Progress®
Kemp® LoadMaster®

Configuration Guide Web User Interface (WUI)

Generated on: 21 April 2026

Copyright

Visit the following page online to see Progress Software Corporation's current Product Documentation Copyright Notice/Trademark Legend: Product Documentation Copyright Notice & Trademarks | Progress

Table of Contents

Chapter 1: Introduction 9

Chapter 2: Home 10

Login Information 12

General Information 12

Virtual Service and Real Server Status 12

WAF Status 13

System Metrics 13

License Information 14

About LoadMaster 14

Chapter 3: Virtual Services 15

Add New 16

View/Modify (Existing HTTP Service) 16

Basic Properties 21

Standard Options 22

QoS/Limiting 32

SSL Properties 33

Advanced Properties 40

Web Application Firewall (WAF) Options 51

Advanced Settings 53

False Positive Analysis 55

WAF Misconfigured Virtual Service Status 57

Edge Security Pack (ESP) Options 57

SMTP Virtual Services and ESP 75

Sub Virtual Services 75

View/Modify (Remote Terminal Service) . . . 79
Real Servers . . . 79
HTTP or HTTPS Protocol Health Checking . . . 84
Binary Data Health Checking . . . 87
Name Server (DNS) Protocol Health Checking . . . 88
Add a Real Server . . . 89
Modify a Real Server . . . 92
Manage Templates . . . 93
Manage SSO Domains . . . 93
Single Sign On Domains . . . 94
Single Sign On Image Sets . . . 107
Cache Configuration . . . 108
Compression Options . . . 109
Kubernetes Settings . . . 109

Chapter 4: Global Balancing . . . 112
Enable/Disable GSLB . . . 112
Manage FQDNs . . . 113
Add an FQDN . . . 114
Add/Modify an FQDN . . . 115
Manage Clusters . . . 123
Add a Cluster . . . 123
Modify a Cluster . . . 123
Delete a Cluster . . . 125
Upgrading GEO Partners . . . 125
Miscellaneous Params . . . 125
Source of Authority . . . 126
Resource Check Parameters . . . 128
Stickiness . . . 130
Location Data Update . . . 130
IP Range Selection Criteria . . . 131
IP Access List Settings . . . 132

Configure DNSSEC 134
GSLB Statistics 136
GEO System Info / Debug 137

Chapter 5: Statistics 140
Real Time Statistics 140
Global 141
Real Servers 143
Virtual Services 146
WAF 148
Client Limits 149
Historical Graphs 150

Chapter 6: SDN Statistics 153
Device Information 155
Path Information 156

Chapter 7: Real Servers 159

Chapter 8: Rules & Checking 161
Content Rules 161
Content Matching Rules 162
Content Matching 162
Add Header 164
Delete Header 165
Replace Header 166
Modify URL 167
Replace String in Response Body 168

Limitations of Using Regular Expressions in the LoadMaster WUI 169
Check Parameters. 170
Service (Health) Check Parameters 170
Adaptive Parameters 171
SDN Adaptive Parameters 172

Chapter 9: Certificates & Security 174
SSL Certificates. 175
Intermediate Certificates 176
ACME Certificates 176
Let's Encrypt Certificates 177
Generate CSR (Certificate Signing Request) 181
Backup/Restore Certs 184
Cipher Sets 185
Remote Access 187
Administrator Access 187
GEO Settings 195
GEO Partners Status 196
WUI Authentication and Authorization 196
Partner Communications 201
Admin WUI Access 202
OCSP Configuration 207
LDAP Configuration. 209
Intrusion Detection Options (IPS/IDS). 211
SSL Options 215

Chapter 10: Web Application Firewall 217
Access Settings. 218
Export Logs 219
Custom Rules 220
False Positives 221

Chapter 11: System Configuration 224

Network Setup 224
Interfaces 225
Host & DNS Configuration 231
Default Gateway 233
Additional Routes 233
Packet Routing Filter 234
VPN Management 235
Route Based VPN 238
HA and Clustering 243
HA Mode 244
Cluster Control 253
QoS/Limiting 256
Global Limits 257
Limiter Options 258
Client Limiting 259
URL Based Limiting 260
System Administration 260
User Management 260
License Management 267
System Reboot 270
Update Software 270
Backup/Restore 275
Date/Time 278
Logging Options 279
System Log Files 280
Extended Log Files 291
Syslog Options 294
SNMP Options 296
Email Options 302
SDN Log Files 304
Troubleshooting 307
Miscellaneous Options 309

WUI Settings 309
L7 Configuration 315
Network Options 321
SDN Configuration 324
Kemp 360 Central Activation Settings 326

Chapter 12: Network Telemetry 327

Chapter 13: Help 330

Chapter 14: References 332

1

Introduction

Progress Kemp products optimize web and application infrastructure as defined by high-availability, high-performance, flexible scalability, security and ease of management. Progress Kemp products maximize the total cost-of-ownership for web infrastructure, while enabling flexible and comprehensive deployment options.

Document Purpose

This document describes the Web User Interface (WUI) of the LoadMaster. It describes in detail how to configure the various features of the LoadMaster using the WUI.

Note: The available menu options in the LoadMaster may vary from the ones described in this document. The features available in a LoadMaster depend on what license is in place. To upgrade a license, please call a Progress Kemp representative.

2

Home

Clicking the **Home** menu option displays the home page which presents a list of basic information regarding the LoadMaster.

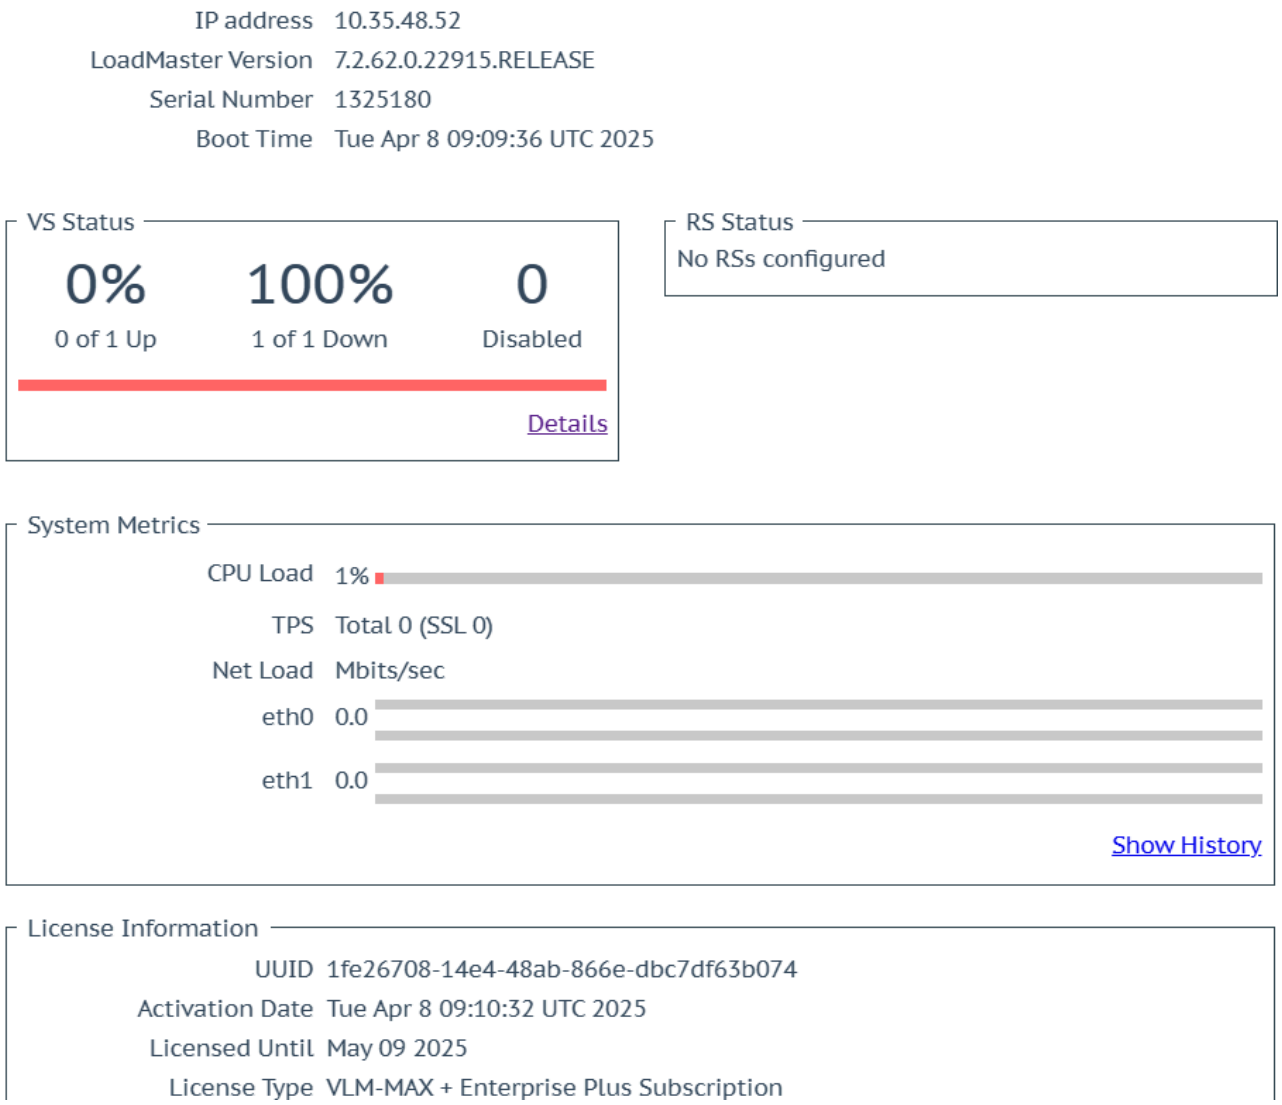

Note: If any of the panels are not displaying information, try resetting the browser to default settings.

Contents

- Login Information
- General Information
- Virtual Service and Real Server Status
- WAF Status
- System Metrics
- License Information
- About LoadMaster

Login Information

After initially logging in to the LoadMaster, if Session Management is enabled - some login information is displayed:

- The last login time and IP address of the current user
- The number of successful logins by the current user in the last 30 days
- The total number of failed login attempts by any user (including unknown usernames) since the last successful login

For further information on Session Management, refer to the Admin WUI Access section.

General Information

IP address: The IP address of the LoadMaster.

LoadMaster Version: The firmware version of the LoadMaster.

Serial Number: The Serial Number of the LoadMaster.

Boot Time: The time of the last server reboot.

Virtual Service and Real Server Status

VS Status

This section displays some monitoring information for the Virtual Services, such as the percentage of Virtual Services that are up and the number of disabled Virtual Services. Clicking the **Details** link will display the **View/Modify Services** screen.

Syslog messages are generated every hour about the number of Virtual Services, SubVSs and Real Servers that are up/down, and so on. Syslog messages are also generated when a status changes.

RS Status

This section displays some monitoring information for the Real Servers, such as the percentage of Real Servers that are up and the number of disabled Real Servers. Clicking the **Details** link will display the **Real Servers** screen.

WAF Status

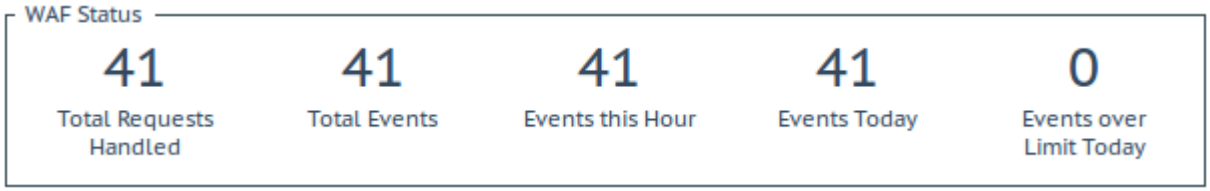

The Web Application Firewall (WAF) Status section is displayed if at least one Virtual Service has WAF enabled. The values shown here are as follows:

- The total number of requests handled by the WAF (shows all requests, whether they were blocked or not). Two requests are recorded for each connection – one incoming and one outgoing request.
- The total number of events handled by the WAF (that is, requests that were blocked)
- The number of events that have happened in the current hour (since xx.00.00)
- The number of events that have happened since midnight (local time)
- The number of times the event counter has gone over the configured warning threshold today. For example, if the threshold is set to 10 and there has been 20 events, this counter is set to 2. The warning threshold is set on a per-Virtual Service basis by filling out the **Hourly Alert Notification Threshold** field in **WAF Options** in the Virtual Service modify screen. For further information, refer to the Web Application Firewall (WAF) Options section.

System Metrics

CPU Load: The percentage of load to the CPU of the LoadMaster appliance.

TPS [conn/s]: The total number of Transactions Per Second and the number of Secure Sockets Layer (SSL) transactions per second.

Note: Only the decrypting transactions are counted here (the re-encrypting transactions are not). If **Reencrypt** is enabled for a Virtual Service, you can roughly double the number. For example, if five Virtual Services are enabled for SSL offloading and only one Virtual Service is enabled for re-encrypting, then double the TPS number for that one Virtual Service and add to the total to get the total value. Doubling only applies to Virtual Services that have re-encryption enabled.

Net Load: Network load in megabits per second, shown for each configured interface. The **Net Load** will only be displayed for interfaces which have been configured.

CPU Temp.: Displays the temperature of the CPU on supported hardware platforms.

The CPU Load and Net Load data is updated every 5 seconds.

These values are only available using SNMP. For further information on **SNMP Options**, refer to the **SNMP Options** section.

License Information

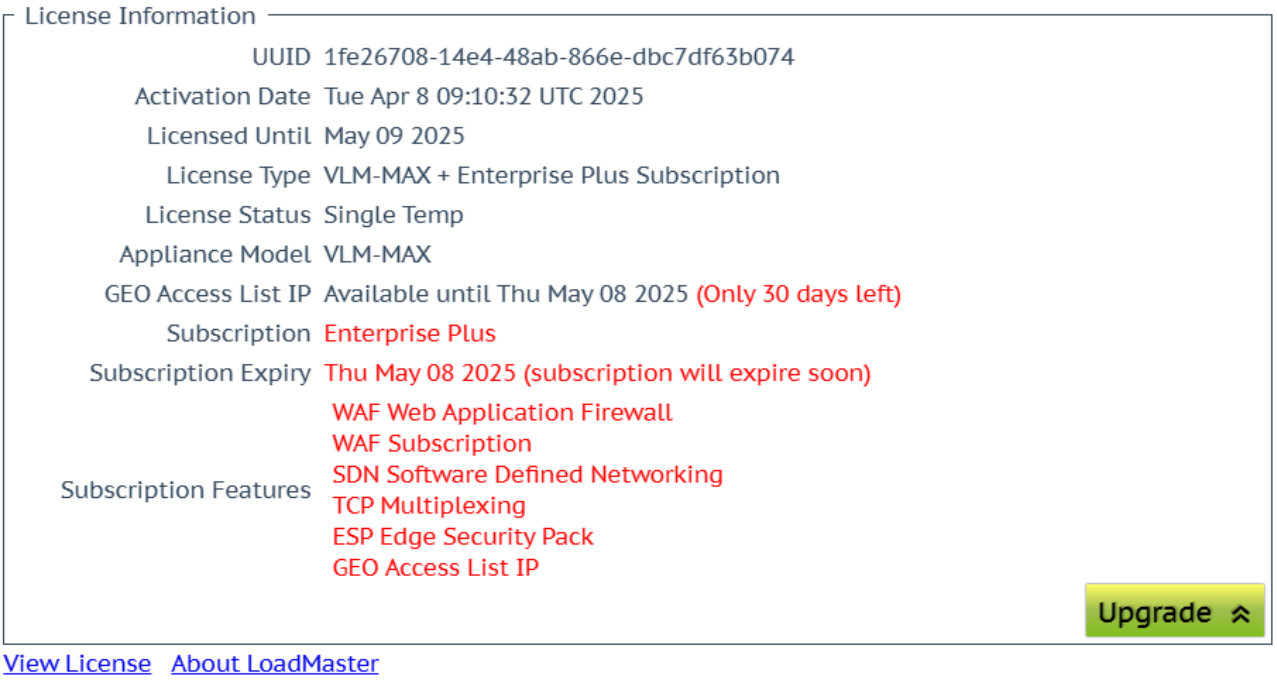

Click the **View License** link to display model, subscription expiry, and subscription feature details, such as the activation date and end date of the LoadMaster license.

Note: If the LoadMaster has a trial or subscription license and it is expired, a message is displayed in the **License Information** section. To renew a subscription, contact Progress Kemp.

Note: LoadMasters that have a permanent license and a support subscription such as Standard, Enterprise, or Enterprise Plus will continue working when the support subscription expires. Only some features like daily reputation data for WAF and GEO may stop working when the support subscription expires.

Note: From 13 days prior to the license or subscription expiry date, a message is displayed on the **Home** screen that states the expiry date. To relicense the appliance or to renew your subscription, contact a Progress Kemp Sales representative.

Upgrade: Upgrade the LoadMaster by buying a license from the Progress Kemp purchase portal.

Note: The Progress Kemp purchase portal is only available for customers in the US.

About LoadMaster

Clicking the **About LoadMaster** link on the LoadMaster User Interface (UI) home page brings you to the current Product Documentation Copyright Notice/Trademark Legend.

3

Virtual Services

Note: From this point onwards, the headings in this document generally correspond to the options in the main menu on the left of the LoadMaster WUI.

Contents

- Add New
- View/Modify (Existing HTTP Service)
- Basic Properties
- Standard Options
- QoS/Limiting
- SSL Properties
- Advanced Properties
- Web Application Firewall (WAF) Options
- Edge Security Pack (ESP) Options
- Sub Virtual Services
- View/Modify (Remote Terminal Service)
- Real Servers
- Manage Templates
- Manage SSO Domains
- Cache Configuration
- Compression Options

- Kubernetes Settings

Add New

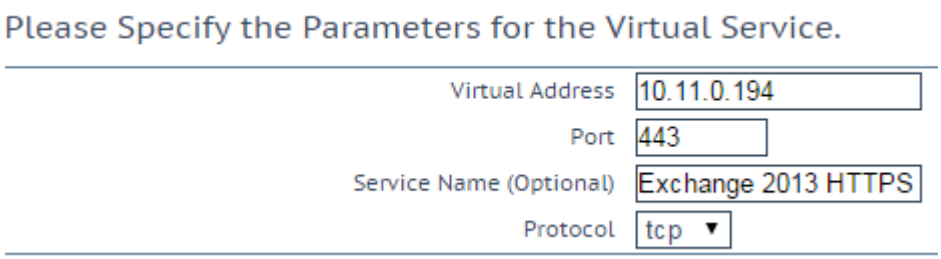

Here the Virtual IP (VIP) address, port, protocol and name are defined. The VIP address, name and port are manually entered into the text boxes and the protocol is selected from the drop-down list.

Note: You cannot use a special character as the first character of the Service Name.

If templates are installed on your machine, a **Use Template** drop-down list is available whereby you can select a template to configure the Virtual Service parameters such as port and protocol.

For further information regarding templates, please refer to the Virtual Services and Templates Feature Description.

View/Modify (Existing HTTP Service)

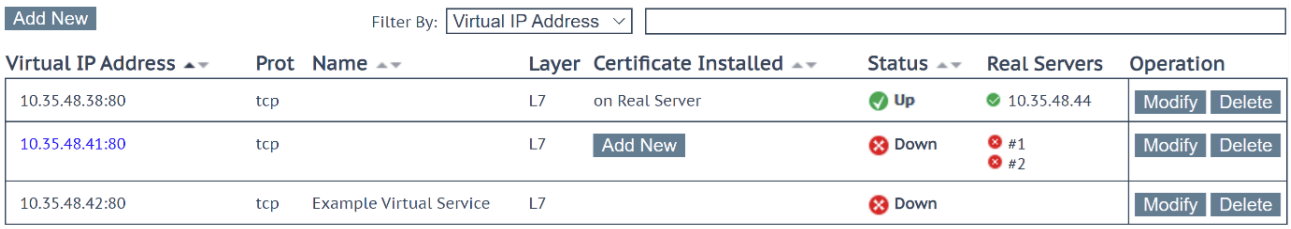

From this screen, you can **Add** or **Modify** a Virtual Service. You can also filter the Virtual Services based on **Virtual IP Address**, **Name**, and **Status** using the filter text box.

Refer to the sections below for further details.

Filter and Sort Virtual Services

You can filter the Virtual Services based on the option selected from the **Filter By** drop-down list. The filter option available are:

- Virtual IP Address

- Name
- Status

By default, the Virtual IP Address option is selected.

If you enter any text in the text box that is contained in the **Virtual IP Address**, **Name**, or **Status** (depending on what is selected in the drop-down), the list is instantly filtered to show results that match. If you enter any text that does not match (based on the option selected in the drop-down), the text box flashes red and the text that does not match is deleted.

The Virtual Services table displays a list of Virtual Services on the LoadMaster, summarizing the main properties of each and giving the options to modify or delete services, or create a new service. You can sort the entries by **Virtual IP Address**, **Name**, **Certificate Installed**, and **Status** by using the up/down arrow options in the table header.

- For **Virtual IP Address**, entries can be sorted first by IP version (IPv4 followed by IPv6), then by IP address (in ascending order) or the opposite.

  Note: The list will not re-sort if the Virtual Service changes while you are viewing the page. Also, the sorting does not apply to SubVSs.

- For **Name**, the order of sorting first shows blank entries, followed by special characters, and the remaining entries are shown alphabetically (or the opposite of this).
- For **Certificate Installed**, the order of sorting is as follows (or the opposite of this):
  - Numeric values
  - Special characters
  - Alphabetical order
  - SSL-based Virtual Services that a certificate must be added to
  - Virtual Services without a certificate (and **SSL Acceleration** is disabled)
- For **Status**, the order of sorting is as follows:

  Note: The meaning of each status is described later in the Virtual Service Statuses section.

  - Up/Unchecked (both are at same level)
  - FailMsgs
  - Redirect
  - Sorry
  - WAF Misconfigured
  - Security Down

- Down
- Disabled

In reverse order:

- Disabled
- Down
- Security Down
- WAF Misconfigured
- Sorry
- Redirect
- FailMsg
- Up/Unchecked (both are at same level)

Note: You can see what sort is currently applied by checking which arrow is darker in color. Only one sort can be applied at any one time. Click one of the lighter color arrows to apply a new sort.

Modify and Delete Virtual Services

Note: **CAUTION** Delete is permanent, there is no UNDO feature. **Use with care.**

Each configured Virtual Service may be changed by clicking the **Modify** button or deleted by clicking the **Delete** button.

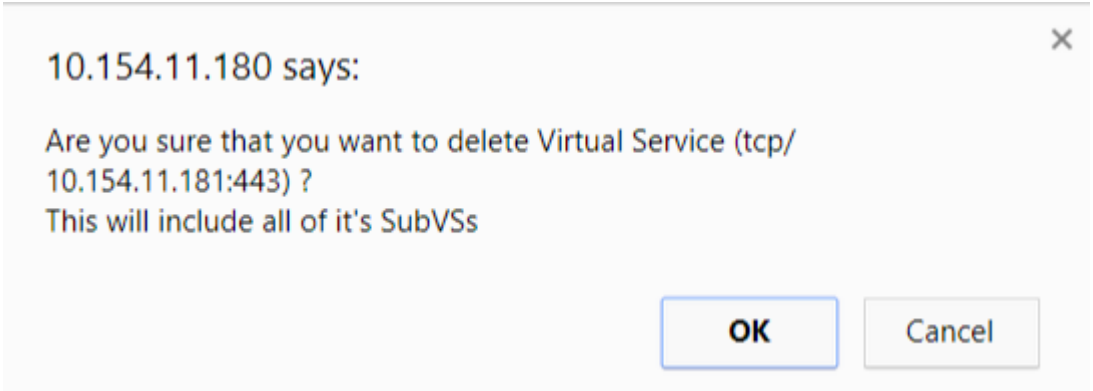

If you try to delete a Virtual Service containing SubVSs, a confirmation warning appears. Click **OK** to confirm the deletion.

Virtual Service Statuses

The status of the Virtual Service is displayed. Health checking is enabled by default when a Virtual Service is created. For further information on health checking, refer to the Real Servers section.

The Virtual Service status may be one of the following:

Status	Description
Up	At least one Real Server is available.
Down	No Real Servers are available.
Sorry	All Real Servers are down and traffic is routed to a separately configured Sorry Server that is not part of the Real Server set, with no health checking.
Disabled	The Virtual Service has been administratively disabled by unticking the **Activate or Deactivate Service** check box in the **Basic Properties** section of the Virtual Service modify screen.
Redirect	A fixed redirect response is in use. One way that redirect Virtual Services can be created is by using the **Add a Port 80 Redirector VS** option in the **Advanced Properties** section. For more information, refer to the Advanced Properties section.
FailMsg	A fixed error message has been configured. A fixed error message can be specified using the Not Available Redirection Handling options. Refer to the Advanced Properties section for more information.
Unchecked	Health checking of the Real Servers has been disabled. All Real Servers are accessed and presumed UP.
Security Down	The LoadMaster is unable to reach the Authentication Server and will prevent access to any Virtual Service which has Edge Security Pack (ESP) enabled for authentication.
WAF Misconfigured	If the Web Application Firewall (WAF) for a particular Virtual Service is misconfigured, for example if there is an issue with a rule file, the status changes to WAF Misconfigured and turns red. If the Virtual Service is in

Status	Description
	this state, all traffic is blocked. WAF can be disabled for that Virtual Service to stop the traffic being blocked (if required) while troubleshooting the problem.

Virtual Service Setting Sections

The image below shows the Virtual Service properties screen. It is composed of several component sections:

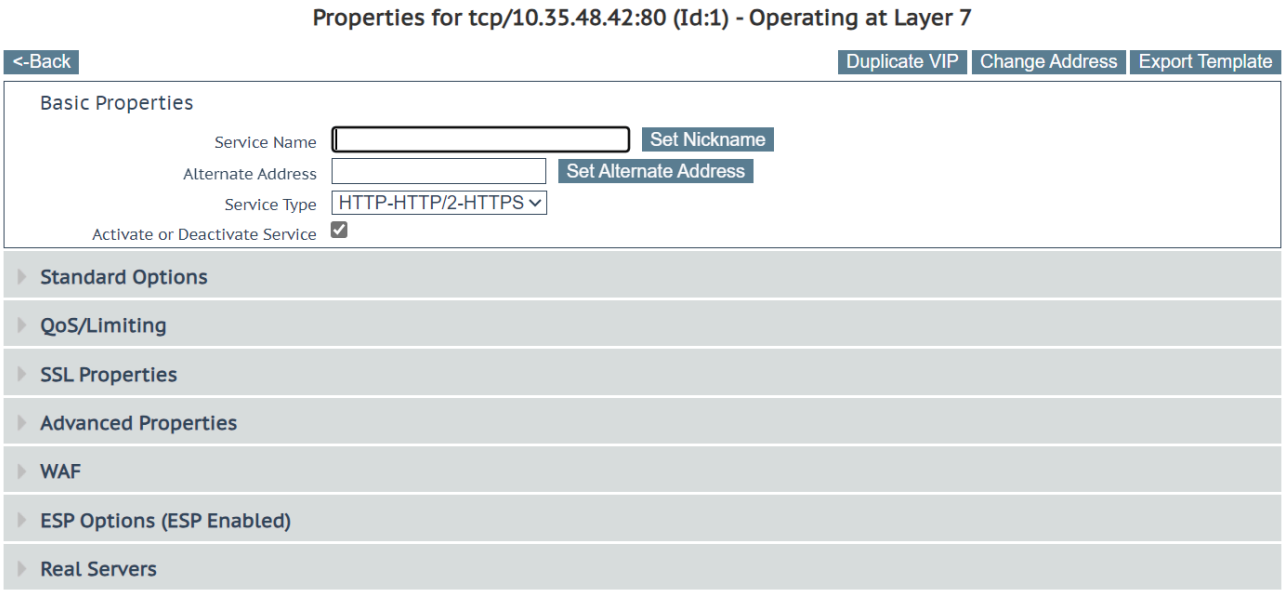

- **Basic Properties** - where the usual and most common attributes are set
- **Standard Options** – the most widely used features of a Virtual Service
- **QoS/Limiting** – where the options relating to the rate limit connections/bandwidth can be set to the LoadMaster and/or from specific client devices.
- **SSL Properties** – if SSL acceleration is being used, it will show Acceleration Enabled and this section of the screen is used to configure the SSL functions
- **Advanced Properties** – the additional features for a Virtual Service
- **WAF** – where the options relating to the Web Application Firewall (WAF) can be set
- **ESP Options** – where the options relating to ESP are set
- **Real Servers/SubVSs** – where Real Servers/SubVSs are assigned to a Virtual Server

The Real Server status is also indicated by an icon to the left of each Real Server. A green icon with a white star in the middle indicates that the **fixed weighting Scheduling Method** is selected. For further details on the possible Real Server statuses, refer to the Real Servers section.

Note: Depending upon the service type, and enabled or disabled features, specific fields and options show in the WUI. The screenshots in this document may not represent every possible configuration.

Basic Properties

There are three buttons adjacent to the **Basic Properties** heading:

Duplicate VIP

This option makes a copy of the Virtual Service, including any related SubVSs. All Virtual Service configuration settings are copied to the duplicate Virtual Service. When this button is clicked, a screen appears where the IP address and port can be specified for the copied Virtual Service.

Change Address

Clicking this button opens a screen where the virtual IP address and port of the Virtual Service can be modified.

Export Template

Export the Virtual Service settings as a template. Templates can be used to create Virtual Services quickly and easily.

Note: When exporting a Virtual Service template in which the Virtual Service uses a custom **Cipher Set**, the LoadMaster on which the template is imported must include the same custom **Cipher Set**.

Virtual Services created from a template will have all of the settings preconfigured based on the settings in the template. The settings in the Virtual Service can then be changed, as needed. For more information on templates, refer to the Virtual Services and Templates, Feature Description.

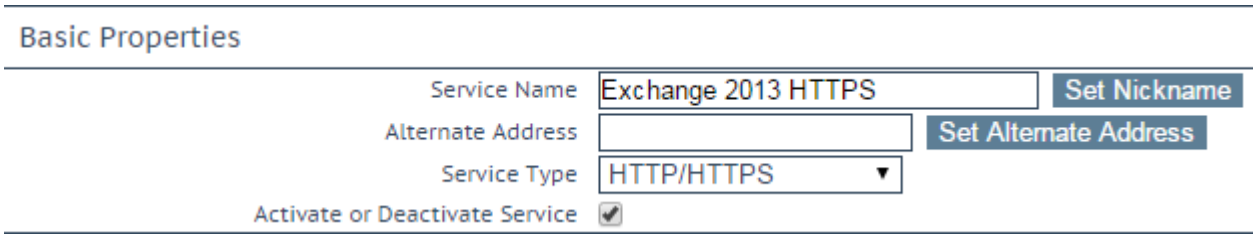

Service Name

This text box allows you to assign a nickname to the Virtual Service being created, or change an existing one.

Note: In addition to the usual alphanumeric characters, the following 'special' characters can be used as part of the Service Name: **. @ - _**

However, you cannot use a special character as the first character of the Service Name.

Alternate Address

This is where, if so desired, you would specify a secondary address in either IPv6 or IPv4 format.

Service Type

Setting the **Service Type** controls the options displayed for the Virtual Service. It's important to make sure the Service Type is set according to the type of application that you are load balancing.

Note: WebSocket Virtual Services must be set to the **Generic** Service Type.

Note: The **HTTP/2 Pass-through** Service Type allows HTTP/2 traffic. HTTP/2 does not support transparency. Subnet originating, alternate source address functionality, caching, compression and body rewrites all work, as required.

Activate or Deactivate Service

This check box gives you the option to activate or deactivate a Virtual Service. The default (active) is selected.

Standard Options

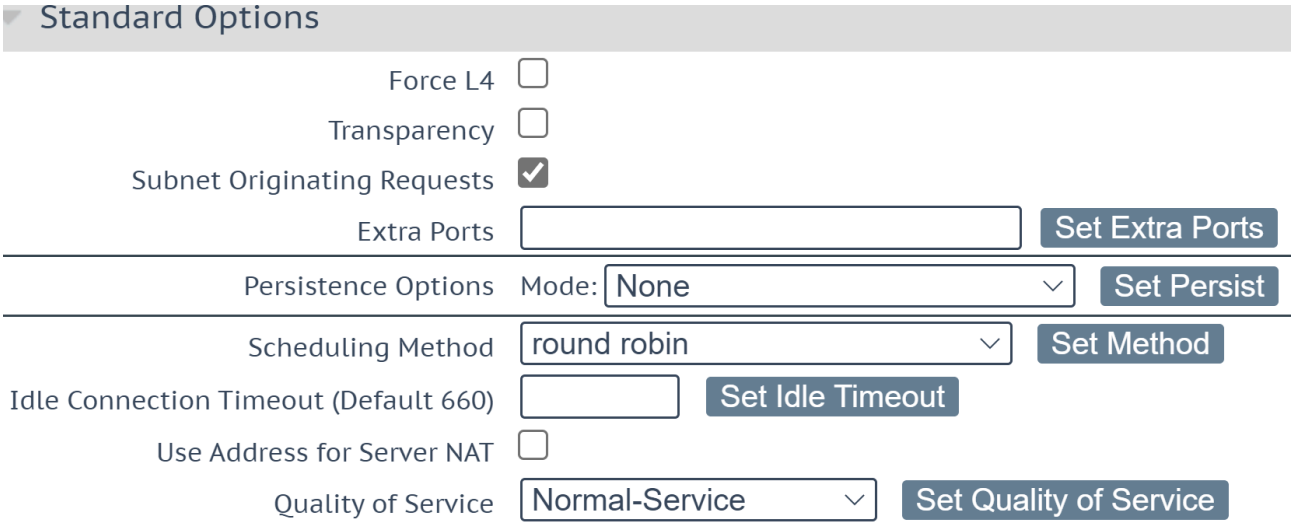

Force L4

Select this check box to force the Virtual Service to run at Layer 4 and not at Layer 7. This is only required in some special circumstances. If in doubt, leave this option unchecked.

L7 Transparency

When using L7, a connection can be transparent. This means the connection arriving at the Real Server appears to come directly from the client. Alternatively, if the connection is not transparent – connections at the Real Server appear to come from the LoadMaster. We recommend keeping transparency disabled in most configurations.

Enabling transparency makes the Virtual Service transparent (no Network Address Translation (NAT)). However, if the client resides on the same subnet as the Virtual IP and Real Servers, then the Virtual Services will automatically NAT the source IP (enabling non-transparency).

If the **Real Servers are local** option is enabled, then the Real Servers are NATed (non-transparent), even if **L7 Transparency** is enabled. This only happens if the Real Server is the originator of the request to the

Virtual Service (and not just answering requests from other clients). For further information on the **Real Servers are local** option, refer to the **L7 Configuration** section.

For further information on transparency in general, refer to the Transparency Feature Description.

Subnet Originating Requests

Note: This option is only available if **Transparency** is disabled.

If **Subnet Originating Requests** is enabled, the source addresses for connections to the Real Servers is the interface address of the LoadMaster. If this option is disabled, the source address is the Virtual Service IP address. If transparency is enabled, the source address is the IP address of the client and the **Subnet Originating Requests** option is ignored.

If the Real Server is on a subnet, and the **Subnet Originating Requests** option is enabled, then the subnet address of the LoadMaster is used as the source IP address.

The Subnet Originating Requests feature was designed for 'local' Real Servers. It works fine for re-encrypt unless the Real Server is non-local and not on the Default Gateway interface. In this case, you can force the local address by setting it in the **Alternate Source Addresses** field. This works for both normal and re-encrypted Virtual Services.

This switch allows control of subnet originating requests on a per-Virtual Service basis. If the global switch (**Subnet Originating Requests** in **System Configuration > Miscellaneous Options > Network Options** in the main menu) is enabled then it is enabled for all Virtual Services.

Note: It is recommended that the **Subnet Originating Requests** option is enabled on a per-Virtual Service basis.

For more information about the global option, refer to the Network Options section.

If the global option is not enabled, it can be controlled on a per-Virtual Service basis.

Note: If this option is switched on for a Virtual Service that has SSL re-encryption enabled, all connections currently using the Virtual Service is terminated.

Extra Ports

You may specify a range of ports, sequential or otherwise, starting with the base port already configured for the Virtual Service. The port numbers are inputted to the field and separated with a space, and the maximum range is 510 ports.

You can enter the extra ports either as port ranges or single ports separated by spaces or comma in whatever order you wish, for example, entering the list **8000-8080, 9002, 80, 8050, 9000** will add the ports 80, 8000 to 8080, 9000 and 9002.

Note: It is not possible to load balance both HTTP and HTTPS on the same Virtual Service. If you need to load balance both HTTP and HTTPS, create two separate Virtual Services or one service for HTTPS and one service to redirect HTTP traffic to HTTPS.

Note: Extra ports cannot be used with SSL re-encryption.

Server Initiating Protocols

By default, the LoadMaster will not initiate a connection with a Real Server until it has received some data from a client. This prohibits certain protocols from working as they need to communicate with the Real Server before transmitting data.

If the Virtual Service uses one of these protocols then select the protocol from the drop-down list to enable it to work correctly.

The protocols that can be selected are:

- SMTP
- SSH
- IMAP4
- MySQL
- POP3
- Other Server Initiating Protocols

Note: The **Server Initiating Protocols** option is not visible when the port specified in the Virtual Service is **80, 8080** or **443**.

Persistence Options

Persistence is setup on a per Virtual Service basis. This section allows you to select whether persistence is enabled for this service, to set the type of persistence and the persistence timeout value.

If persistence is enabled it means that a client connection to a particular Real Server using the LoadMaster is persistent, in other words - the same client will subsequently connect to the same Real Server. The timeout value determines for how long this particular connection is remembered.

The drop-down list gives you the option to select the type of persistence. These are listed and described below.

Note: The persistence modes available for selection will change depending on the settings of the Virtual Service you are editing. For example, for a Virtual Service on port 443, if **SSL Acceleration** is disabled you will only see **Source IP Address** as an option for the **Persistence Mode**.

Source IP Address

The source IP address (of the requesting client) is used as the key for persistency in this case.

Super HTTP

Super HTTP is the recommended method for achieving persistence for HTTP and HTTPS services with the LoadMaster. It creates a unique fingerprint of the client browser and uses that fingerprint to preserve connectivity to the correct Real Server. The fingerprint is based on the values of the **User-Agent** field, if the **User-Agent** value does not contain the **MSRPC** string.

If the **User-Agent** value contains the **MSRPC** string, then the value of the Authorization header is used to achieve persistence.

In the case that the Authorization header is not present, and the **User Agent** value contains the **MSRPC** string, the persist value will simply be blank with a length of zero.

Note: Persistence cookies generated by the LoadMaster for HTTPS Virtual Services include the HttpOnly flag.

Server Cookie

The LoadMaster checks the value of a specially set cookie in the HTTP header. Connections with the same cookie will go to the same Real Server.

Server Cookie or Source IP

If cookie persistence fails, it reverts to source IP-based persistence.

Active Cookie

With Active Cookie persistence, the cookies are generated by the LoadMaster, not the server. When a connection comes into a LoadMaster Virtual Service configured with Active Cookie, the LoadMaster looks for a specific cookie. If that cookie is not there, the LoadMaster inserts it into the HTTP stream with a Set-Cookie directive. Existing cookies are not affected. As with the Server Cookie persistence method, the value for the LoadMaster-generated cookie is unique to each user, allowing the LoadMaster to differentiate between users. A benefit of this method is that no cookies need to be managed or generated by the servers, relieving the burden of server configuration. To gain better dispersion per client connection you can enable the Add Port to Active Cookie feature in the L7 configuration. For further information on this option, refer to the L7 Configuration section.With Active Cookie persistence, the cookie is valid for the session or until the persistence time expires. For example, if using Active Cookie persistence with the persistence timeout set to 10 minutes and the client connects at 2pm, then disconnects and reconnects at 2.05pm – this would reset the persistence timeout value. If the client tries to connect to a Virtual Service after the persistence timeout has expired, they would present the old cookie. The LoadMaster will check its persistence table and see that it does not have a valid entry. The LoadMaster would then generate a new cookie for the client and would update its persistence table.

Note: There may be issues with persistence if the persistence **Mode** set to **Active Cookie** on both a parent and SubVS and the cookie name is the same.

Active Cookie or Source IP

If active cookie persistence fails, it reverts to source-based persistence.

Hash All Cookies

The **Hash All Cookies** method creates a hash of the values of all cookies in the HTTP stream. Cookies with the same value are sent to the same server for each request. If the values change, then the connection is treated as a new connection and the client is allocated to a server according to the load balancing algorithm.

Hash All Cookies or Source IP

Hash All Cookies or Source IP is identical to **Hash All Cookies**, with the additional feature that it will fall back to Source IP persistence in the event no cookies are in the HTTP string.

Super HTTP and Source IP Address

This is the same as super HTTP but it also appends the source IP address to the string, thus improving the distribution of the resulting HASH.

URL Hash

With URL Hash persistence, the LoadMaster will send requests with the same URL to the same server.

HTTP Host Header

With HTTP Host Header persistence, the LoadMaster will send all requests that contain the same value in the HTTP Host: header to the same server.

Hash of HTTP Query Item

This method operates in exactly the same manner as Server Persistence, except that the named item being inspected is a Query Item in the Query String of the URL. All queries with the same Query Item value is sent to the same server.

Selected Header

With Selected Header persistence, the LoadMaster will send all requests that contain the same value in the specified header to the same server.

SSL Session ID

Each session over SSL has its own session ID which can be persisted on.

Note: For this option to appear as a persistence method, the Virtual Service needs to have a Service Type of Generic and SSL acceleration must be disabled.

If a Virtual Service is an SSL service and not offloaded, the LoadMaster cannot meaningfully interact with any of the data in the stream at Layer 7. The reason is, the data is encrypted and the LoadMaster has no way of decrypting it.

If, in the above scenario, a persistence mode that is not based off source IP is required, this is the only other option. When an SSL session is started, it generates a session ID for the connection. This session ID can be used to cause the client to persist to the correct server.

There are some downsides to this however, as most modern browsers regenerate the session ID at very short intervals, basically overwriting it, even if there is a longer interval set on the persist timeout.

UDP Session Initiation Protocol (SIP)

This persistence mode is only available in a UDP Virtual Service when **Force L4** is enabled. SIP uses request and response transactions, similar to HTTP. An initial INVITE request is sent, which contains a number of header fields. These header fields can be used for persistence.

Timeout

For each persistence method, there is a configurable timeout value that determines how long the persistence for each user is honored, selectable from one minute to seven days.

Note: In LoadMaster firmware version 7.2.53, the maximum value of the persistence timeout setting has increased from 7 days to 28 days. You can configure the persistence **Timeout** drop-down list when a persistence **Mode** is selected. If the persistence **Timeout** is set to 4 days or more, a **Refresh Persist** check box appears. This is disabled by default. When **Refresh Persist** is enabled, the persist entries are auto-refreshed each day for long-lived connections.

This timeout clock is started when the initial connection is established. The persistence timeout value is updated if the client reconnects within the timeout period. For example, if the persistence timeout is set to 1 hour and the client starts a connection at 2pm, if the client disconnects and then reconnects before 3pm they will still persist to the same Real Server. Also, the persistence record is updated to reflect this and the persistence countdown timer is reset back to 1 hour for this client.

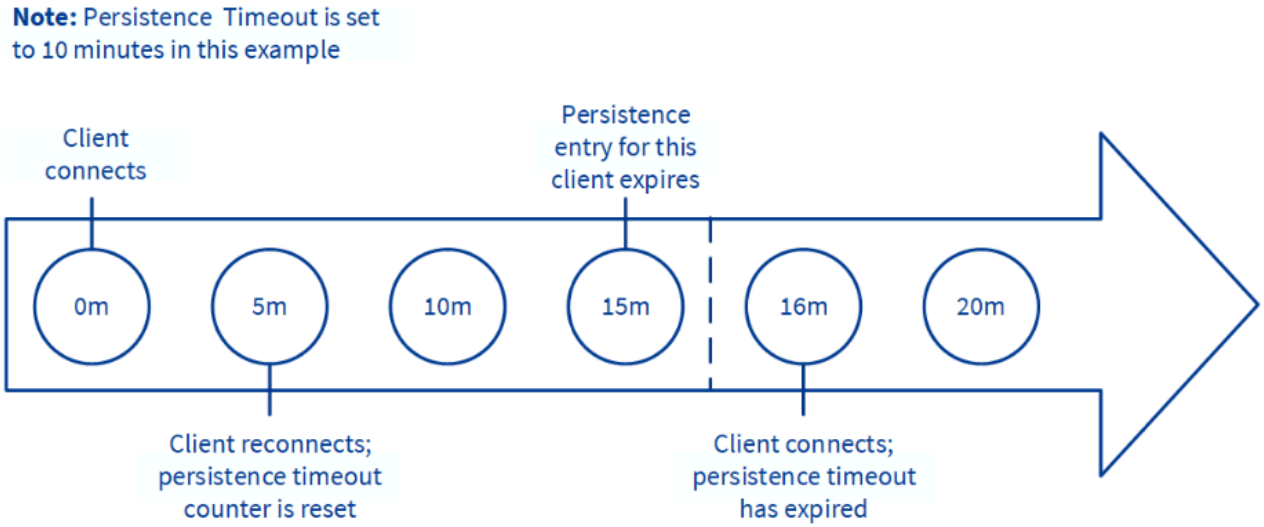

If a client made connections to the Virtual Service repeatedly within the timeout period, the persistence would be honored indefinitely. For instance, given the following scenario:

- Persistence Timeout is set to 10 minutes
- A user makes several requests in the course of 20 minutes, but the time between connections is always less than 1 minute

The request should be sent to the correct Real Server, as long as it is available (that is, passing health checks).

If the active connection goes idle for 20 minutes, then the next connection is counted as a new session, and may be sent to a different server, depending on scheduling. If the connection is opened for more than 10 minutes and the client disconnects and reconnects, the persistence record would have expired, the LoadMaster will create a new persistence entry for that client and possibly send the client to a new Real Server. This is due to the fact that the persistence countdown starts once a connection is established, not at the closing of the connection.

If you are experiencing persistence issues, this may be due to the fact that the persistence timeout is not long enough. If this is not long enough, then the timeout value should be set for a higher amount. In general, matching this value to your server timeout value is recommended.

Header field name

When **UDP Session Initiation Protocol** is selected as the persistence mode is selected sin the LoadMaster, a text box called **Header field name** will appear. The header field that is to be used as the basis for the persistence information should be entered here.

Scheduling Methods

This section allows you to select the method by which the LoadMaster will select a Real Server, for this particular service. The scheduling methods are as follows:

- Round Robin:

Round Robin causes the LoadMaster to assign Real Servers to a session in order, for example the first session connects to Real Server 1, the second to Real Server 2 and so on. There is no bias in the way the Real Servers are assigned.

- Weighted Round Robin:

This method uses the weight property of the Real Servers to determine which Real Servers get preference. The higher the weight a Real Server has, the higher the proportion of connections it will receive.

- Least Connection:

With this method, the current Real Server with the fewest open connections is assigned to the session.

- Weighted Least Connection:

As with **Least Connection**, but with a bias relative to the weight.

- Resource Based (Adaptive):

Adaptive scheduling means that the load on the Real Servers is periodically monitored and that packets are distributed such that load is approximately equal for all machines.

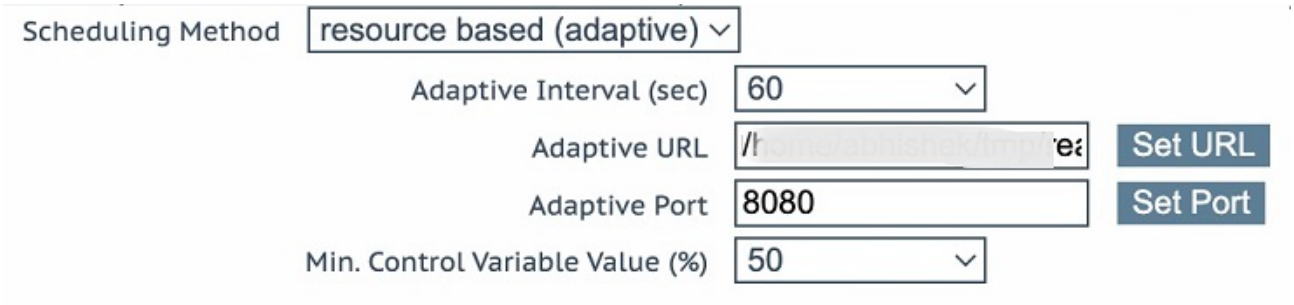

If the **Scheduling Method** is set to **resource based (adaptive)**, the adaptive parameters are displayed. By default, the global values configured in **Rules & Checking > Check Parameters** are used. This is indicated by **Use Global** and **Using Global** in the field values. However, you can change these values for each Virtual Service, if needed. For further details on the adaptive parameters, including a description of each of the fields, refer to the following section: Adaptive Parameters.

Note: When you first enable adaptive scheduling and there is no HTTP server running on the Real Server, the Real Server weight is set to 100 (heavily loaded).

- Resource Based (SDN Adaptive):

A Virtual Service which is using an adaptive scheduling method (whether using SDN or not) can be viewed as a control system. The intent is to achieve an evenly distributed load over the Real Servers and the controller calculates an error value from this (that describes the deviation from the desired even distribution). It also calculates a set of control values (Real Server weights) that are fed back into the system in a way to decrease the error value.

- Fixed Weighting:
  - All new connections, and connections without a valid persistence entry are forwarded to the highest weighted available server. If multiple Real Servers share the same highest weight, connections are forwarded to the first highest-weighted, available Real Server.
  - Current connections are forwarded to their current server until that connection is terminated
  - Connections with a persistence entry are forwarded to their current server until that connection is terminated and the persistence entry expires. Drain time could also apply in this scenario. For further details on drain time, refer to the **Drop at Drain Time End** section of the following topic: L7 Configuration.



When fixed weighting is in use, the Real Server with the higher weight is indicated with a green star icon.

- Weighted Response Time:

Every 15 seconds the LoadMaster measures the time it takes for a response to arrive for a health check probe and uses this time to adjust the weights of the Real Servers accordingly, that is, a faster response time relative to the other Real Servers leads to a higher weight which in turn leads to more traffic sent to that server.

- Source IP Hash:

When using the **source IP hash** scheduling method with Layer 4 (L4) Virtual Services, LoadMaster uses a prepopulated hash table to determine how client traffic is distributed across Real Servers.

For L4 source IP hash scheduling:

- LoadMaster uses a fixed hash table containing 4096 destination slots.
- Each Real Server is assigned a number of slots based on its configured Server weight.
- The weights of all Real Servers associated with the Virtual Service are tallied together to determine how the 4096 slots are distributed.

This prepopulation directly affects how source IP addresses are mapped to Real Servers.

By default, each Real Server is assigned a weight of **1000**.

When multiple Real Servers are configured using the default weight, the total combined weight may not evenly divide into 4096.

In such cases, LoadMaster allocates slots sequentially based on Server weight until all 4096 slots are

assigned. Any remaining slots are assigned starting again from the first Real Server.

As a result, some Real Servers, particularly those added later, may receive significantly fewer connections or no connections at all.

Example 1: Four Real Servers, weight 1000 each

- rs1: 1000 slots
- rs2: 1000 slots
- rs3: 1000 slots
- rs4: 1000 slots
- rs1: remaining 96 slots

Example 2: Five Real Servers, weight 1000 each

- rs1: 1000 slots
- rs2: 1000 slots
- rs3: 1000 slots
- rs4: 1000 slots
- rs5: 96 slots

In the second example, the fifth Real Server receives a very small share of the hash table, which can lead to minimal or no traffic being forwarded to that Server.

When using source IP hash scheduling, ensure that the sum of all Real Server weights is a factor of 4096.

Users should adjust Real Server weights rather than relying on the default value of 1000 based on the number of configured Servers. This ensures an even distribution of hash slots and prevents traffic imbalance.

- URL Hash:

The **URL Hash** method works by creating a hash value based on the object referenced in the client request's URL and the number of Real Servers or SubVSs in the Virtual Service. All requests for a particular URL are sent to the same Real Server/SubVS, unless a Real Server or SubVS is added or removed – in which case all hash values are re-calculated and subsequent traffic is redistributed accordingly. A write always succeeds regardless of any outage (unless everything is down). The URL hash method sends write requests to the next available SubVS when a SubVS is down. For example:

- A Virtual Service has three SubVSs. A write request is received for which there is an existing hash that says to send the write to SubVS 2.
- SubVS 2 is down. The request is sent to SubVS 3.
- If SubVS 3 is down, send to SubVS 1 (in a round-robin fashion).
- When SubVS 2 comes back online, go back to honoring the hash and send future requests to SubVS 2.

Both path style addressing (for example, https://s3.us-west-2.amazonaws.com/mybucket/puppy.jpg) and virtual host style addressing (for example, https://my-bucket.s3.us-west-2.amazonaws.com/puppy.png) are supported.

This scheduling method was developed primarily to support Dell EMC Elastic Cloud Storage (ECS)

applications and efficient use of ECS-based resources, but could also be used to support other workloads where storage efficiency is the primary goal. For Dell ECS deployments, the load traffic is distributed across the Virtual Data Centers (VDCs) in the deployment, each of which are represented on LoadMaster as a SubVS. Within each VDC, the traffic is distributed across the Real Servers in the SubVS.

Idle Connection Timeout (Default 660)

The seconds before an idle connection is closed. Setting it to **0** ensures that the default L7 connection timeout is used. Configuring the **Idle Connection Timeout** in the **Standard Options** section of a Virtual Service sets the value for the Virtual Service. You can also modify the default **Connection Timeout** value by going to **System Configuration > Miscellaneous Options > Network Options**.

Use Address for Server NAT

By default, when the LoadMaster is being used to SNAT Real Servers, the source IP address used on the internet is that of the LoadMaster. The **Use Address for Server NAT** option allows the Real Servers configured on the Virtual Service to use the Virtual Service as the source IP address instead - if the Real Server makes an outbound request using the same port as the Virtual Service. The LoadMaster does not NAT all outbound ports.

Note: The **Use Address for Server NAT** option is most useful for services such as SMTP when the LoadMaster is in a public domain and when the service requires a reverse DNS check to see if the source address sent from the LoadMaster is the same as the Mail Exchanger (MX) record of the sender.

If the Real Servers are configured on more than one Virtual Service that has this option set, the LoadMaster examines the destination port of the server's request and then selects the Virtual Service with a matching port. The LoadMaster then uses this Virtual Service as the source IP address. If no match is found for the port being requested, the IP address of the LoadMaster is used as the source IP address.

Note: The **Use Address for Server NAT** option only works on Virtual Services which are operating on the default gateway. This option is not supported on non-default gateway interfaces.

Quality of Service

The **Quality of Service** drop-down sets a Type of Service (ToS) in the IP header of packets that leave the Virtual Service. This means that the next device or service that deals with the packets will know how to treat and prioritise this traffic. Higher priority packets are sent from the LoadMaster before lower priority packets.

The different options are described below:

- **Normal-Service**: No special priority given to the traffic
- **Minimize-Cost**: Used when data needs to be transferred over a link that has a lower “cost”
- **Maximize-Reliability**: Used when data needs to travel to the destination over a reliable link and with little or no retransmission
- **Maximize-Throughput:** Used when the volume of data transferred during an interval is important, even if the latency over the link is high
- **Minimize-Delay**: Used when the time required (latency) for the packet to reach the destination must be low. This option has the quickest queue of each of the **Quality of Service** choices.
- **Pass Through:** In LoadMaster firmware version 7.2.52, the **Pass Through** value was introduced. When this is selected, connections that contain the Quality Of Service (QOS) flags are passed through to the

Real Server. There are a couple of points to be aware of regarding SubVSs:

- If you select **Pass Through** as the **Quality of Service** on the parent Virtual Service, all SubVSs under the parent Virtual Service will use **Pass Through**. The SubVS will not display the **Quality of Service** field and you will not be able to change the **Quality of Service** value using the Application Programming Interface (API).
- If you select any other option apart from **Pass Through** as the **Quality of Service** on the parent Virtual Service, no SubVS under that Virtual Service will have the **Pass Through** option in the **Quality of Service** drop-down list and you will not be able to set the **Quality of Service** to **Pass Through** using the API.

The ToS values for each option are provided in the following table:

Bits	Decimal	Importance
1000	8th	Minimal delay
0100	4	Maximum throughput
0010	2	Maximum reliability
0001	1	Minimal costs (in the form of money)
0000	0	Normal service

Note: The **Quality of Service** feature only works with Layer 7 traffic. It does not work with Layer 4 traffic.

QoS/Limiting

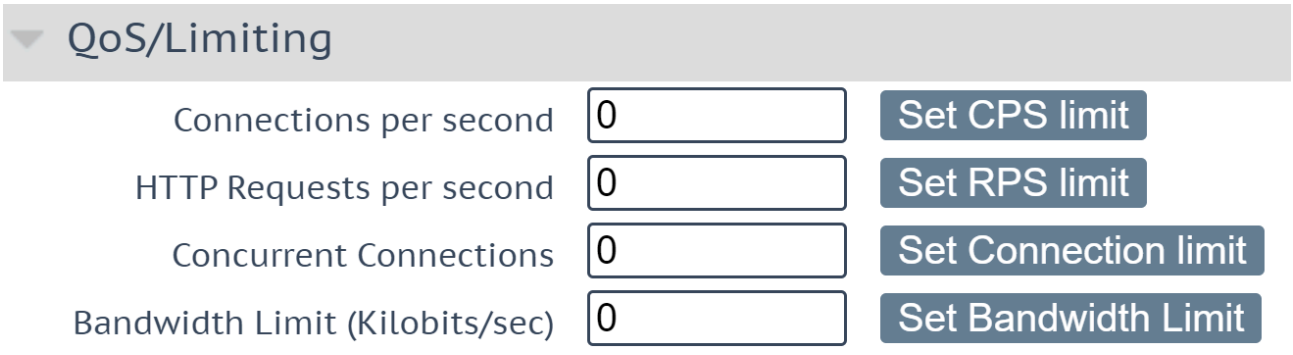

Connections per second: Set the maximum connections per second of this Virtual Service. Setting the limit to 0 disables this option. Valid values are 0 - 100000000.

HTTP Requests per second: Set the maximum HTTP requests per second of this Virtual Service. Setting the limit to 0 disables this option. Valid values are 0 - 1000000.

Note: For Virtual Services created on 443 port, the **HTTP requests per second** option is only available when the **SSL Acceleration** option is enabled.

Concurrent Connections: Set the maximum concurrent connections for this Virtual Service. Setting the limit to 0 disables this option. Valid values are 0 - 1000000.

Bandwidth Limit (Kilobits/sec): Set the maximum bandwidth of this Virtual Service. The value is in Kilobits/second. The minimum value is 16. The maximum value is 99999999. Setting this value to 0 removes any bandwidth limit. This limits all traffic going through the Virtual Service. If a bandwidth limit is set for a Virtual Service, then it is forced to be a Layer7 (L7) service.

SSL Properties

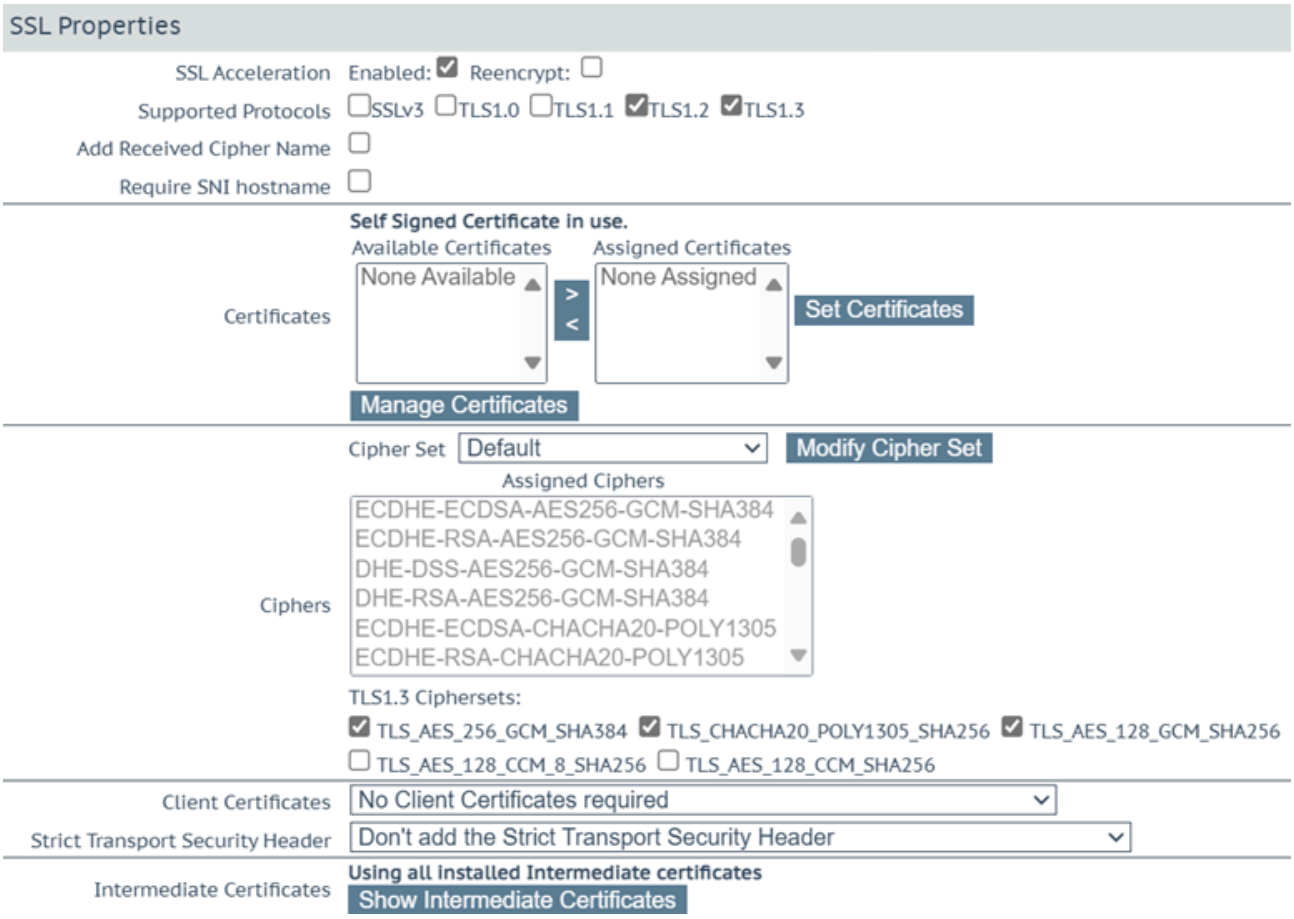

SSL Acceleration

This check box appears when the criteria for SSL Acceleration have been met. Select this check box to activate **SSL Acceleration**.

Enabled: If the **Enabled** check box is selected and there is no certificate for the Virtual Service, you are prompted to install a certificate. You can add a certificate by clicking **Manage Certificates** and importing or adding a certificate.

Reencrypt: Selecting the **Reencrypt** check box re-encrypts the SSL data stream before sending it to the Real Server.

Note: You cannot use **Extra Ports** or **Transparency** with SSL reencryption.

Reversed: Selecting this check box means that the data from the LoadMaster to the Real Server is re-encrypted. The input stream must not be encrypted, for example, the client sends HTTP port 80 traffic to the LoadMaster and the LoadMaster sends HTTPS port 443 traffic to the Real Server. This is only useful in connection with a separate Virtual Service which decrypts SSL traffic then uses this Virtual Service as a Real Service and loops data back to it. In this way, the client to Real Server data path is always encrypted on the wire.

Note: The **Reversed** option is only available when the **Service Type** is set to **Generic** (**Virtual Services > View/Modify Services > Modify > Basic Properties**).

Supported Protocols

The check boxes in the **Supported Protocols** section enable you to specify which protocols are supported by the Virtual Service. By default, TLS1.1, TLS1.2, and TLS1.3 are enabled and SSLv3 and TLS1.0 are disabled.

Note: The **TLS1.3** check box will not be visible if the **OpenSSL Version** setting in **Certificates & Security > SSL Options** is set to **Use older SSL library - no TLS 1.3**.

Starting with version 7.2.37, when re-encryption is enabled, the TLS version that can be negotiated between the LoadMaster and the Real Servers behind it are no longer constrained by the TLS version settings configured on the client side. All TLS versions and ciphers that are supported on the LoadMaster can be negotiated without restriction by Real Servers. In this way, the LoadMaster can, for example, provide strict security for client-side application access and still support server-side connections to legacy servers that only support specific, less secure, TLS versions, and ciphers. This is illustrated in the example below.

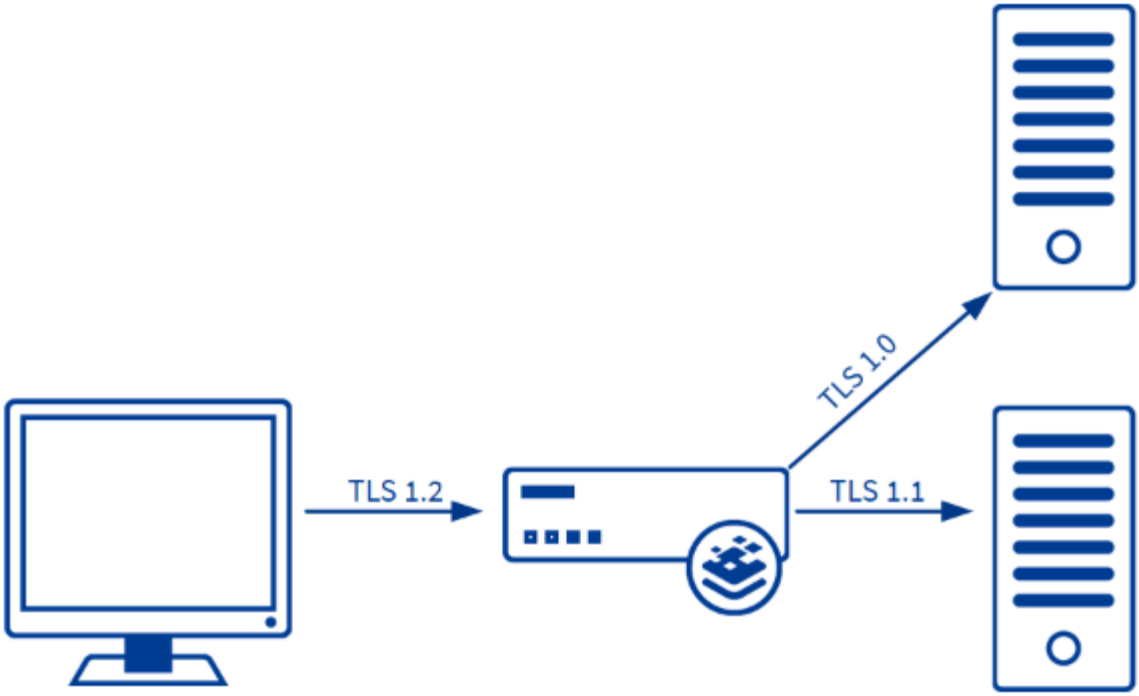

Server connections are only restricted by the configuration of the Real Servers, regardless of the TLS version selected on the client side. Each Real Server can be configured independently of the others. The LoadMaster negotiates connections according to the requirements of each Real Server.

Add Received Cipher Name

In LoadMaster version 7.2.52 and above, a new check box called **Add Received Cipher Name** was added. This option is disabled by default. When this option is enabled, the LoadMaster adds X-SSL headers containing client SSL information such as TLS version, TLS cipher, client certificate serial number, and SNI host as described in below table.

The information contained in these headers can be used in content rules by referencing the appropriate header name in the rule (see the table below). This allows you to make load balancing decisions based on, for example, the cipher used.

This information can also be useful, for example, as you maintain cipher sets over time; it allows you to see which ciphers are being used and can help you plan what ciphers to change or delete in the cipher sets. The **Add Received Cipher Name** check box must be enabled to use the headers in the table below in content rules.

Header	Description	Example Value
X-SSL-Cipher	The cipher used.	X-SSL-Cipher: ECDHE-RSA-AES256-GCM-SHA384
X-SSL-Protocol	The SSL protocol version used.	X-SSL-Protocol: TLSv1.2
X-SSL-Serialid	The Virtual Service certificate serial number.	X-SSL-Serialid: 4900000006A2ABDC165ACEAD55000000000006

Header	Description	Example Value
X-SSL-ClientSerialid	The client certificate serial number.	X-SSL-ClientSerialid: 490000005D6898F3C7E5905361000 10000005D
X-SSL-SNIHost	The value of the received SNI name.	X-SSL-SNIHost: sni.test.com

Require SNI hostname

If require Server Name Indication (SNI) is selected, the hostname is always required to be sent in the TLS client hello message.

When **Require SNI hostname** is disabled, the first certificate is used if a host header match is not found.

When **Require SNI hostname** is enabled, a certificate with a matching common name must be found, otherwise an SSL error is yielded. Wildcard certificates are also supported with SNI.

Note: When using a Subject Alternative Name (SAN) certificate, alternate source names are not matched against the host header.

Note: Wildcard certificates are supported but note that the root domain name is not matched, as per RFC 2459. Only anything to the left of the dot is matched. Additional certificates must be added to match the root domain names. For example, www.kemptechnologies.com is matched until a wildcard of *.kemptechnologies.com. Kemptechnologies.com is not matched.

Note: To send SNI host information in HTTPS health checks, enable **Use HTTP/1.1** in the **Real Servers** section of the relevant Virtual Service(s) and specify a host header. If this is not set, the IP address of the Real Server is used.

Pass through SNI hostname

In LoadMaster firmware version 7.2.52 and above, when this option is enabled and when re-encrypting, the received SNI hostname is passed through as the SNI to be used to connect to the Real Server. If the Virtual Server has a **Reencryption SNI Hostname** set, this overrides the received SNI.

Note: This field is only visible when SSL re-encryption is enabled.

Certificates

Available certificates are listed in the **Available Certificates** select list on the left. To assign or unassign a certificate, select it and click the right or left arrow button. Then click **Set Certificates**. Multiple certificates can be selected by holding **Ctrl** on your keyboard and clicking each required certificate.

Note: There is a limit of 8171 characters when assigning certificates to a Virtual Service using the WUI.

Note: A Virtual Service can be configured using both RSA and ECC certificates. However, if an RSA and an ECC certificate have the same common name, for example, progress.com, the first certificate is preferred. If the ECC certificate is first in the list, and a client does not have an ECC cipher, the connection fails. Conversely, if the RSA certificate is first in the list, and a client does not have an RSA cipher, the connection

fails.

Note: The total number of certificates that you can add to a Virtual Service is 256, but this number may be further limited by the size of the certificate file names used. In LMOS Version 7.2.47 and later releases, the number of characters in each certificate file name and extension (not counting the period between them), plus all spaces used to separate multiple file names, must add up to 8176 characters or less (in earlier releases, the limitation is 1023 characters.)

Clicking **Manage Certificates** brings you to the SSL Certificates screen.

CAUTION: If you click **Manage Certificates** and replace a certificate in the current list, the replaced certificate moves to the bottom of the certificate list upon returning to the Virtual Service page. If the replaced certificate should be higher in the list, you can select all the certificates in the **Assigned Certificates** box on the right, move them back to the **Available Certificates** box on the left, and then re-add them in the required order.

If you add a certificate to the LoadMaster in version 7.2.51 or later (or in 7.2.48.3 LTS or a later LTS version) and then downgrade to 7.2.50 or an earlier version (or 7.2.48.2 LTS or an earlier version) - the certificate will not work. To work around this, create a backup of all SSL certificates before downgrading and then restore the certificates after downgrading (**Certificates & Security > Backup/Restore Certs**). If you forget to take the backup before downgrading: upgrade the firmware again, take the certificate backup, downgrade, and then restore the certificate backup.

Reencryption Client Certificate

With SSL connections, the LoadMaster gets a certificate from the client and also gets a certificate from the server. The LoadMaster transcribes the client certificate in a header and sends the data to the server. The server still expects a certificate. This is why it is preferable to install a pre-authenticated certificate in the LoadMaster.

Reencryption SNI Hostname

In LoadMaster firmware version 7.2.52 and above, it is possible to set a Reencryption SNI Hostname at the SubVS level. If this is set in a SubVS, this overrides the parent Virtual Service value and/or the received SNI value.

Note: This field is only visible when SSL re-encryption is enabled.

Cipher Set

A cipher is an algorithm for performing encryption or decryption.

Each Virtual Service (which has **SSL Acceleration** enabled) has a cipher set assigned to it. This can either be one of the system-defined cipher sets or a user-customized cipher set. You can select system-defined cipher sets to quickly and easily select and apply the relevant ciphers. You can create and modify custom cipher sets by clicking **Modify Cipher Set**.

Note: The selection of TLS1.2 and below protocols ciphers become unavailable when only TLS1.3 protocol is selected.

Ciphers

The **Ciphers** list is read only and displays a list of the currently assigned ciphers. Clicking **Modify Cipher Set** brings you to the **Cipher Set Management** screen. This screen allows you to create new, and modify existing custom cipher sets.

TLS1.3 Cipher Sets

Select the cipher sets for the TLS1.3 protocol to be allowed using any combination of supported ciphers over an SSL-enabled Virtual Service. By default, the following three cipher sets are enabled:

- TLS_AES_256_GCM_SHA384
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_GCM_SHA256

Note: It is not possible to disable all ciphersets.

Note: The TLS1.3 cipher sets only become available when the **TLS1.3** protocol is enabled.

Note: TLS1.3 ciphers are always available regardless of whether the selected cipher set contains any TLS1.3 ciphers, which is how the OpenSSL libraries behave. TLS1.3 ciphers can only be removed from offered ciphers by disabling the TLS1.3 protocol (which is not recommended).

Client Certificates

- **No Client Certificates required:** enables the LoadMaster to accept HTTPS requests from any client. This is the recommended option.

By default the LoadMaster accepts HTTPS requests from any client. Selecting any of the other values below requires all clients to present a valid client certificate. In addition, the LoadMaster also passes information about the certificate to the application.

Note: You should not change this option from the default of **No Client Certificates required**. Only change from the default option if you are sure that all clients that access this service have valid client certificates.

- **Client Certificates required:** requires that all clients forwarding a HTTPS request must present a valid client certificate.
- **Client Certificates and add Headers**: requires that all clients forwarding a HTTPS request must present a valid client certificate. The LoadMaster also passes information about the certificate to the application by adding headers. When a client certificate is used, the **X-SSL-ClientSerialid** header is set.
- The below options send the certificate in its original raw form. The different options let you specify the format that you want to send the certificate in:
  - Client Certificates and pass DER through as SSL-CLIENT-CERT
  - Client Certificates and pass DER through as X-CLIENT-CERT
  - Client Certificates and pass PEM through as SSL-CLIENT-CERT
  - Client Certificates and pass PEM through as X-CLIENT-CERT

If a Virtual Service:

- Has **SSL Acceleration** enabled, and
- Any of the client certificate required options with "pass through as SSL-CLIENT-CERT/X-CLIENT-CERT" is selected in the **Client Certificates** drop-down list, and
- A **Delete Header** rule is applied to that Virtual Service to delete the SSL/X-CLIENT-CERT header field, then

  The **Delete Header** content rule preserves the LoadMaster inserted client certificate header fields and removes any header with that name passed in by the client.

Verify Client using OCSP

Verify (using Online Certificate Status Protocol (OCSP)) that the client certificate is valid.

Note: This option is only visible when ESP is enabled.

Strict Transport Security Header

Enable this option to add the Strict-Transport-Security header to all LoadMaster-generated messages (ESP and error messages). The options in this drop-down list are as follows:

- Don't add the Strict Transport Security Header
- Add the Strict Transport Security Header - no subdomains
- Add the Strict Transport Security Header - include subdomains
- Add the Strict Transport Security Header - no subdomains + preload
- Add the Strict Transport Security Header - include subdomains + preload

Intermediate Certificates

Prior to the **Intermediate Certificates** field being added to the **SSL Properties** section, there was no ability to assign intermediate or root certificates to a Virtual Service. The Certificate Authority (CA) for client certificates was kept in the global certificate store, so the following could occur:

- Client certificates from two different CAs are installed on the LoadMaster
- Client A presents a certificate issued from CA 1 and as a network administrator, you only want them to be able to access Virtual Service 1.
- Client B presents a certificate issued from CA 2 and as a network administrator, you only want them to be able to access Virtual Service 2.
- Because both client certificates are validated against the global LoadMaster trust store, client A is also allowed access to Virtual Service 2 and client B is also allowed access to Virtual Service 1.

The **Intermediate Certificates** field allows you to assign intermediate and root certificates to specific Virtual Services. This provides the ability to restrict access. It also enables control on what client certificates are eligible to be used when connecting to a service which is useful in environments with multiple client certificates signed by multiple authorities. For example, when this is configured correctly for the scenario above - Client A will only have access to Virtual Service 1 and Client B will only have access to Virtual Service 2.

To configure this, follow the steps below:

1. Upload the relevant certificates.
2. Then in the LoadMaster User Interface (UI), go to **Virtual Services > View/Modify Services**.
3. Click **Modify** on the relevant Virtual Service.
4. Expand the **SSL Properties** section.
5. Click **Show Intermediate Certificates**.
6. Select the relevant certificates from the boxes and click the arrows to remove/assign them from/to the Virtual Service.
7. Then, click **Set Intermediate Certificates**.

Note: It is not possible to unassign all certificates from the Virtual Service. If you do not want client certificates to be required - select **No Client Certificates required** in the **Client Certificates** drop-down list.

Advanced Properties

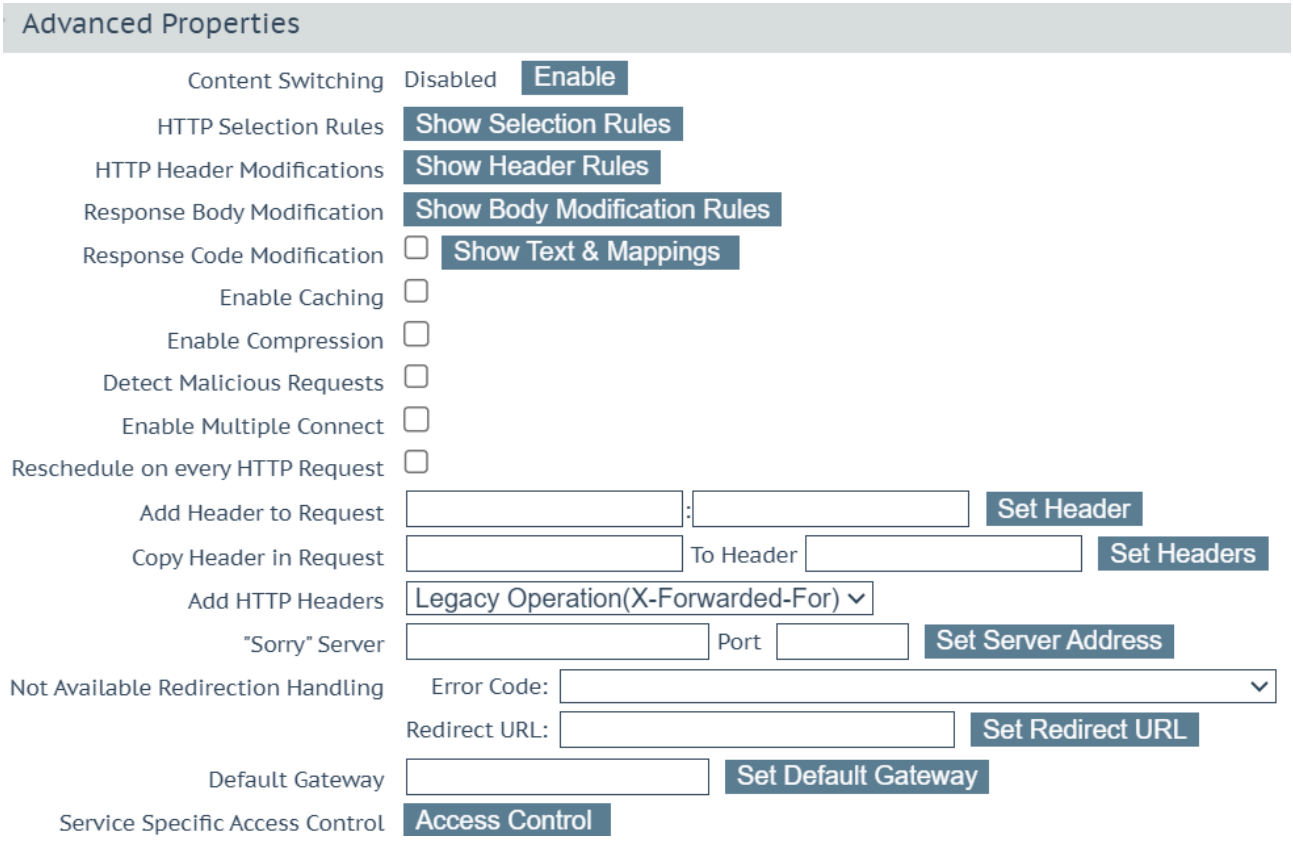

Content Switching

Clicking the **Enable** button, enables rule-based Content Switching on this Virtual Service. The **Enable** button only appears if the Virtual Service has a Real Server or SubVS. Once enabled, rules must be assigned to the various Real Servers. Rules can be attached to Real Server by clicking the **None** button located next the Real Server. Once rules are attached to a Real Server the **None** button will display the count of rules attached.

Rules Precedence

Clicking the **Rules Precedence** button displays the order in which Content Switching rules are applied. This option only appears when Content Switching and rules are assigned to the Real Server(s).

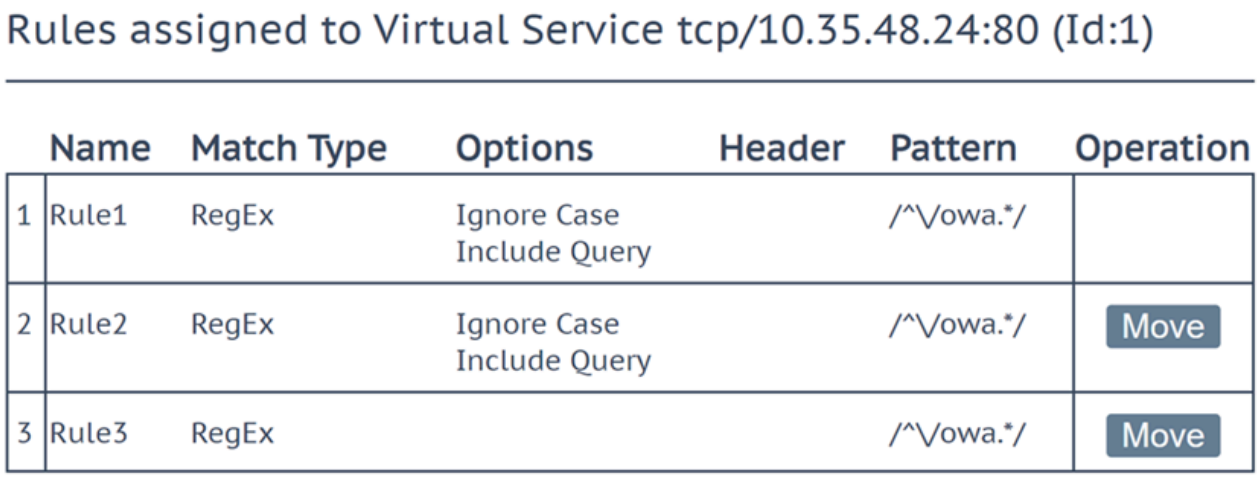

This screen shows the Content Switching rules that are assigned to the Real Servers of the Virtual Services and the order in which they apply. A rule may be promoted in the order of precedence by clicking its corresponding Move button.

In LoadMaster firmware 7.2.52 and above, it is easier to reorder the priority of rules in a Virtual Service - there is a move option that allows you to specify the position to move the rule to.

HTTP Selection Rules

Show the selection rules that are associated with the Virtual Service.

HTTP Header Modifications

Clicking **Show Header Rules** displays the order in which Header Modification rules are implemented. The number of rules (of both request and response type) is displayed on the actual button.

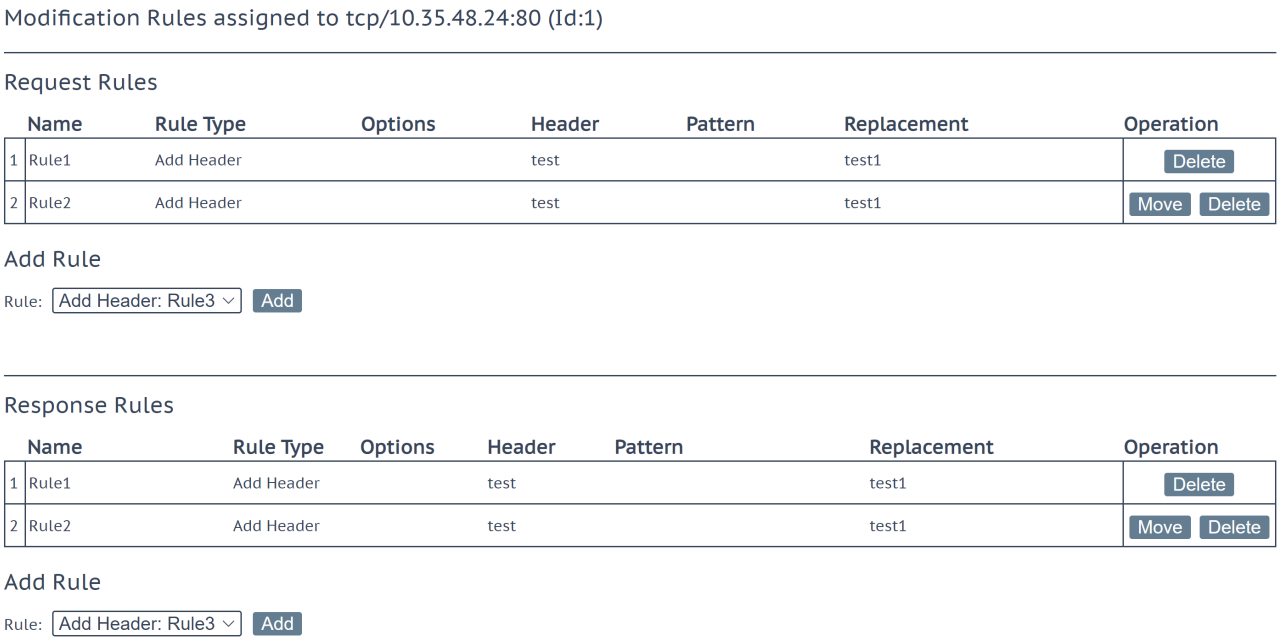

From within the screen you can **Add** and **Delete** Header Modification rules. The order in which the rules are applied can be changed by clicking the **Move** buttons.

In LoadMaster firmware 7.2.52 and above, it is easier to reorder the priority of rules in a Virtual Service - there is a move option that allows you to specify the position to move the rule to. Numbers are also now displayed on the page showing the content rules assigned to a Virtual Service to indicate the priority.

As of LoadMaster firmware version 7.2.51, you can assign URL modification rules to a response by selecting the relevant **Modify Response** rule in the **Response Rules** section.

Response Body Modification

Clicking the **Show Body Modification Rules** button displays the response body modification rules assigned to the Virtual Service. The number of assigned rules is displayed in the button label.

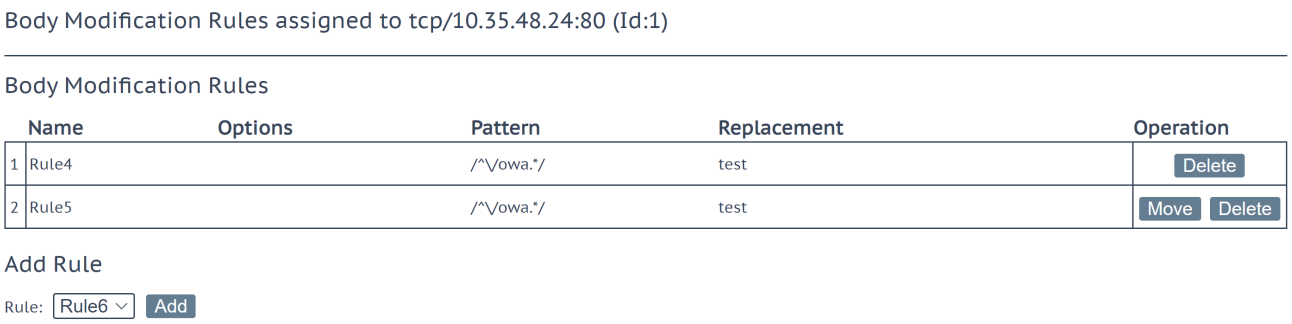

From this screen, you can **Add** and **Delete** response body modification rules to/from the Virtual Service. You can change the order that the rules are applied in by clicking the **Move** button.

In LoadMaster firmware 7.2.52 and above, it is easier to reorder the priority of rules in a Virtual Service - there is a move option that allows you to specify the position to move the rule to. Numbers are also now displayed on the page showing the content rules assigned to a Virtual Service to indicate the priority.

Note: Response body rules are not compatible with Kerberos Constrained Delegation (KCD). If KCD is enabled on a Virtual Service, it is not possible to assign a body rule to it.

Response Code Modification

The response code modification feature allows you to:

- Obfuscate error codes returned by servers
- Provide additional context (text) in server responses

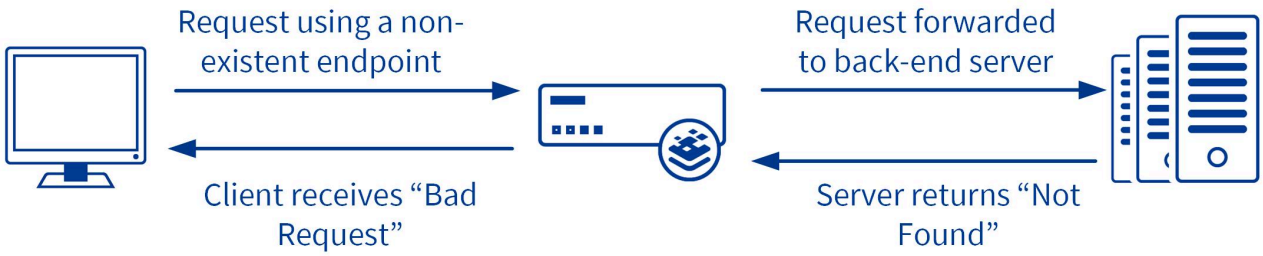

In the above example, the endpoint does not exist. Returning this information to the client could be considered an information leak. Using the response code modification feature, you can modify the response on the way back to the client and change it to "Bad Request", for example.

If **Response Code Modification** is enabled, the error messages specified in the **Response Mappings** list that are coming from the Real Servers are replaced as specified. The text of the message is encoded in JSON and contains a { "status" : <status code>, "message": <message string> }, where status is the response code coming from the Real Server (unless this has been changed in the error response mappings) and the message string is the standard message text for that error message (as defined here: HTTP response status code). No Real Server-generated text or header fields are passed to the client.

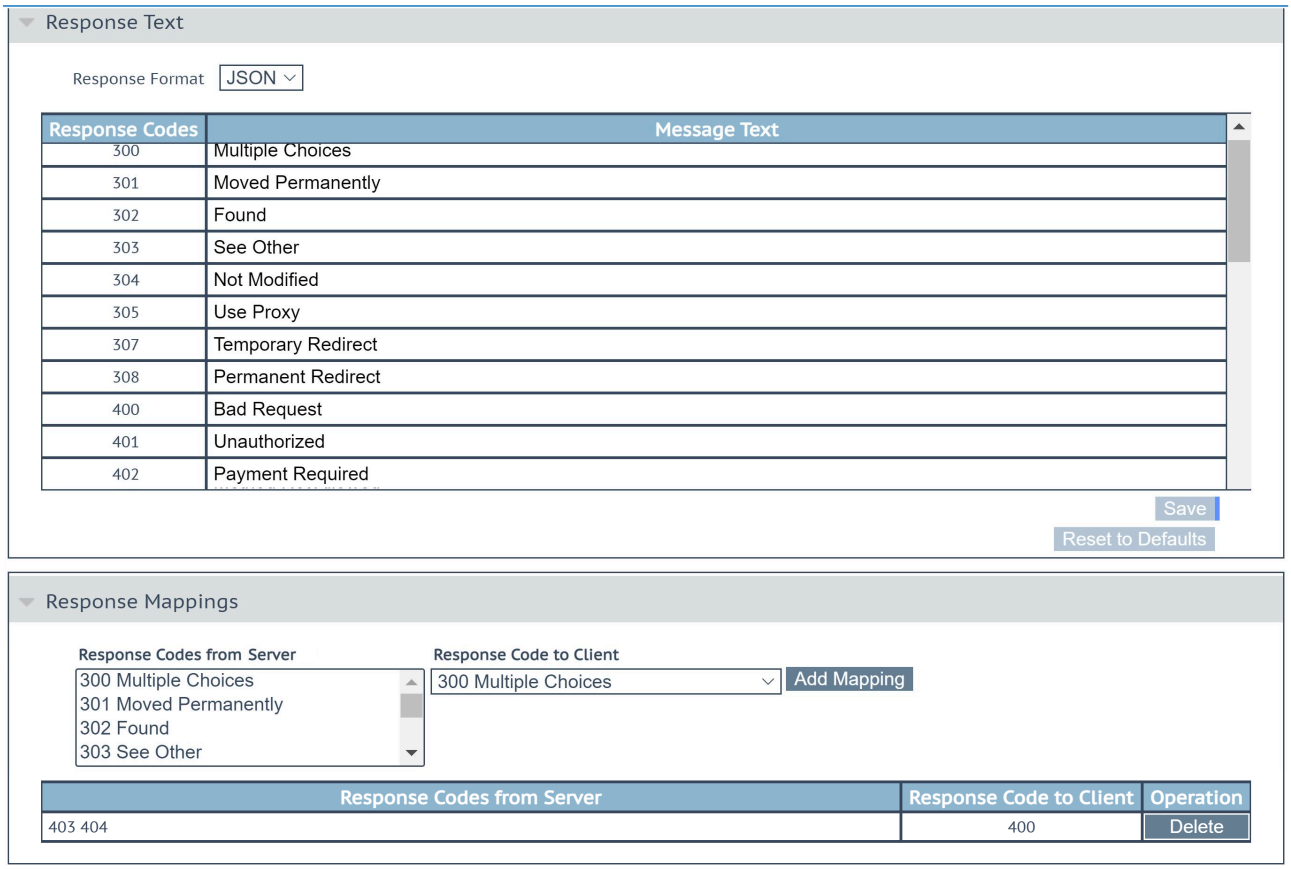

Click **Show Text & Mappings** to open HTTP Response Code Management screen. Clicking **Response Text** displays a table of response codes and corresponding text messages you can select and edit. Click **Save** to save all updated text messages. To restore the default set of messages, click **Reset to Defaults**. You can select the desired response format (JSON/XML) from the **Response Format** drop-down list.

In the **Response Mappings** section, you can specify what response code to return to the client based on the response code from the server. For example, you could map the **404 Not Found** response code from the server to **400 Bad Request** before it is sent back to the client.

This section also displays the list of existing Error Response Code Mappings for a Virtual Service. Each server response code can only be mapped once per Virtual Service. However, you can map multiple different server response codes to the same client response code. For example, you can map errors 403 and 404 to error 400. You can select multiple server response codes at once if required.

Enable HTTP/2 Stack

Enable HTTP/2 client requests to be served by the LoadMaster directly. HTTP/2 requests are made using a secure connection. Please ensure the **SSL Properties** are configured and the **BestPractices Cipher Set** is selected if enabling this option. The **Enable Caching** check box should also be selected to optimize end user experience.

Enable Caching

This option enables caching of static content. This saves valuable Real Server processing power and bandwidth. You can enable caching on HTTP, offloaded HTTPS Virtual Services, and Virtual Services with SSL re-encryption enabled.

Before caching, the LoadMaster checks if there is any reason the content should not be cached. Some reasons the content may not be cached are:

- The extension is on the **File extensions that should not be cached** list in **Virtual Services > Cache Configuration**.
- The content is already cached.
- The status code is not 200, 203, 300, 301, or 410.
- The file was created in the last 60 seconds.

Items are held in the cache for a maximum of 1 hour. This time is not customizable. Every 15 minutes the LoadMaster checks if what is currently in the cache should still be there (it checks the maximum age of the header) and at that point it clears the cache as needed. The cache can be flushed on the browser by refreshing the page (for example, by pressing **Ctrl + F5** on your keyboard). Disabling and re-enabling caching on the Virtual Service flushes the cache for that Virtual Service on the LoadMaster.

Note: Types of files that can be cached may be defined in **Cache Configuration** under the **Virtual Services** menu.

Maximum Cache Usage

If you enable caching on a Virtual Service, you can then limit the amount of cache being used by that service. The LoadMaster reserves 20% of the total system memory for the global cache value. For example, if a LoadMaster has 1GB of RAM - 20% (or 204MB) is reserved for caching and compression.

When you enable caching on a Virtual Service, you then have the option to set **No Limit** (which is the default value) or a percentage usage value. Selecting **No Limit** allows you to use all of the available cache. Alternatively, you can specify a percentage from 1 to 99%. This allows you to break up the 204MB in the example above into a percentage. For example, if you set 20% on the Virtual Service level - the Virtual Service has a total of 40MB available to cache server responses (20% of 204MB = 40.8MB). This leaves approximately 160MB available for other Virtual Services.

Note: This memory is also used by the LoadMaster for content switching and body response rules.

After selecting a percentage, the LoadMaster displays the overall total current usage assigned. This is a combination of the current usage assigned for all Virtual Services. For example, if there are two Virtual Services and on one the **Maximum Cache usage** is set to 10% and on the other it is set to 7%, the **Current usage assigned** value displays 17%. It is possible to over-provision/commit - if you attempt to do this, you will get a message saying **Setting this value will overcommit the Cache memory by N%**.

It is recommended to limit the cache size to prevent unequal use of the cache store. Ensure that the cache maximum usage is adjusted so that each Virtual Service has a percentage of cache to use. If there is no remaining space to be allocated for a cache-enabled Virtual Service, that service will not cache content.

Enable Compression

Files sent from LoadMaster are compressed with Gzip.

Note: If compression is enabled without caching, LoadMaster performance may suffer. When caching and compression are both enabled, caching has priority and therefore the first file is cached but not compressed.

On second and subsequent requests, the file is found in the cache and then it is compressed.

The types of files that can be compressed may be defined in **Compression Options** in the **Virtual Services** menu of the LoadMaster WUI.

Note: Compression is not recommended for files 100MB or greater in size.

Note: More RAM may need to be added to Virtual LoadMasters using the hypervisor to compress large files.

Detect Malicious Requests

The Intrusion Prevention System (IPS) service will provide in-line protection of Real Server(s) by providing real-time mitigation of attacks and isolation of Real Server(s). Intrusion prevention is based on the industry standard SNORT database and provides real-time intrusion alerting.

Note: To get updated or customized rules, please refer to the SNORT website: https://www.snort.org/.

Note: The detection code only handles HTTP-classed rules.

Selecting the **Detect Malicious Requests** check box enables the IPS per HTTP and offloaded HTTPS Virtual Services. There are two options for handling of requests that match a SNORT rule. **Drop Connection**, where a rule match will generate no HTTP response, or **Send Reject**, where a rule match will generate a response to the client of HTTP 400 “Invalid Request”. Both options prevent the request from reaching the Real Server(s).

Enable Multiple Connect

Enabling this option permits the LoadMaster to manage connection handling between the LoadMaster and the Real Servers. Requests from multiple clients are sent over the same TCP connection.

Note: Multiplexing only works for simple HTTP GET operations. The **Enable Multiple Connect** check box will not be available in certain situations, for example if WAF, ESP or SSL Acceleration is enabled.

Reschedule on every HTTP Request

The **Reschedule on every HTTP Request** check box is only available on Virtual Services if the Protocol is set to **tcp** and the **Service Type** is set to **HTTP-HTTP/2-HTTPS**.

If the **Reschedule on every HTTP Request** option is enabled and content switching is not in use, this forces the reselection of a Real Server for each request. By default, the Real Server is selected only once per connection.

Note: You cannot enable the **Reschedule on every HTTP Request** checkbox if the **Enable HTTP/2 Stack** option is enabled.

Note: **Persistence** must be set to **None** (in **Standard Options**) to use the **Reschedule on every HTTP Request** feature.

Port Following

Port following is set when two services need to share persistence records. Typically, this is done for HTTP and HTTPS services so users maintain a server session, regardless of whether they connect securely or not.

If the Real Server for one of the Virtual Services fails, the persistence records for the same Real Server on the other Virtual Service will be cleared.

Port following has several requirements:

- The Virtual Services must have the same set of Real Servers
- The Virtual Service must be using the same persistence options

After meeting these conditions, in the Virtual Service modify screen there will be an option under **Advanced Properties** for **Port Following**. Ensure to set this on both Virtual Services to ensure that port following is done bi-directionally. Port following must be set up bi-directionally to ensure that, regardless of whether the client connects using HTTP or HTTPS, the persistence and session is saved.

Note: Prior to LoadMaster firmware version 7.2.51, it was not possible to configure port following for Virtual Services where the **Service Type** was set to **Generic**. This is possible in version 7.2.51 and above.

For further information, refer to the Port Following Feature Description.

Add Header to Request

Input the key and the value for the extra header that is to be inserted into every request sent to the Real Servers.

Click the **Set Header** button to implement the functionality.

Copy Header in Request

This is the name of the source header field to copy into the new header field before the request is sent to the Real Servers. Enter the name of the header field into which the source header is to be copied in the **To Header** text box.

Add HTTP Headers

This option allows you to select which headers are to be added to the HTTP stream. The options available include:

- Legacy Operation(X-Forwarded-For)
- None
- X-Forwarded-For (+ Via)
- X-Forwarded-For (No Via)
- X-ClientSide (+ Via)
- X-ClientSide (No Via)
- Via Only

HTTP headers let the client and the server pass additional information with an HTTP request or response. Not all Virtual Services can handle HTTP data (header/body), for example, RDP Virtual Services. To improve the performance and to save unnecessary work, not all HTTP Virtual Services need to look at the HTTP data. Simple pass-through Virtual Services do not need to process the HTTP requests at all. If rules are defined on the Virtual Service, a header is added, or some more complicated feature is required, the Virtual Service does need to look at the HTTP data.

In older (legacy) versions of the LoadMaster, the setting for **Add HTTP Headers** was only global and it could not be set on a per Virtual Service basis. You can now update the setting for **Add HTTP Headers** on a specific Virtual Service if required. The default value is **Legacy Operation(X-Forwarded-For)** for backward compatibility reasons.

The **Add HTTP Headers** field is not configurable in the parent Virtual Service if a SubVS exists. We recommend not changing the **Add HTTP Headers** value in the parent Virtual Service if you intend to add a SubVS. If you need to change this in the parent Virtual Service after adding a SubVS, you can either use a content rule or run an API command to change the value, for example:

```
https://<LoadMasterIPAddress>/access/modvs?vs=<VSIndex>&Addvia=5
```

For further details on the RESTful API, including the valid values for the **Addvia** parameter, refer to the RESTful API documentation.

"Sorry" Server

Enter the IP Address and Port number in the applicable fields. If no Real Servers are available, the LoadMaster will forward to a specified IP:PORT, with no checking. The IP address of a "Sorry" Server must be on a network or subnet that is defined on the LoadMaster.

Note: When using a Layer 4 Virtual Service, the "Sorry" Server should be on the same subnet as the Real Server.

Note: When using a Layer 7 Virtual Service, the "Sorry" Server can be on any local network. It is also possible to add a non-local "Sorry" Server. For this, **Transparency** must be disabled, there must be a route to the "Sorry" Server and the **Enable Non-Local Real Servers** option must be enabled (**System Configuration > Miscellaneous Options > Network Options**).

Not Available Redirection Handling

When no Real Servers are available to handle the request you can define the error code and URL that the client should receive.

- **Error Code:** If no Real Servers are available, the LoadMaster can terminate the connection with a HTTP error code. Select the appropriate error code.
- **Redirect URL:** When there are no Real Servers available and an error response is to be sent back to the client, a redirect URL can also be specified. If the string entered in this text box does not include **http://** or **https://** the string is treated as being relative to the current location, so the hostname is added to the string in the redirect. This field also supports the use of wildcards such as **%h** and **%s** which represent the requested hostname and Uniform Resource Identifier (URI) respectively.
- **Error Message:** When no Real Servers are available and an error response is to be sent back to the client, the specified error message is added to the response.

For security reasons, the returned HTML page only returns the text **Document has moved**. No request-

supplied information is returned.

- **Error File:** When no Real Servers are available and an error response is to be sent back to the client, the specified file is added to the response. This enables simple error HTML pages to be sent in response to the specified error.

The maximum size of this error page is 16KB.

"Sorry" Server/Port



In a UDP Virtual Service there is an option to specify a **"Sorry" Server** and **Port**. When there are no Real Servers available to handle the request, this option defines the URL that the client will receive.

Note: The value of the **"Sorry" Server** can only be changed for UDP if the service is not currently using the **"Sorry" Server**.

Add a Port 80 Redirector VS

If no port 80 Virtual Service is configured, one can be created. It will then redirect the client to the URL specified in the **Redirection URL:** field.

Click the **Add HTTP Redirector** button to implement the redirector.

Note: When the **Add HTTP Redirector** button is clicked, a redirect Virtual Service is created and this WUI option disappears from the relevant Virtual Service.

Default Gateway

Specify the Virtual Service-specific gateway to be used to send responses back to the clients. If this is not set, the global default gateway is used.

Click the **Set Default Gateway** button to implement the default gateway. The **Default Gateway** for a Virtual Service is only used for that Virtual Service.

Note: If the global **Use Default Route Only** option is set in **System Configuration > Miscellaneous Options > Network Options**, traffic from Virtual Services that have the **Default Gateway** set is only routed to the interface where the Virtual Service's default route is located. This can allow the LoadMaster to be directly connected to client networks without returning traffic directly using the adjacent interface.

Alternate Default Gateway

This field is only visible if an **Alternate Address** is set for the Virtual Service. You should only use the **Alternate Default Gateway** field if all of the following conditions are met:

- The **Alternate Address** has a different address family to the main Virtual Service address, for example, the Virtual Service address is IPv4 and the **Alternate Address** is IPv6 or the other way around.

- The Virtual Service **Default Gateway** is set.
- A second Virtual Service **Default Gateway** is set for the other address family.

Alternate Source Addresses

If no list is specified, the LoadMaster will use the IP address of the Virtual Service as its local address. Specifying a list of addresses ensures the LoadMaster will use these addresses instead.

Click the **Set Alternate Source Addresses** button to implement the Alternate Source Addresses.

Note: This option is only available if the **Allow connection scaling over 64K Connections** option is enabled in the **L7 Configuration** screen.

Service Specific Access Control

Allows you to change the Virtual Service-specific **Access Control** lists.

Note: If you implement the Access Control Lists option, the **Extra Ports** option will not work correctly.

When using Access Control Lists on a Virtual Service that has the same IP address as an interface (which we do not recommend) the following ports are never blocked for Real Servers on the same network interface accessing the VS as a client:

- 443 (WUI)
- 22 (SSH)
- 53 (DNS)
- 161 (SNMP)

Web Application Firewall (WAF) Options

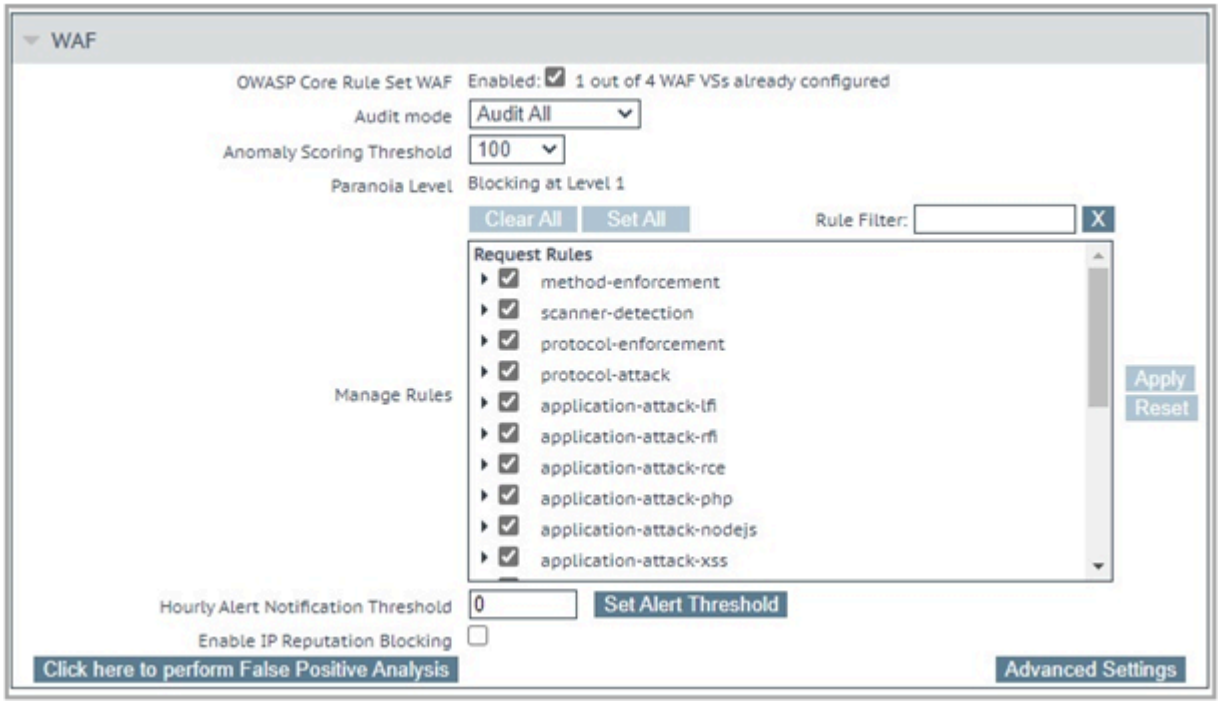

By default, WAF is disabled. To enable WAF, select **Enabled**.

When WAF is enabled for a Virtual Service, the section heading in the Virtual Service options changes from **WAF** to **WAF - Enabled**.

The maximum number of WAF-enabled Virtual Services is the total (unused or available) RAM (in MB)/512 MB. For example: 8 GB/512 MB = 16 WAF-enabled Virtual Services. When the maximum is reached, no additional Virtual Services can be enabled with WAF.

A message displays if there is insufficient memory available to enable WAF.

Note: A message is displayed next to the **Enabled** check box displaying how many WAF-enabled Virtual Services exist and the maximum number of WAF-enabled Virtual Services that can exist. If the maximum number of WAF-enabled Virtual Services is reached, the **Enabled** check box is grayed out.

Audit mode

There are three audit modes:

- **No Audit:** No data is logged.
- **Audit All:** Logs all data through the Virtual Service.
- **Audit Relevant:** Logs data that is of a warning level and higher. This is the default option for this setting.

Note: Selecting the **Audit All** option produces a large amount of log data. We do not recommend selecting

---

the **Audit All** option for normal operation. However, the **Audit All** option can be useful when troubleshooting a specific problem.

Anomaly Scoring Threshold

For each request, every triggered detection raises the anomaly score, most rules having a score of 5. If the cumulative anomaly score per request hits the configured limit, the request will be blocked. The default value is 100 and allowable range is 1 to 10000.

Note: The **Paranoia Level** can be set in **Advanced Settings**, but the value is displayed here for informational purposes.

Manage Rules

Rules are grouped in the **Request Rules** section as per the OWASP numbering system. Rule groups or Individual rules within each ruleset can be enabled/disabled as required. To enable a rule or group of rules, select the relevant check box. If you have previously enabled/disabled rules in that ruleset, within that Virtual Service – the rules retain their previous settings.

Note: Some rules or rule sets may have dependencies on other rules. There is no dependency check in the LoadMaster when rules are disabled - before disabling any rule, be aware of any rule chains or dependencies.

If a user has created custom rules, they can be enabled or disabled within the **Custom Rules** section.

There is a **Run First** check box available for custom rules. If the **Run First** check box is enabled for a custom rule, the rule will be run first, before the OWASP Core Rule Set (CRS). If the Run First check box is disabled for a custom rule, the custom rule runs after the CRS. The **Run First** check box is disabled by default.

In the **Workloads** section there are several workloads available.

When finished making changes, click **Apply**.

Clicking **Reset** reverses any changes that you have made that have not been applied.

To filter rules, enter text into the **Rule Filter** text box and only rules containing that text will be shown. You can select the filtered rules by clicking **Set All** or deselect the filtered rules by clicking **Clear All**. Click **Apply** to apply the changes.

Hourly Alert Notification Threshold

This is the number of incidents per hour before sending an alert. Setting this to **0** disables alerting.

Enable IP Reputation Blocking

In **Web Application Firewall > Access Settings** you can download and install the latest IP reputation file. If **Enable IP Reputation Blocking** is selected for a Virtual Service, client addresses are checked against the IP access list file and are blocked if a match is found.

Contents

- Advanced Settings
- False Positive Analysis
- WAF Misconfigured Virtual Service Status

Advanced Settings

Click the **Advanced Settings** button to configure the advanced OWASP settings.

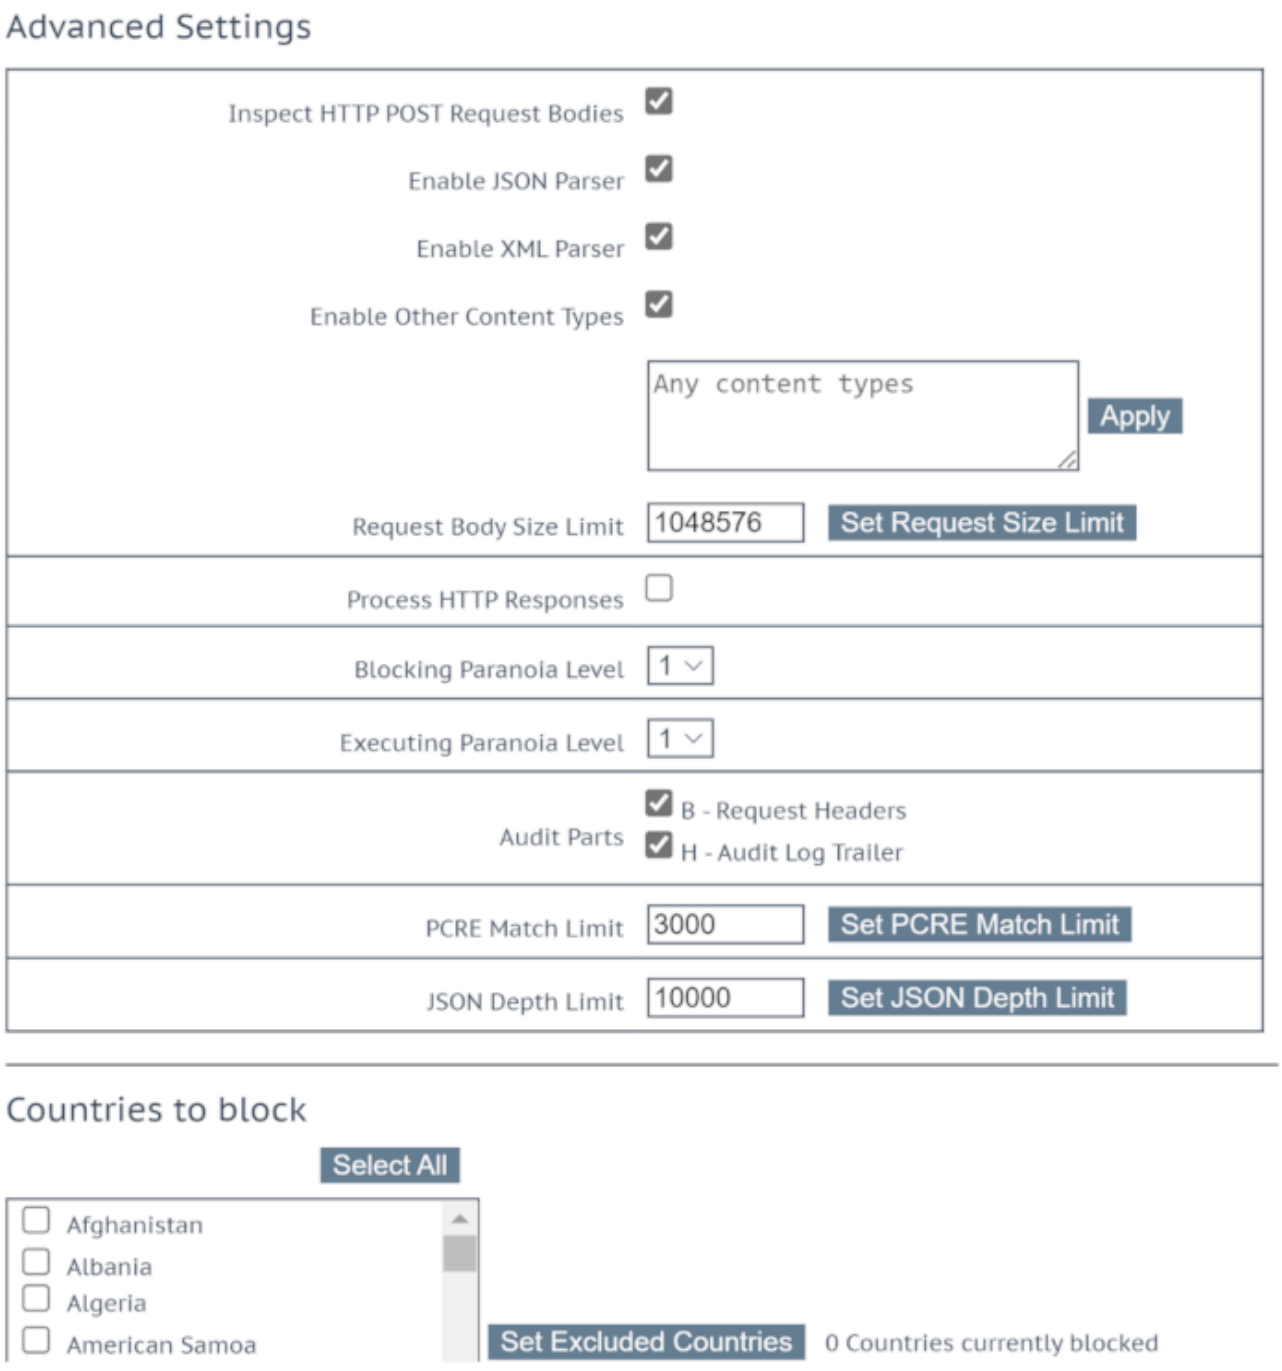

Inspect HTTP POST Request Bodies

This option is disabled by default. If you enable this option, three more check boxes become available that allow you to enable the processing of JavaScript Object Notation (JSON), Extensible Markup Language (XML) requests, and other content types.

Request Body Size Limit

This option allows you to set the maximum size of POST request bodies that the WAF engine will allow through. The default value is 1048576 bytes. The range of valid values is 1024 to 52428800 bytes (50 MB).

CAUTION: Higher values require more memory resources and may impact WAF engine performance.

Process HTTP Responses

Enables checking of the responses from the server to the client.

Enabling the **Process HTTP Responses** option makes two more options, **E - Intended Response Body** and **F- Response Headers**, available in the **Audit Parts** options

---

Note: The processing of response data can be CPU and memory intensive and may impact on performance.

Blocking Paranoia Level

Defines how strictly the ModSecurity engine implements each rule. The default Paranoia Level value is set at 1. With each paranoia level increase, the CRS enables stricter implementations of the rules, giving you a higher level of security. However, higher paranoia levels also increase the possibility of blocking some legitimate traffic due to false positives. If you use higher paranoia levels, you will likely need to add some exclusion rules for certain applications that need to receive complex input patterns.

Executing Paranoia Level

Defines the paranoia level at which the ModSecurity engine checks/verifies the requests coming from the servers. The results of the checks will be logged but the **Executing Paranoia Level** is not used to determine what traffic will be blocked.

Though the **Executing Paranoia Level** can be higher than the **Blocking Paranoia Level**, it cannot be lower. A higher **Executing Paranoia Level** enables users to see which rules would be triggered at a higher Paranoia level without blocking traffic.

Audit Parts:

A single string that contains the sections that are to be entered in the WAF audit log for each request. The supported values are A, B, E, F, H, K, Z, though only the values B, E, F, H can be enabled or disabled.

For further information regarding the Audit Parts, please refer to https://github.com/SpiderLabs/ModSecurity/wiki/ModSecurity-2-Data-Formats

PCRE Match Limit:

This setting sets the maximum iterations that the internal PCRE engine uses before failing a match. A lower value may cause a valid match to fail, whereas a higher value may cause the WAF engine to run slower. The

default value is 10000. The maximum value is 9999999.

JSON Depth Limit

This value sets the maximum depth that will be accepted during JSON parsing. Lower values may cause a valid match to fail. Higher values may cause the WAF engine to run slower. The default value is 10000. The range of valid values is 1000 to 99999.

Workloads

When a workload is selected, the OWASP CRS optimizes the rules to ensure that known false positives are not returned.

Countries to block:

Based on GEO IP information, you can select countries that should not be allowed access. Click the **Select All** button to block the access for all countries or select individual countries from the country list that are to be blocked and click the **Set Excluded Countries** button.

False Positive Analysis

This feature allows users to perform false positive analysis against their applications to obtain enhanced visibility of attacks and fine-tune protection. Click the **Click here to perform False Positive Analysis** button to check False Positives against any virtual service that runs OWASP CRS rules.

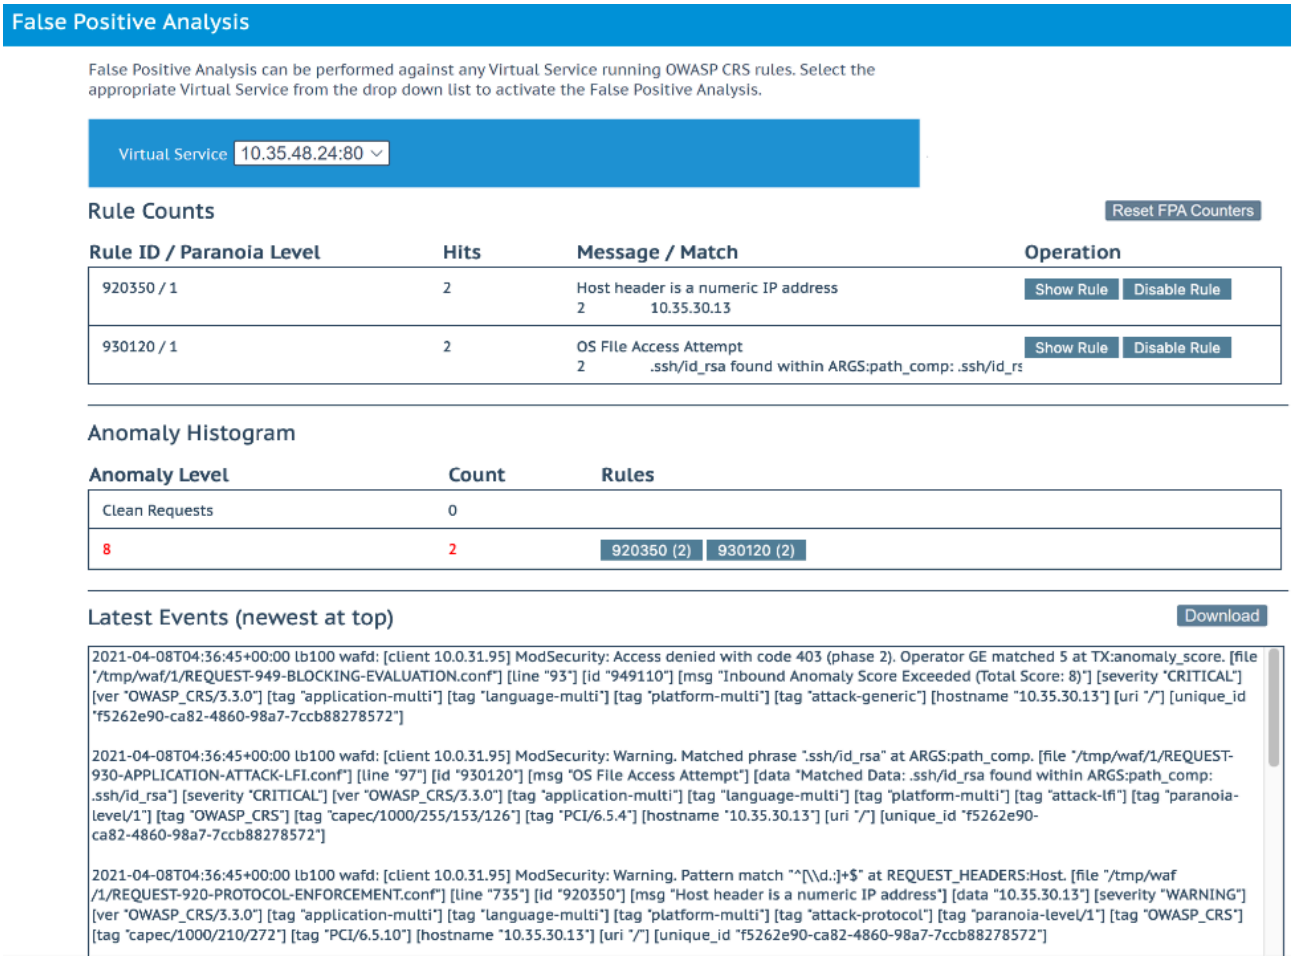

Rule Counts

The Rule Counts section displays information on any rules that are being triggered by requests. It displays the Rule ID, the paranoia level the rule is running under, the number of hits per requests that have triggered the rule and the message or match for the request are displayed for each rule that is triggered.

Clicking the **Show Rule** button in the **Operation** column displays the contents of the rule file associated with the triggered rule. This opens in a separate tab and the URL contains the triggered rule id.

The rule can be disabled by clicking the **Disable Rule** button.

Reset FPA Counter

Reset all False Positive Analysis Counters (Anomaly Histogram and Latest Events) for the virtual service. Clearing the Latest Events does not remove the logs from the LoadMaster, they are still available under **System Configuration > Logging Options > System Log Files > WAF Event Log File**.

Anomaly Histogram

The first row of the **Anomaly Histogram** section displays how many requests have been run without triggering a rule.

Each subsequent row gives details of rules that have been triggered and which are affecting the Anomaly Score. In each row the cumulative Anomaly Score, the number of requests which have triggered the rule and

the rule details are provided

Latest Events (newest at top)

Displays the event details for each rule that is triggered. These messages are in the standard ModSecurity log format and contains the anomaly score, the warning message, the attack state, and the paranoia level.

Download

Click the **Download** button to download the displayed WAF event logs details.

WAF Misconfigured Virtual Service Status



On the View/Modify Services screen in the LoadMaster UI, the Status of each Virtual Service is displayed. If the WAF for a particular Virtual Service is misconfigured (for example, if there is an issue with a rule file), the status changes to WAF Misconfigured and turns to red.

CAUTION: If the Virtual Service is in a WAF Misconfigured state, all traffic stops flowing. WAF can be disabled for that Virtual Service to stop the traffic being blocked, if required, while troubleshooting the problem.

Edge Security Pack (ESP) Options

The ESP feature must be enabled before you can configure these options. To enable the ESP function, please select the **Enable ESP** check box.

Note: ESP is not available on all Support tiers. Refer to the following page for details on what features are included in the different Support tiers: LoadMaster Support Subscriptions

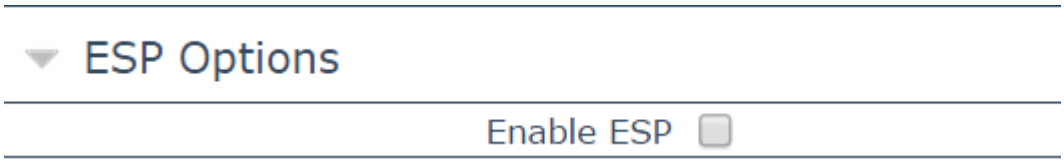

The full **ESP Options** screen will appear.

Note: The ESP feature can only be enabled if the Virtual Service is a HTTP, HTTPS, or SMTP Virtual Service.

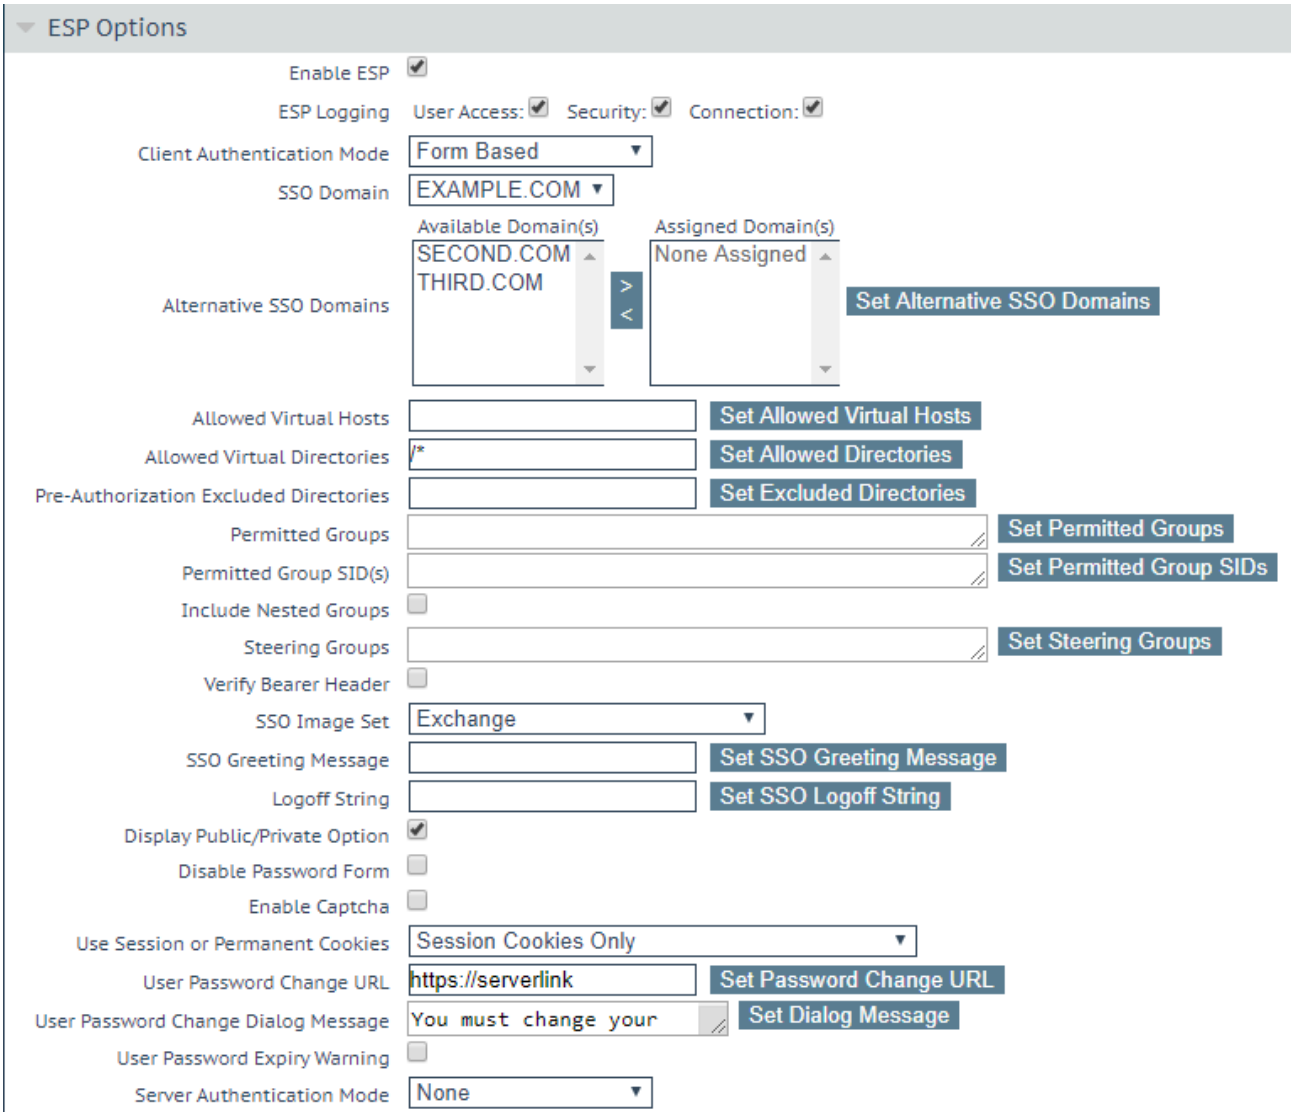

Enable ESP

Enable or disable the ESP feature set by selecting or removing the checkmark from the **Enable ESP** checkbox.

ESP Logging

There are three types of logs stored in relation to the ESP feature. Each of these logs can be enabled or disabled by selecting or deselecting the relevant checkbox. The types of log include:

- **User Access:** logs recording all user logins
- **Security:** logs recording all security alerts
- **Connection:** logs recording each connection

Logs are persistent and can be accessed after a reboot of the LoadMaster. For further information on logs please refer to the Extended Log Files section.

Client Authentication Mode

Specifies how clients attempting to connect to the LoadMaster are authenticated. The following types of methods are available:

- **Delegate to Server:** the authentication is delegated to the server
- **Basic Authentication**: standard Basic Authentication is used
- **Form Based:** clients must enter their user details within a form to be authenticated on the LoadMaster
- **Client Certificate:** clients must present the certificate which is verified against the issuing authority

In LoadMaster firmware version 7.2.53, support was added for **Client Certificate** authentication with no server side authentication. For further details, refer to the ESP Feature Description.

- **NTLM/NTLM-Proxy:** NTLM credentials are based on data obtained during the interactive logon process and consist of a domain name and a user name
- **SAML:** The LoadMaster supports SAML, playing the role of a SAML service provider. The service provider provides secure, gated access to a resource.
- **Pass Post:** In LoadMaster firmware version 7.2.53, a new mode called **Pass Post** was introduced. With this change introduced, users with valid credentials using the Workspace client app can successfully log in (using Single Sign On (SSO)) using POST-based authentication on the client side and **Form Based** Authentication (FBA) on the server side and access is granted to the VDI workspace.
- **OIDC/OAUTH:** Open ID Connect (OIDC) is an authentication protocol based on the OAuth2 protocol used to enable Single Sign On of users across multiple applications via a single Identity Provider. OIDC uses the standardized message flows from OAuth2 to provide identity services.

Note: The remaining fields in the **ESP Options** section will change based on the **Client Authentication Mode** selected.

SSO Domain

Select the Single Sign-On (SSO) Domain within which the Virtual Service is included.

Please refer to the Manage SSO Domains section for further information on configuring SSO Domains. An SSO Domain must be configured to correctly configure the ESP feature.

Note: Only SSO domains with the **Configuration type** of **Inbound Configuration** are shown as options in this **SSO Domain** field.

Alternative SSO Domains

Many organizations use extranets to share information with customers and partners. It is likely that extranet portals will have users from two or more Active Directory domains. Rather than authenticating users from individual domains one at a time, assigning **Alternative SSO Domains** gives the ability to simultaneously authenticate users from two or more domains using one Virtual Service.

Note: This option appears only when more than one domain has been configured and when the **Authentication Protocol** for the SSO domains are set to **LDAP**.

Please refer to the Manage SSO Domains section for further information on configuring **SSO Domains**.

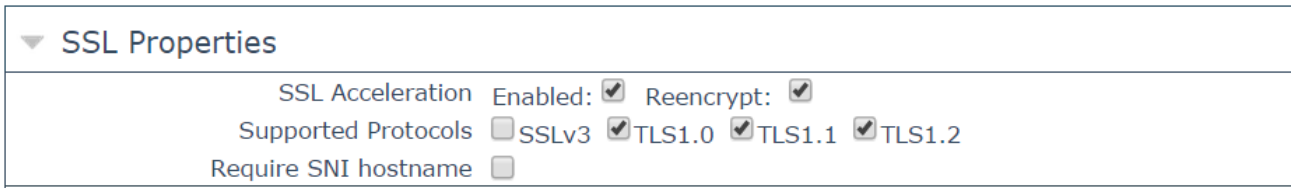

Before enabling ESP, ensure that SSL offloading is configured for the HTTPS Virtual Service.

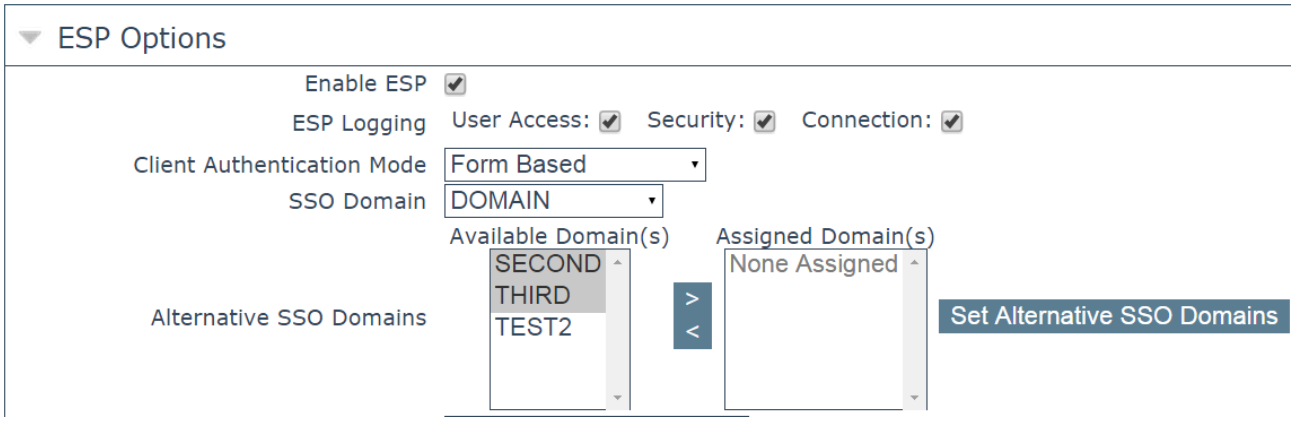

The domain name which appears in the **SSO Domain** drop-down list is the default domain. This is also the domain which is used if only one is configured.

Previously configured alternative domains appear in the **Available Domain(s)** list.

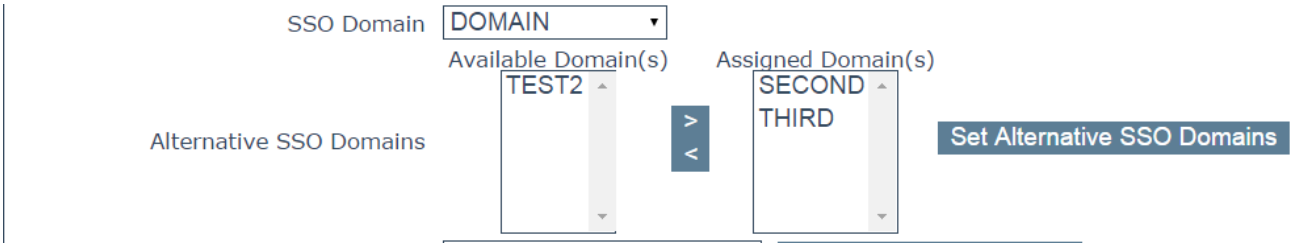

To assign alternative SSO Domains:

1. Highlight each of the domains you wish to assign and click the **>** button.

   Note: An assigned domain is a domain which can be authenticated using a particular Virtual Service.

   Note: All domains which appear as available may be assigned to a Virtual Service.

2. Click the **Set Alternative SSO Domains** button to confirm the updated list of Assigned Domain(s).
3. Choose **Basic Authentication** from the **Server Authentication Mode** drop-down list.

Note: When logging in to a domain using the ESP form, users should enter the name of the SSO Domain if an alternative domain needs to be accessed. If no domain name is entered in the username, users are, by default, logged on the domain entered in the default SSO Domain drop-down list.

To view the status of the Virtual Services, click **Virtual Services** and **View/Modify Services** in the main menu.

A list of the **Virtual Services** displays showing the current status of each service.

If alternative domains are assigned and there is an issue with a particular domain, the affected domain name is indicated in the **Status** column.

Allowed Virtual Hosts

The Virtual Service will only be allowed access to specified virtual hosts. Any virtual hosts that are not specified are blocked.

Enter the virtual host name(s) in the **Allowed Virtual Hosts** field and click the **Set Allowed Virtual Hosts** button to specify the allowed virtual hosts.

If a Virtual Service is configured with non-standard HTTP/HTTPS ports, ensure the Virtual Service WebsiteName:Port is entered in the **Allowed Virtual Hosts** text box to allow the HTTP/HTTPS request to pass the ESP check. For example, if the requested Virtual Service website name is **sample.com** and the non-standard port is **8888**, ensure the **Allowed Virtual Hosts** text box contains **sample.com:8888**. Alternatively, you could use a wildcard symbol (*****) to allow any port number, for example, **sample.com:***.

Multiple domains may be specified within the field allowing many domains to be associated with the Single Sign On Domain.

The use of regular expressions is allowed within this field.

Note: If you use quotes in regular expressions in the LoadMaster WUI, there are limitations.For more information, refer to the section Limitations of Using Regular Expressions in the LoadMaster WUI.

Note: If the **Allowed Virtual Hosts** field is left blank, the Virtual Service is blocked. However, when the **Client Authentication Mode** is set to **Delegate to Server**, leaving this field blank does not block traffic.

Note: When using the **Permitted Groups** field in **ESP Options**, you need to ensure that the SSO domain set here is the directory for the permitted groups. For example, if the **SSO Domain** is set to webmail.example and webmail is not the directory for the permitted groups within example.com, it will not work. Instead, the **SSO Domain** needs to be set to .example.com.

Allowed Virtual Directories

The Virtual Service will only be allowed access to the specified virtual directories, within the allowed virtual hosts. Any virtual directories that are not specified are blocked.

Enter the virtual directory name(s) in the **Allowed Virtual Directories** field and click the **Set Allowed Virtual Directories** button to specify the allowed virtual directories.

The use of regular expressions is allowed within this field.

Note: If the **Allowed Virtual Directories** is left blank, the Virtual Service is blocked. However, when the **Client Authentication Mode** is set to **Delegate to Server**, leaving this field blank does not block traffic.

Pre-Authorization Excluded Directories

Any virtual directories specified within this field will not be pre-authorized on this Virtual Service and are passed directly to the relevant Real Servers.

Permitted Groups

Specify the groups that are allowed to access this Virtual Service. When set, if a user logs in to a service published by this Virtual Service, the user must be a member of at least one of the groups specified. Performance may be impacted if a large number of groups are entered. Groups entered in this field are validated using an LDAP query.

Some guidelines about this field are as follows:

- The group(s) specified must be valid groups on the Active Directory in the SSO domain associated with the Virtual Service. The SSO domain in the LoadMaster must be set to the directory for the groups. For example, if the SSO domain in the LoadMaster is set to webmail.example and webmail is not the directory for the groups, it will not work. Instead, the SSO domain may need to be set to .example.com.
- The group(s) listed must be separated by a semi-colon

Note: A space-separated list does not work because most groups contain a space in the name, for example **IT Users**.

- Do not use the **Domain Users** group because it is a default primary group for new users.
- The following characters are not allowed in permitted group names:**/ : + ***
- The authentication protocol of the SSO domain must be LDAP
- The groups should be specified by name, not by full distinguished name
- Do not enter the same group name in both the **Permitted Groups** and **Steering Groups** fields. This causes a conflict. When you specify a steering group, it is assumed to behave like a permitted group, so you do not need to enter the same group in both the **Permitted Groups** and **Steering Groups** fields.

Permitted Group SID(s)

This field is the equivalent of the **Permitted Groups** field. If specifying permitted groups, you can complete either the **Permitted Groups** field or the **Permitted Groups SID(s)** field (security identifiers).

In the **Permitted Group SID(s)** field you can specify the group SIDs that are allowed to access this Virtual Service. After you type the groups, click **Set Permitted Group SIDs**.

This field allows a list of group SIDs of up to 64 bytes in length (192 characters in the format *NN NN NN*).

Each group is separated by a semi-colon. Spaces are used to separate bytes in certain group SIDs. Here is an example:

S-1-5-21-3763804817-1170992687-1336323834-1151

SIDs can be found by using the **get-adgroup-Identity GroupName** command.

Include Nested Groups

This field relates to the **Permitted Groups** setting. Enable this option to include nested groups in the authentication attempt. If this option is disabled, only users in the top-level group are granted access. If this option is enabled, users in both the top-level and first sub-level group are granted access.

Multi Domain Permitted Groups

In LoadMaster firmware version 7.2.52, a new check box was added to the **ESP Options** section of the Virtual Service modify screen called **Multi Domain Permitted Groups**. This check box is configurable with the following client authentication modes:

- Basic Authentication
- Form Based
- Client Certificate
- NTLM

When **Multi Domain Permitted Groups** is enabled, the LoadMaster checks for permitted group membership within all sub-domains under the top-level domain.

Note: The **Multi Domain Permitted Groups** option works with the **Permitted Groups**, **Permitted Group SID(s)**, and **Include Nested Groups**.

If **Multi Domain Permitted Groups** is disabled, users must be in the same domain or sub-domain that the user profile is defined, or the group check fails.

The **Multi Domain Permitted Groups** option is disabled by default.

The **Include Nested Groups** option works with **Multi Domain Permitted Groups**. For example, if you have group1 in server1 and group2 inside group1 in the same server with different users, those users can be authenticated if both **Include Nested Groups** and **Multi Domain Permitted Groups** are enabled.

Steering Groups

Steering groups can be used to steer client traffic to individual Real Servers in a Virtual Service based on the Active Directory (AD) group membership of users initiating client traffic. An example scenario would be a Virtual Service which has four Real Servers. Two Real Servers could be configured to have a primary association with Active Directory Group 1 and two Real Servers could be configured to have a primary association with AD Group 2. When a user attempts to access the Virtual Service, their group membership will be verified and the information used to steer their request to the appropriate Real Servers. If the Real Servers selected based on group membership are not available, the default behavior is to fall back to the assigned scheduling method for the Virtual Service.

For further information, refer to the ESP Steering Groups Technical Note.

Note: Steering groups are not available if using **Basic Authentication**, **Client Certificate**, or **SAML** authentication.

Note: Do not enter the same group name in both the **Permitted Groups** and **Steering Groups** fields. This causes a conflict. When you specify a steering group, it is assumed to behave like a permitted group, so you do not need to enter the same group in both the **Permitted Groups** and **Steering Groups** fields.

SSO Image Set

This option is only available if **Form Based** is selected as the **Client Authentication Mode**. You can choose

which form to use to gather the Username and Password. There are three form options, **Exchange**, **Blank** and **Dual Factor Authentication**. There are also options to display the form and error messages in other languages.

- Exchange Form

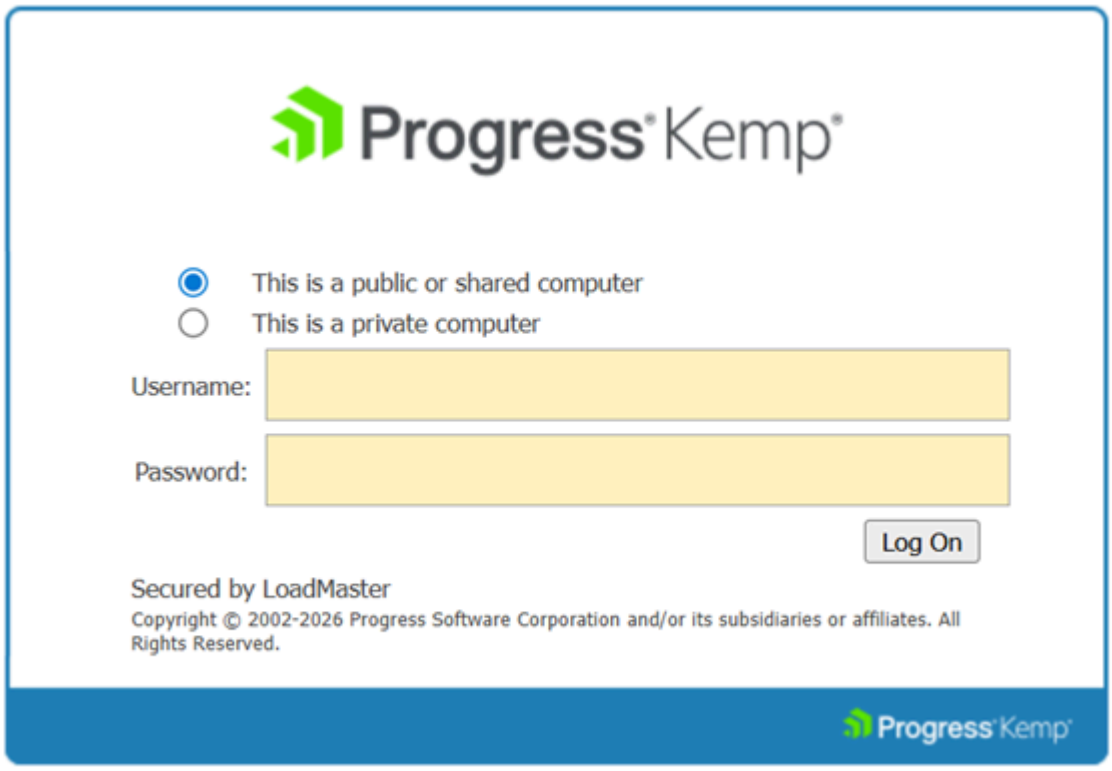

The **Exchange Form** contains the Progress Kemp Logo

- Blank Form

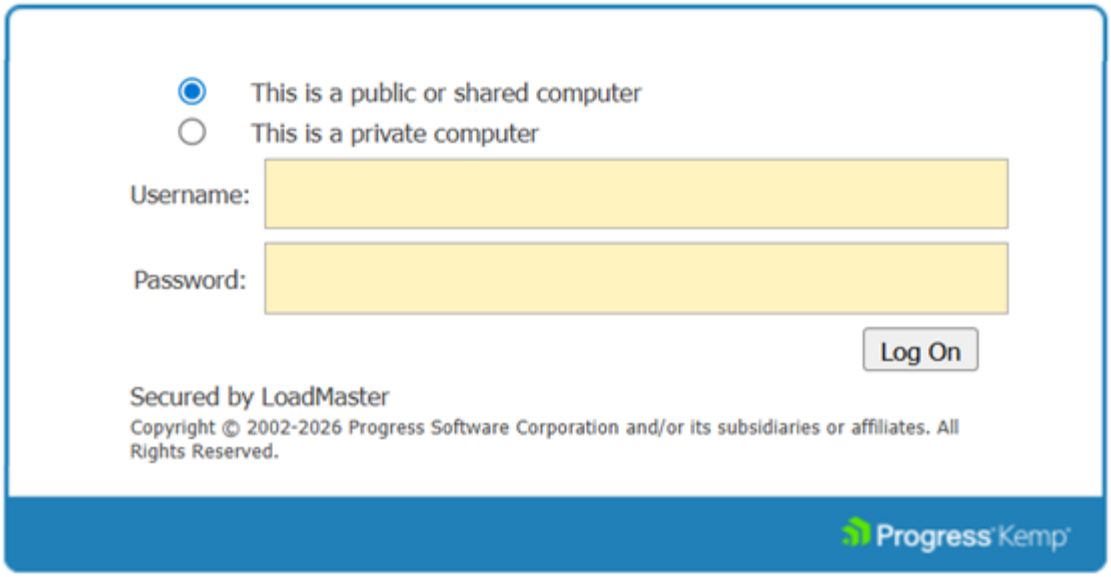

The **Blank Form** does not contain the large Progress Kemp logo.

- Dual Factor Authentication

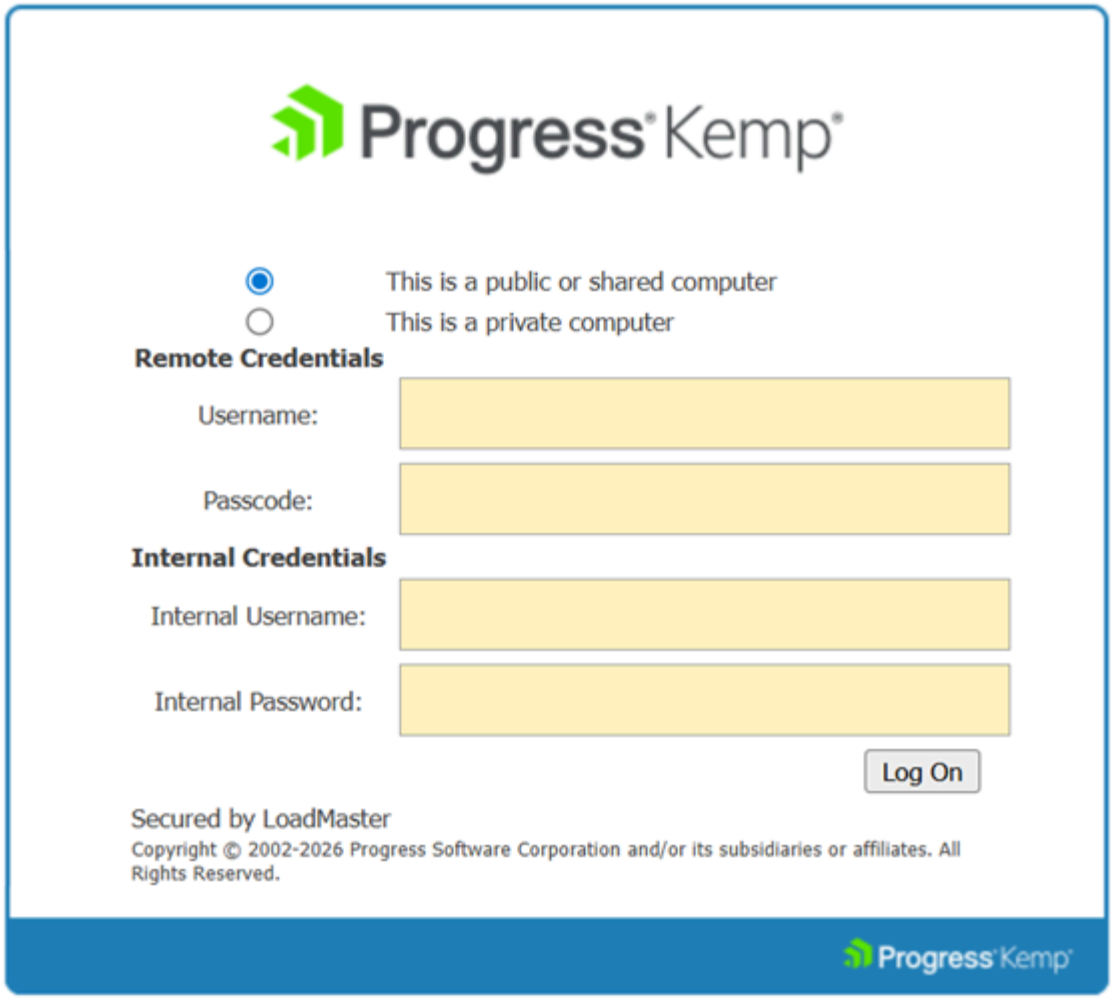

The **Dual Factor Authentication** form contains four fields - two for the remote credentials and two for the internal credentials.

Remote Credentials are credentials that are used to authenticate against remote authentication servers such as RADIUS, before allowing the user to authenticate against Domain Servers such as Active Directory servers.

Internal Credentials are credentials that are used to authenticate against the internal domain servers such as Active Directory Servers.

Note: If the **Authentication Protocol** of the relevant **SSO Domain** is set to **RADIUS and LDAP**, the **SSO Image Set** must be set to **Dual Factor Authentication**.

SSO Greeting Message

This option is only available if **Form Based** is selected as the **Client Authentication Mode.** The login forms can be further customized by adding text. Enter the text that you would like to appear on the form within the **SSO Greeting Message** field and click **Set SSO Greeting Message**. The message can have up to 255 characters.

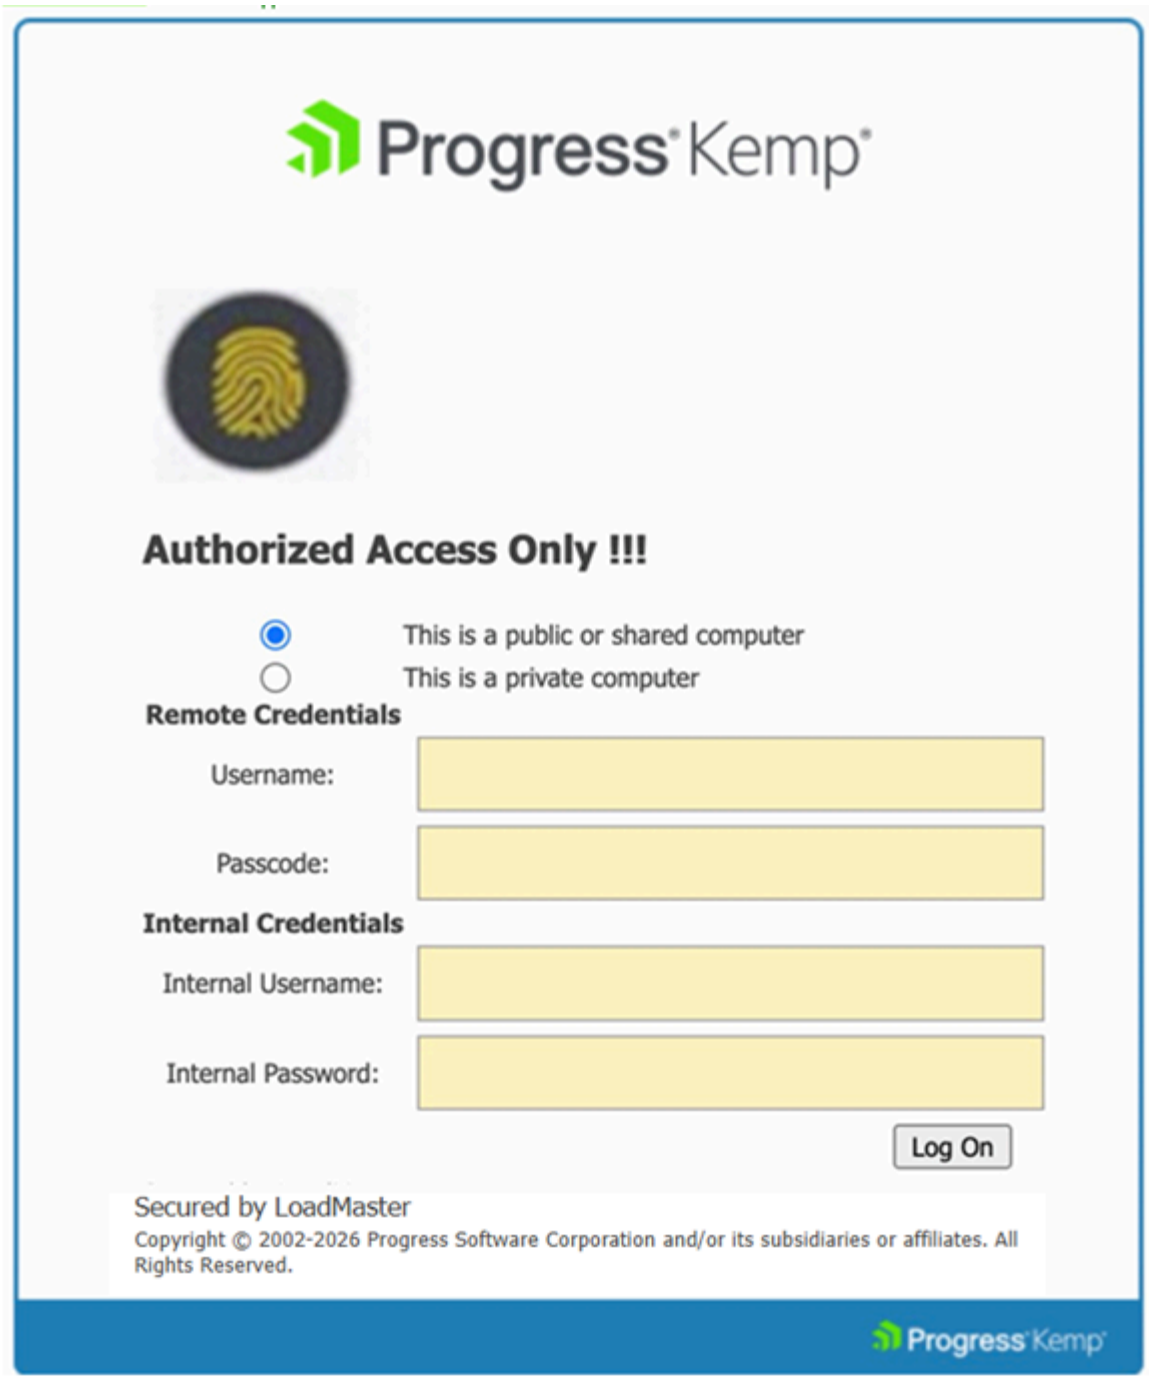

The SSO Greeting Message field accepts HTML code, so you can insert an image if required.

Note: There are several characters that are not supported. These are the grave accent character (` ) and the single quote ('). If a grave accent character is used in the SSO Greeting Message, the character will not display in the output, for example a`b`c becomes abc. If a single quote is used, users will not be able to log in.

Logoff String

This option is only available if **Form Based** or **SAML** is selected as the **Client Authentication Mode.** Normally this field should be left blank. For OWA Virtual Services, the **Logoff String** should be set to **/owa/logoff.owa** or in customized environments, the modified **Logoff String** needs to be specified in this text box. Multiple logoff strings can be entered by using a space-separated list. You can enter up to 255 characters in this field.

Note: If the URL to be matched contains sub-directories before the specified string, the logoff string will not be matched. Therefore, the LoadMaster will not log the user off.

Additional Authentication Header

This option is only available if **SAML** or **OIDC/OAUTH** is selected as the **Client Authentication Mode**. Specify the name of the HTTP header. This header is added to the HTTP request from the LoadMaster to the Real Server and its value is set to the user ID for the authenticated session. You can enter up to 255 characters in this field.

When validating the token from the Identity Provider (IdP), the LoadMaster checks the claims (attributes) in the following sequence:

- preferred_username
- UPN
- unique_name
- email

For example, if the preferred_username is not available then the UPN is used. If the UPN is also not available, then unique_name or email is used for token validation.

Display Public/Private Option

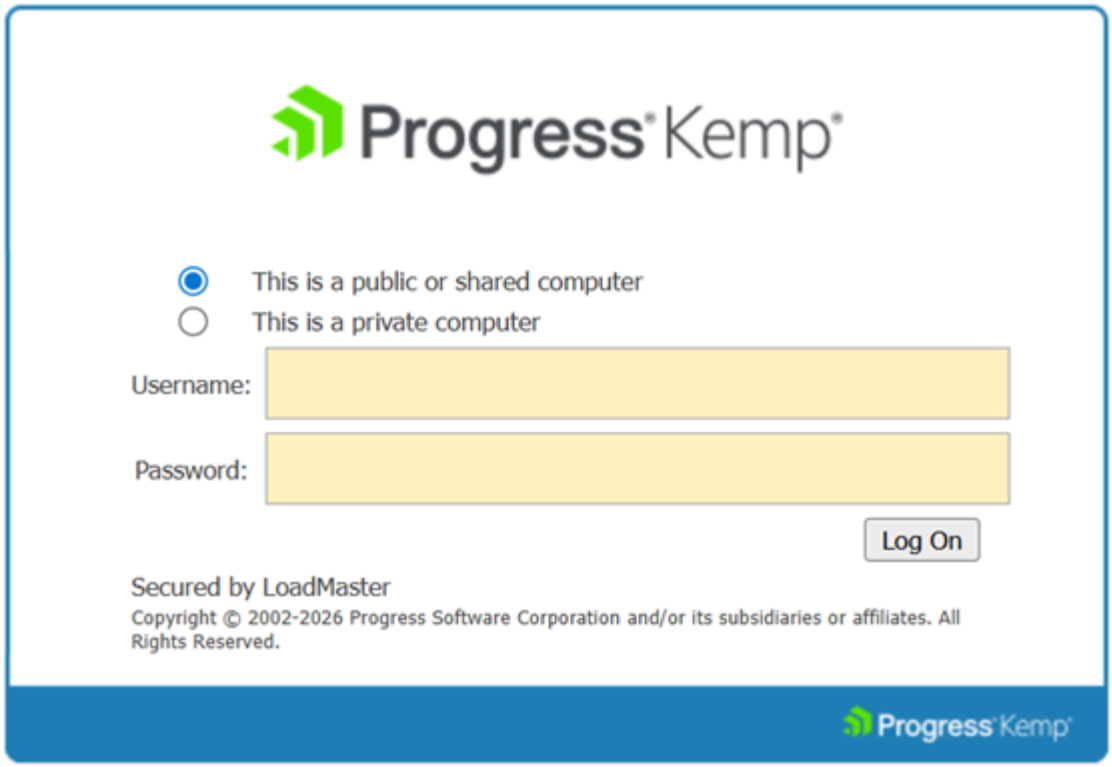

Enabling this check box will display a public/private option on the ESP log in page. Based on the option the user selected on the login form, the **Session timeout** value is set to the value specified for either public or private in the **Manage SSO Domain** screen. If the user selects the private option their username is stored for that session. Refer to the Manage SSO Domains section for more information about these fields.

Disable Password Form

Enabling this option removes the password field from the login page. This may be needed when password validation is not required, for example if using RSA SecurID authentication in a singular fashion. By default, this option is disabled.

Enable Captcha

Select this check box to allow CAPTCHA verification on the login page.

Note: The LoadMaster only supports CAPTCHA v2.

Note: The **Client Authentication Mode** must be set to **Form Based** for the **Enable Captcha** check box to be visible.

CAUTION: All CAPTCHA parameters must be set before it can be used.

CAUTION: Both the LoadMaster and the client machine must be able to access Google for this to work.

Before the CAPTCHA has been correctly answered, the submit button on the login form is disabled. If the user does not submit the form within two minutes of answering the CAPTCHA, the CAPTCHA times out

(Google-specified timeout), and the user must verify a new CAPTCHA (the submit button is disabled until the new CAPTCHA has been verified).

Captcha Public Key

The key that was provided as the public key when you signed up for the CAPTCHA service.

Captcha Private Key

The key that was provided as the private key when you signed up for the CAPTCHA service.

Captcha Access URL

The URL of the service that provides the CAPTCHA challenge. Usually:

www.google.com/recaptcha/api.js

Note: Do not start this URL with https.

Note: Only CAPTCHA V2 is currently supported.

CAPTCHA Verification URL

The URL of the service that verifies the response to the CAPTCHA challenge. Usually:

www.google.com/recaptcha/api/siteverify

Note: Do not start this URL with https.

Note: Only CAPTCHA V2 is currently supported.

Use Session or Permanent Cookies

Three options are available to select for this field:

- **Session Cookies Only:** This is the default and most secure option
- **Permanent Cookies only on Private Computers:** Sends permanent cookies only on private computers
- **Permanent Cookies Always:** Sends permanent cookies in all situations

Note: Permanent cookies only work with Internet Explorer (IE) and IE must be set to accept **Third Party Cookies** and the site must be added to the **Trusted Sites**.

Note: The expiry time of a permanent cookie can be set by configuring the **Session Timeout** fields in the modify SSO screen. The maximum value is 7 days (**604800** seconds).

Specify if the LoadMaster should send session or permanent cookies to the users' browser when logging in.

Note: Permanent cookies should only be used when using single sign on with services that have sessions spanning multiple applications, such as SharePoint.

Cookie SameSite Processing

This option allows the SameSite attribute to be explicitly specified for cookies used by the LoadMaster Edge Security Pack. This influences the way browsers will use cookies across sites and specifically the different behaviour across First and Third party sites. (Cookies that match the domain of the current site, i.e. what's displayed in the browser's address bar, are referred to as first-party cookies. Cookies from domains other than the current site are referred to as third-party cookies.)

The following are the available options to select for this field:

- **SameSite option not added**: This option availability is depend upon the global level configuration i.e. **Default ESP Cookie SameSite Processing** . If the option other than **SameSite Option Not Added** is configured for **Default ESP Cookie SameSite Processing** field (under **System Configuration > Miscellaneous Options > L7 Configuration** page) then only the **SameSite Option Not Added** option will appear in drop-down list for VS.
- **SameSite=None:** Signals that the cookie data can be shared with third parties/external sites (for advertising, embedded content, etc).
- **SameSite=LAX:** Signals that the cookie may be used as a first party cookie but may also be used when accessing the site from an external site via a link clicked by the user.
- **SameSite=Strict:** It is a subset of lax and only enables the cookie be used in first party context and excludes its use when accessing via an incoming link from an external site.
- **System Default**: By default this option is selected and VS will be using the Global level settings.

Note: When a new Virtual Service is created on the LoadMaster and ESP has **Form-based** enabled, then the **Cookie SameSite Processing** option is always configured to **System Default**. When this is the case, theVirtual Service will use the default global settings configured on LoadMaster. When the user changes the configured System Default setting of the Virtual Service to some other option, then the configuration of the selected Virtual Service option overrides the Global SameSite configuration.

User Password Change URL

This is relevant when using client-side forms-based authentication and LDAP. Specify the URL that users can use to change their password, for example https://mail.progressqakcd.net/owa/auth/expiredpassword.aspx?url=/owa/auth.owa

If a user's password has expired, or if they must reset their password, this URL and the **User Password Change Dialog Message** is displayed on the login form.

This URL must be entered in the ESP **Pre-Authorization Excluded Directories** field - this is required to bypass pre-authentication.

If using this expired password functionality in an Exchange 2010 environment:

- The **Pre-Authorization Excluded Directories** must be set to **/owa/auth.owa /owa/auth* /owa/14.3.123.3****. 14.3.123.3 is the sub-path of the Exchange server that must be added to the excluded directories.
- When changing passwords, users cannot use a User Principal Name (UPN) (for example, joebloggs@example.com) in the **Domain\user name** field in the Change Password window, unless Exchange 2010 SP1 RU3 or later is deployed on the Client Access servers.

For further information, refer to the following Microsoft TechNet article: https://technet.microsoft.com/en-us/library/bb684904(v=exchg.141).aspx

User Password Change Dialog Message

This text box is only visible if something is set for the **User Password Change URL** text box. Specify the text to be displayed on the login form when the user must reset their password. Special characters are not permitted in this field.

Note: If the **User Password Change URL** field is set, you must also set the **User Password Change Dialog Message** because both settings are required when a user attempts to change their password. If ESP is enabled on the SubVS level, you must set the password expiry/change configuration on the first SubVS in the SubVS list (for example, if the Exchange template is used this is the **Authentication Proxy** SubVS).

User Password Expiry Warning

By default, SSO users are notified about the number of days before they must change their password. If you disable this option, the password expiry notification will not appear on the login forms.

You can specify the number of days to show the warning before the password is expired. The default value for this field is **15 days**. This field is only visible if the **Client Authentication Mode** is set to **Form Based** and the **User Password Change URL** is set.

The language of the warning text is based on the **SSO Image Set** that is selected (English, French, or Portuguese).

Verify Bearer Header

Enable this check box to verify the authenticity of a JSON Web Token (JWT) included in the Authorizations header of incoming client requests. You can verify authenticity either by checking the validity of the signature on the token (which requires a certificate) or by checking for specific text in the token; or, both.

Note: The **Verify Bearer** Header field (and the two fields detailed below) are only available if the **Client Authentication Mode** is set to **Delegate to Server**.

Bearer Header Validation Certificate

This option is only visible if the **Verify Bearer Header** check box is selected.

Specify the name of the relevant certificate from the **Bearer Header Validation Certificate** drop-down list (this must be first uploaded to the LoadMaster by going to **Certificates & Security > SSL Certificates > Import Certificate**) containing a Public Key used to validate the authenticity of the bearer header token signature. If you are not performing signature validation, the certificate can be set to **None**.

Bearer Header Validation Text

This option is only visible if the **Verify Bearer Header** check box is selected.

You can optionally enter up to 5 comma-separated strings to match against the Audience Claim Field (aud) in the token. If provided, at least one string must match the Audience Claim Field's content or the token is rejected.

Server Authentication Mode

Specifies how the LoadMaster is authenticated by the Real Servers. The following types of methods are available:

- **None:** no client authentication is required
- **Basic Authentication**: standard Basic Authentication is used
- **KCD:** KCD authentication is used
- **Server Token:** On reception and verification of the SAML response, the LoadMaster requests a long-lived token. The LoadMaster then builds a redirection URL with the token specified.

Note: You can only select **Server Token** as the **Server Authentication Mode** if **SAML** is selected as the **Client Authentication Mode**.

- **Form Based:** When **Form Based** authentication is selected, the **Form Authentication Path** field appears.

Note: You can only select **Form Based** as the **Server Authentication Mode** if **Form Based** is selected as the **Client Authentication Mode**.

- When you enter a value in the **Form Authentication Path** field and click the **Set Path** button, the **Form POST Format** and **Post Format Username Only** fields appear. The username and password from the client-side, form-based authentication is injected into the form POST format to build the POST body.
- This feature is predominantly used in Microsoft Exchange deployments and has only been tested with Exchange 2013 and 2016. Therefore, the following strings do not need to be explicitly configured for Exchange 2013/2016. They are used by default in the implementation:
- **Form Authentication Path:** /owa/auth.owa
- **Form POST Format:** destination=%s#authRedirect=true&flags=4&forcedownlevel=0&username=%s&password=%s&passwordText=&isUtf8=1

Note: The **Form POST Format** field only becomes visible when the **Form Authentication Path** is set.

Note: If the deployment is not Exchange, we recommend that the settings are evaluated based on the required interaction with the Real Server and subsequently set appropriately.

POST Format Username Only

Enable this option to send the username only (without the domain part) in the server-side form based authentication POST request.

If **Delegate to Server** is selected as the **Client Authentication Mode**, then **None** is automatically selected as the **Server Authentication mode**. Similarly, if **Basic Authentication** is selected as the **Client Authentication Mode**, then **Basic Authentication** is automatically selected as the **Server Authentication mode**.

When choosing a specific **Client Authentication Mode** protocol, it is important to understand what **Server Authentication Mode** protocols are compatible:

Client Authentication Mode	Default Compatible Server Authentication Mode
Delegate to Server	None
Basic Authentication	Basic Authentication
Form Based	Basic Authentication
	KCD
	Form Based
	None
NTLM	KCD
	None
Client Certificate	KCD
SAML	KCD
SAML	None
SAML	Server Token

Server Side configuration

Note: This option is only visible when the **Server Authentication mode** value is set to **KCD**.

Select the SSO domain for the server side configuration. Only SSO domains which have the **Configuration type** set to **Outbound Configuration** are shown here.

Token Server FQDN

Note: This option is only visible when the **Server Authentication mode** value is set to **Server Token**.

Set the FQDN for the token server. When set, LoadMaster contacts the token server at the given FQDN during sign-on and obtains a permanent access token from that token server. If this parameter is unset, then LoadMaster obtains the token from the Real Server (as in previous releases).

Contents

- SMTP Virtual Services and ESP

SMTP Virtual Services and ESP

If you create an SMTP Virtual Service (with **25** as the port), the ESP feature is available when you select the **Enable ESP** checkbox but with a reduced set of options.



Enable ESP

Enable or disable the ESP feature set by selecting or deselecting the **Enable ESP** checkbox.

Connection Logging

Logging of connections can be enabled or disabled by selecting or deselecting the **Connection Logging** checkbox.

Permitted Domains

All the permitted domains that are allowed to be received by this Virtual Service must be specified here. For example, if you wish the Virtual Service to receive SMTP traffic from **john@progress.com**, then the **progress.com** domain must be specified in this field.

Sub Virtual Services

From within a Virtual Service you can create one or more 'Sub Virtual Services' (SubVS). A SubVS is linked to, and uses the IP address of, the 'parent' Virtual Service. The SubVSs may have different settings (such as health check methods and content rules) to the parent Virtual Service and to each other. This allows the grouping of related Virtual Services, all using the same IP address. This could be useful for certain configurations such as Exchange or Lync which typically are comprised of a number of Virtual Services.

Note: Users with the Virtual Services permission can add a SubVS.

Note: Users with the Real Server permission cannot add a SubVS.

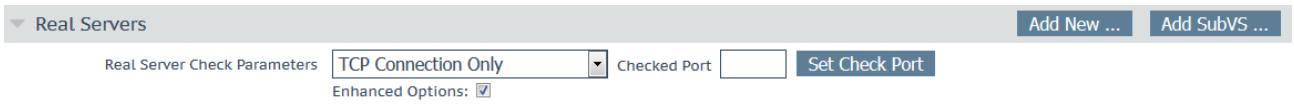

To create a SubVS, within a Virtual Service configuration screen, expand the **Real Servers** section and click the **Add SubVS** button.

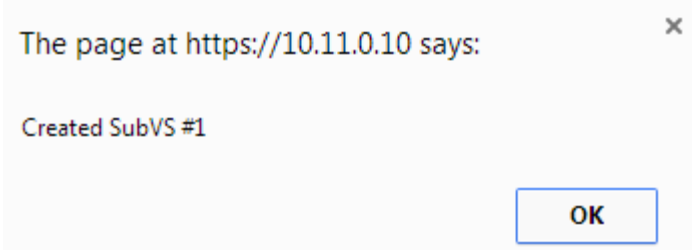

A message appears stating that the SubVS has been created.

Note: You cannot have Real Servers and SubVSs associated with the same Virtual Service. You can however, associate a Real Server with a SubVS.

SubVSs — Add New ...

Id	Name	Weight	Limit	Critical	Status	Operation
1		1	1	☐	Enabled	Disable Modify Delete
2		1000	0	☐	Enabled	Disable Modify Delete

When the SubVS is created, the **Real Servers** section of the Virtual Services configuration screen is replaced with a **SubVSs** section.

All the SubVSs for the Virtual Service are listed here. The **Critical** check box can be enabled to indicate that the SubVS is required for the Virtual Service to be considered available. If a non-critical SubVS is down, the Virtual Service is reported as up and a warning is logged. If a critical SubVS is down, a critical log is generated and the Virtual Service is marked as down. If the email options are configured, an email is sent to the relevant recipients. For further information on the email options, refer to the Email Options section. In all cases, if the Virtual Service is considered to be down and the Virtual Service has a sorry server or an error message configured, these are used.

To modify the SubVS, click the relevant **Modify** button. A configuration screen for the SubVS appears. This contains a subset of the configuration options available for a normal Virtual Service.

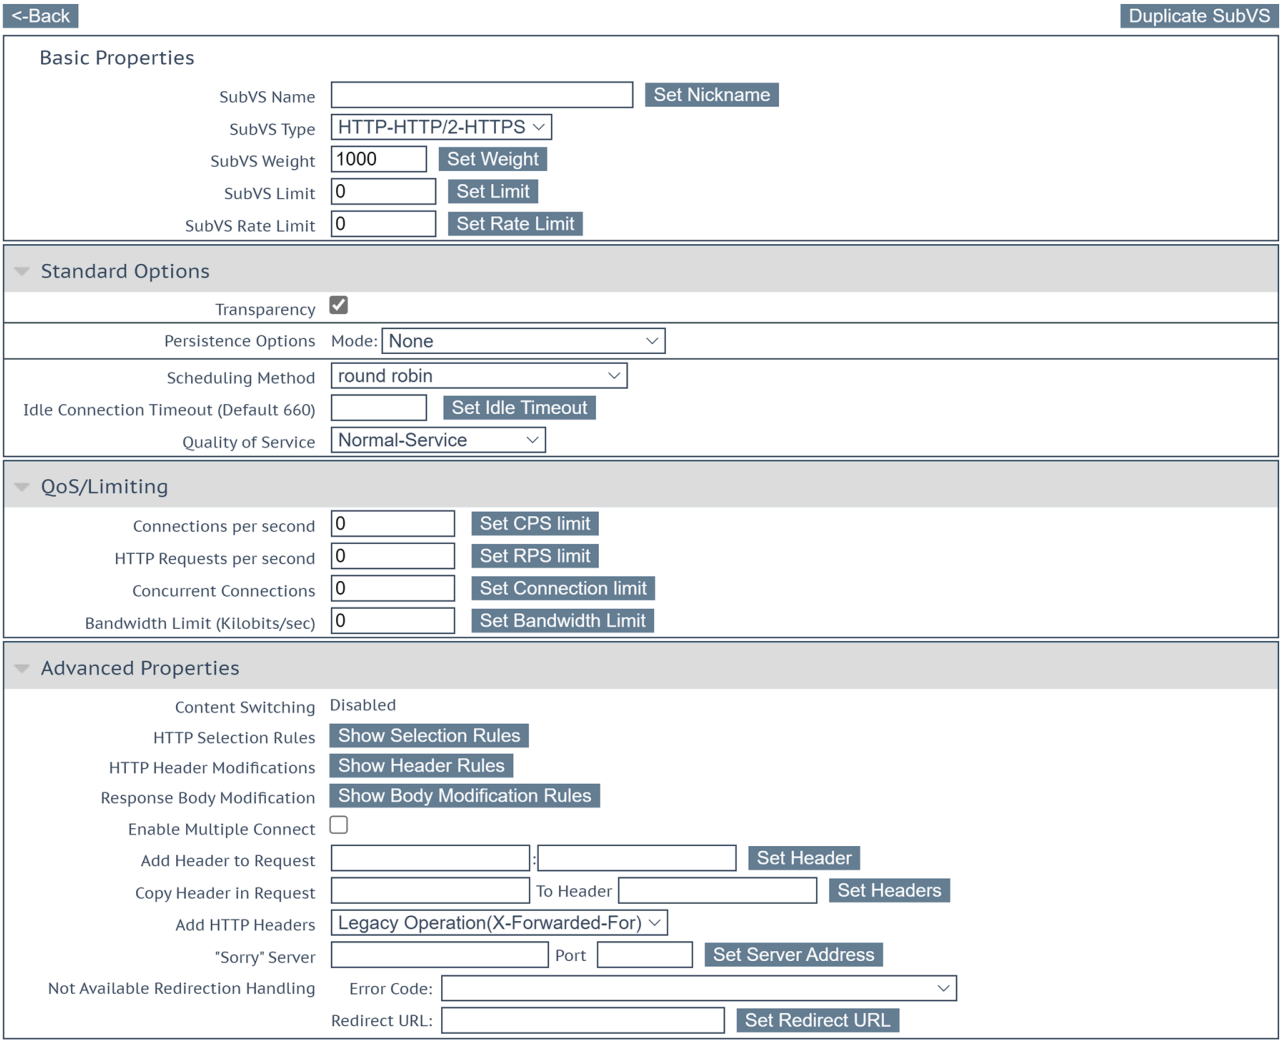

Note: In LoadMaster firmware version 7.2.52 and above, it is possible to set a **Reencryption SNI Hostname** at the SubVS level. If this is set in a SubVS, this overrides the parent Virtual Service value and/or the received SNI value.

The SubVSs can also be modified by clicking the relevant **Modify** button from within the main Virtual Services view. A Virtual Service with SubVSs is colored differently within the Virtual IP address section and the SubVSs are listed in the Real Server section. The SubVS details can be viewed by clicking the 'parent' Virtual Service to expand the view to include the SubVSs.

Duplicate SubVS

You can click **Duplicate SubVS** to create a duplicate SubVS within the same Virtual Service. All SubVS configuration settings are copied to the duplicate SubVS.

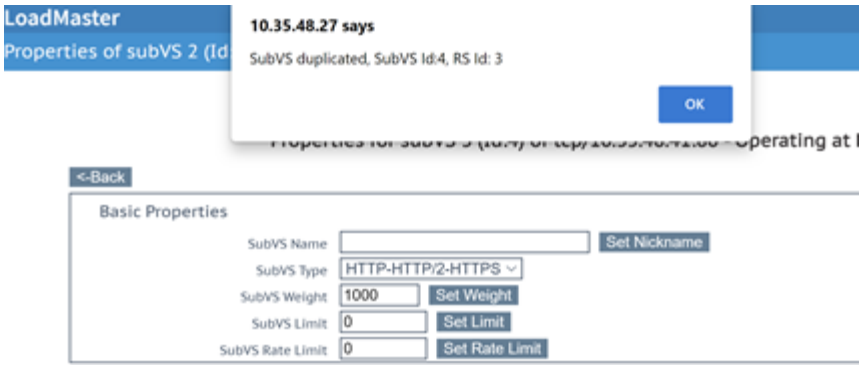

When you click **Duplicate SubVS**, a pop-up message appears like "SubVS duplicated, SubVS Id:4, RS Id: 3". Click **OK** to close the pop-up. A SubVS has both a Virtual Service ID (SubVS Id) and a Real Server ID (RS Id).

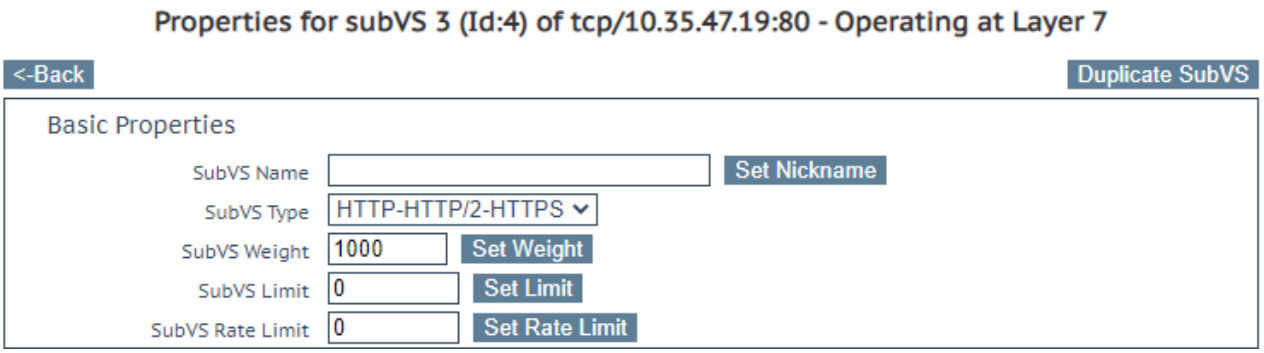

The SubVS Id in the message refers to the Virtual Service ID. You can see this ID in the heading at the top of the SubVS modify screen for the relevant SubVS.

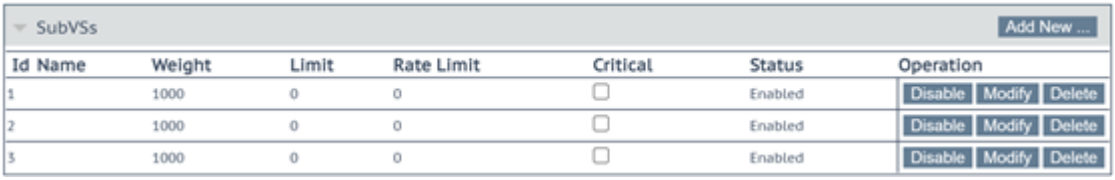

The RS Id refers to the "Real Server" Id. For SubVSs, this Real Server Id is used to identify the SubVS in the SubVSs section of the Virtual Service modify screen.

SubVSs have two additional limit fields in the **Basic Properties** section of the SubVS modify screen:

- **SubVS Limit:** The maximum number of connections that can be sent to this SubVS before it is taken out of rotation from the main Virtual Service. The maximum limit is 1000000.
- **SubVS Rate Limit:** The maximum number of connections per second that can be sent to this SubVS before it is taken out of rotation from the main Virtual Service. The maximum limit is 1000000.

Note: The lowest of the possible limits (global, client, Virtual Service, and SubVS) that is reached first is enforced. Note that the global limit is for all Virtual Services, the Virtual Service limit is for the current Virtual Service with multiple clients, and the client limit can be applicable for all clients, multiple clients, or for a single client - depending on the configuration. Client limits are enforced regardless of what Virtual Service is

---

in use.

Note: If you would like to remove a Virtual Service which contains SubVSs, you must remove the SubVSs first before you are able to delete the main service.

Note: SubVSs may have different ESP configurations than their parent Virtual Service, however care must be taken to ensure that the parent Virtual Service and SubVS ESP options do not conflict.

View/Modify (Remote Terminal Service)

Properties of the Virtual Service include the Generic Type and also provide Remote Terminal specific options.

Persistence

If the terminal servers support a Session Directory, the LoadMaster will use the "routing " supplied by the Session Directory to determine the correct host to connect to. The LoadMaster persistency timeout value is irrelevant here - it is a feature of the Session Directory.

Note: The switch "IP address redirection" in the Session Directory configuration must not be selected for this to work.

Using Session Directory with LoadMaster is optional, in terms of persistence. If the client pre-populates the username and password fields in the initial request, then this value is stored on the LoadMaster. As long as these fields are still populated upon reconnect, the LoadMaster will look up the name and reconnect to the same server as the original connection. The persistence timeout is used to limit the time the information is kept on the LoadMaster.

If using **Terminal-Service or Source IP** mode, then if neither of these two modes succeeds, then the source IP address is used for persistency.

Service Check for the Virtual Service

Only three options are available; **ICMP**, **TCP** and **RDP**. Remote Terminal Protocol (RDP) opens a TCP connection to the Real Server on the Service port (port 3389). The LoadMaster sends an a1110 Code (Connection Request) to the server. If the server sends an a1101 Code (Connection Confirm) then LoadMaster closes the connection and marks the server as active. If the server fails to respond within the configured response time for the configured number of times, or if it responds with a different status code, it is assumed dead.

Real Servers

This section allows you to create a Real Server and lists the Real Servers that are assigned to the Virtual Service. The properties of the Real Servers are summarized and there is also the opportunity to add or delete a Real Server, or modify the properties of a Real Server. When Content Switching is enabled, there is also the opportunity to add rules to, or remove rules from, the Real Server (see Add Rule).

Real Server Check Method

This provides a list of health checks for well-known services, as well as lower level checks for TCP/UDP or ICMP. With the service health checks, the Real Servers are checked for the availability of the selected service. With TCP/UDP the check is simply a connect attempt**.**

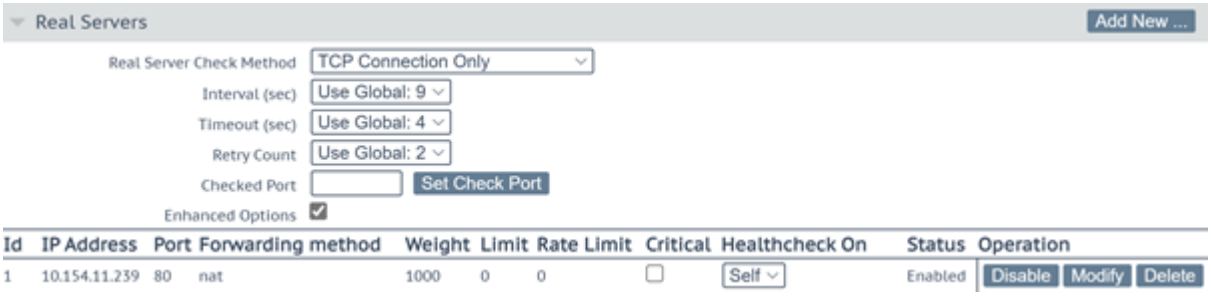

The tables below describe the options that may be used to verify Real Server health. You may also specify a health check port on the Real Server. If none are specified here, it will default to the Real Server port.

When the **HTTP/HTTPS**, **Generic** and **STARTTLS protocols** Service Types are selected, the following health check options are available.

Method	Action
ICMP Ping	An ICMP ping is sent to the Real Server
HTTP	HTTP checking is enabled
HTTPS	HTTPS (SSL) checking is enabled
TCP	A basic TCP connection is checked
Mail	The SMTP (Simple Mail Transfer Protocol) is used
NNTP	The NNTP (Network News Transfer Protocol) is used
FTP	The FTP (File Transfer Protocol) is used
Telnet	The Telnet protocol is used

Method	Action
POP3	The POP3 (Post Office Protocol – mail client protocol) is used
IMAP	The IMAP (Internet Message Access Protocol – mail client protocol) is used
Name Service (DNS) Protocol	The **Name Server (DNS) Protocol** value is only available in the **Real Server Check Method** drop-down list when the Virtual Service **Protocol** is set to **udp**. The LoadMaster performs nslookups against an A record on the server over UDP port 53. If the server successfully responds to the DNS query, the LoadMaster marks it as active. If the server fails to respond within the configured response time for the configured number of times or if it responds unsuccessfully to the A record request, it is assumed down.
Binary Data	Specify a hexadecimal string to send and specify a hexadecimal string to check for in the response
LDAP	Select an LDAP endpoint to use for the health check. The LDAP health check uses the LDAP credentials and protocol specified in the LDAP endpoint. The health check is run against the Real Server IP address and port. The LDAP health check comprises of a LoadMaster connecting to a Real Server and validating the specified user credentials. The health check is performed in two steps: Step 1: Check if the Real Server specified port is up and available. Step 2: Attempt to log in to the Real Server using the LDAP specified credentials. If step 1 and step 2 are true, the health check passes. If step 1 or step 2 fails, the health check fails. For further information on LDAP endpoints, refer to the LDAP Configuration section.

Method	Action
None	No checking performed

When the **Remote Terminal** Service Type is selected the following health check options are available.

Method	Action
ICMP Ping	An ICMP ping is sent to the Real Server
TCP	A basic TCP connection is checked
Remote Terminal Protocol	An RDP Routing Token is passed to the Real Server. This health check supports Network-Level Authentication.
None	No checking performed

Note: For a UDP virtual service, only the **ICMP Ping** and **Name Service (DNS) Protocol** options are available for use.

Check Parameters

In LoadMaster firmware version 7.2.52, the check **Interval**, **Timeout**, and **Retry Count** settings can be configured on each Virtual Service or SubVS. Previously, these were just global settings. You can configure the global settings in **Rules & Checking > Check Parameters**. The global settings are used by default for all Virtual Services.

Interval (seconds): This field specifies the number of seconds that will pass between consecutive health check attempts and defines the length of the 'health check cycle'. To override the global interval, you can select any other value from the drop-down list. The global option cannot be selected if either the **Timeout** or **Retry Count** parameters have been set to a value other than the global setting.

Timeout (seconds): This is the allowed maximum wait time for a connection to be established and for the receipt of a reply to a health check. If no reply is received before the timeout expires, then the connection is re-attempted as long as the Retry Count has not been reached. If any HTTP response is received in response to the health check, then no retries are attempted in the current health check cycle. To override the global timeout value, you can select any other value from the drop-down list.

Retry Count: This specifies the number of times the health check will be re-attempted if the timeout above is reached before receiving a response. To override the global retry count value, you can select any other value

from the drop-down list. The retry count does not apply if any response is received in response to the health check. If, for example, either a 200 or a 404 response is received, then no retries are attempted in the current cycle.

To configure these settings for a specific Virtual Service, expand the **Real Servers** section of the Virtual Service or SubVS modify screen. A **Real Server Check Method** must be selected to see the relevant fields. You can configure these settings to either use the global value, set a specific value within the provided range, or reset to the default value.

If you configure these settings for a parent Virtual Service and then create a SubVS within that Virtual Service, the check values are reset to use the global values.

Enhanced Options

Enabling the **Enhanced Options** check box provides an additional health check option – **Minimum number of RS required for VS to be considered up**. If the **Enhanced Options** check box is disabled (the default), the Virtual Service is considered available if at least one Real Server is available. If the **Enhanced Options** check box is enabled, you can specify the minimum number of Real Servers that must be available to consider the Virtual Service to be available.

Minimum number of RS required for VS to be considered up

Note: This option will only appear if the **Enhanced Options** check box is enabled and if there is more than one Real Server.

Select the minimum number of Real Servers required to be available for the Virtual Service to be considered up.

If less than the minimum number of Real Servers is available, a critical log is generated. If some Real Servers are down but it has not reached the minimum amount specified, a warning is logged. If the email options are configured, an email is sent to the relevant recipients. For further information on the email options, refer to the Email Options section.

Note: Note that the system marks a Virtual Service as down whenever a Real Server that is marked as **Critical** becomes unavailable – even if **Enhanced Options** are enabled and there are more than the specified minimum number of Real Servers still available.

In all cases, if the Virtual Service is considered to be down and the Virtual Service has a sorry server or an error message configured, these are used.

If the minimum number is set to the total number of Real Servers and one of the Real Servers is deleted, the minimum will automatically reduce by one.

When using content rules in a SubVS, the minimum number of Real Servers required has a slightly different meaning. A rule is said to be available and can be matched if and only if the number of available Real Servers with that rule assigned to them is greater than the limit. If the number of available Real Servers is below this limit, the rule can never be matched - the SubVS is marked as down and this is logged appropriately.

If a Real Server on a SubVS is marked as critical – the SubVS is marked as down if that Real Server is down. However, the parent Virtual Service will not be marked down unless that SubVS is marked as critical.

Contents

- HTTP or HTTPS Protocol Health Checking
- Binary Data Health Checking
- Name Server (DNS) Protocol Health Checking
- Add a Real Server
- Modify a Real Server

HTTP or HTTPS Protocol Health Checking

When either the **HTTP Protocol** or **HTTPS Protocol** options are selected a number of extra options are available as described below.

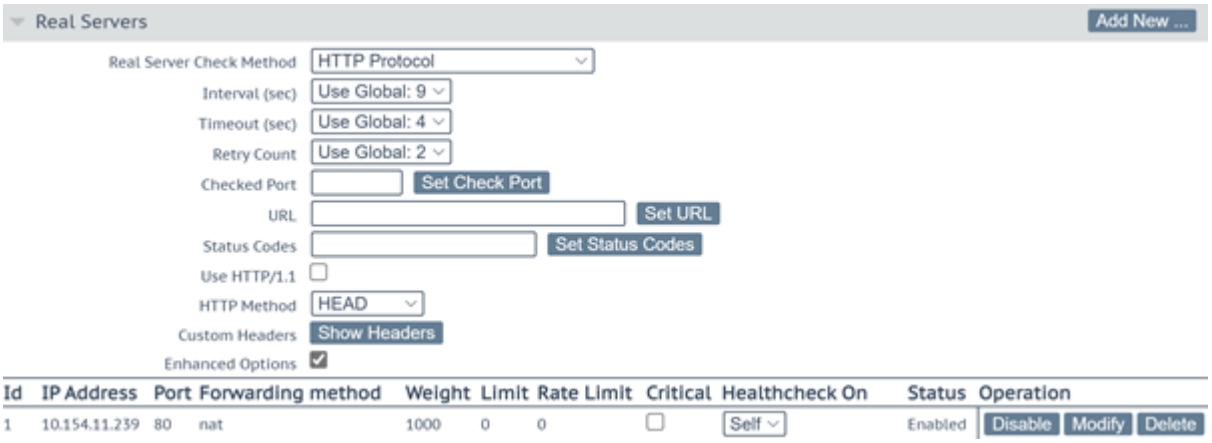

Note: The **post data** option only appears if the **POST HTTP Method** is selected.

Note: The **Reply 200 Pattern** option only appears if either the **POST** or **GET HTTP Method** is selected

URL

By default, the health checker tries to access the URL to determine if the machine is available. A different URL can be specified here.

Status Codes

Health check status codes can be set to override default functionality. Without any **Status Codes** set, the following HTTP status codes are considered to be Up:

- 200-299
- 301
- 302
- 401

Additionally, 2xx status codes are subject to pattern matching the response data, if this is configured. Other codes are considered up without pattern matching, even if it is set.

If custom health check codes are set:

- Check codes may be set to a list of numbers, each from 300-599
- Check codes may be up to 127 characters long, which means 32 valid codes
- Any code in the list is considered to have a health check status of Up
- Configured codes override the default set
  - 2xx codes are always considered up in all cases and are subject to pattern matching, if configured
  - Check codes may be official HTTP status codes, unofficial codes or custom-defined user codes – as long as they fall in the range of 300-599
    - For a list of official HTTP status codes, refer to: https://en.wikipedia.org/wiki/List_of_HTTP_status_codes
    - For a list of unofficial codes, refer to: https://en.wikipedia.org/wiki/List_of_HTTP_status_codes#Unofficial_codes
- Microsoft sub-codes using decimals can be supported, but only by the top-level status code
  - For a list of Microsoft sub-codes using decimals, refer to: https://support.microsoft.com/en-us/kb/943891
  - Sub-codes may not be configured in the Status Codes field – please use the three digit code
  - Sub-codes are grouped by the top-level code

Use HTTP/1.1

By default, the LoadMaster uses HTTP/1.0. However you may opt to use **HTTP/1.1** which will operate more efficiently. When using HTTP/1.1, the health checks are multiplexed to a single connection. This means that more health checks are sent to the server in a single connection which is more efficient from a connection point of view, that is, there is only one connection rather than multiple connections.

Note: Optimization only works on HTTP (not HTTPS) connections.

HTTP/1.1 Host

Note: This field will only be visible if '**Use HTTP/1.1'** is selected.

When using **HTTP/1.1** checking, the Real Servers require a hostname to be supplied in each request. If no value is set, then this value is the IP address of the Virtual Service.

To send SNI host information in HTTPS health checks, please enable **Use HTTP/1.1** in the **Real Servers** section of the relevant Virtual Service(s) and specify a host header. If this is not set, the IP address of the Real Server is used.

HTTP Method

When accessing the health check URL, the system can use either the **HEAD**, **GET** or **POST** method.

In LoadMaster firmware version 7.2.52, support was added for the **OPTIONS** method for HTTP and HTTPS health check types. This specifies that the server will be marked up when the LoadMaster receives a 200 OK in response to an HTTP (or HTTPS) OPTIONS request sent by the LoadMaster.

The OPTIONS HTTP method requests a description of the permitted communication options from the server. A 200 OK response from the server contains a response body which can be optionally searched for specific text in order to provide an additional check. To search the response body, specify the search text in the **Reply 200 Pattern** text box that appears when you select the **OPTIONS** HTTP method. The server will be

marked up if the provided text is found in the response body; otherwise, the server is marked down.

Post Data

This field will only be available if the **HTTP Method** is set to **POST**. When using the **POST** method, up to 2047 characters of POST data can be passed to the server.

Reply 200 Pattern

When using the **GET** or the **POST** method, the contents of the returned response message can be checked. If the response contains the string specified by this Regular Expression, then the machine is determined to be up. The response will have all HTML formatting information removed before the match is performed. Only the first 4K of response data can be matched.

If the response message contains multiple lines, each line is split and is considered as a separate entity. The lines are tested one-by-one (in sequence) against the pattern until one succeeds. For example, if the response is

```
Example

Pattern
```

If the **Reply 200 Pattern** is “/Example/”, it will attempt to match it first against “Example” and succeed.

If the **Reply 200 Pattern** is “/Pattern/”, it will attempt to match it first against “Example” and fail, then move on and try to match it against “Pattern” and succeed, resulting in an UP health check.

The LoadMaster will only check for this phrase if the reply from the server is a 200 code. If the reply is something else, the page is marked as down without checking for the phrase. However, if the reply is a redirect (code 302), the page is not marked as down. This is because the LoadMaster assumes that the phrase will not be present and also it cannot take the service down, as the redirect would then become useless.

An exclamation mark (**!**) at the start of a search pattern negates the outcome of the search. For example, a pattern of 'fail' returns true if the string 'fail' appears in the response; while a pattern of '!fail' returns false if 'fail' appears in the response.

Both Regular Expressions and Perl Compatible Regular Expressions (PCRE) can be used to specify strings. For further information on Regular Expressions and PCRE, please refer to the Content Rules Feature Description.

Note: If you use quotes in regular expressions in the LoadMaster WUI, there are limitations. For more information, refer to the section Limitations of Using Regular Expressions in the LoadMaster WUI.

Custom Headers

Here you can specify up to 4 additional headers/fields which are sent with each health check request. Clicking the **Show Headers** button will show the entry fields. The first field is where you define the key for the custom header that is to be part of the health check request. The second field is the value of the custom header that is to be sent as part of the health check request. Once the information is input, click the **Set Header** button. Each of the headers can be up to a maximum of 20 characters long and the fields can be up to a maximum of 100 characters long. However, the maximum allowed number of characters in total for the 4 header/fields is 256.

The following special characters are allowed in the **Custom Headers** fields:

; . ( ) / + = - _

If a user has specified **HTTP/1.1**, the Host field is sent as before to the Real Server. This can be overridden by specifying a Host entry in the additional headers section. The User-Agent can also be overridden in the same manner. If a Real Server is using adaptive scheduling, the additional headers which are specified in the health check are also sent when getting the adaptive information.

It is possible to perform a health check using an authenticated user: enable **Use HTTP/1.1**, select **HEAD** as the **HTTP Method** and enter the Authorization header with the correctly constructed value. The Authorization field is constructed as follows:

1. The username and password are combined into a string “username:password”.
2. The resulting string is then encoded using the RFC2045-MIME variant of Base64, except not limited to 76 char/line.
3. The authorization method and a space, for example, “Basic “ is then put before the encoded string.

For example, if the user agent uses 'Aladdin' as the username and 'open sesame' as the password then the field is formed as follows:

Authorization: Basic QWxhZGRpbjpvcGVuIHNlc2FtZQ==

To send SNI host information in HTTPS health checks, please enable **Use HTTP/1.1** in the **Real Servers** section of the relevant Virtual Service(s) and specify a host header. If this is not set, the IP address of the Real Server is used.

Rules

If any of the Real Servers have Content Switching rules assigned to them the **Rules** column appears in the Real Servers section. A button with the number of rules assigned to each of the Real Server (or with **None** if there are no rules assigned) is displayed in the **Rules** column.

Clicking the button within the **Rules** column opens the **Rules Management** screen.



From within this screen you can **Add** or **Delete** the rules assigned to a Real Server.

Binary Data Health Checking

When **Binary Data** is selected as the health check method, some other fields are available, as described below.

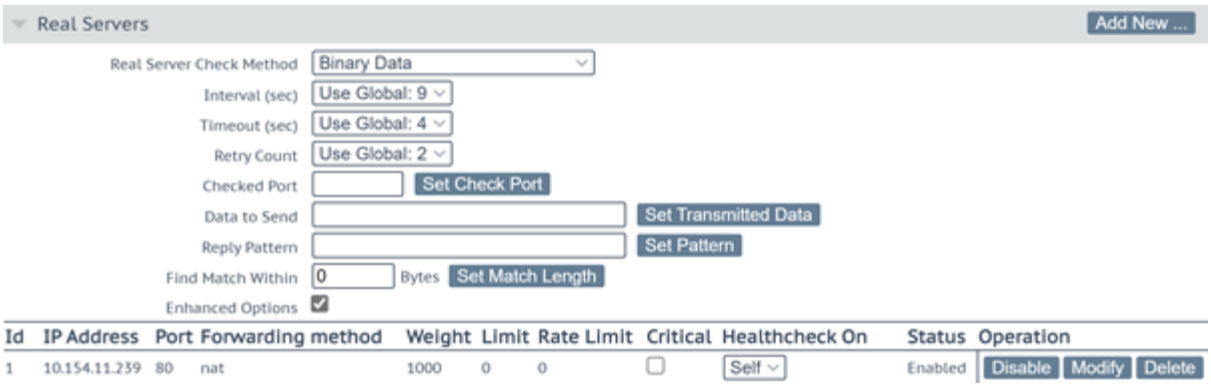

Data to Send

Specify a hexadecimal string to send to the Real Server.

Note: This hexadecimal string must contain an even number of characters.

Reply Pattern

Specify the hexadecimal string which is searched for in the response sent back from the Real Server. If the LoadMaster finds this pattern in the response, the Real Server is considered up. If the string is not found, the Real Server is marked as down.

Note: This hexadecimal string must contain an even number of characters.

Find Match Within

When a response is returned, the LoadMaster will search for the **Reply Pattern** in the response. The LoadMaster will search up to the number of bytes specified in this field for a match.

Setting this to **0** means that the search is not limited. Data is read from the Real Server until a match is found. A maximum of 8 KB is read from the Real Server.

Setting the value to less than the length of the reply string means that the check will act as if the value has been set to **0**, that is, all packets (up to 8 KB) are searched.

Name Server (DNS) Protocol Health Checking

Name Server (DNS) Protocol health checking is only available when using a UDP Virtual Service.

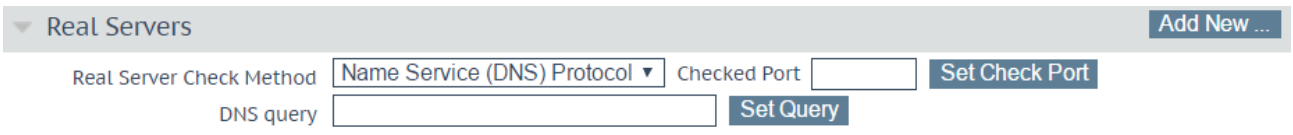

Checked Port

The port to be checked. If there is no port specified, the Real Server port is used.

DNS query

Specify the query string to be requested from the name server. This field has a maximum length of 126 characters.

Add a Real Server

Clicking the **Add New** button brings you to the following screen where the properties of the Real Server are set.

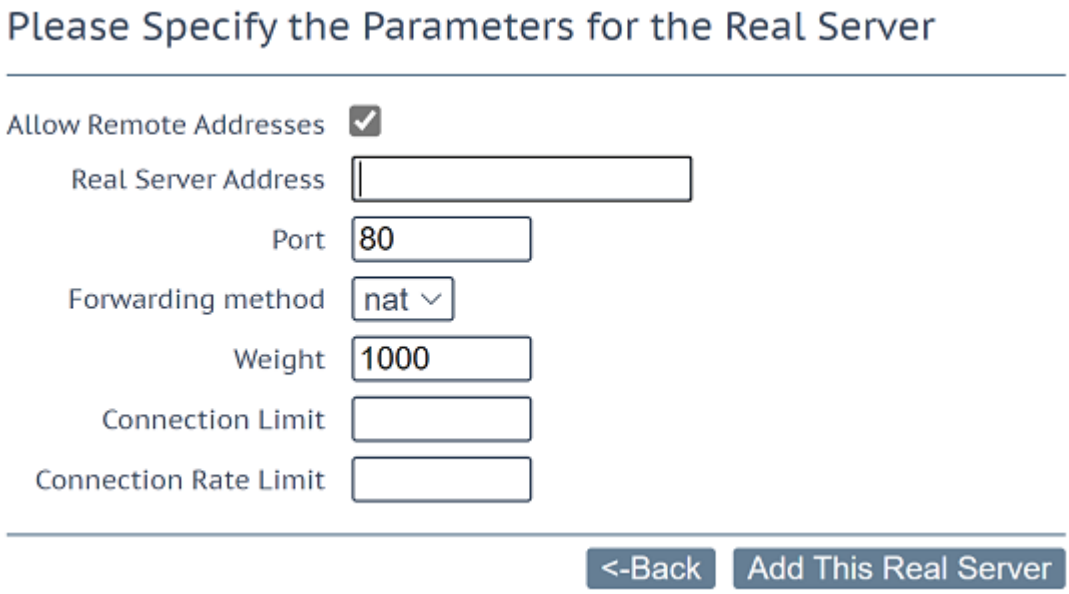

Allow Remote Addresses: By default only Real Servers on local networks can be assigned to a Virtual Service. Enabling this option will allow a non-local Real Server to be assigned to the Virtual Service. This option is enabled by default.

Note: To make the **Allow Remote Addresses** option visible, **Enable Non-Local Real Servers** must be selected (in **System Configuration > Miscellaneous Options > Network Options**). Also, **Transparency** must be disabled in the Virtual Service.

Note: When alternative gateways/non-local Real Servers are set up, health checks are routed through the default gateway.

Real Server Address: The Real Server address. This can either be an IP address or a Fully Qualified Domain Name (FQDN). This is not editable when modifying a Real Server. An FQDN can only be used if a **Nameserver** is configured. The resolved name is listed next to the IP address in parenthesis. For further information, refer to the Host & DNS Configuration section. If an FQDN is used when adding a Real Server – the name is resolved at the time of adding. If it fails to resolve, the Real Server is not created and an error is generated.

You can either type the address of a new Real Server, or select an existing Real Server from the drop-down list provided. The entries before the line in the drop-down list are existing Real Servers. The entries below the line are auto-complete form options. Real Servers already added to this SubVS are not listed in the drop-down list.

Note: This drop-down list does not appear on Safari browsers due to a browser limitation.

Add to all SubVSs

When adding a Real Server to a SubVS, a check box is available that, when selected, adds the Real Server to all SubVSs in that Virtual Service.

Port: The forwarding port of the Real Server. This field is editable, so the port may be altered later if required.

Forwarding Method: Either NAT (Network Address Translation) or Route (Direct) forwarding. The available options are dependent on the other modes selected for the service.

Weight: The Real Server's weight. This is weight of the Real Server, as used by the Weighted Round Robin, Weighted Least Connection and Adaptive scheduling methods. The default initial value for the weight is **1000**, the maximum is **65535**, and the minimum is **1**. It is a good benchmark to give a Real Server a weight relative to its processor speed, for example, if server1 seems to bring four times the power of server2, assign a weight of **4000** to server1 and weight of **1000** to server2.

Connection Limit: The maximum number of open connections that a Real Server will accept before it is taken out of the rotation. This is only available for Layer 7 traffic. The limit stops new connections from being created, but it will allow requests that already have persistent connections to the server. Persistent connections include connections to a Virtual Service using Session Broker Persistence which include a Session Broker cookie as set by the Connection Broker.

A maximum number of 1024 different Real Servers is allowed. For example, if you want two Virtual Services with the same Real Server pool, you could have 1024 Real Servers assigned to each Virtual Service.

Note: For the LoadMaster Exchange, there is a limit of six Real Servers that may be configured.

Click the **Add This Real Server** button and it gets added to the pool.

Connection Rate Limit

As of LoadMaster firmware version 7.2.51, there is a **Connection Rate Limit** field when configuring a Real Server. This allows you to set the maximum number of open Connections Per Second (CPS) that can be sent to a Real Server before it is taken out of rotation. The maximum limit is **100000** (100,000 CPS). If the **Connection Rate Limit** is set to 0 (the default) the feature is disabled, that is, there is no rate limiting.

If rate limiting is enabled, the LoadMaster limits the number of new connections going to a particular Real Server. If the limit is reached, the Real Server is taken out of rotation/selection (that is, removed from the load balancing scheduling algorithm) until the current period is finished. The 'rate limit period' is 0.1 seconds. Any value less than 10 CPS is handled like 10 CPS.

When a Real Server is taken out of rotation, the current connection will then be scheduled to go to a different Real Server in the Virtual Service (or SubVS). This includes new connections with persistence settings to the rate-limited Real Server - these are also sent to another Real Server when the rate limit is exceeded. If no Real Servers are found, the usual rejection methods are used.

No new connections are sent to the rate-limited Real Server until the current 'rate limit period' expires and the Real Server is returned to the load balancing scheduling process.

The rate value for the Real Server also takes slow start into consideration. At Real Server restart, the CPS

limit will get to its full value at the end of the slow start period. For further details on the slow start feature, refer to the heading **Least Connection Slow Start** under section L7 Configuration

Critical

Note: This option will only appear if the **Enhanced Options** check box is enabled. For further information on the **Enhanced Options** check box, refer to the Real Servers section.

In the Real Servers section of the Virtual Service modify screen, there is a **Critical** check box for each of the Real Servers. Enabling this option indicates that the Real Server is required for the Virtual Service to be considered available. The Virtual Service is marked as down if the Real Server has failed or is disabled.

If a Real Server on a SubVS is marked as critical – the SubVS is marked as down if that Real Server is down. However, the parent Virtual Service will not be marked down unless that SubVS is marked as critical.

Note: This option overrides the **Minimum number of RS required for VS to be considered up** field. For example, if the minimum is set to two and only one Real Server is down but that Real Server is set to critical – the Virtual Service is marked as down.

In all cases, if the Virtual Service is considered to be down and the Virtual Service has a sorry server or an error message configured, these are used.

Healthcheck On

Note: This option will only appear if the **Enhanced Options** check box is enabled. For further information on the **Enhanced Options** check box, refer to the Real Servers section.

In the Real Servers section of the Virtual Service modify screen, there is a **Healthcheck On** drop-down list for each of the Real Servers. This allows you to specify what Real Server the health check is based on. This can either be set to **Self** to perform the health check based on that particular Real Server status, or another Real Server can be selected. For example – if Real Server 1 is down, any Real Servers which have their health check based on Real Server 1 will also be marked as down, regardless of their actual Real Server status.

Some points to be aware of are listed below:

- A Real Server can only follow a Real Server and not a SubVS.
- A Real Server can follow a Real Server that is also following a third Real Server. The status of the first two Real Servers will reflect the status of the third Real Server.
- Chains of Real Servers are allowed – but loops are not allowed and cannot be created.
- If a Real Server is deleted (either singly or as part of a Virtual Service), all Real Servers that are following the Real Server are reset to normal behaviour (that is, they will start using the Virtual Service health check options).
- If all Real Servers in a Virtual Service are following Real Servers on a different Virtual Service, the health check parameters for the Virtual Service are not shown on the WUI (because the settings do not affect any Real Servers).
- Disabling the **Enhanced Options** check box will disable all Real Server health check following for that Virtual Service.

Modify a Real Server

When you click the **Modify** button of a Real Server, the following options are available:

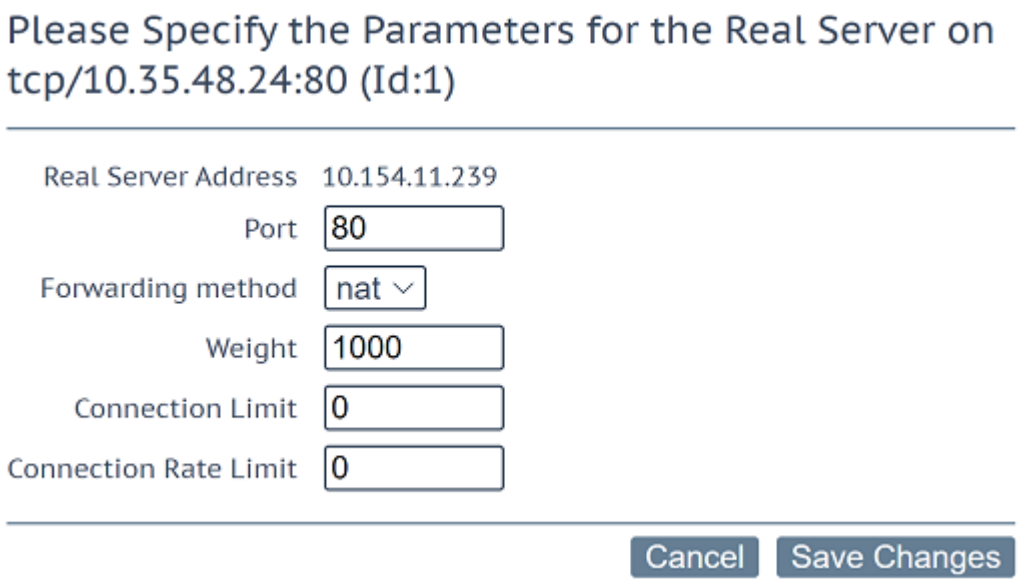

Real Server Address

This field shows the address of the Real Server. This is not an editable field.

Port

This is a field detailing the port on the Real Server that is to be used.

Forwarding Method

This is a field detailing the type of forwarding method to be used. The default is NAT; Direct Server Return can only be used with L4 services.

Weight

When using Weighted Round Robin Scheduling, the weight of a Real Server is used to indicate what relative proportion of traffic should be sent to the server. Servers with higher values will receive more traffic.

Connection Limit

This is the maximum number of open connections that can be sent to the real server before it is taken out of rotation. The maximum limit is 100,000.

Connection Rate Limit

This is the maximum number of open Connections Per Second (CPS) that can be sent to a Real Server before it is taken out of rotation. The maximum limit is 100000 (100,000 CPS).

Manage Templates

Templates make the setting up of Virtual Services easier by automatically creating and configuring the parameters for a Virtual Service. Before a template can be used to configure a Virtual Service, it must be imported and installed on the LoadMaster.

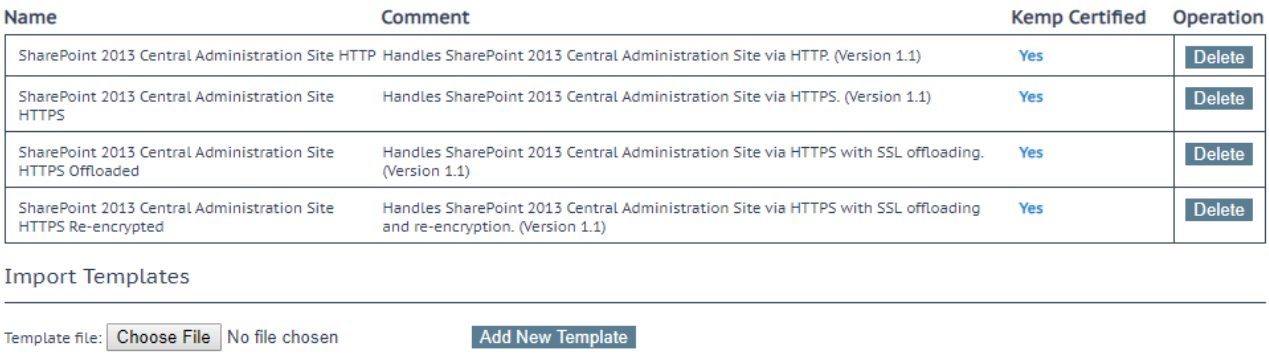

Click the **Choose File** button, select the template you wish to install and click the **Add New Template** button to install the selected template. This template is now available for use when you are adding a new Virtual Service.

Click **Delete** to remove the template.

The **Kemp Certified** column will indicate whether the template was supplied by Progress Kemp or not. If a template is certified, it has been provided by Progress Kemp. If a template is not certified, it might be a template created by you (by exporting a Virtual Service).

For further details on templates, such as how to use a template to create and configure a new Virtual Service and where to obtain Progress Kemp templates, please refer to the Virtual Services and Templates Feature Description.

Manage SSO Domains

Before using the Edge Security Pack (ESP) the user must first set up a Single Sign-On (SSO) Domain on the LoadMaster. The SSO Domain is a logical grouping of Virtual Services which are authenticated by an LDAP server.

Note: The maximum number of SSO domains that are allowed is 128.

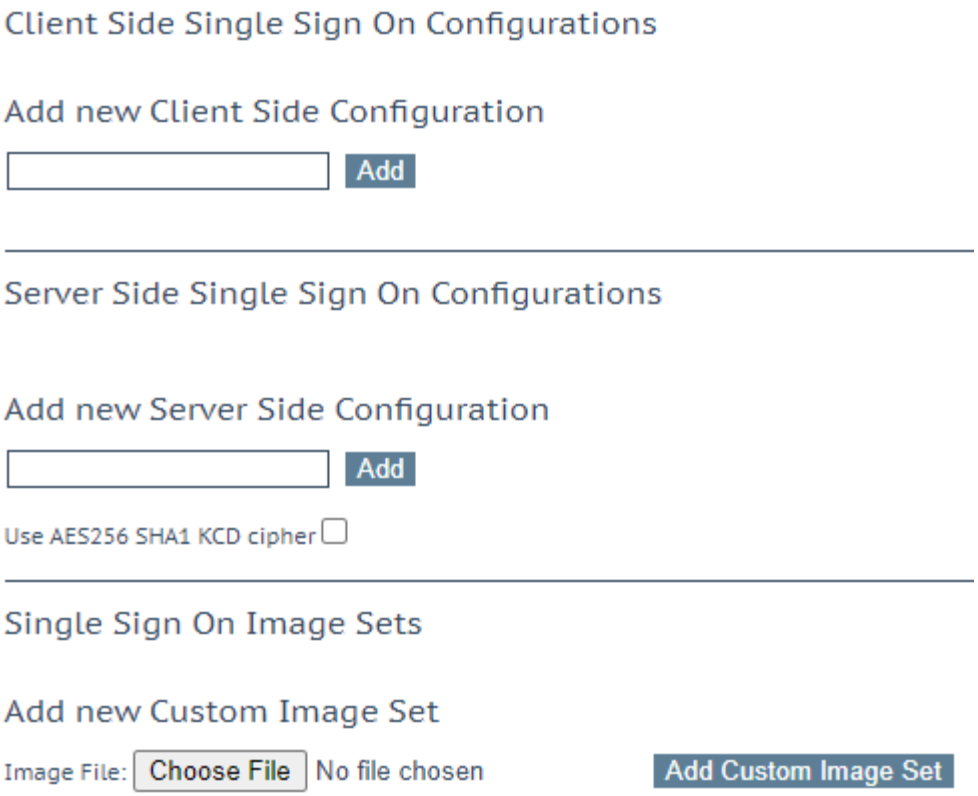

Click the **Manage SSO Domains** menu option to open the **Manage Single Sign On Options** screen.

Contents

- Single Sign On Domains
- Single Sign On Image Sets

Single Sign On Domains

Two types of SSO domains can be created – Client Side and Server Side. Client Side single sign on domains enable the configuration of how the LoadMaster authenticates clients including protocols used and authentication endpoints. A Server Side domain is required if utilizing Kerberos Constrained Delegation for the authentication of connections from the LoadMaster to the servers.

Client Side configurations allow you to set the **Authentication Protocol** to **LDAP**, **RADIUS, RSA-SecurID**, **Certificates**, **RADIUS and LDAP** or **RSA-SecurID and LDAP**.

Note: As of LoadMaster firmware version 7.2.52, RADIUS two-factor and LDAP authentication is supported. Refer to the RADIUS ESP Authentication Feature Description for further details.

Server Side configurations allow you to set the **Authentication Protocol** to **Kerberos Constrained Delegation** (KCD).

To add a new SSO Domain enter the name of the domain in the **Name** field and click **Add**. You can enter up to 64 characters in this field. The name entered here does not need to relate to the allowed hosts within the

Single Sign On Domain.

Note: When using the **Permitted Groups** field in **ESP Options**, you need to ensure that the SSO domain set here is the directory for the permitted groups. For example, if the **SSO Domain** is set to webmail.example and webmail is not the directory for the permitted groups within example.com, it will not work. Instead, the **SSO Domain** needs to be set to .example.com.

Note: If the **Domain/Realm** field is not set, the domain **Name** set when initially adding an SSO domain is used as the **Domain/Realm** name.

Contents

- Client Side (Inbound) SSO Domains
- Client Side (Inbound) SAML SSO Domains
- Client Side (Inbound) OIDC / OAUTH SSO Domains
- RADIUS Two-factor and LDAP Authentication
- Sessions
- Server Side (Outbound) SSO Domains

Client Side (Inbound) SSO Domains

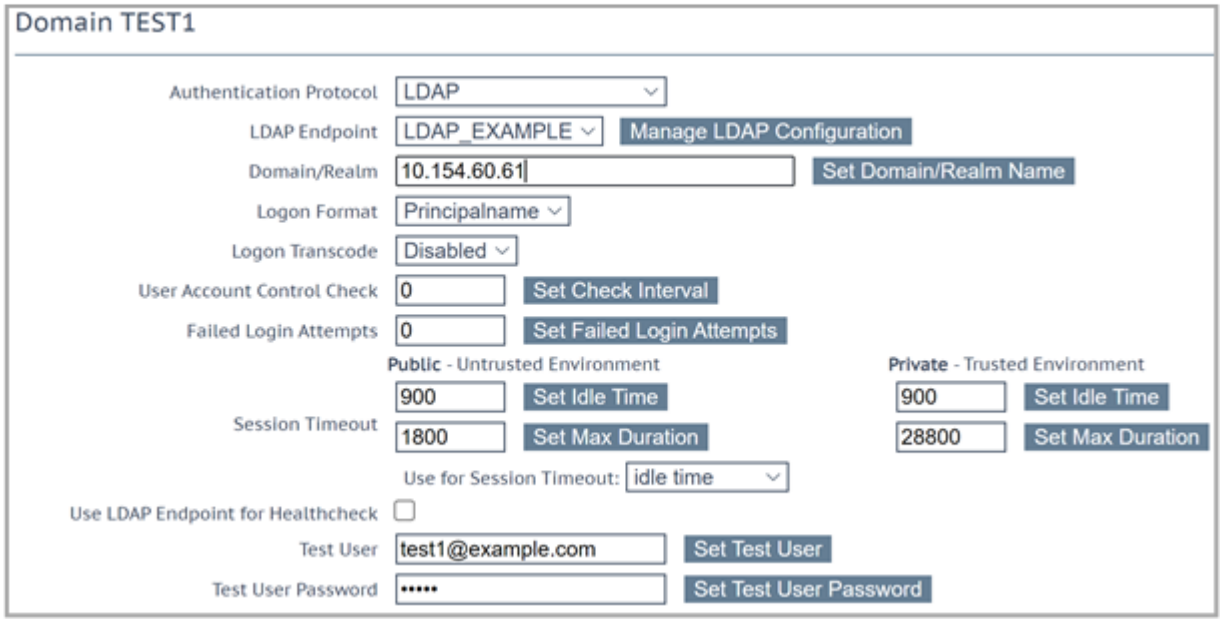

Authentication Protocol

This dropdown allows you to select the transport protocol used to communicate with the authentication server. The options are:

- LDAP
- RADIUS
- RSA-SecurID

- Certificates

Note: If you create an SSO domain with the **Authentication Protocol** set to **Certificates**, ensure to set the **LDAP Protocol** to **LDAPS** in the LDAP endpoint.

- RADIUS and LDAP
- RSA-SecurID and LDAP
- SAML
- OIDC/OAUTH

Note: The fields displayed on this screen will change depending on the **Authentication protocol** selected.

LDAP Endpoint

Select the LDAP endpoint to use. Click the **Manage LDAP Configuration** button to go to the **LDAP Configuration** screen. For further information on LDAP endpoints, refer to the LDAP Configuration section.

Note: This option is only available if the **Authentication Protocol** is set to **LDAP**, **RADIUS and LDAP** or **RSA-SecurID and LDAP**.

RADIUS/ RSA-SecurID Server(s)

Type the IP addresses of the server or servers which are used to authenticate the domain into the server(s) field and click the set server(s) button.

Multiple server addresses can be entered within this text box. Each entry must be separated by a space.

Note: IPv6 is not supported for RADIUS authentication.

RADIUS Shared Secret

The shared secret to be used between the RADIUS server and the LoadMaster (48 character limit).

Note: This field is only available if the **Authentication Protocol** is set to **RADIUS** or **RADIUS and LDAP**.

Send NAS Identifier

If this check box is disabled (default), a Network Access Server (NAS) identifier is not sent to the RADIUS server. If it is enabled, a NAS identifier string is sent to the RADIUS server. By default, this is the hostname. Alternatively, if a value is specified in the **RADIUS NAS Identifier** text box, this value is used as the NAS identifier. If the NAS identifier cannot be added, the RADIUS access request is still processed.

Note: This field is only available if the **Authentication Protocol** is set to **RADIUS** or **RADIUS and LDAP**.

Sending the NAS identifier serves two purposes:

- It helps to classify the device type that is sending the request as opposed to simply sending the host IP

address, which makes troubleshooting and consuming logs easier.

- It enables customized authentication responses to be sent back from the server based on the identifier.

RADIUS NAS Identifier

If the **Send NAS Identifier** check box is selected, the **RADIUS NAS Identifier** field is shown. When specified, this value is used as the NAS identifier. Otherwise, the hostname is used as the NAS identifier. If the NAS identifier cannot be added, the RADIUS access request is still processed.

Note: This field is only available if the **Authentication Protocol** is set to **RADIUS** or **RADIUS and LDAP** and the **Send NAS Identifier** check box is enabled.

Select Certificate to User Mapping

This option is only available when the **Authentication Protocol** is set to **Certificates**. The **Select Certificate to User Mapping** field has the following values:

- User Principal Name (default value)
- Subject
- Issuer and Subject
- Issuer and Serial Number

Note: In LoadMaster firmware version 7.2.53, support for Personal Identity Verification (PIV) smart card authentication was added. Refer to the ESP Feature Description for further details.

If this option is enabled and the check fails, the login attempt will fail. If this option is not enabled, only a valid client certificate (with the username in the SubjectAltName (SAN)) is required to log in, even if the altSecurityIdentities attribute for the user is not present or not matching.

For more information, refer to the **Kerberos Constrained Delegation Feature Description**.

Allow fallback to check Common Name

Enabling this option allows a fallback to check the Common Name (CN) in the certificate when the SAN is not available.

Note: This field only appears when the **Authentication Protocol** is set to **Certificates**.

Domain/Realm

The login domain to be used. This is also used with the logon format to construct the normalized username, for example;

- **Principalname:** <username>@<domain>
- **Username:** <domain>\<username>

Note: If the **Domain/Realm** field is not set, the **Domain** name set when initially adding an SSO domain is

---

used as the **Domain/Realm** name.

RSA Authentication Manager Config File

Note: This option is only available when the **Authentication Protocol** is set to **RSA-SecurID**.

This file needs to be exported from the RSA Authentication Manager.

Note: For more information on the RSA authentication method, including how to configure it, refer to the **RSA Two Factor Authentication Feature Description**.

RSA Node Secret File

Note: This option is only available when the **Authentication Protocol** is set to **RSA-SecurID**.

A node secret must be generated and exported in the RSA Authentication Manager.

Note: It is not possible to upload the RSA node secret file until the RSA Authentication Manager configuration file is uploaded. The node secret file is dependent on the configuration file.

Logon Format

This drop-down list allows you to specify the format of the login information that the client has to enter.

Note: The options available vary depending upon which **Authentication Protocol** is selected.

Not Specified: The username will have no normalization applied to it - it is taken as it is typed.

Principalname: Selecting this as the **Logon format** means that the client does not need to enter the domain when logging in, for example **username@domain**. The SSO domain added in the corresponding text box is used as the domain in this case.

Note: When using **RADIUS** as the **Authentication protocol** the value in this SSO domain field must exactly match for the login to work. It is case sensitive.

Username: Selecting this as the **Logon format** means that the client needs to enter the domain and username, for example **domain\username**.

Username Only: Selecting this as the **Logon Format** means that the text entered is normalized to the username only (the domain is removed).

Note: The **Username Only** option is only available for the **RADIUS** and **RSA-SecurID** protocols.

Logon Format (Phase 2 Real Server)

Specify the logon string format used to authenticate to the Real Server.

The **Logon Format (Phase 2 Real Server)** field only appears if the **Authentication Protocol** is set to one of the following options:

- RADIUS
- RSA-SecurID

Logon Format (Phase 2 LDAP)

Specify the logon string format used to authenticate to LDAP.

The **Logon Format (Phase 2 LDAP)** field only appears if the **Authentication Protocol** is set to one of the following options:

- RADIUS and LDAP
- RSA-SecurID and LDAP

Logon Transcode

Enable or disable the transcode of logon credentials, from ISO-8859-1 to UTF-8, when required.

If this option is disabled, log in using the format that the client dictates. If this option is enabled, check if the client uses UTF-8. If the client does not use UTF-8, use ISO-8859-1.

User Account Control Check

If the UAC check interval value is set to 0 minutes (default value), then UAC is not performed periodically for users after successful login.

When you specify an interval value in the range of 1 to 300 minutes, the periodic UAC check is performed per user for the requests received after the interval expiry.

The UAC detects:

- Unknown users
- Disabled accounts
- Locked accounts
- Expired passwords on accounts

Extended ESP user logs provide the results of the UAC check. Additional information is logged for the user such as start session time, total duration, protocol information, KCD information, and blocked user events.

The check may occur on new connection establishment or as part of existing sessions. The **msDS-User-Account-Control-Computed** and **userAccountControl** attributes are used to determine the UAC status.

Failed Login Attempts

The maximum number of consecutive failed login attempts before the user is locked out. Valid values range from **0** to **99**. Setting this to **0** means that users will never be locked out.

Note: When a user is locked out, all existing logins for that user are terminated, along with future logins.

Reset Failed Login Attempt Counter after

When this time (in seconds) has elapsed after a failed authentication attempt (without any new attempts) the failed login attempts counter is reset to **0**. Valid values for this text box range from **60** to **86400**. This value must be less than the **Unblock timeout** value.

Unblock timeout

The time (in seconds) before a blocked account is automatically unblocked, that is, unblocked without administrator intervention. Valid values for this text box range from **60** to **86400**. This value must be greater than the **Reset Failed Login Attempt Counter after** value.

Session timeout

The **idle time** and **max duration** values can be set here for trusted (private) and untrusted (public) environments. The value that is used is dependent on whether the user selects public or private on their login form. Also, either **max duration** or **idle time** can be specified as the value to use.

Idle time: The maximum idle time of the session in seconds, that is, idle timeout.

Max duration: The max duration of the session in seconds, that is, session timeout.

Note: Valid values for these fields range from **60** to **604800** (seconds).

Use for Session Timeout: A switch to select the session timeout behaviour (**max duration** or **idle time**).

Note: The underlying network traffic may keep the session active, even if there is no obvious user interaction.

Use LDAP Endpoint for Healthcheck

Select this check box to use the LDAP endpoint administrator username and password for health checking. If this is enabled, the **Test User** and **Test User Password** textboxes will not be available.

For more information on LDAP endpoints, refer to the LDAP Configuration section.

Note: This option is only available for the following protocols; **LDAP**, **Certificates**, **RADIUS and LDAP** and **RSA-SecurID and LDAP**.

Test User and **Test User Password**

In these two fields, enter credentials of a user account for your SSO Domain. The LoadMaster will use this information in a health check of the Authentication Server. This health check is performed every 20 seconds.

Client Side (Inbound) SAML SSO Domains

The fields vary when the **Authentication Protocol** is set to **SAML**. The SAML-specific fields are described below.

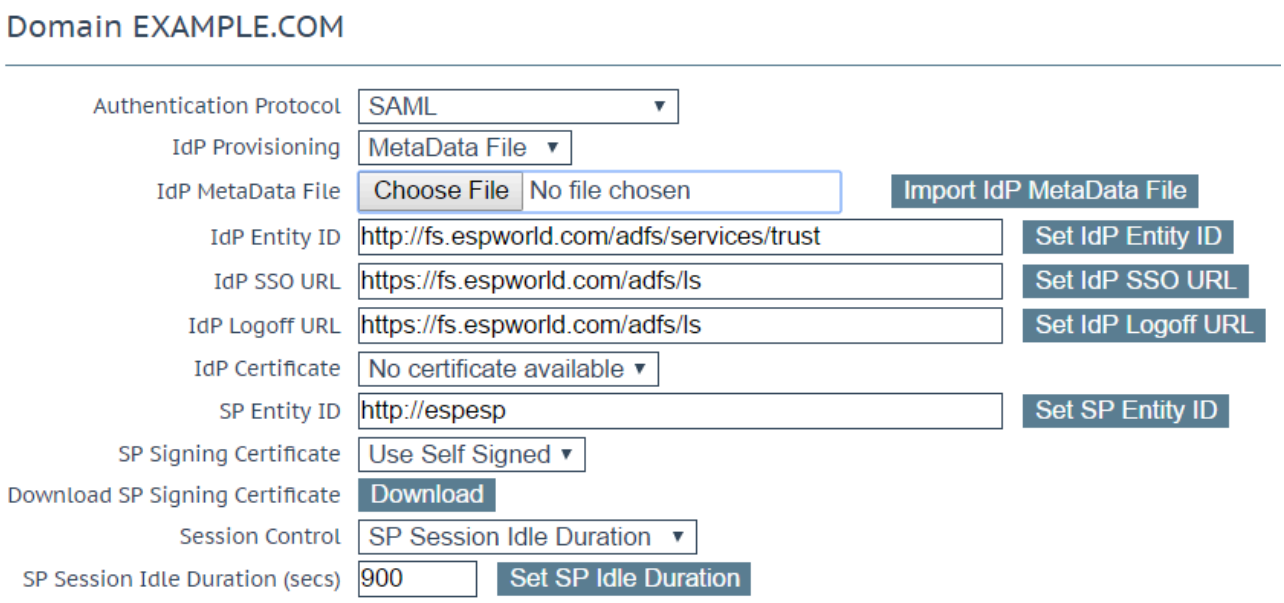

IdP Provisioning

The **Manual** option enables you to manually input details into the IdP fields.

The **MetaData File** option allows you to upload an **IdP MetaData File**. This simplifies the configuration of the IdP attributes, including the **IdP Entity ID**, **IdP SSO URL** and **IdP Logoff URL**. The metadata file can be downloaded from the IdP.

IdP Metadata File

This field is only visible if the **IdP Provisioning** field is set to **MetaData File**. To upload the file - click **Browse**, navigate to and select the relevant file and click **Import IdP MetaData File**.

IdP Entity ID

Specify the IdP entity identifier. The maximum number of characters permitted in this field is 255.

IdP SSO URL

Specify the IdP SSO URL. The maximum number of characters permitted in this field is 255.

IdP Logoff URL

Specify the IdP logoff URL. The maximum number of characters permitted in this field is 255.

IdP Certificate

The **IdP Certificate** is very important in terms of verification of the assertions that must be contained in the SAML response that is received from the IdP. Without the certificate, verification cannot proceed.

IdP Certificate Match

If this option is enabled, the IdP certificate assigned must match the certificate in the IdP SAML response.

SP Entity ID

This is an identifier that is shared to enable the IdP to understand, accept and have knowledge of the entity when request messages are sent from the LoadMaster. This must correlate to the identifier of the relying party on the AD FS server. The maximum number of characters permitted in this field is 255.

SP Signing Certificate

It is optional to sign requests that are sent in the context of logon. Currently, the LoadMaster does not sign those requests.

In the context of log off requests – it is mandatory and these requests must be signed. This is to avoid any spoofing and to provide extra security in relation to log off functionality. This ensures that users are not being hacked and not being logged off unnecessarily.

In the **SP Signing Certificate** drop-down list, you can choose to use a self-signed certificate or third party certificate to perform the signing.

Download SP Signing Certificate

If using a self-signed certificate, click **Download** to download the certificate. This certificate must be installed on the IdP server (for example AD FS) to be added to the relying party signature.

The AD FS server requires this certificate for use of the public key to verify the signatures that the LoadMaster generates.

Session Control

Select the relevant session control option. The available options are:

- SP Session Idle Duration
- SP Session Max Duration
- IdP Session Max Duration

The IdP maximum duration value cannot be set in the LoadMaster. The value is taken from the IdP protocol. If the value is not already set in the IdP authentication response, the default value of 30 minutes is assigned as the IdP maximum duration.

SP Session Idle Duration

Specify the session idle duration (in seconds). This field is only visible if **SP Session Idle Duration** is set as the **Session Control** option.

SP Session Max Duration

Specify the maximum duration of the session (in seconds). This field is only visible if **SP Session Max Duration** is set as the **Session Control** option.

Client Side (Inbound) OIDC / OAUTH SSO Domains

The fields vary when the **Authentication Protocol** is set to **OIDC / OAUTH**. The OIDC-specific fields are described below.

Application ID

Enter the Application (client) Identifier. The maximum number of characters permitted in this field is 255.

Redirect URI

Specify the redirect Uniform Resource Identifier (URI) or URIs (reply URLs). You can enter multiple URIs separated by a space. A maximum of 255 characters can be specified in the **Redirect URI** text box. Once a value is set for this field, you cannot unset it. For further details on the logic used when the **Redirect URI** field is set, refer to the OIDC Feature Description.

Authorization Endpoint URL

Enter the OAuth 2.0 authorization endpoint URL of the application. The maximum number of characters permitted in this field is 255.

Token Endpoint URL

Specify the OAuth 2.0 Token End Point URL of the application. The maximum number of characters permitted in this field is 255.

Logoff URL

Specify the Logout URL of the application. The maximum number of characters permitted in this field is 255.

Application Secret

Specify the value of the Client Secret of the application.

Session Control

Select the **Session Control**:

- **Session Idle Duration**
- **Session Max Duration**

Session Idle Duration/Session Max Duration

Specify the idle or max duration for the session (depending on what is selected for **Session Control**.

RADIUS Two-factor and LDAP Authentication

As of LoadMaster firmware version 7.2.52, RADIUS two-factor and LDAP authentication is supported. To configure this:

1. Select **RADIUS and LDAP** as the **Authentication Protocol** when adding or modifying a client-side Single Sign On (SSO) domain in **Virtual Services > Manage SSO**. If the RADIUS server is configured to use two-factor authentication, the LoadMaster will detect this automatically and perform RADIUS two-factor authentication.
2. Set the LDAP Endpoint and RADIUS Server(s) for this SSO domain.
3. Select **Exchange** or **Blank** as the **SSO Image Set** in the **ESP Options** section of the Virtual Service **Modify** screen.
4. Set the other parameters as appropriate for your configuration.

Sessions

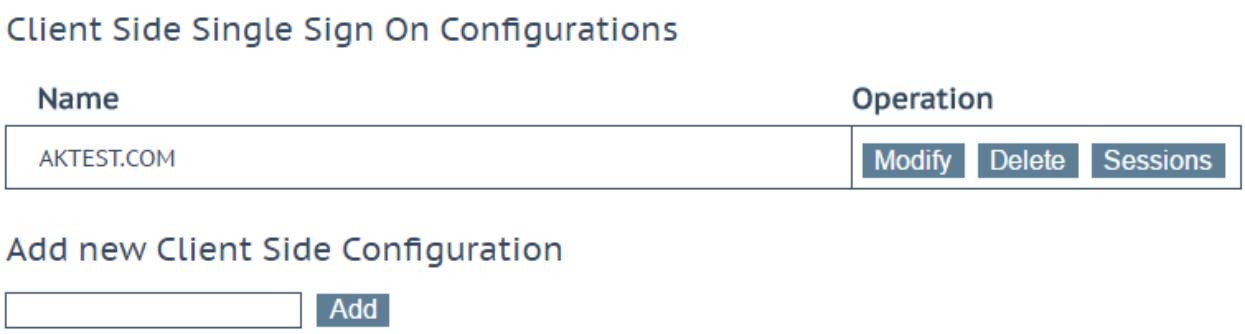

Clicking the **Sessions** button, for a client-side SSO domain, opens a screen listing the current open sessions on that domain.

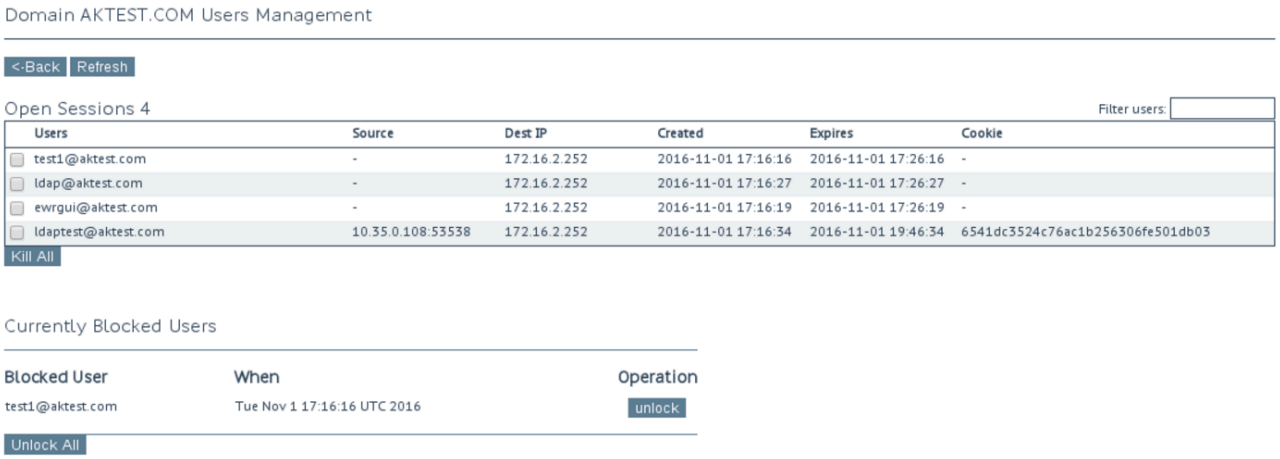

You can filter the list by entering a search term in the **Filter users** text box.

The following information is provided about each session:

- **Users:** The username/domain of the client.
- **Source:** The client (host) IP address and source port.
- **Dest IP:** The destination IP address of the connection.
- **Created:** The date and time the connection was created.
- **Expires:** The date and time that the connection expires.
- **Cookie:** The cookie used in the connection.

Clicking the **Kill All** button kills all open sessions (flushes the SSO cache).

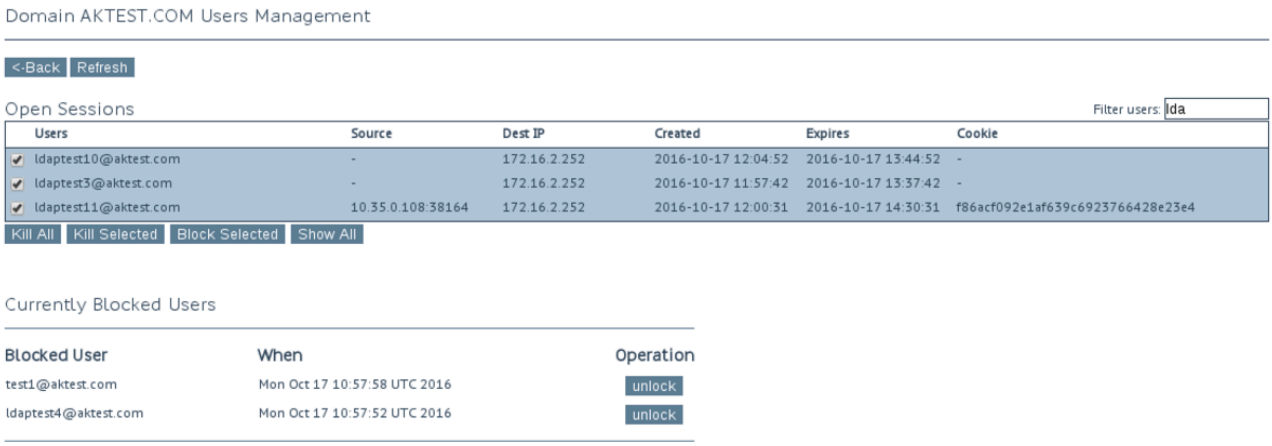

Selecting one or more sessions provides some further options:

- Kill Selected
- Block Selected
- Show All

Logs are added to the audit log for every kill session operation. For example:

- Kill 'non-cookie' session log:Nov 9 16:47:31 LM ssomgr: Deleted a session tester@aktest.com:- for domain AKTEST.COM
- Kill 'cookie' session log:Nov 9 16:47:31 LM ssomgr: Deleted a session ldaptest@aktest.com:420cf78373643b3c0171d95c757e7bf3 for domain AKTEST.COM
- Kill all domain sessions log:Nov 9 16:48:46 LM ssomgr: Deleted all domain AKTEST.COM user sessions

Currently Blocked Users

This section displays a list of users who are currently blocked and it also shows the date and time that the block occurred. It is possible to remove the block by clicking the **unlock** button in the **Operation** drop-down list.

Note: Different formats of the same username are treated as the same username, for example, **administrator@progress.net**, **progress\administrator** and **progress.net\administrator** are all treated as one username.

Server Side (Outbound) SSO Domains

When using Kerberos Constrained Delegation as the Server Side Authentication Protocol it is a requirement to create a Server Side SSO domain. This contains all the configuration required to complete authentication on the LoadMaster to server connection.



In the **Server Side Single Sign On Configurations** section of the **Manage SSO** screen, on creating the Server Side SSO, you may choose to utilize **Use AES256 SHA1 KCD cipher** by selecting the checkbox (by default the RC4 cipher is used). While enabling/disabling the **Use AES256 SHA1 KCD cipher** checkbox, a pop-up message appears that says "The KCD cipher set has been changed. Click OK to activate immediately, which will stop existing sessions. Click Cancel to activate later when the Kerberos key table expires.". If you click **OK**, the SSO cache is flushed which stops existing sessions. If you click **Cancel**, the change activates later when the Kerberos key table expires.

Note: If you change the value of the Use **AES256 SHA1 KCD** cipher check box using the Chrome browser and you navigate to a different tab while the pop-up message is displayed, the pop-up disappears and the value of the check box gets set. Because you did not click **OK** on the pop-up, the SSO cache is not flushed.

To add a new server-side SSO, enter the name of the SSO configuration and click **Add**.

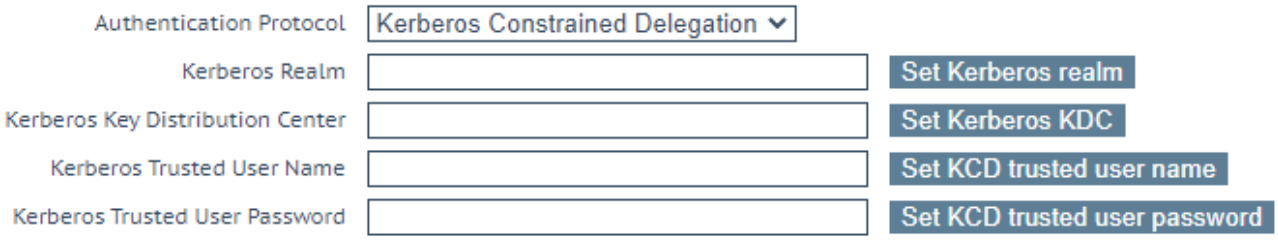

Authentication Protocol

This dropdown allows you to select the transport protocol used to communicate with the authentication server. The only option available for outbound (server side) configurations is **Kerberos Constrained Delegation** (KCD).

For further information on KCD, please refer to the KCD Feature Description.

Kerberos Realm

The address of the Kerberos Realm.

Note: Colons, slashes, and double quotes are not allowed in this field.

Note: This field only supports one address.

Kerberos Key Distribution Center (KDC)

The host name or IP address of the Kerberos Key Distribution Center. The KDC is a network service that supplies session tickets and temporary session keys to users and computers within an Active Directory domain.

When you configure a server-side Kerberos Constrained Delegation (KCD) Single Sign On (SSO) domain, you specify details for the domain. As of LoadMaster firmware version 7.2.51, you can specify two Kerberos Key Distribution Centers (KDCs) separated by a space. This provides a backup in case the active KDC becomes unavailable. Prior to version 7.2.51, you could only specify one KDC.

The first KDC you enter becomes active until it fails. KDC availability is checked and if the KDC fails to respond successfully three times, or if it times out for five seconds, the active KDC is switched. There is no automatic fail-back functionality - the second KDC will be active until it becomes unavailable. To switch back to the first KDC if a failover has occurred and the first KDC becomes available again, clear the SSOMGR cache by going to **System Configuration > Logging Options > System Log Files > Flush SSO Cache**.

When two KDCs are specified, the active Kerberos KDC is shown underneath the **Kerberos Key Distribution Center** field.

Note: If you enter more than one KDC, the username and password must be the same for both KDCs.

Note: Double and single quotes are not allowed in the **Kerberos Key Distribution Center** field.

Kerberos Trusted User Name

Before configuring the LoadMaster, a user must be created and trusted in the Windows domain (Active Directory). This user should also be set to use delegation. This trusted administrator user account is used to get tickets on behalf of users and services when a password is not provided. The user name of this trusted user should be entered in this text box.

Note: Double and single quotes are not allowed in this field.

Kerberos Trusted User Password

The password of the Kerberos trusted user.

Single Sign On Image Sets

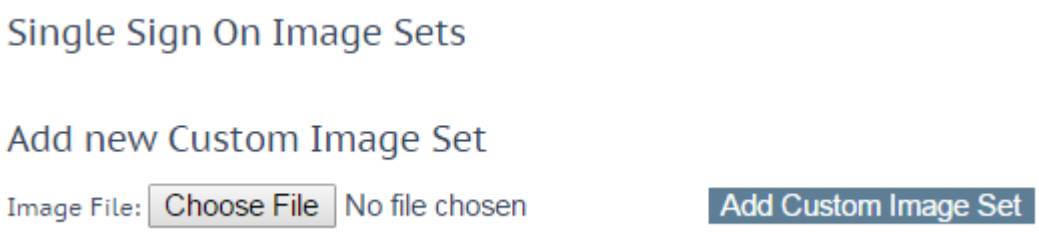

To upload a new image set, click **Choose File**, browse to and select the file and click **Add Custom Image Set**. After adding the file, the supplied image set(s) are listed on this page. It will also be available to select in the **SSO Image Set** drop-down list in the **ESP Options** section of the Virtual Service modify screen.

For more information on SSO image sets, including information on how the .tar file should be structured, refer to the Custom Authentication Form Technical Note.

Cache Configuration

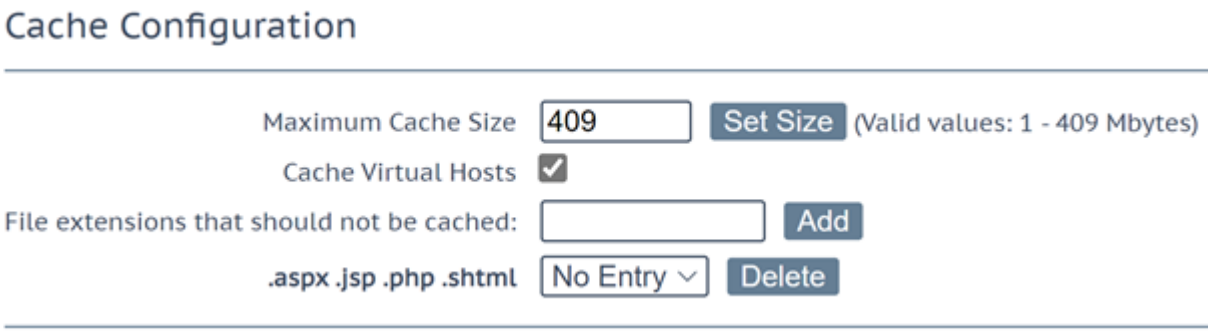

Maximum Cache Size

This defines how much memory can be utilized by the cache in megabytes. The **Maximum Cache Size** defines how much of the main memory should be assigned to the cache. It can never be more than one fifth of the total memory of the machine. Assigning more memory for the cache will reduce the amount of memory available for connections and persist entries. In a system that is correctly configured, there should be enough memory for a full cache and all connections that the system is expected to handle. If this is not the case, the system could run out of memory.

Cache Virtual Hosts

This option is enabled by default. When this option is disabled, the cache presumes there is only one virtual host supported on the Real Server. Enabling this option allows the cache to support multiple virtual hosts which have different content. If a request is cached for one Virtual Service, it will not be cached a second time if the request is sent to another Virtual Service. This is based on the host header of the request.

File extensions that should not be cached

The File extensions that should not be cached field and the label below it specifies a list of file types that should not be cached.

The default extensions that are not cached are .aspx, .jsp, .php, and .shtml. You can delete these from the "no cache" list by selecting them and clicking **Delete**.

You can add other file extensions to the **File extensions that should not be cached** field.

Items are held in the cache for 15 minutes (or until the cache is flushed by refreshing the page (for example, by pressing **Ctrl + F5** on your keyboard).

Compression Options

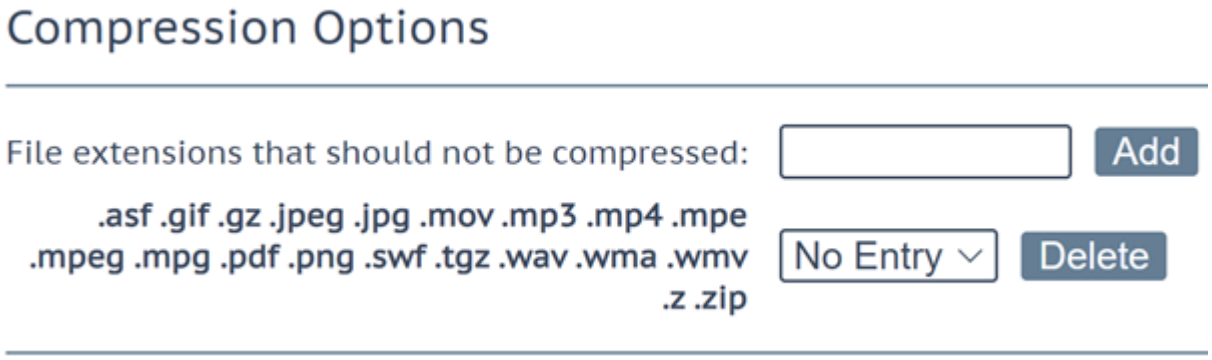

File extensions that should not be compressed

The **File extensions that should not be compressed** field and the label below it specifies a list of file types that should not be compressed. The default extensions that are not compressed are .asf .gif .gz .jpeg .jpg .mov .mp3 .mp4 .mpe .mpeg .mpg .pdf .png .swf .tgz .wav .wma .wmv .z .zip. You can delete these from the "no compress" list by selecting them and clicking **Delete**.

Kubernetes Settings

The Ingress Controller is not installed on the LoadMaster by default. You can easily install it by following these steps:

1. In the LoadMaster User Interface (UI), go to **Virtual Services > Kubernetes Settings**.
2. Click **Install**.
3. Wait for the installation to complete and click **OK** on the confirmation message.
4. Reboot the LoadMaster to activate all required add-ons - **System Configuration > System Administration > System Reboot > Reboot**.

After rebooting, you can use the **Kubernetes Settings** configuration page to enable the LoadMaster Kubernetes integration.

Note: For LoadMasters deployed in AWS or Azure in a High Availability pair, please ensure that the add-ons are installed on both devices.

Note: LoadMaster users (apart from the default admin **bal** user) must be assigned the **All Permissions** option in their user permissions to be permitted access to modify Kubernetes settings in the LoadMaster.

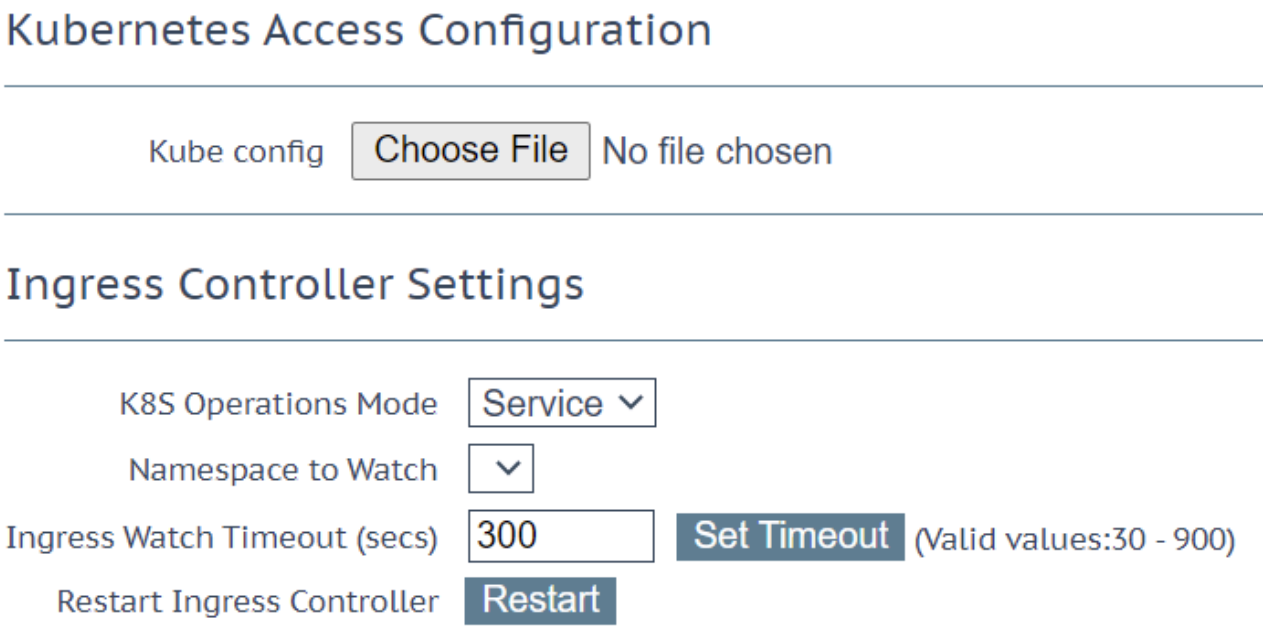

Kube Config

This allows the LoadMaster to communicate with Kubernetes.

The default location for the Kube Config file is **~/.kube/config**, for example, if you are using the Azure Cloud Shell you can access this using **/home/<YourName>/.kube/config** using the **Download File** option at the top of the Cloud Shell window.

Note: If using Minikube, it is recommended to embed the certificates used for authentication inside the kubeconfig file. This can be done with the following command in Minikube:**minikube config set embed-certs true**

Note: If the backup of one LoadMaster is restored to another LoadMaster, it is recommended to upload the **Kube Config** file separately to the restored LoadMaster.

Once the **Kube Config** file is successfully installed, some information is populated in the **Contexts** section. The **Name**, **Cluster**, and **User** are shown.

K8S Operations Mode

Determines the Ingress Controller mode of operation (Ingress or Service). "K8S" is used in the Virtual Service list to indicate a Virtual Service that is under Kubernetes control. For further details on each mode, including instructions on how to configure each mode in Kubernetes, refer to the Progress Kemp Ingress Controller for Kubernetes Feature Description.

Namespace to Watch

The K8s namespace to watch. All namespaces are watched if this field is unset. You can only filter on either a single name space or all name spaces. There is no capability to select multiple name spaces at this time.

Ingress Watch Timeout (secs)

The Ingress Controller watch timeout (in seconds). Valid values range from 30 - 900.

When configured correctly, details on the Kubernetes Nodes and relevant objects (Ingress Objects with an ingress class specification of "kempLB" and Service Objects labeled "kempLB:Enabled") are shown at the bottom of the screen.

For further details, refer to the Progress Kemp Ingress Controller for Kubernetes Feature Description.

4

Global Balancing

Note: This menu option may not be available in your configuration. These features are part of the GSLB Feature Pack and are enabled based on the license that has been applied to the LoadMaster. If you would like to have these options available, contact Progress Kemp to upgrade your license.

Contents

- Enable/Disable GSLB
- Manage FQDNs
- Manage Clusters
- Miscellaneous Params
- IP Range Selection Criteria
- IP Access List Settings
- Configure DNSSEC
- GSLB Statistics
- GEO System Info / Debug

Enable/Disable GSLB

Click this menu option to either enable or disable GEO features. When GEO is enabled, the **Packet Routing Filter** is enabled by default and cannot be changed. When GEO is disabled, it is possible to either enable or disable the **Packet Routing Filter** in **System Configuration > Access Control > Packet Filter**.

Manage FQDNs

A Fully Qualified Domain Name (FQDN), sometimes also referred to as an absolute domain name, is a domain name that specifies its exact location in the tree hierarchy of the Domain Name System (DNS). It specifies all domain levels, including the top-level domain and the root zone. A fully qualified domain name is distinguished by its lack of ambiguity: it can only be interpreted in one way. The DNS root domain is unnamed, which is expressed by the empty label, resulting in an FQDN ending with the dot character.

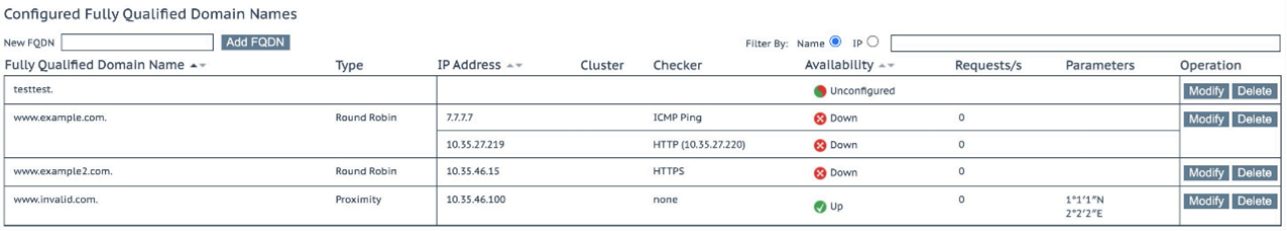

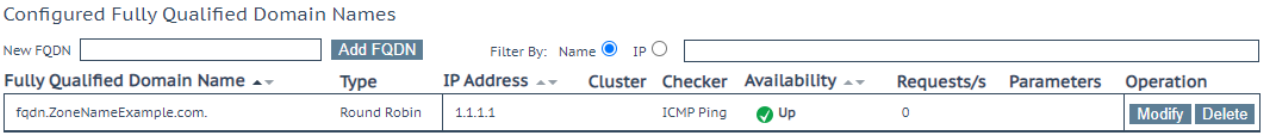

From this screen, you can **Add** or **Modify** an FQDN. You can also filter the FQDNs based on FQDN Name or IP Address using the filter text box.

Filter By Name

By default, radio button for FQDN Name is selected and when you enter any FDQN Name in the search text box, it displays the results for related FQDN entries.

Filter By IP

If you select the radio button for FQDN IP and enter any FQDN IP in the search text box, it displays the results for FQDN entries which have a related IP address.

If you enter any text in the filter text box that is contained in the name/IP, the list will be instantly filtered to show the matching results. If you enter some text in the filter text box that is not in any name/IP (depending on the radio button selected) the text box flashes red and the text that does not match is deleted.

In the FQDN table, you can sort the FQDN entries available for FQDN Name, IP Address, and Availability options using up/down arrow on FQDN table header.

- For FQDN Name, entries can be sorted in alphabetical order or reverse alphabetical order.
- For IP Address, entries can be sorted first by IP version (IPv4 followed by IPv6), then by IP address or reverse alphabetical order.
- For Availability, entries are sorted by the Down status first, followed by the remaining FQDNs only when there are Up and Down status FQDNs available in the list. If there are multiple FQDNs present in the list with multiple statuses like Up, Down, Disabled, and Unconfigured then sorting occurs in the following order:
    - Unconfigured

- Disabled
- Down
- Up

In reverse order:

- Up
- Down
- Disabled
- Unconfigured

Contents

- Add an FQDN
- Add/Modify an FQDN

Add an FQDN



New FQDN

The FQDN name, for example www.example.com. Wildcards are supported, for example ***.example1.com** matches anything that ends in **.example1.com**.

Note: As of LoadMaster version 7.2.60.1, underscores are supported in this field.

Note: The FQDN can be a maximum of 63 characters long. Otherwise, a pop-up message appears that says "FQDN must be 63 characters or less".

Add/Modify an FQDN

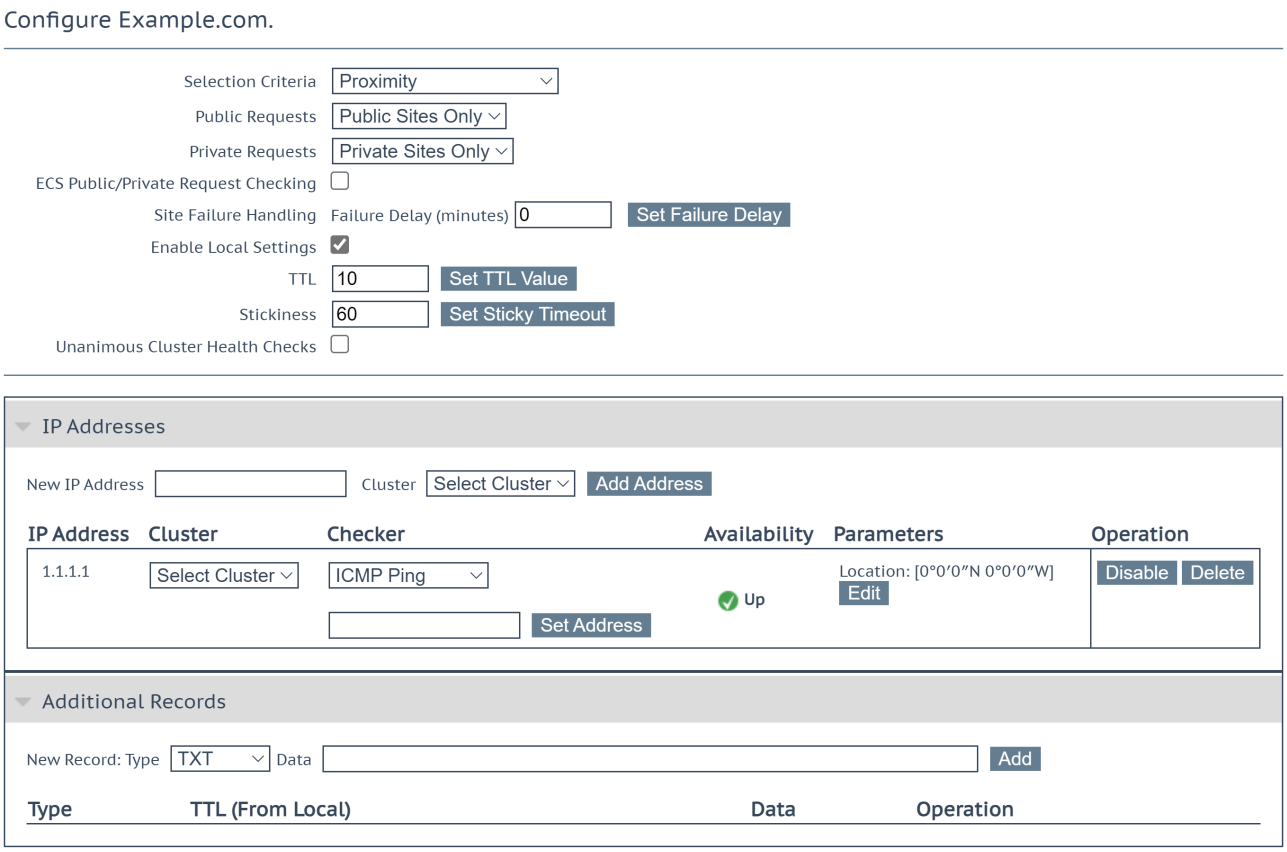

Selection Criteria

The selection criterion used to distribute the resolution requests can be selected from this drop-down list. The Selection Criteria available are:

- **Round Robin** - traffic distributed sequentially across the server farm (cluster), that is, the available servers.
- **Weighted Round Robin** – Incoming requests are distributed across the cluster in a sequential manner, while taking account of a static “weighting” that can be pre-assigned per server.
- **Fixed Weighting** - the highest weight Real Server is used only when other Real Server(s) are given lower weight values. If multiple Real Servers share the same highest weight, the system distributes traffic among them in a round-robin manner until a server with a higher weight becomes available.
- **Real Server Load** - LoadMaster contains logic which checks the state of the servers at regular intervals and independently of the configured weighting.
- **Proximity** – traffic is distributed to the closest site to the client. When using **Proximity** scheduling, new public sites are automatically mapped to geographic coordinates based on the GEO database. New private sites are mapped to 0º0'0" and function as expected. This coordinate should be overridden with accurate values to ensure correct balancing. The position of the client is determined by their IP address.

Note: In LoadMaster firmware version 7.2.52, a bug was fixed which had previously caused GEO location coordinates to be changed after creating or modifying an FQDN when using **Proximity** as the **Selection**

Criteria. This issue no longer occurs in versions 7.2.52 and above. However, if this issue occurred on a version previous to 7.2.52 and you upgrade to 7.2.52 or above, the coordinates do not get automatically fixed. Therefore, if the incorrect coordinates are already in the LoadMaster, you must manually correct them.

- **Location Based** - traffic is distributed to the closest site to the client. The positioning of the sites is set by inputting the location of the site (country or continent) during setup. The position of the client is determined by their IP address. If there is more than one site with the same country code, requests are distributed in a round robin fashion to each of the sites.
- **All Available** – returns all possible healthy targets for an A, AAAA or ANY query request. The contents of the returned list is also controlled by the **Public Requests** and **Private Requests** settings:
- For **Public Sites Only** the list can only contain public addresses. Likewise, for **Private Sites Only** the list can only contain private addresses.
- For **Prefer Public** the list only contains public addresses, unless no public addresses are available – in which case the list contains private addresses (if any are available). Likewise, for **Prefer Private** the list only contains private addresses, unless no private addresses are available – in which case the list contains public addresses (if any are available).
- For **All Sites** the list contains all available addresses

The purpose of this is to provide a list of preferred addresses, if they are available. Otherwise, provide a list of non-preferred addresses as a failback measure for improved availability.

Fail Over

The **Fail Over** option is only available when the **Selection Criteria** is set to **Location Based**. When the **Fail Over** option is enabled, if a request comes from a specific region and the target is down, the connection will fail over and be answered with the next level in the hierarchy. If this is not available, the connection is answered by the nearest (by proximity) target. If this is not possible, the target with the lowest requests are picked. For example, if a request from Ireland is received, but the site assigned to **Ireland** is unavailable, a site assigned to **Europe** is selected. If the site assigned to Europe is also unavailable, a site assigned to **Everywhere** is selected. If this too is unavailable, the site with the lowest requests of the available sites in the same continent is selected using the round robin method. The **Fail Over** setting affects all targets.

Public Requests & Private Requests

The **Isolate Public/Private Sites** setting has been enhanced in version 7.1-30. The checkbox has been migrated to two separate dropdown menus to allow more granular control of DNS responses. Existing behavior has been preserved and is migrated from your current setting, ensuring that no change in DNS responses is experienced.

These new settings allow administrators finer control of DNS responses to configured FQDNs. Administrators may selectively respond with public or private sites based on whether the client is from a public or private IP. For example, administrators may wish to allow only private clients to be sent to private sites.

The following table outlines settings and their configurable values:

Setting	Value	Client Type	Site Types Allowed
Public Requests			
	Public Only	Public	Public
	Prefer Public	Public	Public, Private if no public
	Prefer Private	Public	Private, Public if no private
	All Sites	Public	Private and Public
Private Requests			
	Private Only	Private	Private
	Prefer Private	Private	Private, Public if no private
	Prefer Public	Private	Public, Private if no public
	All Sites	Private	Private and Public

Note: Note that exposing private IP address information to public queries in this way may result in exposed network details. Select this setting at your own risk.

ECS Public/Private Request Checking

In LoadMaster firmware version 7.2.58, a new check box called **ECS Public/Private Request Checking** was added to the modify FQDN screen.

By default, the **ECS Public/Private Request Checking** check box is disabled. When disabled, the device uses the client request's source IP address to determine if the request is public or private. When enabled, the device uses the EDNS Client Subnet (ECS) value instead (if one is received) to determine if the request is public or private. This option becomes inactive in either of the following conditions:

- The **EDNS Client Subnet (ECS)** option (in **Global Balancing > Miscellaneous Params**) is disabled.
- The **Public Requests** and **Private Requests** values are both set to **All Sites**.

The following table summarizes the behavior depending on the configuration:

EDNS Client Subnet (ECS) Setting	Public/Private Requests Settings	Public/Private Behavior	ECS Public/Private Request Checking Effect
Disabled	Any	As in previous releases without the ECS option enabled, the source IP is checked.	This new option is ignored.
On	Public/private != **All Sites**	ECS is ignored; the source IP is checked.	If this new option is enabled, this overrides the behavior described to the left and ECS is used.

EDNS Client Subnet (ECS) Setting	Public/Private Requests Settings	Public/Private Behavior	ECS Public/Private Request Checking Effect
On	Public/private = **All Sites**	ECS is checked; source IP is ignored	This new option is ignored.

Site Failure Handling

The default is for failover to occur automatically. However, in certain circumstances, for example in a multi-site Exchange 2010 configuration, this may not be optimal and different behaviour may be required. **Failure Delay** is set in minutes. If a **Failure Delay** is set, a new option called **Site Recovery Mode** becomes available.

Site Recovery Mode

The **Site Recovery Mode** options only become available when a **Failure Delay** is set. If an endpoint fails the health check it is marked as "failing". After the **Failure Delay** time ends, the endpoint is marked as down. The **Site Recovery Mode** determines what recovery options are implemented when a failed site recovers.

- **Automatic** - when the endpoint recovers (passes the health check), the endpoint is marked as up and the LoadMaster automatically makes the endpoint available for use
- **Manual** - upon failure, the endpoint is administratively disabled and is not available for use until you click **Enable** for the relevant endpoint and the health check is passing.

Enable Local Settings

Selecting this option will display two additional fields – **TTL** and **Stickiness**. These can be set on a per-FQDN basis or globally. To set them for an FQDN – enable local settings and configure them as needed. The per-FQDN settings will default to the value of the global settings when the FQDN is created.

TTL

The Time To Live (TTL) value dictates how long the reply from the GEO LoadMaster can be cached by other DNS servers or client devices. The time interval is defined in seconds. This value should be as practically low as possible. The default value for this field is 10. Valid values range from 1 to 86400.

Stickiness

'Stickiness', also known as persistence, is the property that enables all name resolution requests from an individual client to be sent to the same resources until a specified period of time has elapsed. For further information on Stickiness, refer to the GEO Sticky DNS Feature Description.

Unanimous Cluster Health Checks

If this option is enabled, if any IP addresses fail health checking - other FQDN IP addresses which belong to the same cluster are marked as down. When **Unanimous Cluster Health Checks** is enabled, the IP addresses which belong to the same cluster within a specific FQDN are either all up or all down. For example, **example.com** has addresses 172.21.58.101, 172.21.58.102 and 172.21.58.103 which all belong to cluster **cl58**:

- If 172.21.58.101 fails, the unanimous policy forces 172.21.58.102 and 172.21.58.103 down as well.
- When 172.21.58.101 comes back, the unanimous policy brings back 172.21.58.102 and 172.21.58.103 along with it.

So, at any given time – either all three addresses are available or all three addresses are down.

The same approach applies for site failure mode with manual recovery. Manual recovery causes a failed address to be disabled, so the administrator can re-enable it after fixing the problem. When **Unanimous Cluster Health Checks** is enabled, all three addresses are disabled.

The unanimous policy ignores disabled addresses. So, if you know that an address is down, and for whatever reason you want to continue using the other addresses that belong to the same cluster, you can disable the failed address and the unanimous policy will not force down the other addresses with it.

When **Unanimous Cluster Health Checks** are enabled, some configuration changes may cause FQDN addresses to be forced down or brought back up. For example, if an address is forced down and you remove it from the cluster while the unanimous policy is in effect, the address should come back up. Similarly, if you add an address to a cluster where the unanimous policy is in effect and one of the addresses is down, the new address should be forced down. This change may not occur immediately, but it should happen the next time health checking occurs.

If there are addresses with the **Checker** set to **None** combined with addresses that have health checking configured – addresses with no health checking will not be forced down, but they can be forcibly disabled if the **Site Recovery Mode** is set to **Manual**. For example, say there are three addresses:

- 172.21.58.101 with a **Checker** of **Cluster Checks**
- 172.21.58.102 with a **Checker** of **Cluster Checks**
- 172.21.58.103 with a **Checker** of **None**

If site failure handling is off or automatic, the failure of 172.21.58.101 causes 172.21.58.102 to be forced down, but 172.21.58.103 remains up. The rationale is that if you do not want health checking on 172.21.58.103 then it should remain up.

However, if the **Site Recovery Mode** is set to **Manual**, failure of 172.21.58.101 causes both 172.21.58.102 and 172.21.58.103 to be disabled, along with 172.21.58.101. For site recovery – all addresses are disabled, even the ones with no health checking configured. This is to keep traffic away from the problem data center until the system administrators fix it. This does not conflict with having addresses with no health checking because you can have an address that is up but disabled.

IP Addresses

To create an IP address for an FQDN, enter the IP address of the domain in the **New IP address** text box.

Note: In LoadMaster firmware versions prior to 7.2.57, there is an entry limit of 64 IP addresses per FQDN.

Note: The entry limit was increased to 256 IP addresses in LoadMaster firmware version 7.2.57. If you attempt to add more than this, you get an error message saying **Limit of 256 IP Addresses has been reached for the FQDN**. It is recommended that the LoadMaster should have sufficient memory space (minimum 8GB of RAM) to configure the maximum of 256 IP addresses per FQDN.

Note: We do not recommend downgrading the LoadMaster firmware version to a previous version where the endpoint limit of IP addresses per FQDN is smaller than the number of IP addresses in your configured FQDNs.

Cluster

If needed, the cluster containing the IP address can be selected.

Checker

This defines the type of health checking that is performed. The options include:

- **None:** This implies that no health check is performed to check the health status of the machine (IP address) associated to the current FQDN.
- **ICMP Ping:** This tests the health status by pinging the IP address.
- **TCP Connect:** This will test the health by trying to connect to the IP address on a specified port.
- **Cluster Checks:** When this is selected, the health status check is performed using the method associated with the selected cluster.

Note: If there are no clusters configured, the **Cluster Checks** value is grayed out.

- **HTTP/HTTPS:** In LoadMaster firmware version 7.2.53, support was added to perform Layer7 (L7) HTTP and HTTPS health checks on back-end servers within GEO "sites" that are not handled from the LoadMaster for application delivery. In other words, site health determination can be enhanced directly from GEO by checking the health of back-end servers that are not being health-checked by LoadMaster.

Note: HTTP/1.1 is supported (HTTP/1.0 is not supported).

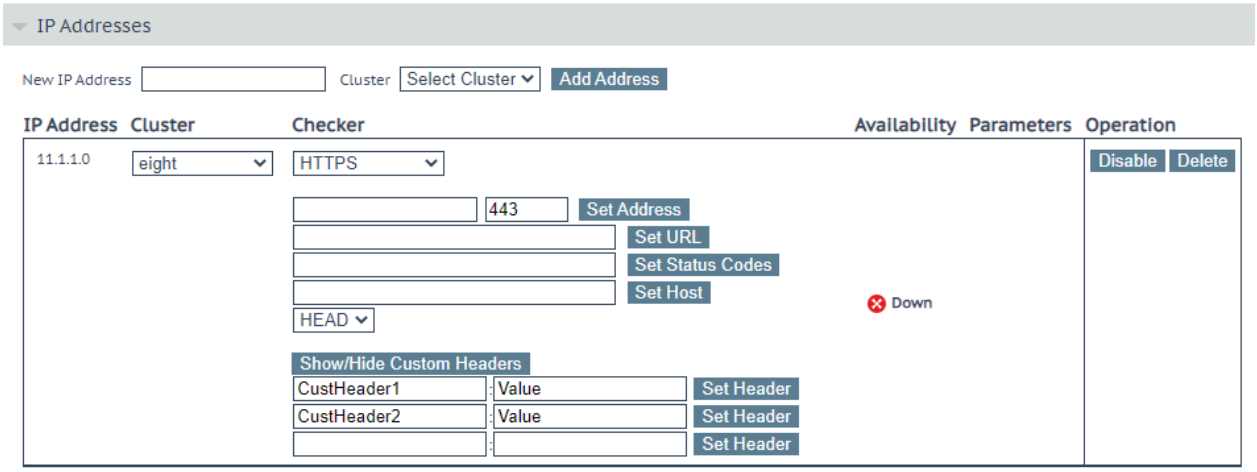

The available options that appear when **HTTP** or **HTTPS** is selected as the **Checker** are as follows:

- **Address:** Set the address and port to use to health check the IP address. The default port is **80** when **HTTP** is selected and **443** when **HTTPS** is selected.
- **URL:** By default, the health checker tries to access the URL forward slash (/) to determine if the machine is available. You can specify a different URL here.
- The URL must begin with a forward slash (/).
- The URL cannot contain **http:** or **https:**.

- The URL can be a maximum of 127 characters.
- If the URL is left blank, a forward slash (/) is sent by default.
- **Status Codes:** A space-separated list of HTTP and HTTPS status codes that should be treated as successful when received from the server.
- Codes must be between 300-599.
- There is a maximum of 32 codes.
- There is a limit of 127 characters.
- **Host:** A hostname can be supplied in the request to the server. If this is not set, the server address is sent as the host.
- There is a limit of 127 characters.
- Allowed characters: alphanumerics and **-._:**
- **Method:** When accessing the health check URL, the system can use the **GET**, the **POST**, or the **HEAD** method. If **POST** is selected, another field appears to set the POST data. Up to 2047 characters of POST data can be passed to the server.
- **Show/Hide Custom Headers:** When you click the **Show/Hide Custom Headers** option, it displays the custom headers that are to be part of the health check request.
- **Header:** Add a header to be part of the health check request.
  - There is a limit of 255 characters for the total check headers string. The first text box allows you to specify a key for the custom header. The second text box allows you to specify a value for the custom header. The value has a maximum length of 128 characters.
  - Allowed characters for header keys: alphanumerics and **_-**
  - Allowed characters for header values: alphanumerics and **_-.+=;()**

Note: You can specify a maximum of four custom headers.

The status of the health check displays in the **Availability** column.

Note: It is not possible to override the GEO health checks on a per-FQDN basis (like you can with Virtual Services).

Note: The **Check Parameters** in **Rules & Checking > Check Parameters** do not apply to GEO. Only the settings in **Global Balancing > Miscellaneous Params** apply to GEO.

For further information regarding health checks, refer to the GEO Feature Description**.**

Parameters

The parameters for the Selection Criteria are described and can be changed within this section. The parameters differ depending on the **Selection Criteria** in use, as described below:

- **Round Robin** – no parameters available
- **Weighted Round Robin** – the weight of the IP address can be set by changing the value in the **Weight** text box and clicking the **Set Weight** button
- **Fixed Weighting** – the weight of the IP address can be set in the **Weight** text box
- **Real Server Load** – the weight of the IP address can be set in the **Weight** text box and the Virtual Service which is measured can be chosen from the **Mapping** field
- **Proximity** – the physical location of the IP address can be set by clicking the **Edit** button
- **Location Based** – the locations associated with the IP address can be set by clicking the **Show Locations** button

Delete IP address

An IP address can be deleted by clicking the **Delete** button in the **Operation** column of the relevant IP address.



Additional Records

In LoadMaster firmware version 7.2.53, a new **Additional Records** section was added to allow you to configure records for a specific FQDN. You can add, modify, or delete a additional TXT, CNAME, and MX records to an FQDN. These record types allow you to communicate domain resources to clients:

- **TXT:** A TXT (text) record is essentially unformatted data that can be used for almost any purpose, but typically contains information to be consumed by clients to classify a domain in some way, provide details about a domain, or specify resources available within a domain.

- **CNAME:** A CNAME record points a DNS name (such as www.example.com) to another DNS name (such as lb.example.com). This is typically used to define a website alias.

  If you select **CNAME** from the **Type** drop-down list, the **NAME** and **RDATA** fields become available. By default, the RDATA field is set with the FQDN name unless some other value is entered in the field. You can add an FQDN that is unrelated to the FQDN to which the records are being added, if desired.

- **MX:** A mail exchanger (MX) record specifies the mail server responsible for accepting email messages on behalf of a domain name.

Note: Underscores are supported in the **Data** field for each of the record types.

For more information and instructions, refer to the GEO Feature Description.

Delete FQDN

An FQDN can be deleted by clicking the **Delete** button at the bottom of the **Modify** (Configure) **FQDN** screen.

Manage Clusters

GEO clusters is a feature mainly used inside data centers. Health checks are performed on a machine (IP address) associated to a specific FQDN, using the containing cluster server, rather than the machine itself.

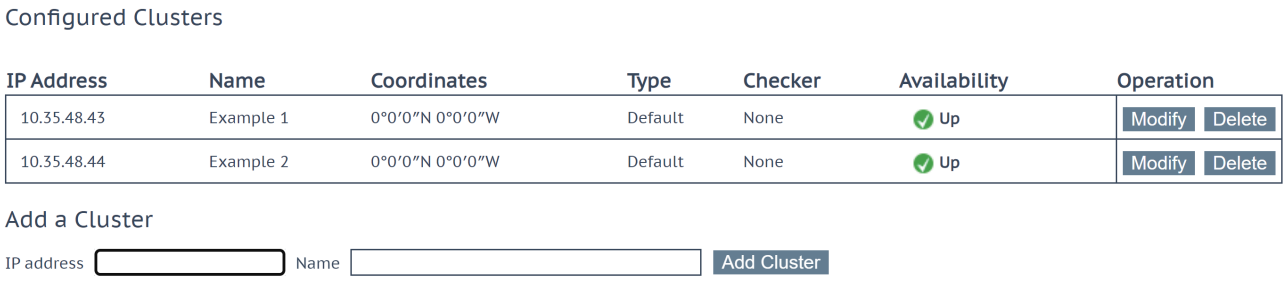

In the **Manage Clusters** screen there are options to **Add**, **Modify** and **Delete** clusters.

Contents

- Add a Cluster
- Modify a Cluster
- Delete a Cluster
- Upgrading GEO Partners

Add a Cluster

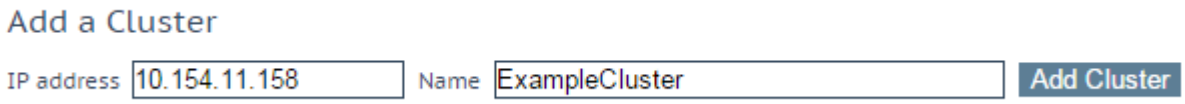

When adding a cluster, there are 2 text boxes to fill out:

- **IP address** – the IP address of the cluster
- **Name** – the name of the cluster. This name can be used to identify the cluster while in other screens.

Modify a Cluster

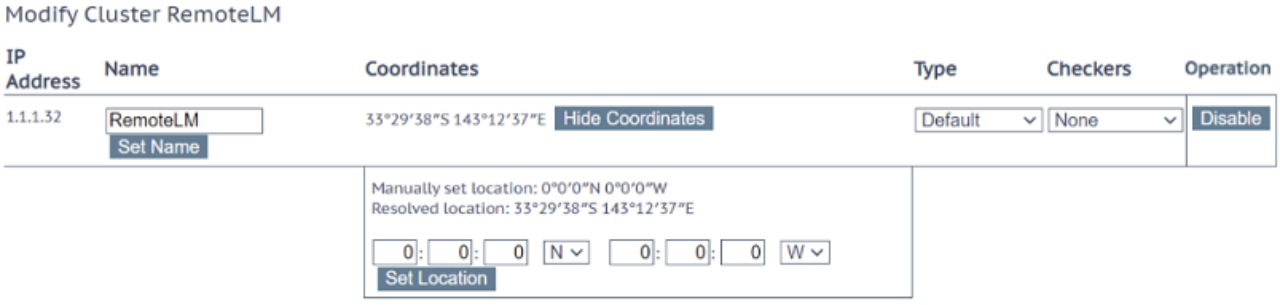

Name

The name of the cluster.

Location

If needed, the **Show Locations** button can be clicked to enter the latitude and longitude of the location of the IP address.

Type

The cluster type can be **Default**, **Remote LM** or **Local LM**:

- **Default:** When the type of cluster is set to **Default**, the check is performed against the cluster using one of the following three available health checks:
- **None:** No health check is performed. Therefore, the machine always appears to be up.
- **ICMP Ping:** The health check is performed by pinging against the cluster IP address.
- **TCP Connect:** The health check is performed by connecting to the cluster IP address on the port specified.
- **Local LM:** When **Local LM** is selected as the **Type**, the **Checkers** field is automatically set to **Not Needed**. This is because the health check is not necessary because the cluster is the local machine.
- **Remote LM:** The health check for this type of cluster is **Implicit** (it is performed using SSH).

When using **Real Server Load** as the **Selection Criteria**, and the cluster **Type** is set to **Local LM** or **Remote LM**, a drop-down list appears called **Mapping Menu**. The **Mapping Menu** drop-down list displays a list of Virtual Service names (where available) and Virtual Service IP addresses from that LoadMaster. If there are multiple Virtual Services using the same port, there is an entry in the drop-down list for the Virtual Service with up to five ports shown and other individual entries with each port. Select the Virtual IP address and port (or ports) that is associated with this mapping. If a Virtual Service with multiple ports is selected, the health check checks all Virtual Services with the same IP address as the one selected. If one of them is in an “Up” status, the FQDN shows as “Up”. The port does not come into consideration. If a Virtual Service with one port is selected, the health check only checks against the health of that Virtual Service when updating the health of the FQDN.

The only difference between **Remote LM** and **Local LM** is that it saves a TCP connection because it gets the information locally and not over TCP. Otherwise, the functionality is the same.

Checkers

The health check method used to check the status of the cluster.

If the **Type** is set to **Default** the health check methods available are **ICMP Ping** and **TCP Connect**.

If **Remote LM** or **Local LM** is selected as the **Type**, the **Checkers** dropdown list is unavailable.

Disable

If needed, a cluster can be disabled by clicking the **Disable** button in the **Operation** column.

Delete a Cluster

To delete a cluster, click the **Delete** button in the **Operation** column of the relevant cluster.

Note: Use the **Delete** function with caution. There is no way to undo this deletion.

Upgrading GEO Partners

When upgrading GEO partners, it is strongly recommended that all nodes are upgraded at the same time. Since GEO partners operate in active-active mode, upgrading at the same time ensures that consistent behavior is experienced across all nodes.

If you must operate a GEO partner with mixed versions, be sure to make all changes from the most recent version. This prevents configuration loss due to incompatible configurations. Additionally, changing configuration options not present in older versions will result in disparate behavior.

Miscellaneous Params

A description of the sections and fields in the **Miscellaneous Params** menu option are below.

Contents

- Source of Authority
- Resource Check Parameters
- Stickiness
- Location Data Update

Source of Authority

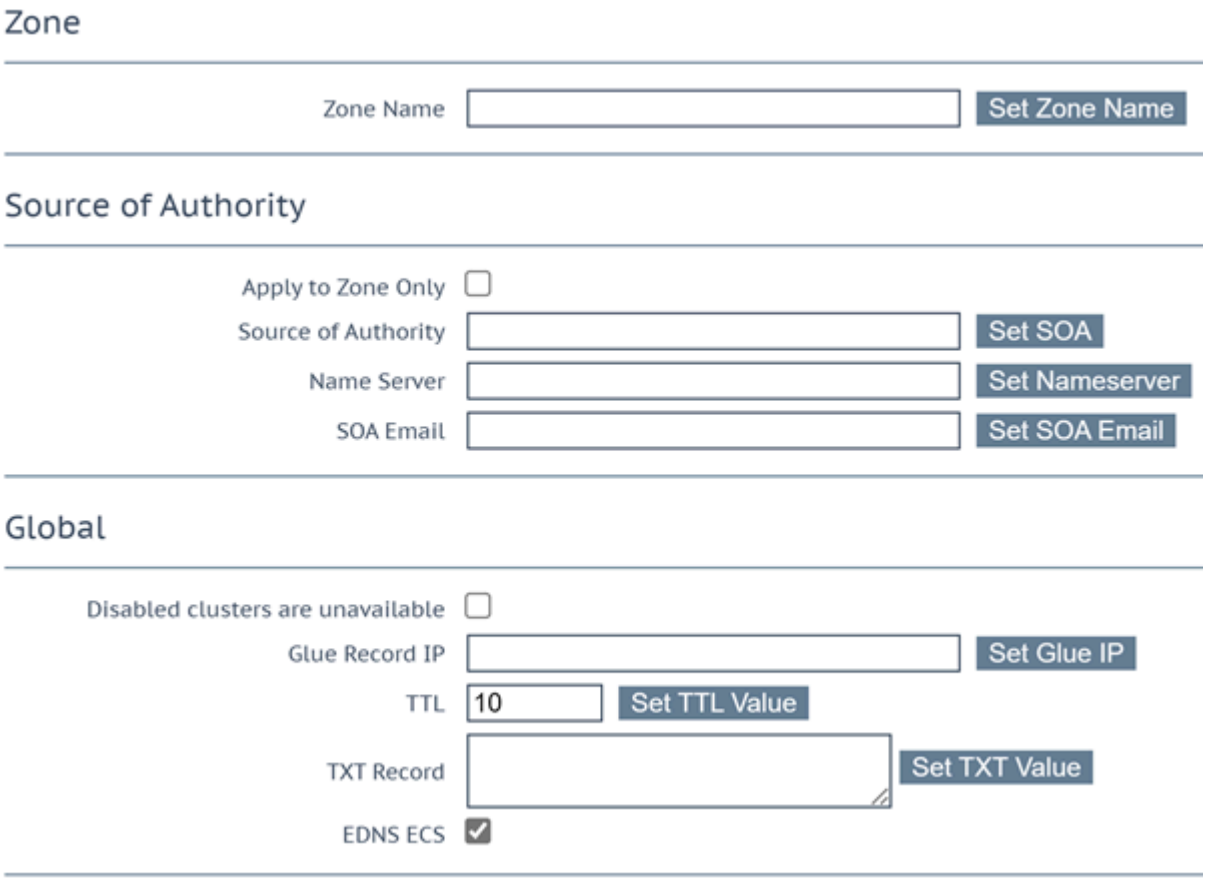

Zone Name

Enter the zone name to use. A zone name is necessary for DNSSEC configurations. All FQDNs within the zone are signed using the provided key. All FQDNs outside the zone continue to work but the responses are unsigned.

Note: As of LoadMaster version 7.2.60.1, underscores are supported in the **Zone Name**, **Source of Authority**, **Name Server**, and **SOA Email** fields.

Note: In LoadMaster firmware version 7.2.52, the **Zone Name** field moved to the new **Zone** section and a new **Apply to Zone Only** check box was added to the **Source of Authority** section. If this option is enabled, the Source of Authority (SOA) parameters are applied only to the zone. If it is disabled, the SOA parameters apply to all Fully Qualified Domain Names (FQDNs). The Apply to Zone Only option is disabled by default.

Source of Authority

This is defined in RFC 1035. The SOA defines global parameters for the zone (domain). There is only one SOA record allowed in a zone file.

Name Server

The Name Server is defined as the forward DNS entry configured in the Top Level DNS, written as a Fully-Qualified Domain Name (FQDN and ends with a dot), for example **lm1.example.com**.

The maximum number of characters permitted in this field is 1022.

If there is more than one Name Server, for example in a HA configuration, then you would add the second Name Server in the field also, separated by a blank space, for example **lm1.example.com lm2.example.com**.

SOA Email

This textbox is used to publish a mail address of a person or role account dealing with this zone with the "@" converted to a ".". The best practice is to define (and maintain) a dedicated mail alias, for example "hostmaster" [RFC 2142] for DNS operations, for example **hostmaster@example.com**.

Disabled clusters are unavailable

In LoadMaster firmware version 7.2.53, a check box named **Disabled clusters are unavailable** was introduced. This option is disabled by default. When it is enabled, requests to the cluster are dropped if a GEO cluster is disabled.

Glue Record IP

In LoadMaster firmware version 7.2.52, a new text box was introduced called **Glue Record IP** which allows you to set the IP address of the name server to return in additional records in a DNS response. Both IPv4 and IPv6 addresses are supported.

Glue records appear in query results if the following conditions are met:

- A **Zone** has been configured
- At least one **Name Server** has been configured and it belongs to the zone
- The FQDN being queried is part of that **Zone**

For each name server that belongs to the zone, its glue record is returned.

Query failures occur when a glue record is to be returned, but it is not present on the system.

You can currently add glue records to the LoadMaster in the following ways:

- In **Global Balancing > Miscellaneous Params** in the **Glue Record IP** field you can add a single glue record IP address and it will always correspond to the first name server configured in the **Name Server** field. The **Glue Record IP** field defaults to 0.0.0.0 if it is not set. The first Name Server is always guaranteed to have a glue record (0.0.0.0 if not set).
- If more than one name server is configured that require glue records, a workaround is to add the name server (or servers) as an FQDN in **Manage FQDNs** and ensure that it has at least one IP address configured. This FQDN then acts as a glue record from BIND's point of view.

Note: If you change the configuration to a scenario where a SERVFAIL occurs due to the lack of a glue record, the issue may not manifest itself until GEO is restarted because BIND appears to be able to mask the issue by caching previous glue records.

TTL

The Time To Live (TTL) value dictates how long the reply from the GEO LoadMaster can be cached by other DNS servers or client devices. This value should be as practically low as possible. The default value for this field is 10. The valid range for this field is between 1 and 86400. The time interval is defined in seconds.

TXT Record

In LoadMaster firmware version 7.2.52, support for the TXT (Text) record type was added to the GEO functionality. A TXT (text) record is unformatted data that can be used for almost any purpose, but typically contains information to be consumed by clients to classify a domain in some way, provide details about a domain, or specify resources available within a domain.

The configured record is returned for a TXT request on any Fully Qualified Domain Name (FQDN).

A maximum of 127 characters can be entered as the global TXT record.

The following are currently not supported:

- Multiple strings within the record: Some DNS providers allow you to put multiple strings in one entry using quotes, for example, "string 1" "string 2" "string 3". Progress Kemp does not currently allow this, so you can only have "string 1".
- Non-ASCII characters.

EDNS Client Subnet (ECS)

In LoadMaster firmware version 7.2.57, a checkbox named **EDNS Client Subnet (ECS)** was introduced. ECS is an option in the Extension Mechanisms for DNS that allows a recursive DNS resolver to specify the subnetwork when making a DNS query on behalf of the host or client.

By default, the **EDNS Client Subnet (ECS)** checkbox is enabled on new installations but if you are upgrading a LoadMaster that previously used GEO functionality then the option is disabled. When enabled, the ECS field (if included in the request) is used to determine the client location. When disabled, this field is ignored.

Note: It is possible that a client that has EDNS enabled may be sending requests through recursive DNS servers that are not enabled for EDNS. That means the client subnet will be dropped from the UDP packet that is received by the LoadMaster. If the EDNS client subnet information is removed it will not be available for the LoadMaster to act on.

Resource Check Parameters

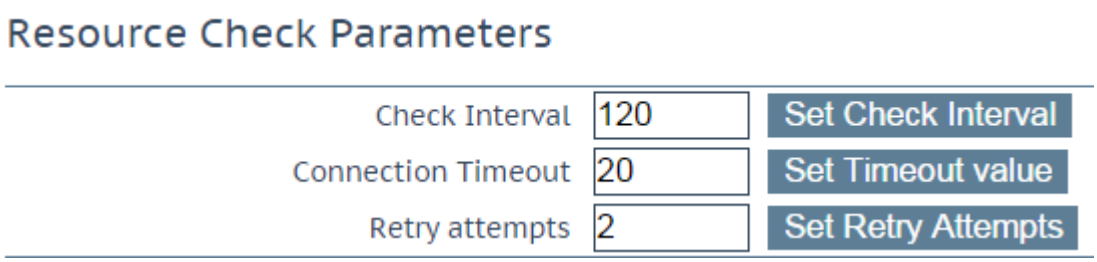

Check Interval (seconds)

The number of seconds between health check attempts. This defines the length of the 'health check cycle' and applies to both GEO clusters and FQDNs. The valid range for this field is between 9 and 3600. The default value is 120.

Note: The interval value must be greater than the timeout value multiplied by the retry value (Interval > Timeout * Retry + 1). This is to ensure that the next health check does not start before the previous one completes.

Note: If the timeout or retry values are increased to a value that breaks this rule, the interval value is automatically increased.

Connection Timeout (seconds)

This is the allowed maximum wait time for a connection to be established and for the receipt of a reply to a health check. If no reply is received before the timeout expires, then the connection is re-attempted as long as the Retry Count has not been reached. If any HTTP response is received in response to the health check, then no retries are attempted in the current health check cycle. The valid range for this field is between 4 and 60. The default value is 20.

Retry Attempts

This specifies the number of times the health check will be re-attempted if the timeout above is reached before receiving a response. The default retry attempts is 2. The retry count does not apply if any response is received in response to the health check. If, for example, either a 200 or a 404 response is received, then no retries are attempted in the current cycle.

The maximum detection window for failed clusters of FQDNs is the **Check Interval** + (**Connection Timeout** * (**Retry attempts** + 1)). On average, the maximum time is half that.

The timeline diagram below illustrates what happens from the time a resource IP is added or enabled, to when it goes down and then comes back up again:

1. When a resource IP is enabled/created, an ICMP request is sent by the LoadMaster to the resource IP. Assuming it responds, the resource is marked UP.
2. After 120 seconds has elapsed (the default **Check Interval**), an ICMP request is sent to the resource IP. If 20 seconds (the default **Connection Timeout**) elapses and the IP fails to respond, the LoadMaster will send up to two additional requests (the default **Retry Attempts**) and wait for 20 seconds between each. If all three of these requests receive no response, then the resource is marked down, and the **Check Interval** timer is reset.
3. After 120 seconds elapses, the LoadMaster attempts to send an ICMP request to the resource IP. If the resource has now come back up and responds before the **Connection Timeout** elapses, the LoadMaster marks it UP and resets the **Check Interval** timer.

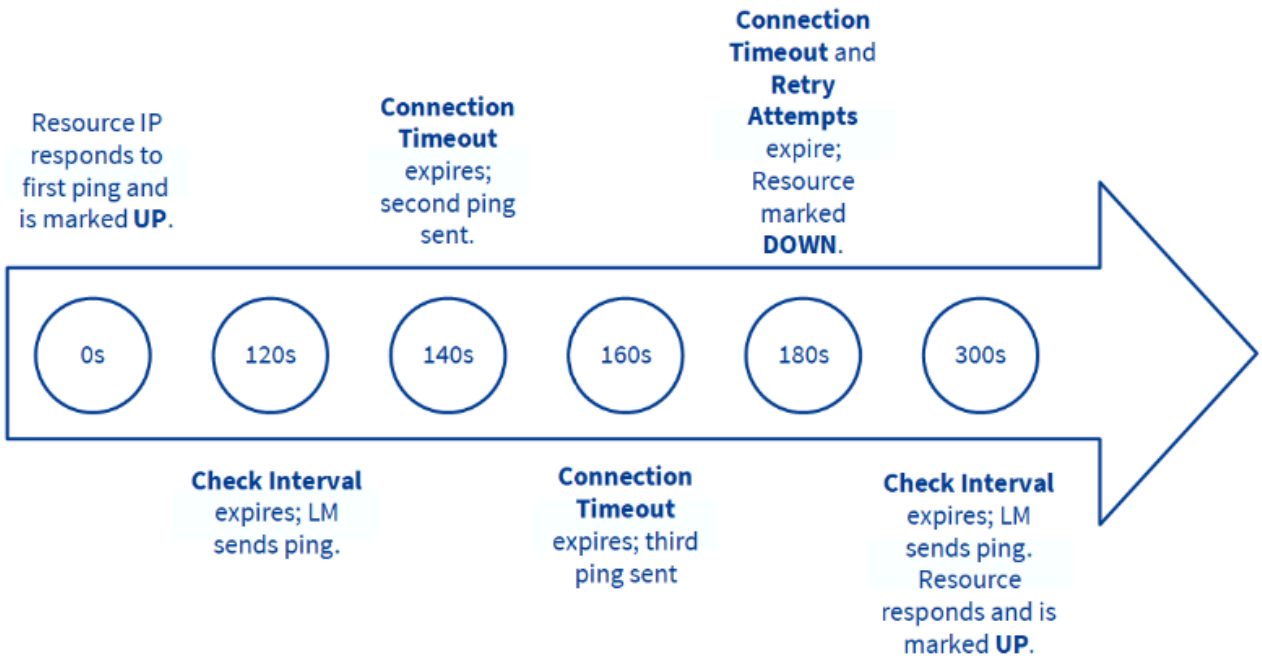

Stickiness

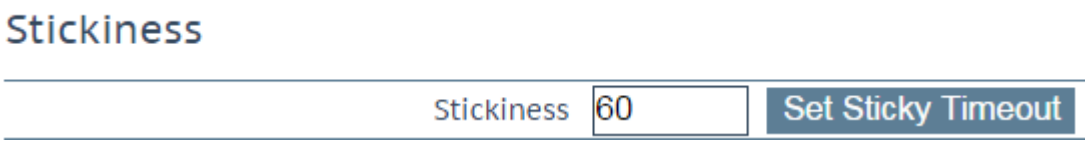

'Stickiness', also known as Global Persistence, is the property that enables all name resolution requests from an individual client to be sent to the same resources until a specified period of time (in seconds) has elapsed. For further information on **Stickiness**, refer to the GEO Sticky DNS Feature Description.

Location Data Update

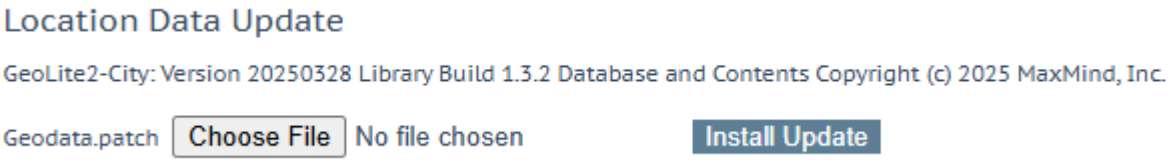

The location patch contains the geographically-encoded IP to location data. Data files can be obtained directly from Progress Kemp using normal support channels. These files are a repackaged distribution of MaxMind; the GeoIP database.

Note: The legacy MaxMind GeoLite database is only supported on LoadMaster version 7.2.44 and below. The new MaxMind GeoLite2 database is only supported on LoadMaster version 7.2.45 and above.

To download the latest database file, go to: GEO IP Database.

Note: Keeping the GEO IP database up-to-date is essential to ensure accurate GEO location functionality.

IP Range Selection Criteria



This section allows a new IP address range to be defined.

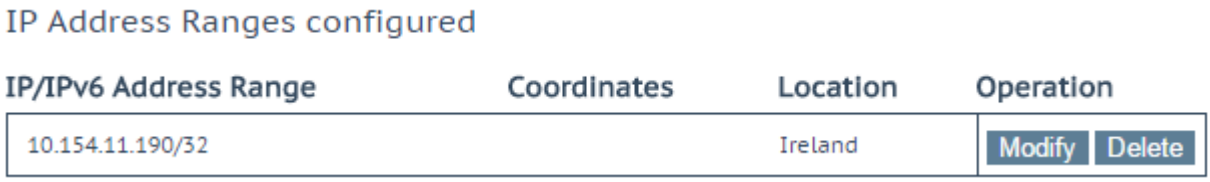

After adding an address, clicking **Modify** will open the modify settings screen. It is also possible to delete a range after it has been added.

Note: You cannot delete IP address ranges if they are configured in an FQDN.

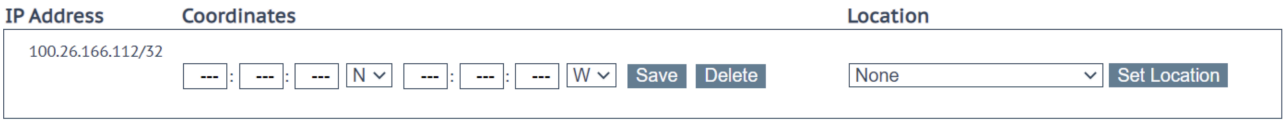

This section allows the definition of up to 64 IP ranges per data center.

IP Address

Specify an IP address or network. Valid entries here are either a single IP, for example **192.168.0.1**, or a network in Classless Inter-Domain Routing (CIDR) format, for example **192.168.0.0/24**.

Coordinates

Specify the latitude and longitude of the location.

Location

Specify the location to be assigned to the address.

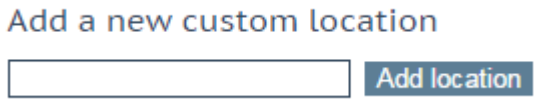

Add Custom Location

This section allows you to add a custom location.



Existing custom locations can also be modified and deleted in this section.

IP Access List Settings

It is possible to download access-list rules from Progress Kemp to block access from IP addresses that are on the access-list. An allow list can be manually specified that will override the access-list.

Note: This is a licensable feature. If you cannot see these options, or if any fields are grayed out, please contact Progress Kemp to upgrade your license.

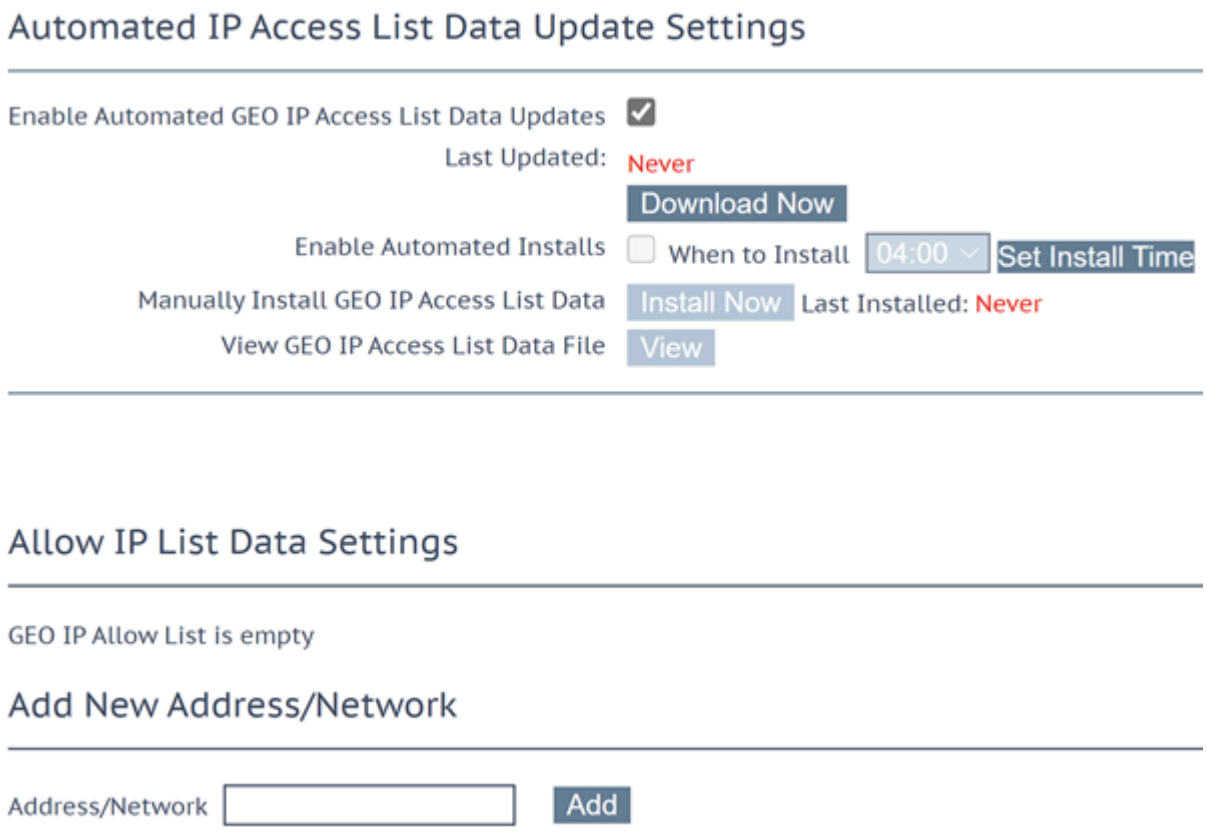

Enable Automated GEO IP Access List Data Updates

If this option is enabled, updates to the GEO IP access-list are downloaded daily. By default, this option is disabled.

Last Updated

The date when the last updates were downloaded is displayed. If the GEO access-list data is more than 7 days old, a message appears to inform you.

Download Now

Click this button to download the updates now.

Note: If you click **Download Now** to download the GEO Access Control List (ACL) when there is a routing issue such as missing DNS on the LoadMaster then the following error message is displayed:

Cannot update the GEO IP Access list. This may be due to a networking issue, an expired support contract, or some other connectivity issue. Please contact support for assistance.

Enable Automated Installs

Select this check box to enable the automatic daily install of updated rules at the specified time.

When to Install

Select the hour at which to install the updates every day.

Note: When the **Enable Auto-Save** option (under **System Configuration > Miscellaneous Options > WUI Settings**) is enabled and you select a new hour value, the value gets auto-updated.

When the **Enable Auto-Save** option is disabled, you must click **Set Install Time** after selecting a new hour value to update the installation time. If the hour value in the drop-down list has not been changed from its previous value and you click **Set Install Time**, the following error message displays:

"Value has not been changed. If auto-save is disabled, select a new value and press the set button again."

Manually Install GEO IP Access List Data

This button enables you to manually install the updates. This section also displays when the updates were last installed. If the GEO access-list data is not updated for more than 7 days, a message appears to inform you.

View GEO IP Access List Data File

Clicking the **View** button displays the current GEO IP access-list data file.

Allow IP List Data Settings

This section displays the IP addresses that are currently on the allow list.

Add New Address/Network

In this section, new addresses and networks can be added to the allow list. The allow list overrides the access-list.

Configure DNSSEC

Before you can configure DNSSEC, a zone must be defined. To define a zone, go to **Global Balancing > Miscellaneous Params** and specify a **Zone Name**.

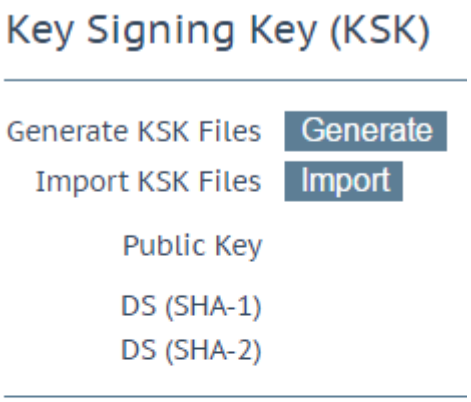



After the zone name is defined, the Key Signing Keys (KSKs) must be configured. You have two choices - you can either:

- Import the KSK files by clicking **Import** and browsing to the file locations.
- Generate the KSK files by clicking **Generate**

Note:

If you have GEO partners and want to use DNSSEC, you must generate the KSK files outside of the LoadMaster using the BIND **dnssec-keygen** command and import them onto each GEO partner separately, for example:

```
dnssec-keygen -a RSASHA256 -f KSK -b 2048 -n ZONE <zone_name>
```

Then, import the generated KSK files onto each GEO LoadMaster separately.



On the generate screen, select the cryptographic **Algorithm** and **Key Size**.

The following algorithms are supported:

- NSEC3RSASHA1
- RSASHA256
- RSASHA512

The default is RSASHA256.

The supported key sizes are 1024, 2048 and 4096 bits. The default is 2048.



After the KSK files have been generated/imported, the DNSSEC screen shows the KSK details and gives you an option to delete the KSK files.

The final step is to enable DNSSEC by selecting the check box.

GSLB Statistics

The **GSLB Statistics** screen (accessible from under the **Global Balancing** section in the main menu) is a centralized snapshot of the GSLB components that enable site resiliency and hybrid traffic distribution.



GSLB Service Status

Boot time	Thu, 13 Mar 2025 06:38:17 GMT
Last configuration	Thu, 13 Mar 2025 06:38:16 GMT

FQDN Statistics

Fully Qualified Domain Name	IP Address	Requests/s	Total
www.[redacted]	1.2.3.4	0	17

Queries

Type	Requests
A	11
AAAA	10
ANY	6

DNS Request Information

Type	Description	Requests
Requestv4	IPv4 Requests Received.	17
Requestv6	IPv6 Requests Received.	10
ReqEdns0	Requests with DNS Extension Mechanisms Received.	7
ReqTCP	TCP requests received.	6
Response	DNS Responses Sent.	27
RespEDNS0	DNS Responses with DNS Extension Mechanisms Sent.	7
QrySuccess	DNS Queries resulted in a successful answer.	17
QryAuthAns	DNS Queries resulted in authoritative answer.	27
QryNxrrset	DNS Queries resulted in NOERROR responses with no data.	10
QryUDP	UDP queries received.	21

The following sections display on the **GSLB Statistics** screen:

- **GSLB Service Status:** Displays the **Boot time** (the start time of the bind device) and the **Last configuration** (the date and time when the configuration was last modified).
- **FQDN statistics:** Displays the FQDN configuration with IP address information.
- **Queries:** Displays the different DNS query types received.
- **DNS Request Information:** Displays the type of DNS requests with a description and request count.

Note: If you disable and re-enable GSLB, the GSLB statistics reset to zero.

GEO System Info / Debug



GEO Information

This section provides information about the state of the GEO system, including:

- Config: Displays the last updated date (in format YYYY-MM-DD) and time, version, and total number of FQDNs and IP addresses
- Services: Displays the nameserver, healthcheck, persistence, and GEO process monitor details
- Remote machines: Displays details about configured GEO partners and remote GEO clients
- GEO file system: Displays details about the GEO file system including the number of forward zones, number of reverse zones, and number of DNS views details
- GEO IP Database: Displays the MaxMind database version

If you click the **Refresh** button, the data displayed in the **GEO Information** section gets refreshed.

GEO Debug Options

Restart GEO Services

Clicking **Restart GEO Services** restarts the nameserver, health check, and stickiness services.

GEO Process Monitor

When the **GEO Process Monitor** is enabled, three primary services (nameserver, healthchecks, and persistence) are monitored. If one or more of these services shut down, the LoadMaster attempts to restart them. By default, the **GEO Process Monitor** option is enabled.

The **GEO Process Monitor** setting is synced across High Availability (HA) units. The GEO Process Monitor setting is not synced across GEO partner units.

Automatic Partner Syncing

Enable or disable automatic synchronization of configuration data from this unit to all partners. Disabling the **Automatic Partner Syncing** option may be useful when debugging complex environments. The manual sync button on the page still functions when the automatic sync is disabled.

Sync Partners - Debug Mode

Clicking **Debug Sync** pushes the GEO configuration to the configured partners and logs the connection debug information in the message log.

Retrieve All Remote Cluster VS' - Debug Mode

Clicking **Debug Retrieve** gets the remote Virtual Service data from all configured LoadMaster clusters and logs the connection debug information in the message log.

Enable SSHD debug logs

Clicking **Enable SSHD debug logs** enables SSHD debug logs to debug the incoming connections from remote machines.

Enable GEO Query Logging

Enable bind debug trace logs for GEO.

Note: GEO Query Logging should only be enabled when troubleshooting because it will impact the performance of the LoadMaster.

Open warning log

Open the warning log in a new window. The warning logs also contain any SSHD debug logs if they are enabled on this page.

Open nameserver log

Open the name server log in a new window.

5

Statistics

Refer to the following sections for details about the **Real Time Statistics** and **Historical Graphs** in the LoadMaster.

Contents

- Real Time Statistics
- Historical Graphs

Real Time Statistics

Shows the activity for the LoadMasters within the system (**Global**), the **Real Servers**, the **Virtual Services**the **WAF**and the **Client Limits**.

Contents

- Global
- Real Servers
- Virtual Services
- WAF
- Client Limits

Global

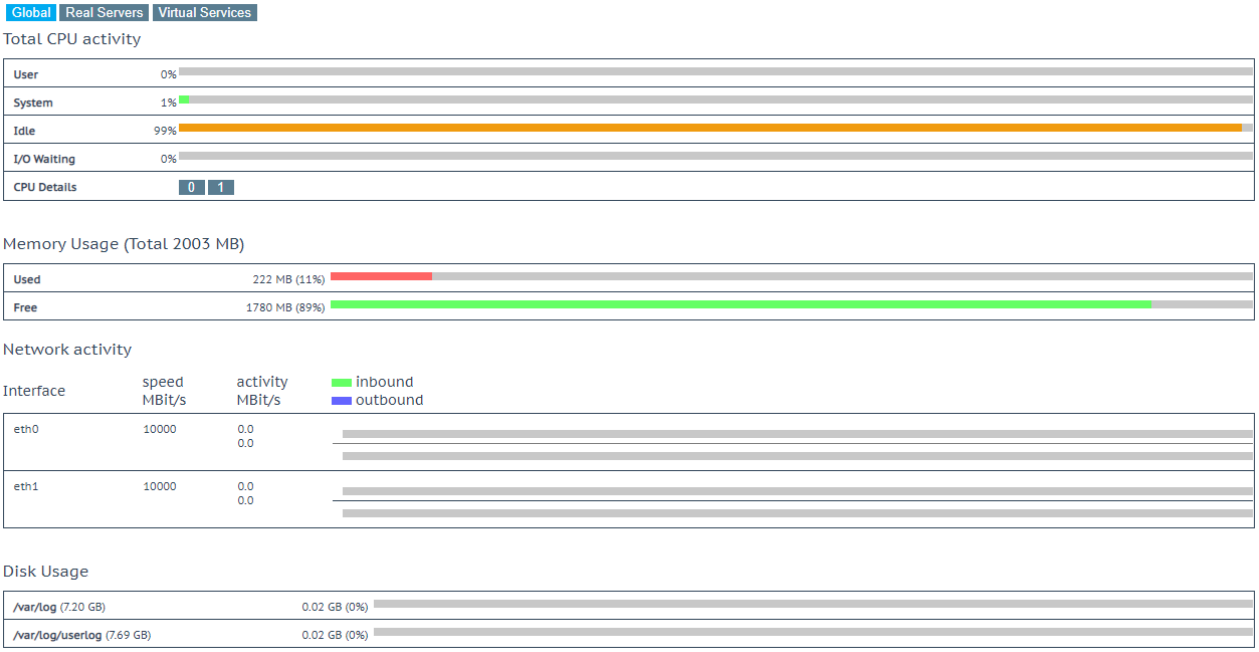

Total CPU Activity

This table displays the following CPU utilization information for a given LoadMaster:

Statistic	Description
User	The percentage of the CPU spent processing in user mode
System	The percentage of the CPU spent processing in system mode
Idle	The percentage of CPU which is idle
I/O Waiting	The percentage of the CPU spent waiting for I/O to complete

The sum of these four percentages equals 100%.

Core Temperatures: The temperature for each CPU core is displayed for LoadMaster hardware appliances.

Temperature will not show on a Virtual LoadMaster statistics screen.

These values are only available using SNMP. For further information on **SNMP Options**, refer to the **SNMP Options** section.

CPU Details: To get statistics for an individual CPU, click the relevant number button in **CPU Details**.

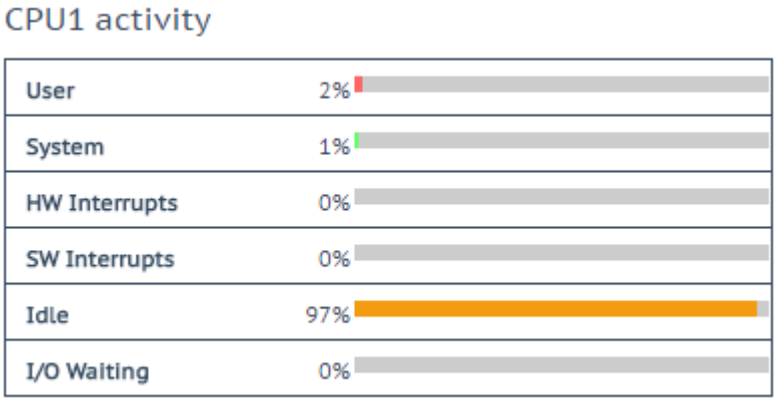

The CPU details screen has two additional statistics displayed - **HW Interrupts** and **SW Interrupts**.

Memory usage

This bar graph shows the percentage (MB) amount of memory in use and the amount of memory free.

Network activity

These bar graphs show the current network throughput on each interface.

Note: For Elastic Network Adapter (ENA) driver interfaces, the speed displayed behaves similar to bonded interfaces in that the speed of the overall interface is the aggregate of the number of separate logical interfaces that have been defined on that interface.

Disk Usage

This section provides an indication of the percentage (GB) used/free of the log partitions. Color-coding is used to highlight different usage levels:

- 0% to 50%: green
- 50% to 90%: orange
- 90% to 100%: red

Real Servers

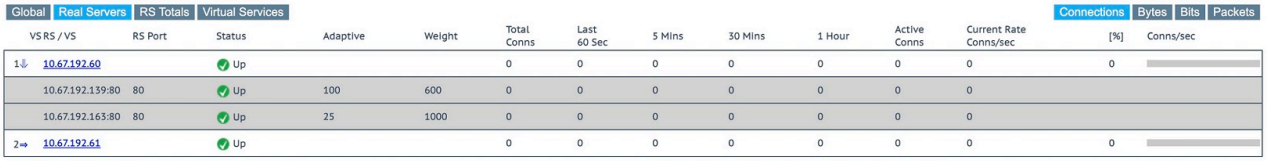

These graphs display the connections, bytes, bits or packets, depending on choice. The buttons in the top right of the page toggle which values are displayed. The values displayed for the Real Server comprise of the values for all the Virtual Services accessing the Real Server.

If the Real Server has been assigned to more than one Virtual Service, you can view the statistics for each Real Server by Virtual Service by clicking the arrow (⇒) to the right of the number in the first column. This expands the view to show the statistics for each Virtual Service on the Real Server.

Note: Because of the way that encrypted services are implemented, it is not possible to view the packet statistics on an encrypted Virtual Service.

VS: The **VS** column is automatically populated based on a DNS lookup.

RS/VS: This column displays the IP address of the Real Servers and Virtual Services.

Clicking the links in the **RS/VS:** column displays another screen containing a number of statistics specific to that Real Server or Virtual Service. For further details, refer to the section: Real Server or Virtual Service Statistics.

RS Port: Displays the port of the Real Server.

Status: This shows the status of the Real Server.

Adaptive: This will only be displayed if an adaptive scheduling method has been selected for a Virtual Service. This column will display the adaptive value.

Weight: This will only be displayed if the scheduling method is set to **resource based (SDN adaptive)** in a Virtual Service. The information which is gathered from the controller determines what the **Adaptive** value is set to. As the adaptive value goes up, the weight of the Real Server goes down. If all adaptive values are the same, all weights are the same. When the adaptive values are different the weights will change. The weight of the Real Servers determines where traffic is sent. If a Real Server is configured in multiple Virtual Services, two numbers are displayed for the weight - the first shows the average of the current weights over all Virtual Services that the Real Server is configured in. The second shows the number of Virtual Services that the Real Server is configured in. For example, a **Weight** of **972/2** means that the average weight of a Real Server which is configured in two Virtual Services is 972.

Total Conns: The total number of connections made.

Note: For Layer 4 UDP connections - the connection count always shows as 0.

Last 60 Sec: The total number of connections in the last 60 seconds.

5 Mins: The total number of connections in the last 5 minutes.

30 Mins: The total number of connections in the last 30 minutes.

1 Hour: The total number of connections in the last hour.

Active Conns: The total number of connections that are currently active.

Note: When using ESP, all connections going through the login process are counted as active connections for the Virtual Service. They are not counted as active connections for the Real Server because they are not actual connections to the Real Server. The WUI page displays the number of active connections associated with the Real Servers, while SNMP displays the number of active connections for the Virtual Service. The API returns the same value as displayed in the WUI. Without ESP, these values are identical. When using ESP, the Virtual Service counts can be much higher than the final counts going to the Real Servers, due to the above reason.

Current Rate Conns/sec: The current rate of connections per second.

[%]: The percentage of connections per second.

Conns/sec: A graphical representations of the connections per second.

System Total Conns: This row displays totals for each of the columns.

RS Totals	Conns
Total	0
Last 60 Sec	0
5 Mins	0
30 Mins	0
1 Hour	0
Active Conns	0
Current rate - Conns/sec	0

Clicking **RS Totals** displays another screen containing total Real Server connection statistics.

Contents

- Real Server or Virtual Service Statistics

Real Server or Virtual Service Statistics

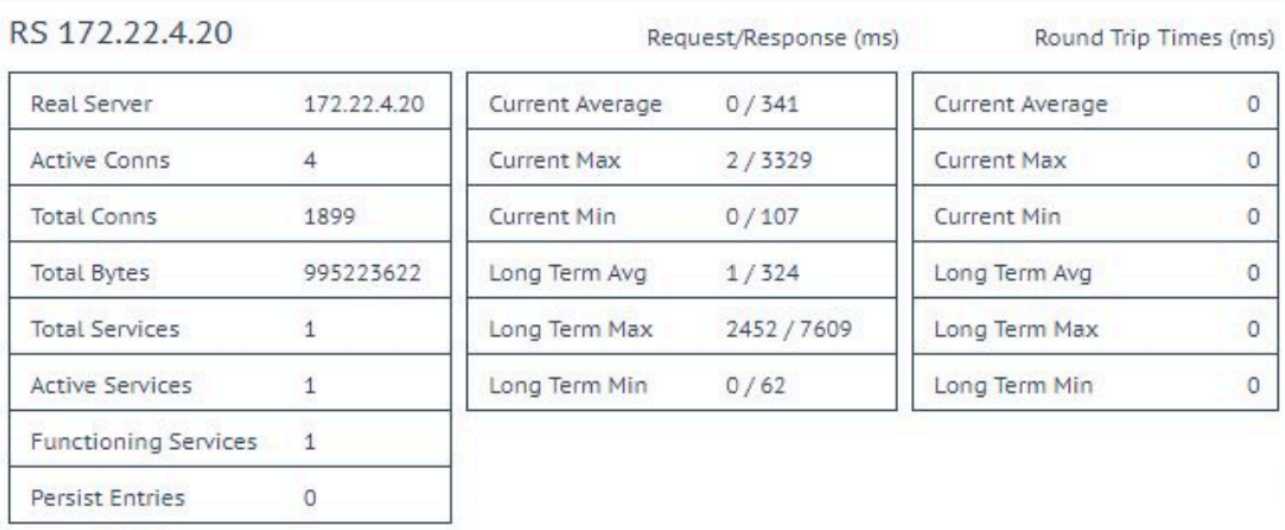

RS 172.22.4.20

Real Server	172.22.4.20
Active Conns	4
Total Conns	1899
Total Bytes	995223622
Total Services	1
Active Services	1
Functioning Services	1
Persist Entries	0

Request/Response (ms)

Current Average	0 / 341
Current Max	2 / 3329
Current Min	0 / 107
Long Term Avg	1 / 324
Long Term Max	2452 / 7609
Long Term Min	0 / 62

Round Trip Times (ms)

Current Average	0
Current Max	0
Current Min	0
Long Term Avg	0
Long Term Max	0
Long Term Min	0

The above screen displays a number of statistics specific to the Real Server or Virtual Service.

Round trip times between the LoadMaster and Real Server are displayed. The performance of the Real Servers is monitored with the following two measurements:

- How long the Real Server takes to respond to the request (first byte).
- How long the Real Server takes to send all the data (last byte) - response time.

In the above graphic the **Request/Response Current Max** figures are **2 / 3329**. This means that it took the Real Server a maximum of 2ms to respond to the request (first byte) and a max of 3329ms to send all the data and this is the measurement of the response time.

Note: Round Trip Time (RTT) to the server cannot be measured if using re-encrypt.

Note: Request/Response times are not applicable if the service being used is not a HTTP-HTTP/2-HTTPS service.

Note: Clicking **Reset Statistics** in the **Debug Options** screen (**System Configuration > System Administration > System Log Files > Debug Options**) clears all the values.

For each RTT or request/response time, the following values are measured:

- **Current Average:** The average time (in milliseconds (ms)) over the last 5 seconds (or since the last statistics reset).
- **Current Max:** The current maximum time (in ms) - in the last 5 seconds (or since the last statistics reset).
- **Current Min:** The current minimum time (in ms) - in the last 5 seconds (or since the last statistics reset).
- **Long Term Avg:** The average time over the whole period since the Virtual Service started handling traffic (or since the last statistics reset).
- **Long Term Max:** The all-time maximum value (in ms) since the Virtual Service started handling traffic (or since the last statistics reset).

- **Long Term Min:** The all-time minimum value (in ms) (or since the last statistics reset).

Values are only present if values have been recorded. For example, no **Current Max** displays if no traffic has passed through the Virtual Service in the last five seconds or if the Virtual Service cannot monitor the specific values.

Note: The **Current Average**, **Current Max**, **Current Min**, **Long Term Avg**, **Long Term Max**, and **Long Term Min** for Virtual Services and Real Servers are also available over SNMP.

Virtual Services

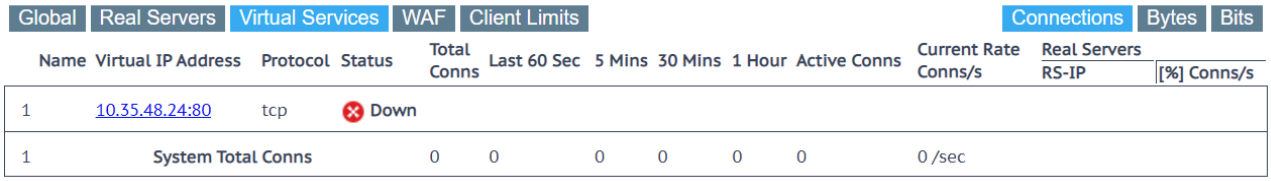

These graphs display the connections, bytes, bits or packets, depending on choice. The buttons in the top right of the page toggle which values are displayed. The percentage of distribution across the Virtual Service's Real Servers are displayed.

Name: The name of the Virtual Service.

Virtual IP Address: The IP address and port of the Virtual Service.

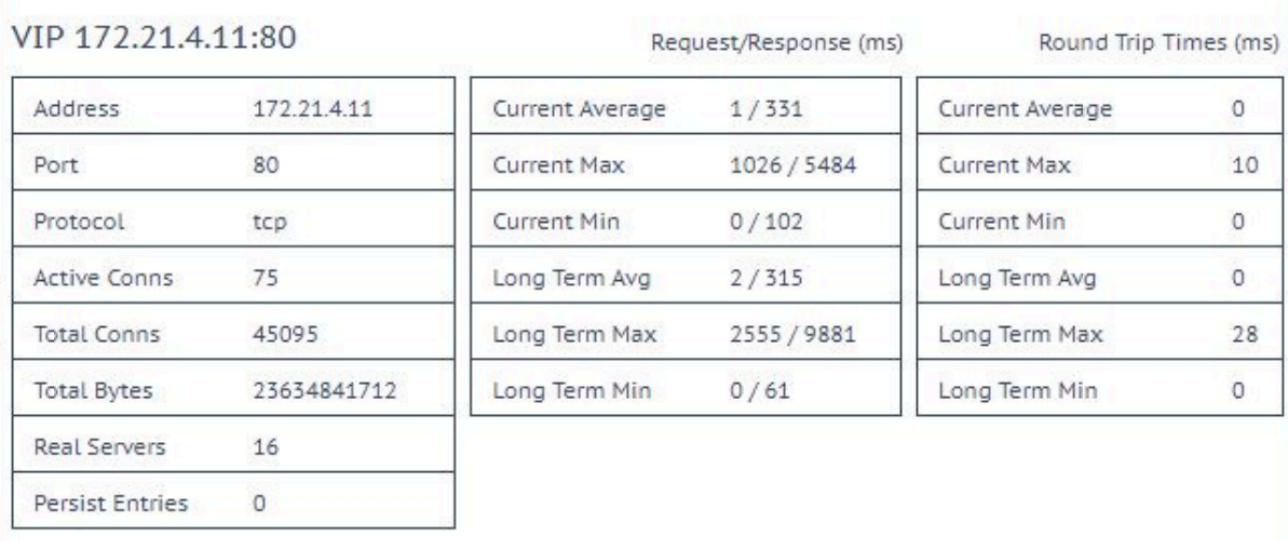

Clicking the links in the **Virtual IP Address** column will display another screen containing a number of statistics specific to that Virtual Service.

Round trip times between the client and LoadMaster are displayed. The performance of the Virtual Service is monitored with the following two measurements:

- How long the Virtual Service takes to respond to the request (first byte).

- How long the Virtual Service takes to send all the data (last byte) - response time.

In the above graphic the **Request/Response Current Max** figures are **1026 / 5484**. This means that it took the Real Server a maximum of 1026ms to respond to the request (first byte) and a max of 5484ms to send all the data and this is the measurement of the response time.

The following limiting statistics are also available in the Virtual Service statistics:

- Conns/Sec Blocked
- Req/Sec Blocked
- MaxConns Blocked

Note: Clicking **Reset Statistics** in the **Debug Options** screen (**System Configuration > System Administration > System Log Files > Debug Options**) clears all the values.

For each RTT or request/response time, the following values are measured:

- **Current Average:** The average time (in milliseconds (ms)) over the last 5 seconds (or since the last statistics reset).
- **Current Max:** The current maximum time (in ms) - in the last 5 seconds (or since the last statistics reset).
- **Current Min:** The current minimum time (in ms) - in the last 5 seconds (or since the last statistics reset).
- **Long Term Avg:** The average time over the whole period since the Virtual Service started handling traffic (or since the last statistics reset).
- **Long Term Max:** The all-time maximum value (in ms) since the Virtual Service started handling traffic (or since the last statistics reset).
- **Long Term Min:** The all-time minimum value (in ms) (or since the last statistics reset).

Values are only present if values have been recorded. For example, no **Current Max** displays if no traffic has passed through the Virtual Service in the last five seconds or if the Virtual Service cannot monitor the specific values.

Address: The IP address of the Virtual Service.

Protocol: The protocol of the Virtual Service. This will either be **tcp** or **udp**.

Active Conns: The total number of connections that are currently active. Here are some notes about how the active connections are calculated:

- The value reported in the WUI and API is an average calculated over five seconds.
- Because the number is an average and not a scalar data point, it does not represent "now" (the number of active connections at a specific point in time).
- Because of the sampling/average method used to report this statistic, the actual number of active connections can be severely under-counted if the incoming traffic stream contains a lot of short-lived connections.

Note: When using ESP, all connections going through the login process are counted as active connections for the Virtual Service. They are not counted as active connections for the Real Server because they are not

---

actual connections to the Real Server. The WUI page displays the number of active connections associated with the Real Servers, while SNMP displays the number of active connections for the Virtual Service. The API returns the same value as displayed in the WUI. Without ESP, these values are identical. When using ESP, the Virtual Service counts can be much higher than the final counts going to the Real Servers, due to the above reason.

Total Conns: The total number of connections made.

Total Bytes: The total number of bytes transmitted.

Real Servers: The total number of Real Servers in this Virtual Service.

Persist Entries: The total number of persistence entries made.

WAF: The status, along with the other WAF statistics below, are displayed if WAF is enabled on the Virtual Service.

Requests: The total number of requests handled by the WAF (shows all requests, whether they were blocked or not). Two requests are recorded for each connection – one incoming and one outgoing request.

Incidents: The total number of events handled by the WAF (that is, requests that were blocked).

Incidents/Hour: The number of events that have happened in the current hour (since xx.00.00).

Incidents/Day: The number of events that have happened since midnight (local time).

Incidents/Dayover: The number of times the event counter has gone over the configured warning threshold today. For example, if the threshold is set to 10 and there has been 20 events, this counter is set to 2. The warning threshold is set on a per-Virtual Service basis by filling out the **Hourly Alert Notification Threshold** field in **WAF Options** in the Virtual Service modify screen. For further information, refer to the Web Application Firewall (WAF) Options section.

System Total Conns: This row displays totals for each of the columns.

WAF

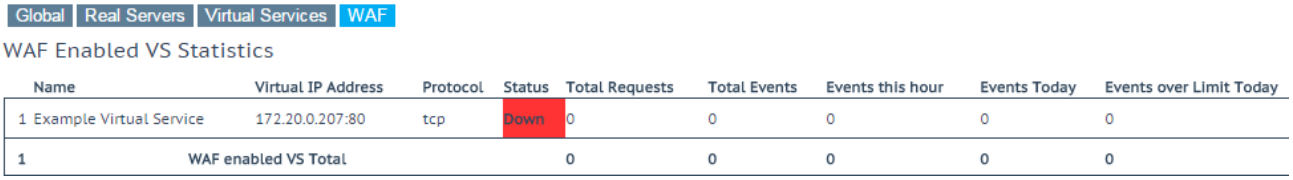

These statistics refresh every 5 to 6 seconds. The following items are displayed on this screen:

Count: The left-most column displays the total number of WAF-enabled Virtual Services.

Name: The name of the WAF-enabled Virtual Service.

Virtual IP Address: The IP address and port of the Virtual Service.

Protocol: The protocol of the Virtual Service (tcp or udp).

Status: The status of the Virtual Service. For information on each of the possible statuses, refer to the View/Modify (Existing HTTP Service) section.

Total Requests: The total number of requests handled by the WAF (shows all requests, whether they were blocked or not). Two requests are recorded for each connection – one incoming and one outgoing request.

Total Events: The total number of events handled by the WAF (requests that were blocked).

Events this hour: The number of events that have happened in the current hour (since xx.00.00).

Events Today: The number of events that have happened since midnight (local time).

Events over Limit Today: The number of times the event counter has gone over the configured warning threshold today. For example, if the threshold is set to 10 and there has been 20 events, this counter is set to 2. The warning threshold is set on a per-Virtual Service basis by filling out the **Hourly Alert Notification Threshold** field in **WAF Options** in the Virtual Service modify screen. For further information, refer to the Web Application Firewall (WAF) Options section.

Client Limits



The **Client Limits** button (under **Statistics > Real TIme Statistics**) is only displayed if there is at least one client limit enabled in the **System Configuration >QoS/Limiting** screen.

Statistics are only generated if the **Generate Limiter Statistics** check box is enabled in **System Configuration > QoS/Limiting > Limiter Options**. Statistics are updated every 10 seconds. There are buttons on the right of the **Client Limits** statistics screen where you can select different pages for **Total Connections** and **Bandwidth Usage**.

Note: These buttons are only displayed if the corresponding client limits are set in **System Configuration >QoS/Limiting**.

The top 10 clients are displayed for the **Last 30 seconds**, **Last 5 minutes**, and **Last 30 minutes**. There are separate columns to show the number of **Ok** and **Blocked** connections. Based on these insights, you can configure specific rate controls for specific client IP addresses.

Historical Graphs

The **Historical Graphs** screen provides a graphical representation of the LoadMaster statistics. These configurable graphs provide a visual indication of the traffic that is being processed by the LoadMaster.

Note: In some cases, after upgrading the LoadMaster firmware from version 7.1.35 to a newer firmware version, historical graphs may not display. To fix this, reset the statistic counters (**System Configuration > Extended Log Files > System Log Files > Debug Options > Reset Statistics**).

There are graphs for the network activity on each interface. There is also an option to view graphs for the overall and individual Virtual Services and the overall and individual Real Servers.

The time granularity can be specified by selecting one of the **hour, day, month, quarter** or **year** options.

In the case of the network activity on the interface graphs, you can choose which type of measurement unit you wish to use by selecting one of the **Packet, Bits** or **Bytes** options.

For the Virtual Services and Real Servers graphs you can choose which type of measurement unit you wish to use by selecting one of the **Connections, Bits** or **Bytes** options.

You can configure which Virtual Service statistics are being displayed by clicking the configuration icon: in the **Virtual Services panel**. This opens the Virtual Services configuration window.

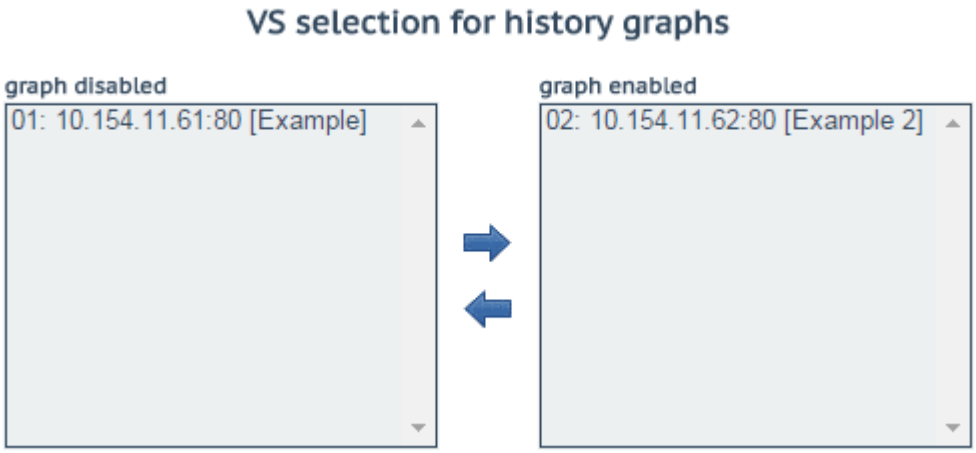

From here, Virtual Services can be added or removed from the statistics display.

You can disable these graphs by disabling the **Enable Historical Graphs** check box in the WUI Settings screen.

A maximum of five Virtual Services can be displayed at the same time.

To close the dialog and apply any changes, please ensure to click the  button within the window itself.

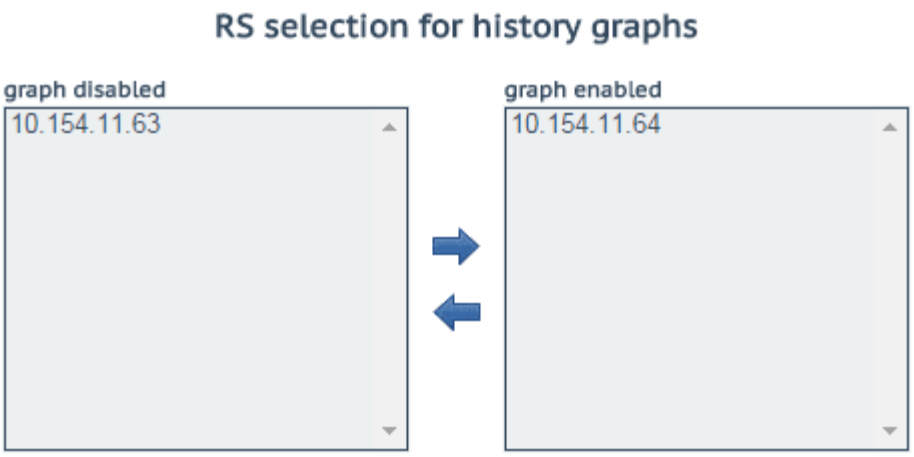

You can configure which Real Server statistics are being displayed by clicking the configuration icon, in the **Real Servers** panel. This opens the Real Servers configuration dialog in a separate window.

From here, Real Servers can be added or removed from the statistics display.

A maximum of five Real Servers can be displayed at the same time.

To close the dialog and apply any changes, please ensure you click the button within the window itself.

By default, only the statistics for the Virtual Services and Real Servers displayed on the Statistics page are gathered and stored. To view statistics for all Virtual Services and Real Servers, enable the **Collect All Statistics** option in **System Configuration > Miscellaneous Options > WUI Settings**.

Note: This option is disabled by default because collecting statistics for a large number of Virtual Services and Real Servers can cause CPU utilization to become very high.

The graphs in the LoadMaster WUI are auto-scaling and are shown using SI magnitude units. The graph will show the prefix of the scaling factor used so the absolute value can be calculated if needed.

The possible scaling factors and their prefixes are listed in the table below.

Symbol	Prefix	Factor
P	peta	10^15
T	tera	10^12
G	giga	10^9

Symbol	Prefix	Factor
M	mega	10^6
k	kilo	10^3
m	milli	10^(-3)
µ	micro	10^(-6)

To calculate the absolute “real” value, take the value shown in the graph and multiply it by the scaling value.

Example

A value of 200 is shown in the connections per second graph with a scaling factor of “m”. As listed in the table above, “m” stands for “milli”. Therefore, to find the absolute value of connections per second for that time – the value of 200 needs to be multiplied by a factor of 10^(-3):

- 10^(-3) = 0.001
- 200 x 0.001 = 0.2 connections per second

This calculation shows that there is less than one connection per second and, due to the fact that the connection rate is so low, if the graph shows the absolute number of connections it will just be a straight line at zero and offers no useful information.

6

SDN Statistics

To view the SDN statistics, go to **Statistics > SDN Statistics** in the main menu of the LoadMaster WUI.

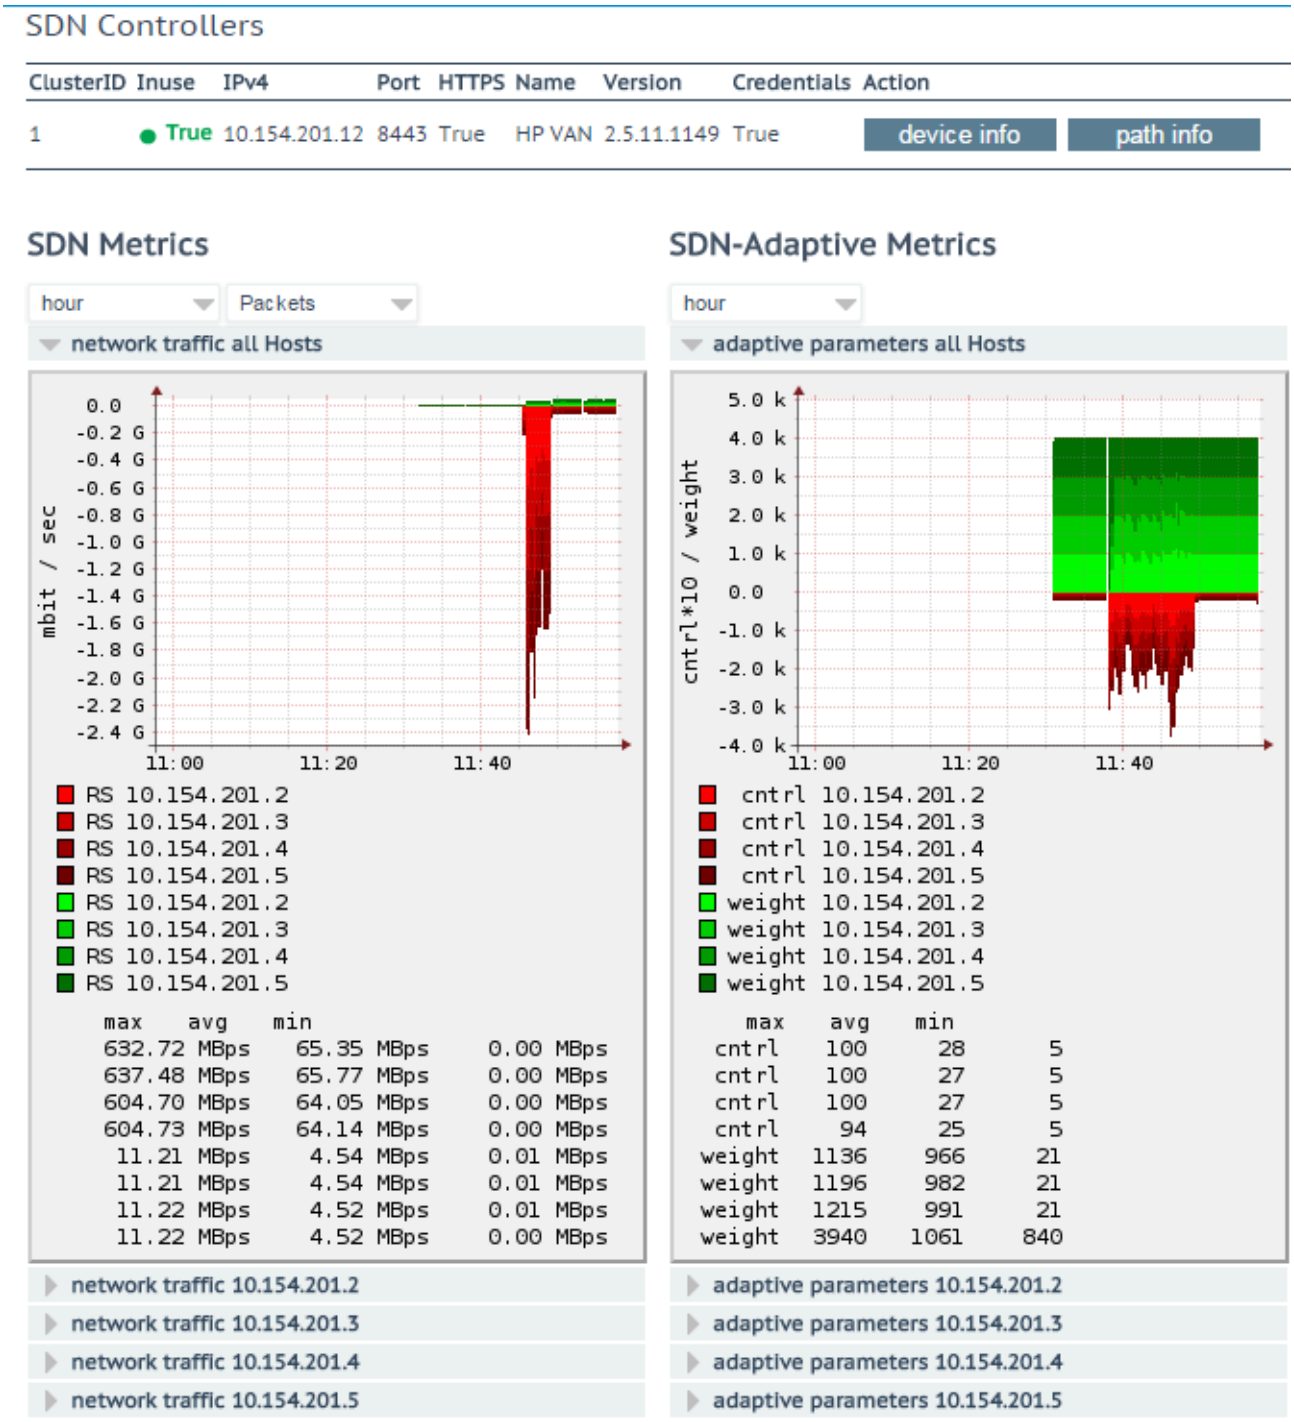

Note: The **Name**, **Version** and **Credentials** are displayed if the LoadMaster has successfully connected to the SDN Controller.

Statistics section

Note: Statistics will not be displayed unless the SDN Controller has been added and is communicating with the LoadMaster. If the **Name**, **Version** and **Credentials** are not displaying it means that the LoadMaster is

not connected to the SDN Controller. This could mean that the configuration is not correct, or the SDN Controller is down.

Two types of statistics are displayed on this screen - network traffic and adaptive parameters:

- Network traffic - this can display the number of bits and bytes transferred per second for each of the Real Servers. The maximum, average and minimum number of bits/bytes per second are shown.
- Adaptive parameters - this displays the adaptive value (**ctrl**) and the weight. As the adaptive value goes up, the weight of the Real Server goes down.

Contents

- Device Information

Device Information

UID	Name	Type
00:00:54:9f:35:1c:c5:30	ovsbr0	Default OpenFlow Switch
00:00:66:52:10:5f:fb:45	ovsbr1	Default OpenFlow Switch

Information about switches on a controller which has OpenFlow enabled can be viewed by clicking the **device info** button.

UID	Name	Type	Vendor	Product
00:00:54:9f:35:1c:c5:30	ovsbr0	Default OpenFlow Switch	Nicira, Inc.	Open vSwitch

Interface Info

ID	Name	State	Mac
id=0x1	Name:eno1	State:[UP]	Mac:54:9f:35:1c:c5:30
id=0x4	Name:vnet2	State:[UP]	Mac:fe:54:00:bc:1b:c3
id=0x7	Name:vnet1	State:[UP]	Mac:fe:54:00:8d:73:9b
id=0x8	Name:vnet7	State:[UP]	Mac:fe:54:00:b1:4b:3b
id=0xa	Name:patch-ovsbr0	State:[UP]	Mac:7e:6d:ac:6b:9f:11
id=0xb	Name:patch-ovsbr3	State:[UP]	Mac:2a:32:8c:e7:4c:5b
id=0xfffffffe	Name:ovsbr0	State:[UP]	Mac:54:9f:35:1c:c5:30

Node Info

ID	VID	Port	Mac
10.154.50.25	0	1	00:0c:29:b1:96:46
10.154.120.62	0	1	00:50:56:b8:13:45
10.154.190.197	0	1	00:50:56:b8:4d:7d
10.154.30.80	0	1	00:0c:29:64:83:1b
10.154.190.104	0	1	00:50:56:b8:e7:31
10.154.190.172	0	1	00:0c:29:91:e6:9d
10.154.190.137	0	1	00:0c:29:d7:aa:5e
10.154.25.30	0	1	00:50:56:b8:b4:5d
10.154.190.145	0	1	00:50:56:b8:54:d5
10.154.120.115	0	1	00:50:56:b8:19:67
10.154.190.111	0	1	00:50:56:b8:e8:08
10.154.190.120	0	1	00:50:56:b8:ee:39
10.154.190.157	0	1	00:50:56:b8:97:f6
10.154.190.126	0	1	80:3f:5d:08:92:d6
10.154.0.3	0	1	20:0c:c8:49:f6:4c
10.154.190.152	0	1	00:0c:29:54:e8:2b
10.154.190.174	0	1	00:50:56:b8:b7:2e
10.154.190.115	0	1	00:50:56:b8:7e:6b
10.154.50.61	0	1	00:50:56:b8:a5:00
10.154.190.151	0	1	00:50:56:b8:1b:67
10.154.190.118	0	1	00:50:56:b8:b7:5c
10.154.190.128	0	1	00:50:56:b8:d4:84
10.154.75.25	0	1	00:50:56:b8:0c:3f
10.154.25.102	0	1	00:50:56:b8:70:8c
10.154.190.190	0	1	00:10:f3:38:4a:e4
10.89.0.44	0	1	00:0c:29:56:ad:2f
10.154.190.150	0	1	00:0c:29:2b:d7:ac
10.154.50.167	0	1	00:0c:29:24:2e:49
10.154.30.81	0	1	00:0c:29:a1:6a:3b

Further information can be seen by clicking the plus (**+**) button to expand each of the devices.

Contents

- Path Information

Path Information

Path information can be viewed by clicking the **path info** button.

Note: The LoadMaster and the SDN controller need to be directly connected for the path information to be

displayed.

To view a graphical representation of the path, click the **=>** or **<=** icon in the **Dir** column for the relevant path.

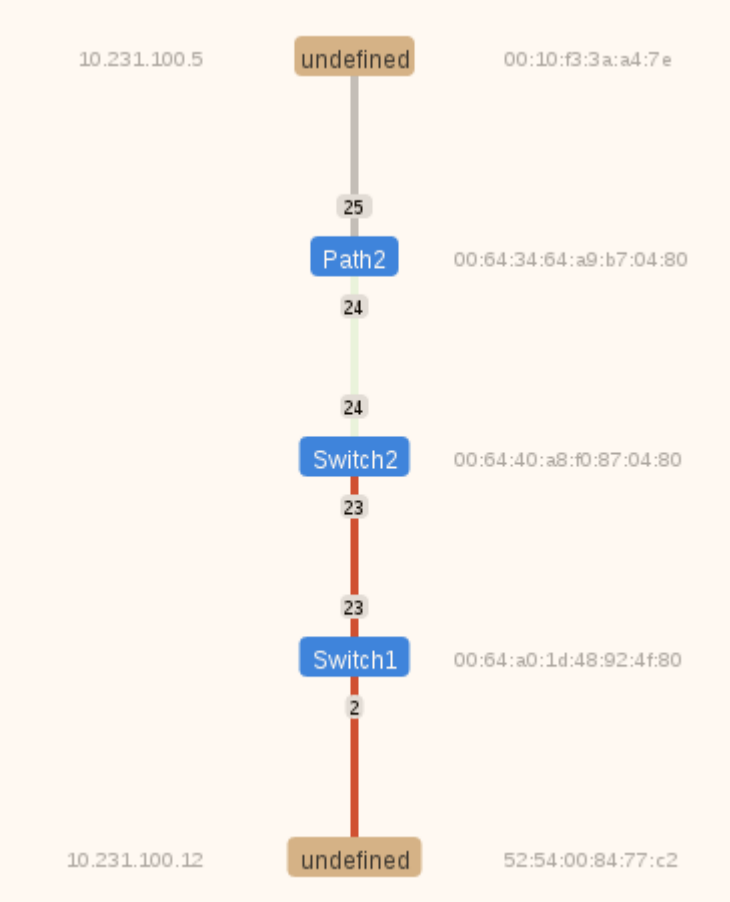

This screen will display the LoadMaster, Real Server and any switches in between. The LoadMaster and Real Server are represented in brown. The LoadMaster is at the top and the Real Server is at the bottom.

The switches are represented in blue. The switch name will appear in the blue boxes if the SDN Controller picks it up.

The Data Path Identifier (DPID) of each switch on the network is displayed on the right of the switches. The DPID is how the controller identifies the different switches.

The Media Access Control (MAC) address of the LoadMaster and Real Server are displayed to the right of those devices. The IP address of the LoadMaster and Real Server will also be displayed on the left.

The colour of the paths are explained below:

- Light green: Traffic is idle and the link is healthy.
- Red: The path is congested with traffic.
- Grey: The path between the LoadMaster and initial switch is shown as grey.

So, in the example screenshot above - the path between the **Path2** and **Switch2** switches is healthy but the

paths between **Switch2** and **Switch1** and the Real Server are congested.

The colour of the path may change as the path gets more or less congested. There is an array of red colours that can be displayed - the darker the red colour is, the more congestion is on the path.

7

Real Servers

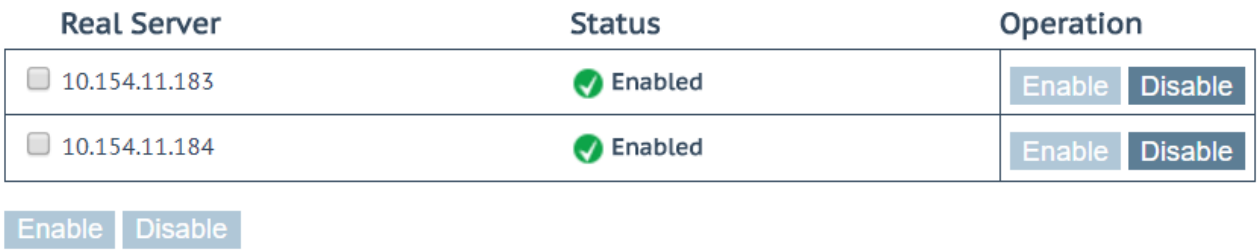

This screen shows the current status of the Real Servers and gives the option to **Disable** or **Enable** each Real Server. Each Real Server has corresponding buttons to disable (take an online server offline) and enable the Real Server. You can also enable or disable multiple Real Servers at the same time by selecting the relevant Real Servers, and clicking the relevant button at the bottom.

The Real Server status may be in one of the following statuses:

Status	Description
Enabled	The Real Server is enabled.
Disabled	The Real Server is disabled.

Status	Description
	The Real Server passed a single health check but failed another health check. The **Ambiguous** status does not show in the Real Servers screen but it does show in the **Virtual Services > View/Modify Services** and **Statistics > Real Time Statistics > Real Servers** screens.
	The Real Server is enabled on one Virtual Service but disabled on another Virtual Service.

Note: **Caution** Disabling a Real Server will disable it for all Virtual Services configured to use it. If it is the only Real Server available (that is, the last one) the Virtual Service is effectively down and will not pass any traffic.

Note: When you enable or disable a Real Server on the **Real Servers** screen, it gets enabled or disabled for all Virtual Services and SubVSs (regardless of what Virtual Services and SubVSs the Real Server was previously enabled or disabled for).

Real Servers that have DNS names assigned to them appear above/below Real Servers without DNS names. You can sort the list of Real Servers by clicking the **Real Server** or **Status** column headings.

8

Rules & Checking

Refer to the following sections for details on the WUI options relating to **Content Rules** and **Check Parameters**.

Contents

- Content Rules
- Check Parameters

Content Rules

Refer to the following sections for details on the LoadMaster WUI options relating to **Content Rules**.

Contents

- Content Matching Rules
- Content Matching
- Add Header
- Delete Header
- Replace Header
- Modify URL
- Replace String in Response Body
- Limitations of Using Regular Expressions in the LoadMaster WUI

Content Matching Rules

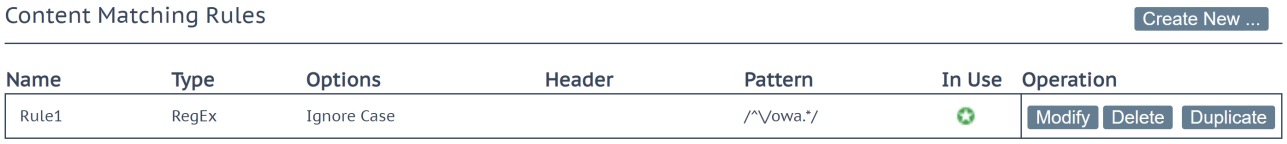

This screen shows rules that have been configured and gives the option to **Modify** or **Delete**.

In LoadMaster firmware 7.2.52 and above, it is possible to **Duplicate** content rules. There is also an **In Use** column on the Content Rules page that indicates if the content rule is in use or not:

- The star icon means the content rule is not assigned to any Virtual Services.



- The tick icon means the content rule is assigned to at least one Virtual Service. The number of assigned Virtual Services is displayed next to the tick icon. Hover over the tick icon to get details about the Virtual Services this content rule is assigned to. The hover text only displays the first 20 assigned Virtual Services.

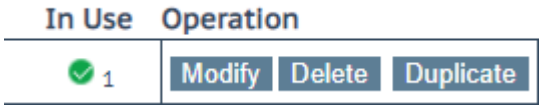

To define a new rule, click the **Create New** button. You must give the rule a name.

The **Rule Name** must be alphanumeric, unique, and cannot contain spaces. Rules are case sensitive, thus two different rules can exist in the form of **Rule1** and **rule1**. It is not possible to name a content rule **default**.

The options that are available depend on the **Rule Type** that you select. The available rules are as follows:

Rule Types:

- **Content Matching:** matches the content of the header or body
- **Add Header:** adds a header according to the rule
- **Delete Header:** deletes the header according to the rule
- **Replace Header:** replaces the header according to the rule
- **Modify URL:** changes the URL according to the rule
- **Replace String in Response Body:** replaces text in the body according to the rule

For further information on configuring rules, please refer to the Content Rules Feature Description document.

Content Matching

When the **Rule Type** selected is **Content Matching** the following describes the options available.

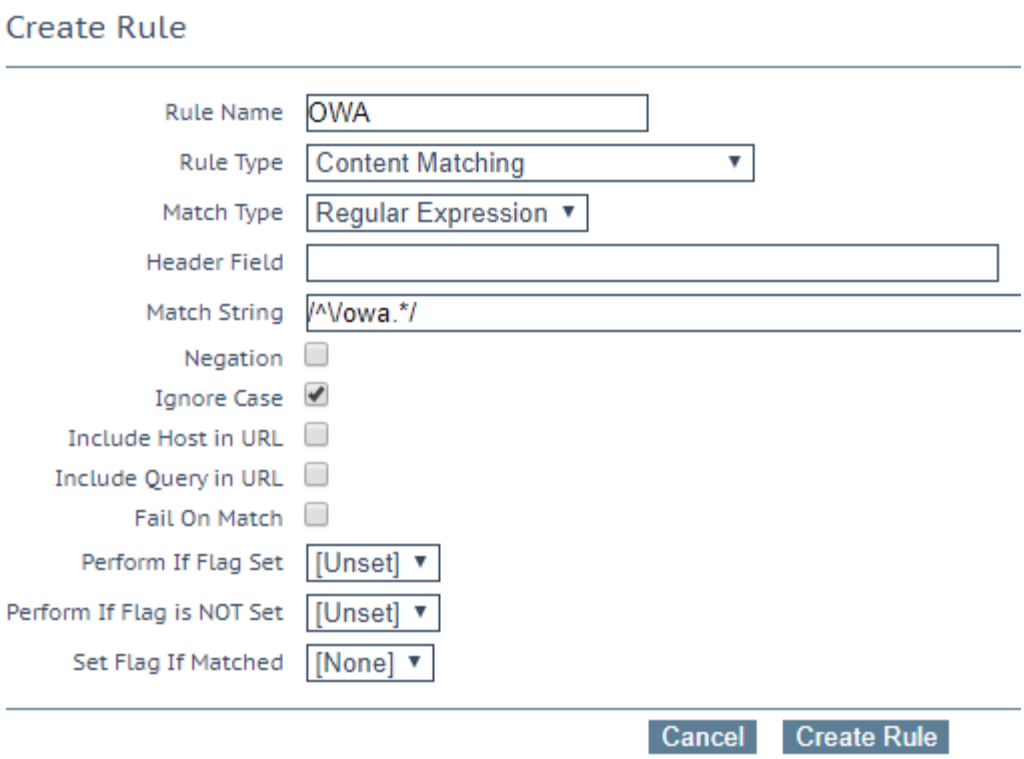

Rule Name

The name of the rule.

Match Type:

- **Regular Expression:** compares the header to the rule

Note: If you use quotes in regular expressions in the LoadMaster WUI, there are limitations. For more information, refer to the section Limitations of Using Regular Expressions in the LoadMaster WUI.

- **Prefix:** compares the prefix of the header according to the rule
- **Postfix:** compares the postfix of the header according to the rule

Header Field

The header field name must be matched. If no header field name is set, the default is to match the string within the URL.

Rules can be matched based on the Source IP of the client by entering **src-ip** within the **Header Field** text box. The header field is populated by the source IP of the client.

Similarly, rules can also be matched based on the HTTP Method used, for example GET, POST or HEAD. The methods that are to be matched should be written in uppercase.

The body of a request can also be matched by typing **body** in the **Header Field** text box.

Match String

Input the pattern that is to be matched. Both Regular Expressions and PCRE are supported. The maximum number of characters allowed is 250.

For further information on Regular Expressions and PCRE, please refer to the Content Rules Feature Description.

Negation

Invert the sense of the match.

Ignore Case

Ignore case when comparing strings.

Include Host in URL

Prepend the hostname to request URL before performing the match.

Include Query in URL

Append the query string to the URL before performing a match.

Fail On Match

If this rule is matched, then always fail to connect.

Perform If Flag Set

Only try to execute this rule if the specified flag is set.

Perform If Flag is NOT Set

Only try to execute this rule if the specified flag is not set.

Set Flag If Matched

If the rule is successfully matched, set the specified flag.

Using the **Perform If Flag Set** and **Set Flag If Matched** options, it is possible to make rules dependent on each other, that is, only execute a particular rule if another rule has been successfully matched. For more detailed instructions on ‘chaining’ rules, please refer to the Content Rules Feature Description.

Add Header

When the **Rule Type** selected is **Add Header** the following describes the options available.

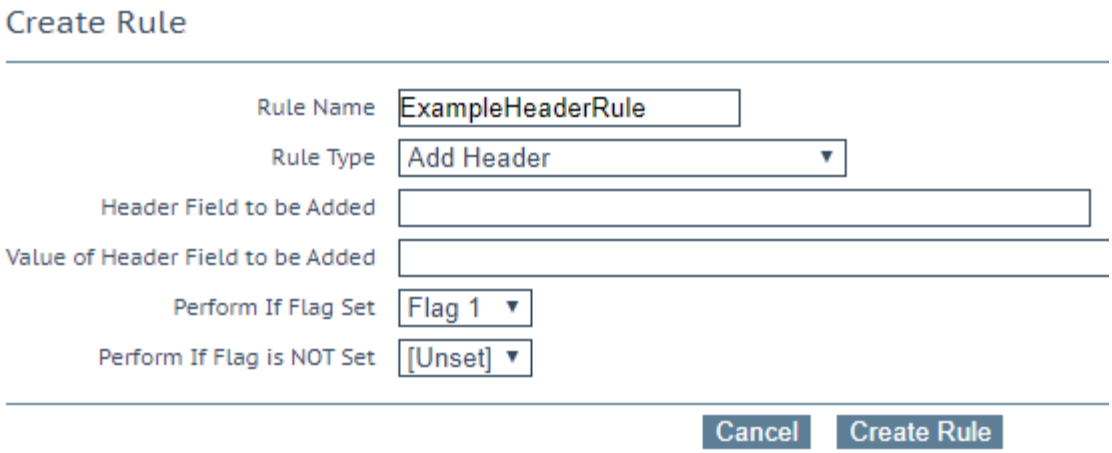

Rule Name

This is a text box to enter the name of the rule.

Header Field to be Added

This is a text box to enter the name of the header field to be added.

Value of Header Field to be Added

Enter the value of the header field to be added. You can enter a maximum of 255 characters in this field.

Perform If Flag Set

Only execute this rule if the specified flag is set.

The flag is set by a different rule. For further information on flags please refer to the Content Matching section.

Perform If Flag is NOT Set

Only try to execute this rule if the specified flag is not set.

Delete Header

When the **Rule Type** selected is **Delete Header** the following describes the options available.



Rule Name

This is a textbox to enter the name of the rule.

Header Field to be Deleted

This is for a text box to enter the name of the header field to be deleted.

Perform If Flag Set

Only execute this rule if the specified flag is set.

The flag will have been set by a different rule. For a more detailed explanation of flags please refer to the Content Matching section.

Perform If Flag is NOT Set

Only try to execute this rule if the specified flag is not set.

Replace Header

When the **Rule Type** selected is **Replace Header** the following describes the options available.

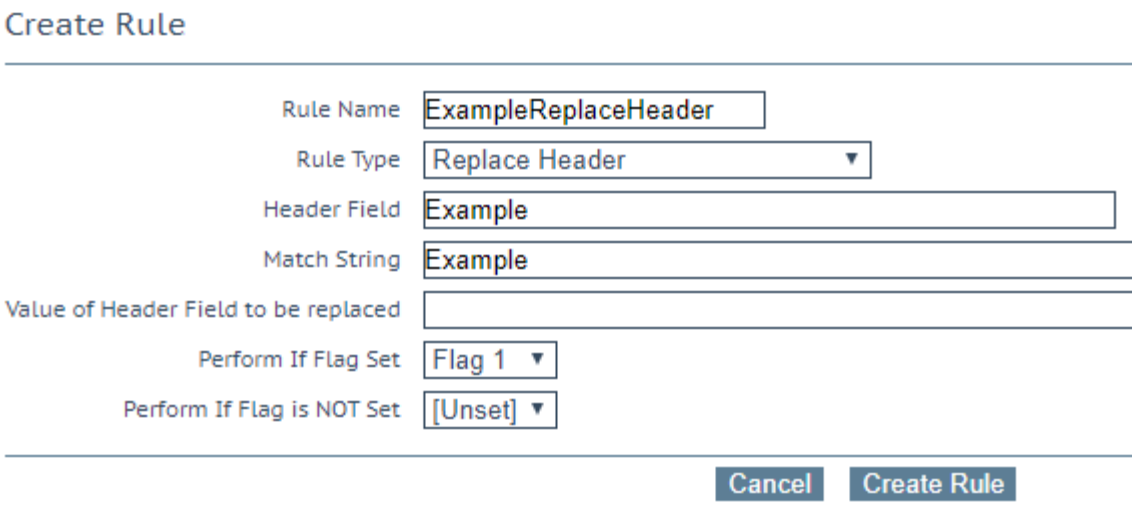

Rule Name

This is for a textbox to enter the name of the rule.

Header Field

This is for a textbox to enter the header name field where the substitution should take place.

Match String

The pattern that is to be matched.

Value of Header Field to be replaced

This is for a textbox to enter the value of the header field to be replaced.

Perform If Flag Set

Only execute this rule if the specified flag is set.

The flag is set by a different rule. For a more detailed explanation of flags please refer to the Content Matching section.

Perform If Flag is NOT Set

Only try to execute this rule if the specified flag is not set.

Modify URL

When the **Rule Type** selected is **Modify URL** the following describes the options available.

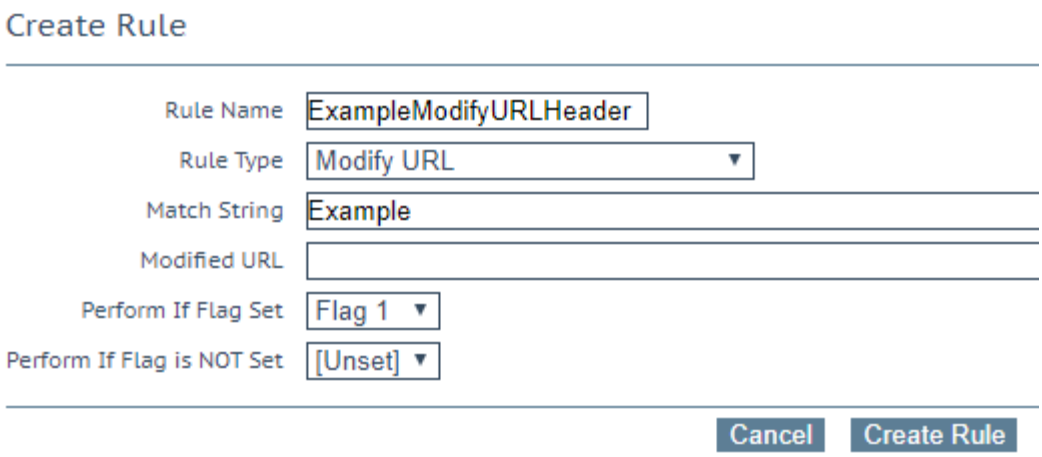

Rule Name

This is for a textbox to enter the name of the rule.

Match String

This is a textbox to enter the pattern that is to be matched.

Modified URL

This is a textbox to enter the URL that is to be modified.

Perform If Flag Set

Only execute this rule if the specified flag is set.

The flag is set by a different rule. For a more detailed explanation of flags please refer to the Content Matching section.

Perform If Flag is NOT Set

Only try to execute this rule if the specified flag is not set.

Replace String in Response Body

When the **Rule Type** selected is **Replace String in Response Body**, the following options available.

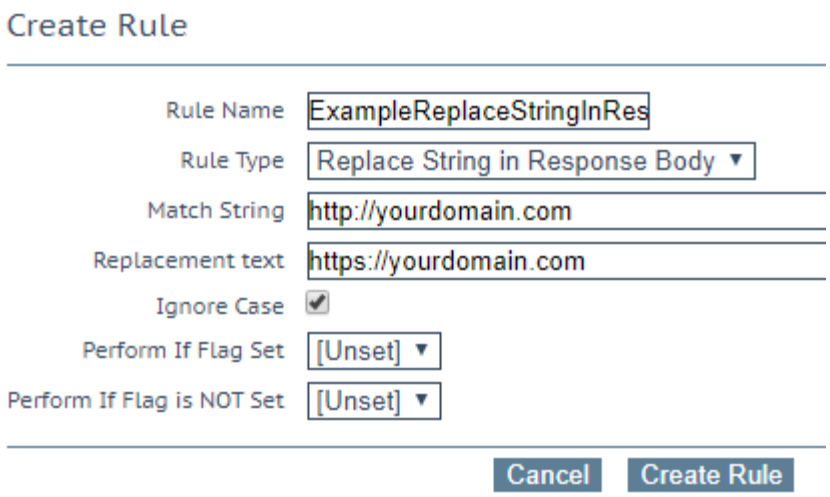

Rule Name

The name of the rule. The rule name must be unique.

Match String

The string to match.

Replacement text

The replacement string.

Ignore Case

Enable this check box to ignore the case of the strings when comparing.

Perform If Flag Set

Only execute this rule if the specified flag is set.

The flag will have been set by a different rule.

Perform If Flag is NOT Set

Only try to execute this rule if the specified flag is not set.

Limitations of Using Regular Expressions in the LoadMaster WUI

When using regular expressions in the LoadMaster WUI, you must use an even number of quotes (single or double). The quotes must also nest correctly, for example, if you use single quotes within double quotes, the single quotes must be matched inside the double-quotes. To use a single " (double-quote) character in a regex, use **\22** instead (or **\27** for a single quote).

If you want to use an uneven number of quotes in a regular expression, use the API instead of the WUI.

For example, trying to set the following **Match String** in the WUI results in an error that says **Please specify a pattern to be matched**:

/\<img([^\>\/]*)\ssrc\=\"([^\"]*)\"([^\>\/]*)\/?>/

However, it is possible to set this using the API, for example:

```
/access/
addrule?name=Example&pattern=/\<img([^\>\/]*)\ssrc\=\"([^\"]*)\"([^\>\/]*)\/?>/
```

Check Parameters

To access the **Check Parameters** screen, go to **Rules & Checking > Check Parameters** in the main menu of the LoadMaster WUI. The **Check Parameters** screen has two sections - **Service Check Parameters** and either **Adaptive Parameters** or **SDN Adaptive Parameters**, depending on the **Scheduling Method** selected in the Virtual Services. If the **Scheduling Method** is set to **resource based (adaptive)**, the **Adaptive Parameters** section is displayed. If the **Scheduling Method** is set to **resource based (SDN adaptive)**, the **SDN Adaptive Parameters** section is displayed.

Refer to the relevant section below to find out more information.

Contents

- Service (Health) Check Parameters
- Adaptive Parameters
- SDN Adaptive Parameters

Service (Health) Check Parameters

The LoadMaster utilizes Layer 3, Layer 4 and Layer 7 health checks to monitor the availability of the Real Servers and the Virtual Services.

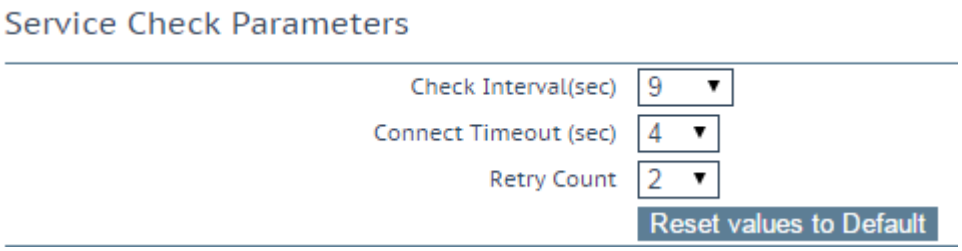

Check Interval(seconds)

This field specifies the number of seconds that will pass between consecutive health check attempts and defines the length of the 'health check cycle'.

Recommended and default value: 9 seconds

Valid values range from the <*mininterval*> (9) to the <*maxinterval*> (901).

The <*mininterval*> is **Retry Count** * **Connect Timeout (seconds)** + 1, that is, a maximum value of 9 by default.

The <*maxinterval*> is 901 [because that is what 60 (maximum **Connect Timeout (seconds)**) * 15 (maximum **Retry Count**) + 1 is].

In the WUI, if the value of **Check Interval** is over 120 (because it is forced to this value by setting the **Connect Timeout (seconds)** and **Retry Count**), it cannot be changed by modifying the **Check Interval(seconds)** drop-down list. To change it, configure the other two options. Otherwise, the maximum value that the interval can be set to is 120.

Connect Timeout (seconds)

This is the allowed maximum wait time for a connection to be established and for the receipt of a reply to a health check. If no reply is received before the timeout expires, then the connection is re-attempted as long as the Retry Count has not been reached. If any HTTP response is received in response to the health check, then no retries are attempted in the current health check cycle.

Default value:4 seconds

Valid values range from **4** to **60**.

Retry Count

This specifies the number of times the health check will be re-attempted if the timeout above is reached before receiving a response. To override the global retry count value, you can select any other value from the drop-down list. The retry count does not apply if any response is received in response to the health check. If, for example, either a 200 or a 404 response is received, then no retries are attempted in the current cycle.

Default value: 2

Valid values range from **2** to **15**.

In LoadMaster firmware version 7.2.52, the check **Interval**, **Timeout**, and **Retry Count** settings can be configured on each Virtual Service or SubVS. Previously, these were just global settings. For further details, refer to the section Real Servers.

Adaptive Parameters

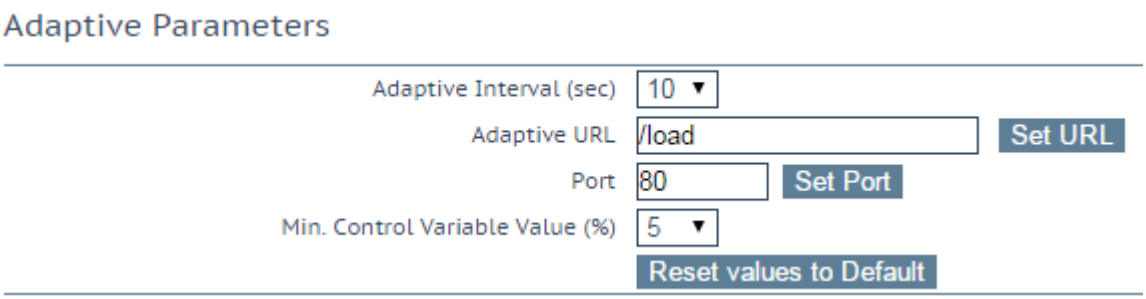

The global adaptive parameters are set in **Rules & Checking > Check Parameters**. By default, these values are used when the **Scheduling Method** of a Virtual Service is set to **resource based (adaptive)**. However, you can configure these adaptive parameters on a per-Virtual Service basis, if needed (in **Virtual Services > View/Modify Services > Modify > Standard Options**).

Adaptive Interval (sec)

This is the interval, in seconds, at which the LoadMaster checks the load on the servers. A low value means the LoadMaster is very sensitive to load, but this comes at a cost of extra load on the LoadMaster itself. **10** seconds is a good starting value. This value must not be less than the HTTP checking interval.

Adaptive URL

The Adaptive method retrieves load information from the servers using HTTP inquiry. This URL specifies the resource where the load information of the servers is stored. This resource can be either a file or program (for example Adaptive Agent) that delivers this information. The standard location is **/load**. It is the servers' job to provide the current load data in this file in ASCII format. In doing so, the following must be considered:

An ASCII file containing a value in the range of 0 to 100 in the first line where: 0=idle and 100=overloaded. As the number increases, that is, the server becomes more heavily loaded, the LoadMaster will pass less traffic to that server. Hence, it 'adapts' to the server loading.

If the server becomes 101% or 102% loaded, a message is added to the logs.

The file is set to "**/load**" by default.

The file must be accessible using HTTP.

The URL must be the same for all servers that are to be supported by the adaptive method.

Note: This feature is not only of interest for HTTP-based Virtual Services, but for all Services. HTTP is merely used as the transport method for extracting the application-specific load information from the Real Server.

Port

This value specifies the port number of the HTTP daemon on the servers. The default value is **80**.

Note: Port **443** will not work.

Min. Control Variable Value (%)

This value specifies a threshold below which the balancer will switch to static weight-based scheduling, that is, normal Weighted Round Robin. The value is a percentage of the maximum load (0-50). The default is **5**.

SDN Adaptive Parameters

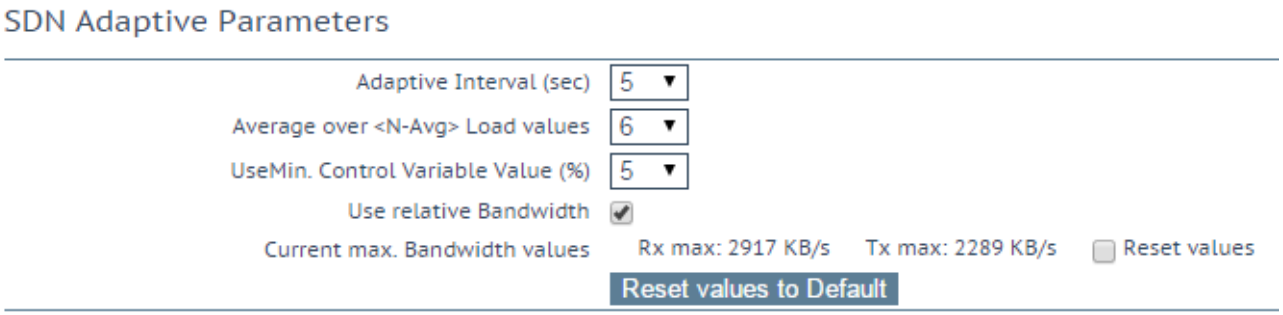

Adaptive Interval (sec)

When using SDN-adaptive scheduling, the SDN controller is polled to retrieve the loading values for the Real Server. This field value specifies how often this occurs.

Average over <N-Avg> Load values

Use this value to dampen fluctuations in the system.

UseMin. Control Variable Value (%)

Anything below the value set here is considered idle traffic and it does not affect the adaptive value (which is displayed on the Real Servers **Statistics** screen), for example - in the screenshot above anything below 5% is considered idle.

Use relative Bandwidth

Use the maximum load observed on the link as link bandwidth. We recommend enabling this option.

Current max. Bandwidth values

This section displays the current received and transmitted maximum bandwidth values.

Reset values

This checkbox can be used to reset the current max. bandwidth values.

9

Certificates & Security

The sections below describe the various screens in the **Certificates & Security** section of the LoadMaster WUI.

Contents

- SSL Certificates
- Intermediate Certificates
- ACME Certificates
- Generate CSR (Certificate Signing Request)
- Backup/Restore Certs
- Cipher Sets
- Remote Access
- Admin WUI Access
- OCSP Configuration
- LDAP Configuration
- Intrusion Detection Options (IPS/IDS)
- SSL Options

SSL Certificates

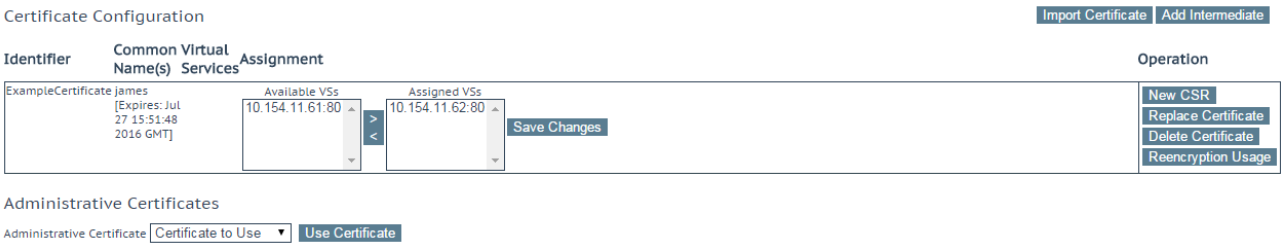

Shown above is the **Manage Certificates** screen. Details about the various options on this screen are below:

Import Certificate – to import the certificate with a chosen filename.

Add Intermediate – refer to the Intermediate Certificates section for further information.

Identifier – is the name given to the certificate at the time it was created.

Common Name(s) – is the FQDN (Fully Qualified Domain Name) for the site.

Virtual Services – the Virtual Service with which the certificate is associated.

Assignment – lists of available and assigned Virtual Services

Operations –

- **New CSR** – generates a new Certificate Signing Request (CSR) based on the current certificate.

Note: If the certificate has Subject Alternative Names (SANs), generating a CSR in this way will not add the SANs. Instead, generate the CSR manually. For further information on this, refer to the Generate CSR (Certificate Signing Request) section.

- **Replace Certificate** – updates or replaces the certificate stored in this file.
- **Delete Certificate** – deletes the relevant certificate.

Note: You cannot delete or replace Let's Encrypt/DigiCert certificates from the **SSL Certificates** screen. You can only delete or replace Let's Encrypt/DigiCert certificates from **Certificates & Security > ACME Certificates**. The **Replace Certificate** and **Delete Certificate** buttons are grayed out on the **SSL Certificates** screen for Let's Encrypt/DigiCert certificates.

- **Reencryption Usage** – display the Virtual Services that are using this certificate as a client certificate when re-encrypting.

Administrative Certificates – the certificate you want to use, if any, for the administrative interface.

The LoadMaster supports key sizes higher than 2048 bit. However, increasing the key size reduces the SSL Transactions Per Second (TPS) performance non-linearly. That means that performance with a 4096 bit key will drop substantially (by at least a power of four) compared to a 2048 bit key. To achieve the same performance with larger keys, more powerful hardware is needed.

However, as indicated by the National Institute of Standards and Technology (NIST); 2048 bit keys have a security lifetime until 2030:

"In many cases, a variety of key sizes may be available for an algorithm. For some of the algorithms (e.g., public key algorithms, such as RSA), the use of larger key sizes than are required may impact operations, e.g., larger keys may take longer to generate or longer to process the data. However, the use of key sizes that are too small may not provide adequate security."

Intermediate Certificates

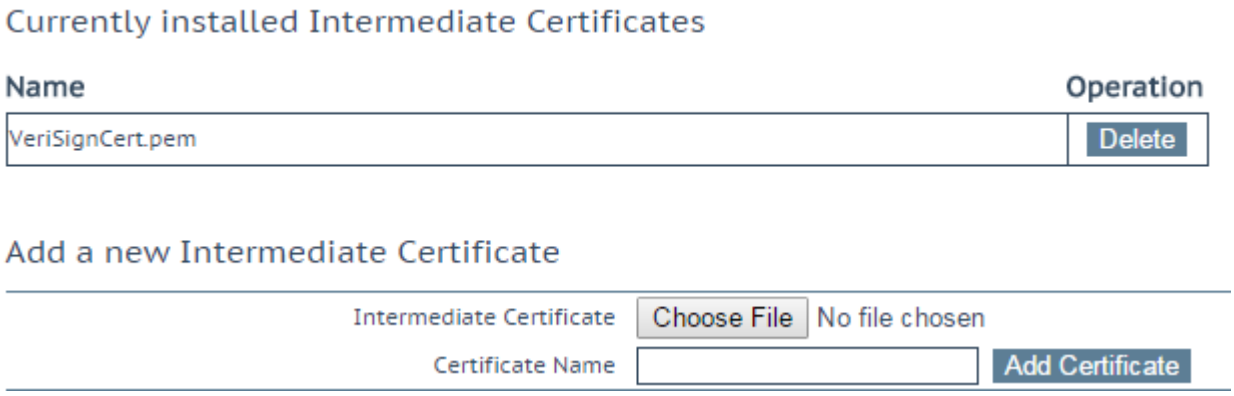

This screen shows a list of the installed intermediate certificates and the name assigned to them.

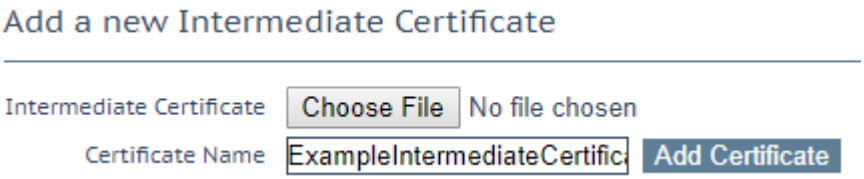

If you already have a certificate, or you have received one from a CSR, you can install the certificate by clicking the **Choose File** button. Navigate to and select the certificate and then enter the desired **Certificate Name**. The name can only contain alpha characters with a maximum of 32 characters.

Uploading several consecutive intermediate certificates within a single piece of text, as practiced by some certificate vendors such as GoDaddy, is allowed. The uploaded file is split into the individual certificates.

ACME Certificates

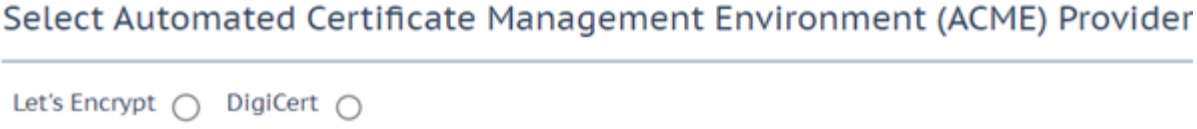

The LoadMaster provides an option of two Automated Certificate Management Environment (ACME) providers:

- Let's Encrypt
- DigiCert

Refer to the relevant section below for further details.

For detailed information related to DigiCert, refer to the DigiCert Feature Description.

Contents

- Let's Encrypt Certificates

Let's Encrypt Certificates

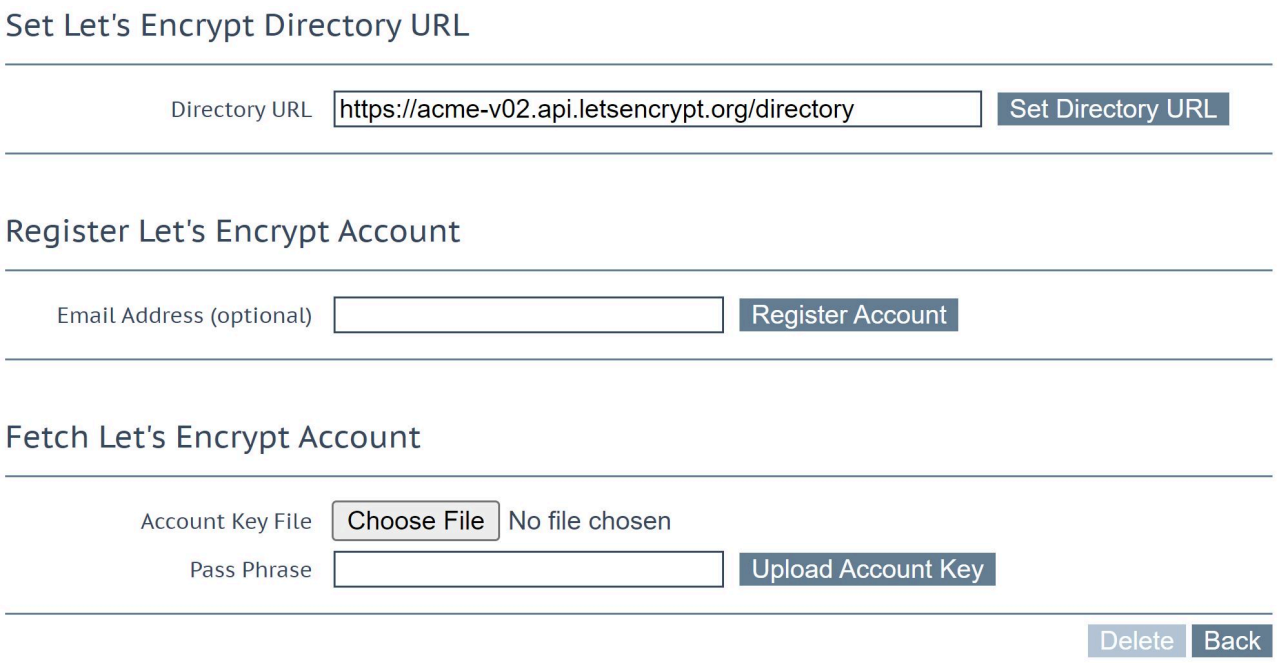

Directory URL: Enter the URL of the Automated Certificate Management Environment (ACME) server in the **Directory URL** field and click **Set Directory URL**. The default URL is the Let's Encrypt production ACME server: **https://acme-v02.api.letsencrypt.org/directory**. This can be changed as needed. The LoadMaster supports API version 2 of the ACME protocol.

Email Address (optional): You can register for Let's Encrypt account by optionally entering your **Email Address** and clicking **Register Account**.

Account Key File: If you already have an existing Let's Encrypt account, you can upload the **Account Key File** by clicking the **Choose File** button. Navigate to and select the key file. You can retrieve the account key file from other ACME clients that you registered the account with (like Certbot).

Pass Phrase: Enter the passphrase associated with the certificate and click **Upload Account Key** to link to your existing account.

Once you have successfully registered or linked to your existing Let's Encrypt account, the **Manage Let's Encrypt Certificates** screen appears.

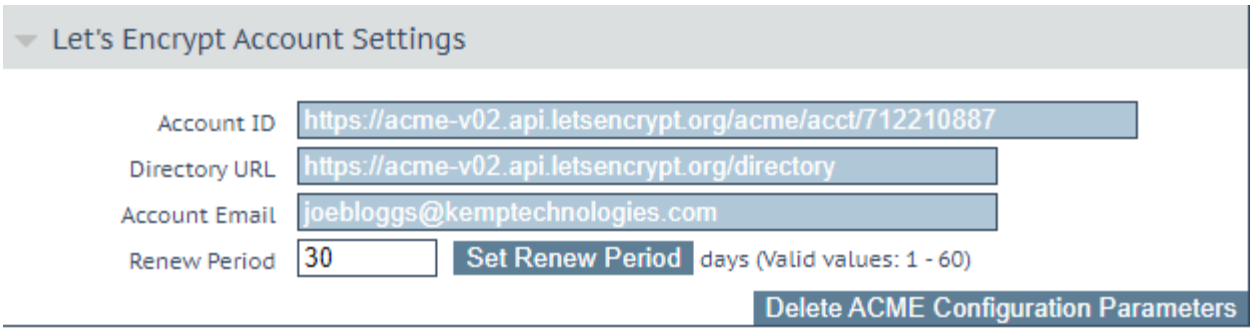

Renew Period

Let's Encrypt certificates are valid for 90 days. The **Renew Period** value specifies how many days in advance of certificate expiry you would like the certificate to be renewed. The **Renew Period** is an account-wide setting. Per-certificate renewal periods are not supported at this time.

The **Renew Period** is set to 30 days by default. Let's Encrypt recommends renewing certificates 30 days before expiry. Valid values for the **Renew Period** field range from 1 to 60 (days). The old certificates are replaced and assigned to the HTTPS Virtual Service when the renewal is successful.

You can click **Delete ACME Configuration Parameters** to remove the ACME account settings.

For more information and instructions, refer to the Let's Encrypt Feature Description.

Request New Certificate

Click **Request New Certificate** to request a new certificate from the Let's Encrypt CA.

All fields on the **Request a New Certificate** screen are optional except for **Certificate Identifier** and **Common Name** (and you must select a Virtual Service next to the **Common Name** field).

Certificate Identifier: Enter a unique identifier. The **Certificate Identifier** value must be unique for all certificates on the LoadMaster.

Common Name: Enter the FQDN of your web server. This is case sensitive. Certificates are only issued to valid hosting domains that you have control over. Select the Virtual Service that is used for this domain. This will be used for the validation challenge to prove ownership of the domain.

Note: A HTTP/HTTPS Layer 7 Virtual Service must be already configured with the ability to add a SubVS (so it should not have any Real Servers added to the parent Virtual Service - but if there are existing SubVSs they can have Real Servers attached). For instructions on how to convert an existing Virtual Service with Real Servers attached to one with SubVSs with Real Servers attached, refer to the Let's Encrypt Feature Description.

Note: A HTTP Redirect Virtual Service must be configured to redirect all port 80 requests to 443 because Let's Encrypt communicates on port 80 to perform the HTTP-01 challenge.

Note: All valid Virtual Services that meet the criteria are listed in the drop-down list.

---

2 Letter Country Code: Optionally enter the two-letter country code. For a list of valid country codes, refer to the following page: https://www.digicert.com/kb/ssl-certificate-country-codes.htm. If using Let's Encrypt, the **2 Letter Country Code** to **Email Address** fields are truncated.

State/Province: Optionally enter the state or province to include in the certificate. Enter the full name, for example **New York** (not NY).

City: Optionally enter the city to include in the certificate.

Company: Optionally enter the name of the company to include in the certificate.

Organization: Optionally enter the department or organizational unit that should be contacted regarding this certificate.

Email Address: Optionally enter the email address of the person or organization that should be contacted regarding this certificate.

Generate Elliptic Curve Request: Optionally enable or disable this option. If this is enabled, an Elliptic Curve request is generated instead of an RSA request.

Key Size: Select the algorithm size from the drop-down list. If you are generating an Elliptic Curve (EC) request, the Key Size drop-down is grayed out. The default size of 256 Bits is used for EC requests. If you are generating an RSA request, you can specify the Key Size.

SAN/UCC Names: Enter the Subject Alternate Name (SAN). This must be a valid domain. You can specify up to 10 SANs.

For every SAN you must select a HTTP/HTTPS Layer 7 Virtual Service (you can use the same Virtual Service). For each SAN you must prove your authority to the Let's Encrypt server. A HTTP/HTTPS Virtual Service must be already configured with the ability to add a SubVS (so it should not have any Real Servers added to the parent Virtual Service - but if there are existing SubVSs they can have Real Servers attached). For instructions on how to convert an existing Virtual Service with Real Servers attached to one with SubVSs with Real Servers attached, refer to the Let's Encrypt Feature Description.

Request Certificate: When you are finished setting the relevant fields, click **Request Certificate** to create a new certificate request using the specified data.

A list of issued certificates and related details are displayed at the bottom of the **Let's Encrypt Certs** screen. The **HTTP Challenge VS(s)** column lists the Virtual Service (or Services) that were used for the HTTP challenge. These are not the Virtual Services that the certificates are assigned to.

Once the certificate is issued successfully, it will be listed in **Certificates & Security > SSL Certificates**. You can then assign it to any HTTPS Virtual Service or use it as an administrative certificate.

Note: When manually assigning a new certificate to a Virtual Service for the first time, the Virtual Service will restart so we recommend doing this outside of working hours.

When Let's Encrypt certificates are renewed, the Virtual Services that have the certificate assigned will be automatically updated with the renewed certificate.

Note: Automatic renewal and updating of certificates is seamless and does not affect Virtual Service traffic.

Certificates are automatically renewed at the number of days specified in the **Renew Period** before the expiry date of each certificate. You can manually renew the certificate by clicking **Renew Certificate**.

You can also delete a certificate associated with the domain by clicking **Delete Certificate**.

Note: If the certificate is used (for example if it is assigned in a Virtual Service or used as an administrative certificate) the **Delete Certificate** button is grayed out.

You cannot delete or replace Let's Encrypt certificates from the **SSL Certificates** screen. You can only delete or replace Let's Encrypt certificates from the **Let's Encrypt Certs** screen. The **Replace Certificate** and **Delete Certificate** buttons are grayed out on the **SSL Certificates** screen for Let's Encrypt certificates.

Note: If you downgrade the LoadMaster from version 7.2.53 (or above) to 7.2.52 (or below), any Let's Encrypt certificates that exist at the time of downgrade are preserved in the downgraded system so that Virtual Service connectivity is not inadvertently affected by the downgrade. However, the automatic certificate management functionality is not available in earlier versions. These certificates are listed on the **SSL Certificates** page and can be deleted after the downgrade, if needed.

How to Request a Wildcard Certificate

Click **Request New Certificate** to request a wildcard certificate from the Let's Encrypt Certificate Authority (CA).

The **Common Name** field supports the use of a wildcard character. For example, ***.example1.com** matches anything that ends in **.example1.com**. Once you enter a **Common Name** beginning with *., the DNS provider **Select DNS API** drop-down becomes available.

All fields on the **Request a New Certificate** screen are optional except for **Certificate Identifier** and **Common Name**. You must select a Virtual Service and DNS provider (including related credential parameters) next to the **Common Name** field.

For wildcard certificate validation, the DNS-01 challenge type is used. This requires the addition and removal of temporary DNS records. For automatic DNS record updates during wildcard name validation, you must select your DNS provider from the **Select DNS API** drop-down list. You must also set the access credential parameters for the selected DNS provider. The fields to fill out vary depending on the selected DNS provider. A list of DNS providers is as follows:

- Progress-LM-GEO
- CloudFlare
- GoDaddy.com
- Azure-DNS
- DNSMadeEasy
- DigitalOcean
- NS1.com
- Amazom-Route53

When you are finished setting the relevant fields, click **Request Certificate** to create a new certificate request using the specified data. It can take approximately 25 seconds to generate the certificate request. If the request fails, you must fill out the form again.

A list of issued certificates and related details are displayed at the bottom of the **Manage ACME Certificates** screen.

Generate CSR (Certificate Signing Request)

If you do not have a certificate, you may complete the Certificate Signing Request (CSR) form and click the **Create CSR** button. CSRs generated by the LoadMaster use SHA256.

Note: If **Self-Signed Certificate Handling** is set to **EC certs with an EC signature** (in **Certificates & Security > Remote Access**), CSR generation is restricted to the administrative (**bal**) user only. If **Self-Signed Certificate Handling** is set to a different value, all users (regardless of their permissions) can generate CSRs.



2 Letter Country Code (ex. US)

The 2 letter country code that should be included in the certificate, for example **US** should be entered for the United States.

State/Province (Entire Name – New York, not NY)

The state which should be included in the certificate. Enter the full name here, for example **New York**, not NY.

City

The name of the city that should be included in the certificate.

Company

The name of the company which should be included in the certificate.

Organization (e.g., Marketing,Finance,Sales)

The department or organizational unit that should be included in the certificate.

Common Name

The Fully Qualified Domain Name (FQDN) for your web server.

Email Address

The email address of the responsible person or organization that should be contacted regarding this certificate.

SAN/UCC Names

A space-separated list of alternate names.

Generate Elliptical Curve Request

By default, CSRs generated by the LoadMaster request an RSA-encrypted key. If you enable the **Generate Elliptical Curve Request** option, the LoadMaster instead requests an EC (Elliptical Curve) key. Smaller EC key sizes generally provide the same cryptographic strength as much larger RSA key sizes. EC keys are becoming increasingly common because of both the reduced storage footprint in addition to reduced processing resources required.

Display Private Key

This new option (introduced in LoadMaster firmware version 7.2.52 and LTS version 7.2.48.3) appears only when the **Certificates & Security > Remote Access > Self-Signed Certificate Handling** option is set to **EC certs with an EC signature** which means that an elliptical curve cipher is used for both the certificate and the digital signature.

Once the above option is selected, a **Display Private Key** check box appears on the **Certificates & Security > Generate CSR** WUI page.

- When **Display Private Key** is disabled (the default), the private key is not displayed in the WUI after the CSR is created. The unsigned CSR is downloaded by the user as in previous releases. Once it is signed by a Certificate Authority, the user uploads the signed certificate to the LoadMaster - the difference from previous releases being that the user does not have to also upload the private key, since LoadMaster maintains it internally when **Display Private Key** is disabled. If the saved private key matches the new certificate, the certificate gets imported and the saved private key is deleted. The stored private key is not encrypted but there is no access to it from the outside and it cannot be seen or displayed.
- When **Display Private Key** is enabled, the LoadMaster behaves as in previous releases: the private key is displayed to the user and must be uploaded to LoadMaster along with the private key.

There is only one private key per machine and it is not shared between High Availability (HA) pairs. This means the newly-generated certificate must be installed on the machine that the CSR was generated on.

Alter clicking the **Create CSR** button, the following screen appears:

The following is your 2048 bit *unsigned* certificate request. Copy the following, in its entirety, and send it to your trusted certificate authority

```
-----BEGIN CERTIFICATE REQUEST-----
MIIC9zCCAd8CAQAwgbExCzAJBgNVBAYTAlVTMREwDwYDVQQIEwhOZXcgWW9yazER
MA8GA1UEBxMITmV3IFlvcmsxGjAYBgNVBAoTEUtFTVAgVGVjaG5vbG9naWVzMR0w
GwYDVQQLExRLbm93bGVkZ2UgTWFuYWdlbWVudDEUMBIGA1UEAxMLRXhhbXBsZS5j
b20xKzApBgkqhkiG9w0BCQEWHGpibG9nZ3NAa2VtcHRlY2hub2xvZ2llcy5jb20w
ggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQC+ohZjEwKEQT3jd6y9gN7k
Snu8E0T8bhA1LuGCD5mN++uC+3Vm4r5m6g5pVS16RF4QaRqkuiaekz5QPWqMV06b
yxveeIhoq1HPVphPOEHBHd1iotC4SLoRJ6/A0vWdlRIjlJVJfe7ka6S60xaVgAog
6lVohNoDtC2RHJOwFvawBhEZh2YzzpuoPSmDoZRnuX8qD9DZN1c9sSKn3YjomY5O
2KRyJmFEII98N8sMmiPATvXYZZCrTUifu2nwfpR9ogx7KVyK7Mi/73P41zDJdN4T
1GM0FMxYehg9bNXL27wkUek4994izLpyrv4whSc9QCbfd1BXz6IdxuFbpMJbMdVx
AgMBAAGgADANBgkqhkiG9w0BAQsFAAOCAQEANRw07oaxj+B6/t+KTMHTVWzzXFDF
79HHQj7ROFtqkw+FfijKEAfBhfNAfOpmRQEC6tWySb70K1acBn2fCI2lr9stsUUC
bq+w4Xl/crsVs+mc+veQ+p3R3zHlNPU1mZ6sofOQUi1E8NbCRUtdZ+6ixxLZL0ah
Y7aN9Ipn5qy2sT/yfYHao4rJWuzLXuKaphqyc1JNWvPkFI/4tDbrdD5rgPZfCdDY
PDOxuN2g6244Htfkn9ZCqfkatGyTI9qVnPsidqapKUAVZ4Zk1j+W7zNFGmw2cXK5
Ff97URaPLwEI+VQrVlbaJgN3/eMzLrvDB/OFD2LCv+9xk+KhAPsiDwvxJQ==
-----END CERTIFICATE REQUEST-----
```

The following is your private key. Copy the following, in its entirety, and save as a .key file. Do this using a text editor such as Notepad or VI (Do not use Microsoft Word - extra characters will be added making the key unusable). Key will later be used during the certificate upload process. **DO NOT** lose or distribute this file!

```
-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQEAvqIWYxMChEE943esvYDe5Ep7vBNE/G4QNS7hgg+Zjfvrgvt1
ZuK+ZuoOaVUtekReEGkapLomnpM+UD1qjFdOm8sb3niIaKtRz1aYTzhBwR3dYqLQ
uEi6ESevwNL1nZUSI5SVSX3u5GukutMWlYAKIOpVaITaA7QtkRyTsBb2sAYRGYdm
M86bqD0pg6GUZ7l/Kg/Q2TdXPbEip92I6JmOTtikciZhRCCPfDfLDJojwE712GWQ
q01In7tp8H6UfaIMeylciuzIv+9z+NcwyXTeE9RjNBTMWHoYPWzVy9u8JFHpOPfe
Isy6cq7+MIUnPUAm33dQV8+iHcbhW6TCWzHVcQIDAQABAoIBAQCt/fLA6pDZdVKv
UoNvUzgc1X6p4kyMuUhBwlBBDUvxs4T5P9mf1kRCWk5dBUlE1zGjeMrAnsaw5WNy
iRu+i9FLkM4W95xJLFS3ESpi483gHQn7BO/Lw1VQYxCexe03rt+nae337eEkyrrH
afKq8PpNoJPjmZ4C02jjkVma1trBPLHBhJOzJ/oT5QtpDu0W+I5ysZriUUo1IOPi
1VzkEl1T08oqZRTJSqIbx12akk3C9QCuA/F+BiGF6Tn76epHmPYGuYykoaAZcjAV
H9ryfkANHtz3B/sRza5lfRmqzTmokeox3sayhf35x6rU68xGSWN5qCr76lRJRx7U
4bjoPxehAoGBAPr+B5lVQyuQ0Gih5fysbqX2suDX2SEM1m55Ts+xuKrog7kc36xY
xTivObfZFuE6ERQhxmGjuD8ZsVhN6giL5PMSDnvFmIL3vg4ja90zAxHKgoR2kpph
IuGfT0UOf/3+ZSTUjfLr/OEZD9uiVRBPpHeH58iWtZJ2YqmqJzMV0193AoGBAMJv
xFK1RZG7MMVXQ1JFYrk+C5A5VG80VVdYh0K+XNv6ThSHk1XqOrrIkcXzhY1qU14o
IuaSqO5+BAsbmJgx9LZlCE5xqHqHtl934WFF4GlBNcBhP9UR6ApnAtQwinWA+8k0
Ii/kAOkRAyAa2ENcT4gF/UdM38lhoid7QSw2B7xXAoGBAIJZs7Caa0wQ5WuxyTO0
ibJ/sN68uvNDK4osThXngrSgFOjqae+kGqkZt6wXfp5x/bSq5dCHqoR633Ow4z6V
CM6ELilxsYczCu1kz/wNJibzOV16ByFOGUN77Ts8EJTkrbq2+RGUJbzxux6h6/OQ
qSW621F9k8cA3LSovbr2NtR5AoGAYDI7x0+346nhL0FFJWb+uPdhcTFr/Li/oD9E
bFkSSCNGjhGla1Q/SjoBJRaedKCuLl9dJQZaXeQqy/QTQvk0QSkrOuQwnq6WJBWD
hES2Cl0g4tU6Z4g8bSkZ1TFOz2PJLnqEj30Wlji8ex3M8UaycnHEJYp7DX8oYrAw
RldU7HECgYBXd4o2+E6pNLiy7uoXXCyIZdHqapMt+MAaiFmg5cCggXbnbY3ftuxH
LDpMa6kZ/Yz10x2Uuji0QXvuh2wL1HlGCB+wJ8GgBI85FtIzaFht70WdR2HzhXY2
m1/Rl5hgtsEBdLLDg9DEN27Pr8LnTtF+7RfRFFVDWbOeDVlm+sqigQ==
-----END RSA PRIVATE KEY-----
```

The top part of the screen should be copied and pasted into a plain text file and sent to the Certificate Authority of your choice. They will validate the information and return a validated certificate.

The lower part of the screen is your private key and should be kept in a safe place. This key should not be disseminated as you will need it to use the certificate. Copy and paste the private key into a plain text file (do not use an application such as Microsoft Word) and keep the file safe.

Backup/Restore Certs

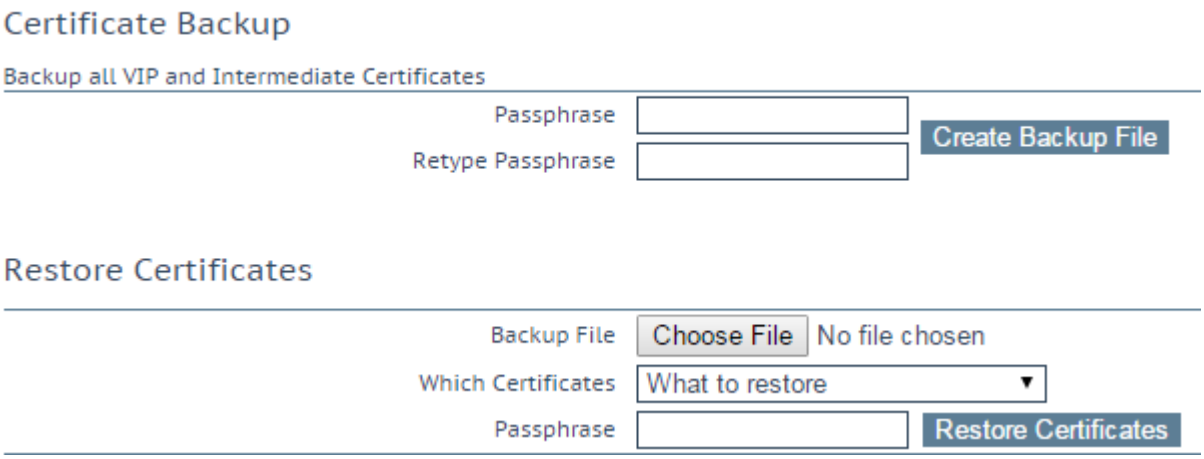

Backup all VIP and Intermediate Certificates: When backing up certificates, you are prompted to enter a mandatory passphrase (password) twice. The parameters of the passphrase are that it must be alpha-numeric and it is case sensitive with a maximum of 64 characters.

Note: This passphrase is a mandatory requirement to restore a certificate. You cannot restore a certificate without the passphrase. If you forget it, there is no way to restore the certificate.

Backup File: select the certificate backup file

Which Certificates: select which certificates you wish to restore

Passphrase: enter the passphrase associated with the certificate backup file

Cipher Sets

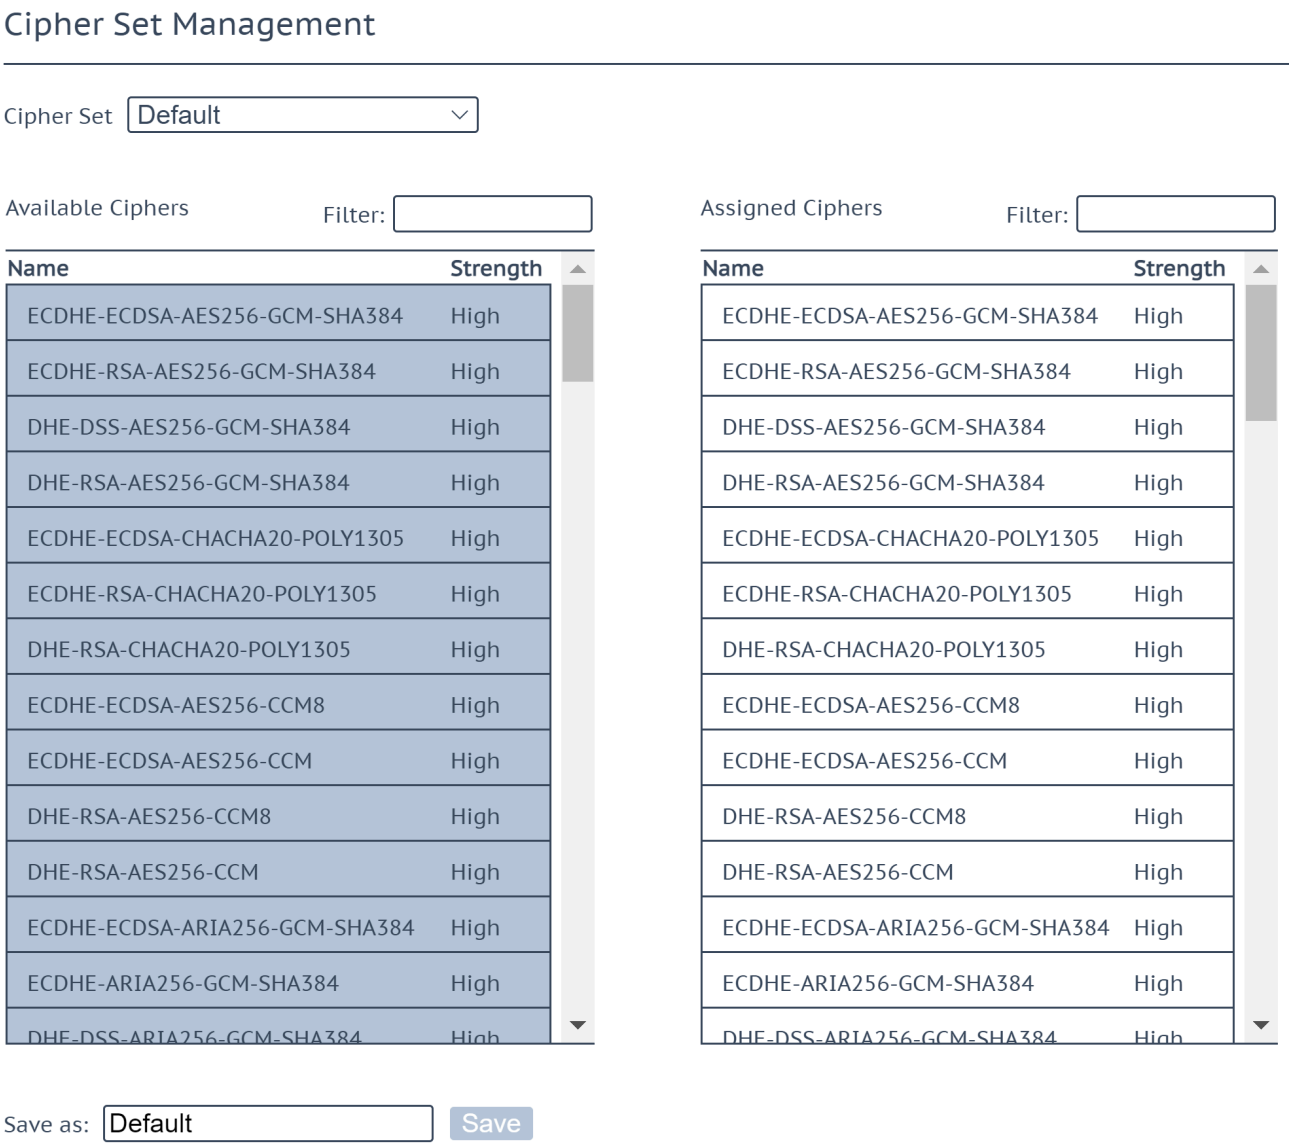

Cipher Set

Select the cipher set to view/modify.

The system-defined cipher sets are as follows:

- **Default:** The cipher set that is configured on the LoadMaster on a fresh installation. This cipher set is geared towards backwards compatibility with previous releases of the LoadMaster.
- **Default_NoRc4:** A more secure version of the default set that does not contain any RC4 ciphers, which are considered to be insecure on modern networks.
- **BestPractices:** This is the recommended cipher set to use on the LoadMaster and it is updated occasionally to reflect the current industry best practices. It does not include older and legacy cipher sets which may be required by older browser and application deployments. The last update to the

BestPractices set was made in LoadMaster version 7.2.60.0. Please see the LoadMaster Release Notes for more information.

- **Intermediate_compatibility:** This cipher set includes some ciphers that are required by older browser and service implementations that are still seen in the field.
- **Backward_compatibility:** This cipher set provides maximum backward compatibility for clients back to Windows XP/IE6 at the risk of using less secure ciphers.

Note: The **Backward_compatibility** cipher set should be used as a last resort only.

- **WUI:** This is the default cipher set used by the administrative user interface. It can be changed by using the controls under **Certificates & Security > Admin WUI Access**.
- **FIPS:** This set contains only ciphers that conform to Federal Information Processing Standards (FIPS) 140-2 level 1 standard and should be used only in those deployments that require it.
- **Legacy:** This cipher set is provided solely for upgrade compatibility for legacy LoadMaster firmware versions (v7.0-10 and previous). After upgrade to a modern version of LoadMaster, it is recommended to choose a more secure cipher set.
- **Null_Ciphers:** This cipher set contains what are called 'null ciphers', which do not provide any cryptographic protection, but rather depend on the application to provide it. In general, use these ciphers only if required by the application and if that application provides independent cryptographic protection.
- **ECDSA_Default:** This cipher set includes only cipher sets that use elliptical curve cryptography and is recommended for those deployments that require EC cryptography.
- **ECSDA_BestPractices:** This is a modified version of the ECDSA_Default set that includes only those ciphers that conform to the Common Criteria standards.

To find out what ciphers are in each cipher set, go to **Certificates & Security > Cipher Sets**. Select the relevant **Cipher Set**.

Note: Progress Kemp reserves the right to change the contents of these cipher sets at any time in response to changes in industry security standards and best practices.

Two lists are displayed – **Available Ciphers** and **Assigned Ciphers**. These lists can be filtered by typing some text into the **Filter** text boxes provided. The **Filter** text boxes will only allow you to enter valid text which is contained in the cipher names, for example **ECDHE**. If invalid text is entered, the text box will turn red and the invalid text is deleted.

Ciphers can be dragged and dropped to/from the **Available** and **Assigned** lists as needed. Ciphers which are already assigned will appear grayed out in the **Available Ciphers** list.

Changes cannot be made to a pre-configured cipher set. However, you can start with a pre-configured cipher set – make any changes as needed and then save the cipher set with a new custom name. Enter the new name in the **Save as** text box and click the **Save** button. Custom cipher sets can be used across different Virtual Services and can be assigned as the WUI cipher set.

Note: When you create or modify a custom cipher set, the LoadMaster restarts the SSL services to apply the updated configuration. As part of this process, all affected Virtual Services are temporarily stopped and restarted.

This behavior is expected and can cause a brief interruption to client connections while the configuration change is applied. To minimize impact in production environments, we recommend:

- Creating and testing new cipher sets on a non-production LoadMaster first.
- Scheduling cipher set changes during a maintenance window.

It is not possible to delete pre-configured cipher sets. However, custom cipher sets can be deleted by selecting the relevant custom cipher set and clicking the **Delete Cipher set** button.

Remote Access

The sections below describe the different areas within the **Remote Access** screen in the LoadMaster WUI.

Contents

- Administrator Access
- GEO Settings
- GEO Partners Status
- WUI Authentication and Authorization
- Partner Communications

Administrator Access

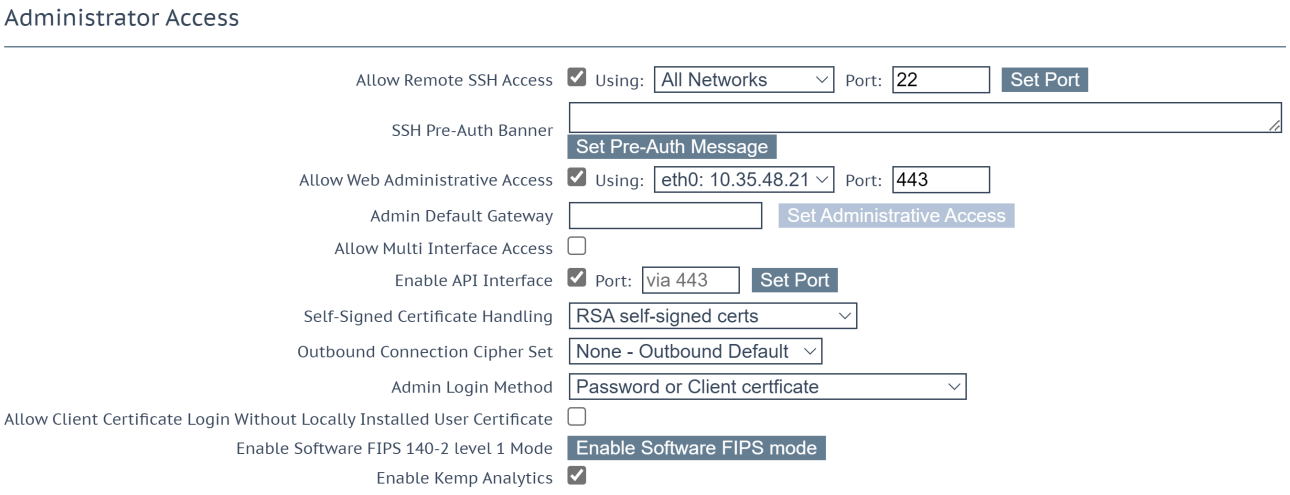

Allow Remote SSH Access

You can limit the network from which clients can connect to the SSH administrative interface on LoadMaster.

Using

Specify which addresses that remote administrative SSH access to the LoadMaster is allowed.

Note: Only the 'bal' user has rights to access the LoadMaster using SSH.

Note: As of LoadMaster firmware version 7.2.48.4 Long Term Support (LTS) and 7.2.53, RSA keys are no longer supported for SSH access to the LoadMaster. If you are using RSA keys for SSH access and you are upgrading to one of these versions (or above) you must move to another key type. There are two key types available as an alternative to the RSA keys: ecdsa-sha2-nistp384 and ssh-ed25519

Port

Specify the port used to access the LoadMaster using the SSH protocol.

SSH Pre-Auth Banner

Set the SSH pre-authentication banner, which is displayed before the login prompt when logging in using SSH. This field accepts up to 5,000 characters.

Allow Web Administrative Access

Selecting this check box allows administrative web access to the LoadMaster. Disabling this option will stop access upon the next reboot. Click **Set Administrative Access** to apply any changes to this field.

Note: Disabling web access is not recommended.

Note: It is critical to the security of your appliance that you use a dedicated network interface for management traffic; refer to the LoadMaster Hardening Technical Note for further details.

Using

Specify the addresses that administrative web access is to be permitted. Click **Set Administrative Access** to apply any changes to this field. You need to reconnect to the WUI using the new address after the change is applied.

Port

Specify the port used to access the administrative web interface. Click **Set Administrative Access** to apply any changes to this field. You need to reconnect to the WUI using the new port after the change is applied.

Admin Default Gateway

When administering the LoadMaster from a non-default interface, this option allows the User to specify a different default gateway for administrative traffic only. Click **Set Administrative Access** to apply any changes to this field.

Allow Multi Interface Access

Enabling this option allows the WUI to be accessed from multiple interfaces. When this option is enabled, a new option appears in each of the interface screens (**System Configuration > eth<n>**) called **Allow Administrative WUI Access**. When both of these options are enabled, the WUI can be accessed from the IP address of the relevant interface(s) and any **Additional addresses** configured for that interface. Click **Set Administrative Access** to apply any changes to this field.

Note: The certificate used by default to secure WUI connections specifies the initial WUI IP address, and so will not work for WUI connections on other interfaces. If you enable the WUI on multiple interfaces, you will need to install a wildcard certificate for the WUI. For more information on certificates, refer to the SSL Accelerated Services, Feature Description.

Note: Enabling the WUI on multiple interfaces can have a performance impact on the system. There is a maximum of 64 network interfaces that can be tracked. There are a maximum of 1024 total addresses where the system will listen on.

RADIUS Server

Here you can enter the address of the RADIUS server that is to be used to validate user access to the LoadMaster. To use a RADIUS server, you have to specify the **Shared Secret**.

A **Shared Secret** is a text string that serves as a password between the LoadMaster and the RADIUS server.

The **Revalidation Interval** specifies how often a user should be revalidated by the RADIUS server.

RADIUS Server Configuration

To configure RADIUS to work correctly with the LoadMaster, authentication must be configured on the RADIUS server and the RADIUS Reply-Message must be mapped to LoadMaster permissions.

The Reply-Message values correspond to LoadMaster permissions as shown in the table below.

Reply-Message	LoadMaster Permission
real	Real Servers
vs	Virtual Services
rules	Rules
backup	System Backup
certs	Certificate Creation
cert3	Intermediate Certificates

Reply-Message	LoadMaster Permission
certbackup	Certificate Backup
users	User Administration
geo	GEO Configuration

The values in the Reply-Message should map to the user permissions page in the WUI as per Figure 119, with the exception of “All Permissions”:

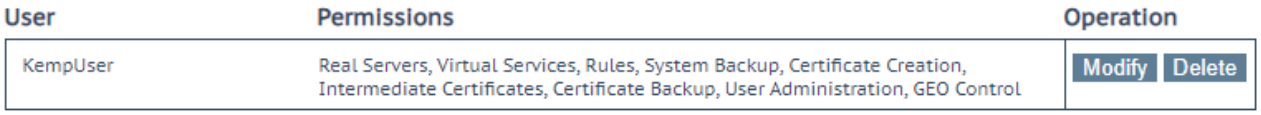

To configure the Linux FreeRADIUS server, please insert the text below into the /etc/freeradius/users file in the sections indicated within the file. The example below is to configure permissions for the user ‘LMUSER’.

LMUSER Cleartext-Password := "1fourall"Reply-Message = "real,vs,rules,backup,certs,cert3,certbackup,users"

The /etc/freeradius/clients.conf file must also be configured to include the LoadMaster IP address. This file lists the IP addresses that are allowed to contact RADIUS.

Note: When Session Management is enabled, the **RADIUS Server** options are not available within this screen. Please refer to the WUI Authentication and Authorization section for further information on how to configure RADIUS Server when Session Management is enabled.

Enable API Interface

Enables/disables the RESTful Application Program Interface (API). You can also specify the port used to access the API interface. If the port is unset, you can access the API over the web interface port. Some things to note are as follows:

- If you try to use the API on a port other than one on which its running, the LoadMaster returns a HTML 404 (not found) response.
- If you try to use the WUI on the port configured specifically for the API, an unreadable page and/or 404 response is displayed (depending on the browser used).
- You can set the API port value to an empty string which will unset the value. If the API port is not set, the WUI port is used.

Self-Signed Certificate Handling

Select the type of self-signed certificates that the system will use. The options are described below:

- **RSA self-signed certs:** By default, these are RSA certificates that are signed with the Progress Kemp RSA root certificate.
- **EC certs with a RSA signature:** The LoadMaster can generate an EC certificate also signed by the original RSA Progress Kemp root certificate.
- **EC certs with an EC signature:** The LoadMaster can generate an EC certificate signed by the Progress Kemp EC root certificate. In this mode, any CSRs generated will also be EC.

Note: If **Self-Signed Certificate Handling** is set to **EC certs with an EC signature**, CSR generation is restricted to the administrative (**bal**) user only. If **Self-Signed Certificate Handling** is set to a different value, all users (regardless of their permissions) can generate CSRs.

Note: If **Self-Signed Certificate Handling** is set to an EC mode and Random Number Generation (RNG) fails (for example, if the hardware does not support it like on legacy systems) a message will display on the home screen saying **Could not start CC mode - system disabled** and the WUI is unusable. This also generates a critical log message saying **Cannot initialize RNG, CC mode disabled** and an authlog saying **Failed to start RNG, CC mode not started**.To get out of this mode, you must use the isetup menu (either using the console or SSH). Navigate to **Local Admin > Web Address > Confirm switch out of CC mode**. This option is only shown when the LoadMaster is in this state. This allows the system to work as usual (but not in Common Criteria (CC) mode).

You should not switch from **RSA self-signed certs** to **EC certs with an EC signature** directly. If you do this, connections will fail because there is no EC Progress Kemp Certificate Authority (CA) certificate. To work around , you must first switch from **RSA self-signed certs** to **EC certs with a RSA signature**.

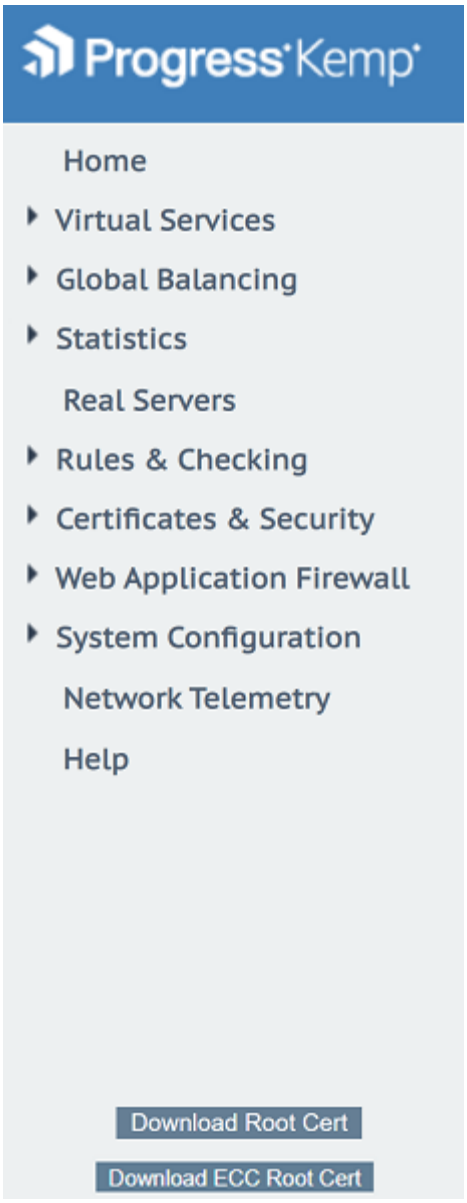

Then, download the new EC Progress Kemp CA certificate by clicking **Download ECC Root Cert** in the bottom-right of the WUI under the main menu after refreshing the page. After you have downloaded the certificate, you can switch to **EC certs with an EC signature** with no loss of connection.

Outbound Connection Cipher Set

This option allows you to select a pre-defined cipher set to use for all outbound connections, including:

- Remote logging (syslog)
- Email notifications
- LDAP authentication
- OCSP certificate validation

- Re-encrypted client requests
- HTTPS health checks

The default setting is **None - Outbound Default**, which means there is no specific cipher set chosen. However, the outbound connection is encrypted and the LoadMaster defaults to all ciphers available for the agreed TLS protocol.

This is global for all outbound connections. For information on each of the cipher sets available, refer to the Cipher Sets section.

Note: The LoadMaster applies Online Certificate Status Protocol (OCSP) stapling (if enabled) to verify certificates for all outbound connections originated by LoadMaster, except for re-encrypted connections to Real Servers.

Admin Login Method

Note: This option will only appear if Session Management is enabled. For further information on Session Management, refer to the Admin WUI Access section or the User Management Feature Description**.**

Specify the login option for access to the LoadMaster WUI. The following options are available:

Note: The **Pre-Auth Click Through Banner** in the **Admin WUI Access** screen must be set for all **Admin Login Method** options to be made available.

- **Password Only Access (default):** This option provides access using the username and password only – there is no access using client certificates.
- **Password or Client certificate:** The user can log in using either the username/password or using a valid client certificate. If a valid client certificate is in place, the username and password is not required.The client will be asked for a certificate. If a client certificate is supplied, the LoadMaster will check for a match. The LoadMaster checks if the certificate is a match with one of the local certificates, or checks if the Subject Alternative Name (SAN) or Common Name (CN) of the certificate is a match. The SAN is used in preference to the CN when performing a match. If there is a match, the user is allowed access to the LoadMaster. This works both using the API and user interface.An invalid certificate will not allow access.If no client certificate is supplied, the LoadMaster will expect that a username and password is supplied (for the API) or will ask the user to enter a password using the standard WUI login page.
- **Client certificate required:** Access is only allowed with the use of a client certificate. It is not possible to log in using the username and password. SSH access is not affected by this (only the **bal** user can log in using SSH).
- **Client certificate required (Verify via OCSP):** This is the same as the **Client certificate required** option, but the client certificate is verified using an OCSP service. The OCSP Server Settings must be configured for this to work. For further information on the OCSP Server Settings, refer to the Cipher Sets section.

Note: In LoadMaster firmware version 7.2.53 and above, the OCSP server settings do not need to be configured in the LoadMaster if the certificate has an Authority Information Access (AIA) extension. The LoadMaster attempts to connect with the provided AIA. For further details on the functionality introduced, refer to the section WUI Authentication and Authorization.

Some points to note regarding the client certificate methods are below:

- The **bal** user does not have a client certificate. Therefore, it is not possible to log into the LoadMaster as **bal** using the **Client certificate required** methods. However, a non-**bal** user can be created and granted **All Permissions**. This will allow the same functionality as the **bal** user.
- There is no log out option for users that are logged in to the WUI using client certificates, as it is not possible to log out (if the user did log out the next access would automatically log them back in again). The session is terminated when the page is closed, or when the browser is restarted.

For further information on client certificate WUI authentication, including step-by-step instructions on how to configure it, please refer to the User Management Feature Description.

Allow Client Certificate Login Without Locally Installed User Certificate

Enabling this option allows client certificate logins for local users even if the client certificate has been deleted from the LoadMaster. By default, this option is disabled. Legacy local certificate login is not secure. Only enable this option if necessary. The **Allow Client Certificate Login Without Locally Installed User Certificate** option is only visible if one of the **Client certificate** values is selected for the **Admin Login Method**. To set the **Admin Login Method** to a **Client certificate** option, you must set a **Pre-Auth Click Through Banner** in **Certificates & Security > Admin WUI Access**.

Note: When enabling this option, a confirmation warning appears. Click **OK** to confirm.

Enable Software FIPS 140-2 level 1 Mode

Note: FIPS mode cannot be enabled if Session Management is disabled. For further information on Session Management, refer to the Admin WUI Access section.

Switch to FIPS 140-2 level 1 certified mode for this LoadMaster. The LoadMaster must be rebooted to activate.

Note: Before switching to FIPS mode, you must export your certificates (**Certificates & Security > Backup/ Restore Certs**). Then, after switching to FIPS mode, reset the UI certificates, restore your certificates and then use the correct certificate that you want for the UI.

Note: A number of warnings will appear before enabling FIPS. If FIPS is enabled on a LoadMaster, it cannot easily be disabled. If FIPS has been enabled and you want to disable it, please contact Progress Kemp Support.



When a LoadMaster is in FIPS level 1 mode - **FIPS-1** will appear in the top-right of the LoadMaster WUI.

FIPS level 1 has a different set of ciphers to a non-FIPS LoadMaster. There is a **Default** cipher set and there are no other system-defined cipher sets to choose from.

Note: If FIPS is enabled, you cannot use RADIUS authentication.

Enable Kemp Analytics

Enables statistical and usage data to be sent to Progress Kemp for analysis. This data is strictly about product usage, enabled capabilities, and statistics. No sensitive user data, or traffic of any kind is either collected or communicated. For more information, visit the following page: Kemp Analytics Disclosure and Usage.

GEO Settings

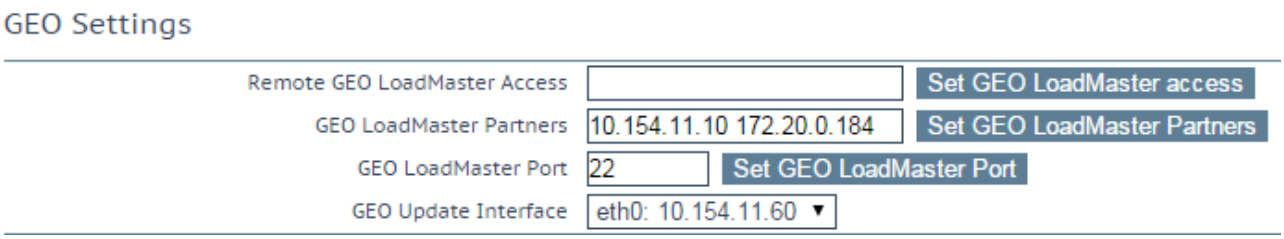

Remote GEO LoadMaster Access

A space-separated list of the IP addresses of GEO LoadMasters that are permitted to retrieve Virtual Service status information from this LoadMaster. When in HA mode, only the shared address needs to be entered.

Note: The **Set GEO LoadMaster access** and **Set GEO LoadMaster Partners** buttons will be greyed out if the **Partner Shared Secret** is not set. For further details, refer to the following section: Partner Communications.

GEO LoadMaster Partners

Note: GEO functionality comes as part of the GSLB Feature Pack and is enabled based on the license that has been applied to the LoadMaster. If you would like to get the GSLB Feature pack, contact Progress Kemp to upgrade your license.

A space-separated list of the IP addresses of GEO LoadMasters that will keep their DNS configuration synchronized with this GEO LoadMaster.

Note: Before partnering GEO LoadMasters, a backup should be taken of the relevant GEO LoadMaster which has the correct/preferred configuration. This backup should then be restored to the other LoadMasters that will be partnered with the original LoadMaster. For more information and step-by-step instructions, refer to the GEO, Feature Description.

Note: Up to 64 GEO HA partner addresses can be added.

GEO LoadMaster Port

The management port used to communicate with all GEO LoadMaster partners.

GEO update interface

The network interface and IP address used to communicate with all GEO LoadMaster partners.

GEO Partners Status

Note: This section is only visible when GEO partners have been set.



A GEO partner status of **Green** indicates the two partners can see each other.

A GEO partner status of **Red** indicates the LoadMasters cannot communicate. The reasons for this include (among other possibilities); one of the partners is powered down, there may be a power outage or a cable may be disconnected.

If there is a failure to update the GEO partner, the logs display an error message saying the GEO update to the partner failed. The message displays the IP address of the partner.

WUI Authentication and Authorization

WUI Authorization Options

Click the **WUI Authorization Options** button on the **Remote Access** screen to display the **WUI Authentication and Authorization** screen. This option is only available when Session Management is enabled.

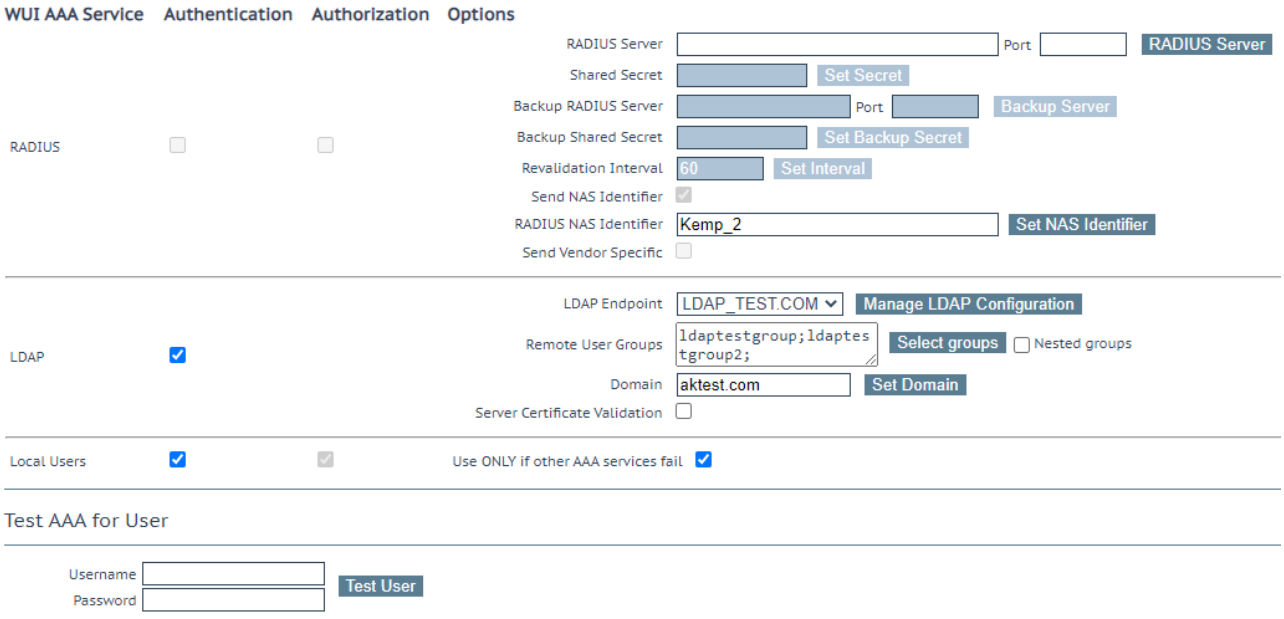

The **WUI Authentication and Authorization** screen enables the administration of the available authentication (login) and authorization (allowed permissions) options relating to LoadMaster WUI logins.

Authentication

Users must be authenticated before logging on to the LoadMaster. The LoadMaster allows authentication of users to be performed using the RADIUS and LDAP authentication methods as well as Local User authentication.

When the **RADIUS**, **LDAP**, and **Local Users** authentication methods are selected, the LoadMaster attempts to authenticate users using the authentication methods in the following order:

1. RADIUS
2. LDAP
3. Local Users

For example, if the RADIUS server is not available then the LDAP server is used. If the LDAP server is also not available, then Local User authentication methods are used.

If neither RADIUS nor LDAP authentication methods are selected, then the Local User authentication method is selected by default.

The consequence of this ordering is that when a local user (that is using a local password) logs in to the LoadMaster and they do not exist in the RADIUS or LDAP configuration, error messages appear in the log for the RADIUS and LDAP authentication checks, even though local password authentication succeeds.

Authorization

LoadMaster allows the users to be authorized by either RADIUS, LDAP, or using Local User authorization. The user's authorization decides what level of permissions the user has and what functions on the LoadMaster they are allowed to perform.

Note: The **RADIUS Authentication** check box must be enabled to use the **RADIUS Authorization** method. Authentication is for access (to ensure the user has a valid username and password) and authorization is used for permissions.

When all authorization methods are selected, the LoadMaster checks authorization in the following order:

1. RADIUS
2. LDAP
3. Local Users

For example, if the RADIUS server is not available then the LDAP server is used. If the LDAP server is also not available, then Local User authorization methods are used.

If neither RADIUS nor LDAP authentication/authorization methods are selected, then the Local User authentication method is selected by default.

You must configure the RADIUS/LDAP server that you are using to authorize the same user permissions that appear on the WUI's user permissions page (with the exception of 'All Permissions'). The Reply-Message returned by the RADIUS server indicates the permissions it is allowing. On a Linux system, the message looks similar to the following:

LMUSER Cleartext-Password := "1fourall"Reply-Message = "real,vs,rules,backup,certs,cert3,certbackup,users"

The preceding example is of a RADIUS user configuration on a RADIUS server deployed on a Linux system. The LoadMaster determines the user's permissions from the "Reply-Message" (the permissions are similar to

the ones for a local WUI user on the LoadMaster).

Note: The **bal** user is always authenticated and authorized using the Local User authentication and authorization methods. Disabling Local User authentication does not lock out the **bal** user. **Bal** is an admin/super user and is allowed to log in to the LoadMaster WUI even if Local User Authentication is disabled on the LoadMaster.

RADIUS Server Configuration

RADIUS Server

The IP address and Port of the RADIUS Server that is to be used to authenticate user WUI access to the LoadMaster.

Shared Secret

This input field is for the Shared Secret of the RADIUS Server.

A Shared Secret is a text string that serves as a password between the LoadMaster and the RADIUS server.

Backup RADIUS Server

The IP address and Port of the backup RADIUS Server that is to be used to authenticate user WUI access to the LoadMaster. This server will be used in case of failure of the main RADIUS Server.

Backup Shared Secret

This text box is to enter the Shared Secret of the backup RADUS Server.

Revalidation Interval

Specifies how often a user should be revalidated by the RADIUS server.

Send NAS Identifier

If this check box is disabled (default), a NAS identifier is not sent to the RADIUS server. If it is enabled, a Network Access Server (NAS) identifier string is sent to the RADIUS server. By default, this is the hostname. Alternatively, if a value is specified in the **RADIUS NAS Identifier** text box, this value is used as the NAS identifier. If the NAS identifier cannot be added, the RADIUS access request is still processed.

Sending the NAS identifier serves two purposes:

- It helps to classify the device type that is sending the request as opposed to simply sending the host IP address which makes troubleshooting and consuming logs easier.
- It enables customized authentication responses to be sent back from the server based on the identifier.

Send Vendor Specific

In LoadMaster firmware version 7.2.51 and above, there is a check box called **Send Vendor Specific** in the User Interface (UI) when a **RADIUS Server** is set. When the **Send Vendor Specific** check box is enabled and a user is logging into the LoadMaster UI using RADIUS authentication with Cisco Access Control Server (ACS) or Identity Services Engine (ISE), the LoadMaster sends an Attribute Value Pair (AVP) to the server as part of the login request which contains Progress Kemp's vendor ID. The server can use this AVP upon receipt to identify the LoadMaster device. The format and requirements for this attribute are in Section 5.26

of RFC 2865.

The Progress Kemp vendor ID is 12196.

RADIUS NAS Identifier

If the **Send NAS Identifier** check box is selected, the **RADIUS NAS Identifier** field is shown. When specified, this value is used as the NAS identifier. Otherwise, the hostname is used as the NAS identifier. If the NAS identifier cannot be added, the RADIUS access request is still processed.

LDAP Endpoint

Select the relevant **LDAP Endpoint** to use. Click the **Manage LDAP Configuration** button to go to the **LDAP Configuration** screen. For further information on LDAP endpoints, refer to the LDAP Configuration section.

In LoadMaster firmware version 7.2.53, support for PIV smart card authentication was added. As a result, a new **Select Certificate to User Mapping** drop-down list was added to the **Certificates & Security > Remote Access > WUI Authorization Options** screen. This field has the following values:

- User Principal Name (default value)
- Subject
- Issuer and Subject
- Issuer and Serial Number

Some configuration caveats are below:

- **Session Management** must be enabled (**Certificates & Security > Admin WUI Access**) to see the **WUI Authorization Options** button.
- The **Admin Login Method** in **Certificates & Security > Remote Access** must be set to a **Client certificate** method to see the new **Select Certificate to User Mapping** drop-down list.
- The Pre-Auth Click Through Banner must be set in **Certificates & Security > Admin WUI Access** before you can select a **Client certificate** method as the **Admin Login Method** in **Certificates & Security > Remote Access**.
- After a certificate is revoked, the certificate fails authentication. However, sometimes it remains in the cache so to make it fail instantly ensure to use the **Flush OCSPD Cache** option in **System Configuration > System Administration > Logging Options > Debug Options**.
- If the LDAP query returns more than one match, the login fails.
- If the Authority Information Access (AIA) is present in the certificate, the LoadMaster attempts to connect with the provided AIA. If this does not work, it tries to connect with the local server.
- If the LoadMaster cannot get the status of the server configured in the certificate AIA, the LoadMaster does not fail back to the local server.
- If the certificate cannot be validated because the server is unavailable, there is an option in **Certificates & Security > OCSP Configuration** called **Allow Access on Server** Failure where you can decide if you want to pass the authentication or not. Enabling this check box treats an OCSP server connection failure

or timeout as if the OCSP server has returned a valid response. That is, the client certificate is treated as valid.

If client users are being authenticated with client certificates, the Common Name (CN) is normalized to lowercase. Therefore, the associated local user entries (with no password), which may be required for permissions, should be in lowercase also.

Remote User Groups

Any remote user groups that are selected are displayed here. To select, clear, or order the groups, click **Select groups**.

Note: It is important to select and apply the group, or groups. If there are no groups selected, no group checking is performed and remote users can log in without a group check.

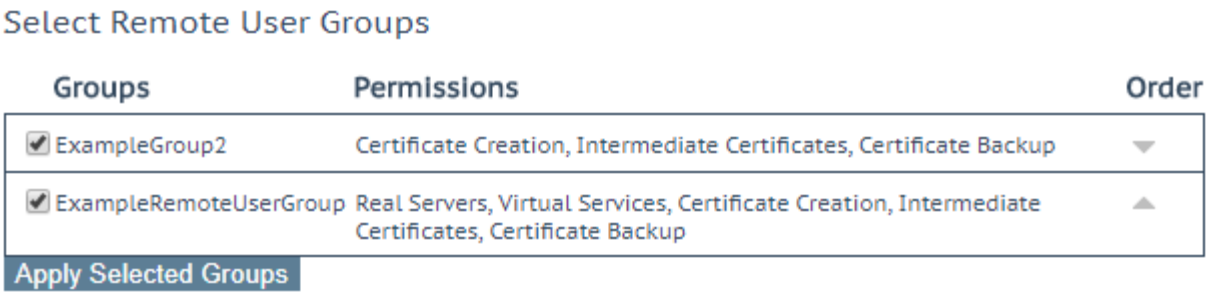

The groups displayed on this screen are taken from the remote user groups set up in **System Configuration > System Administration > User Management**. For more information, refer to the User Management section.

When a user logs in, a check of the user groups on the Active Directory is performed if all the following conditions are met:

- If LDAP WUI Authentication is enabled
- A list of groups is defined
- The user logging in is not locally defined or the **Local Users** option is disabled

You can change the order of the groups on this screen. The first group is checked first. On the first group match, access is enabled and no further groups are checked. If no groups are matched, user access fails and an appropriate log is reported in the syslog. If the user logs in using the group check, the matched group permissions are granted.

Nested Groups

You can enable or disable user nested groups on the WUI Authentication and Authorization screen by using the **Nested groups** check box.

Domain

Specify the domain to use if no domain is provided in the username when group WUI authentication is in use. It is always used as the domain for group search if the Windows logon is used in the format *prefix\username*.

Note: The **Domain** field only appears if a group, or groups, are assigned.

Server Certificate Validation

This check box only appears if **StartTLS** or **LDAPS** is the the **LDAP Protocol** for the selected **LDAP Endpoint**.

Note: There is a known issue that **Server Certificate Validation** does not work with **StartTLS**.

When **Server Certificate Validation** is enabled, it ensures that the host name or IP address that was used to initiate the secure connection resides in the Certificate Subject or Subject Alternative Names (SAN) of the certificate.

Server Certificate Validation is disabled by default.

Local Users Configuration

Use ONLY if other AAA services fail

When selected, the Local Users authentication and authorization methods are used only if the RADIUS and/or LDAP authentication and authorization services fail to respond/time out.

Test AAA for User

To test a user's credentials, enter their username and password in the **Username** and **Password** fields and click the **Test User** button.

A message appears to inform you whether the user is validated or not. This is a useful utility to check a user's credentials without having to log in or out.

Partner Communications



The **Partner Shared Secret** is required to secure communications between partner devices and must be enabled on all High Availability (HA) partners, all LoadMasters in a cluster, and all GEO partners. The **Partner Shared Secret** must be the same on:

- Both units in a HA setup
- All units in a LoadMaster cluster
- All GEO partners
- All remote GEO machines that retrieve Virtual Services from this device

When an incoming shared secret does not match the local **Partner Shared Secret** (including if only one side is providing a shared secret), a warn-level log message is recorded that says **Unauthorized Remote Machine connection from <ClientIPAddress>** and the connection fails.

You can find this secret in the following locations:

- Regular/shared Web User Interface (WUI): **Certificates & Security > Remote Access**
- Local WUI (of a configured HA or cluster unit): **Local Administration > Remote Access**

This secret must have a minimum of 8 and a maximum of 127 characters. The following characters are supported:

- Numeric: 0-9
- Uppercase alphabetic: A-Z
- Lowercase alphabetic: a-z
- Special characters: !"#$%&()*+,-./:;<=>?[\~]^_@`{|}

Admin WUI Access

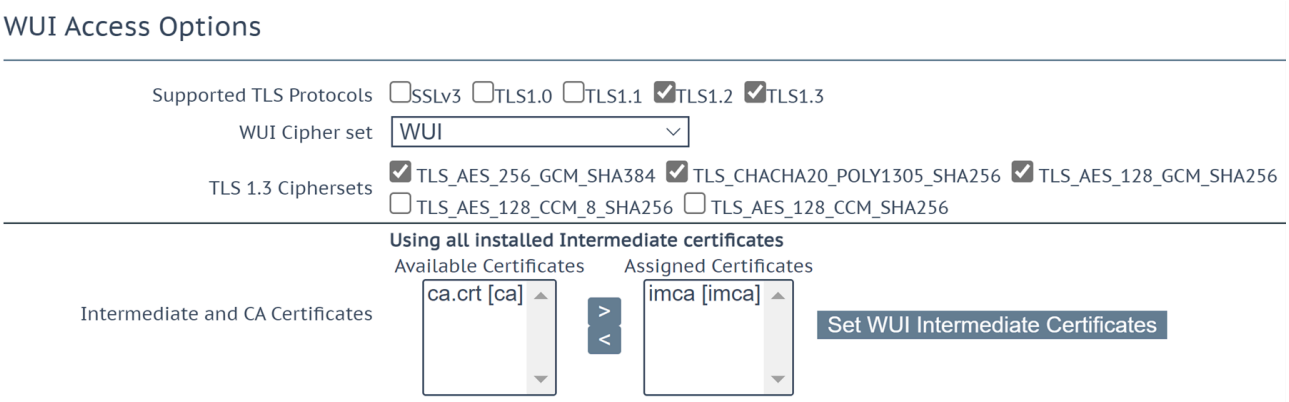

Supported TLS Protocols

Check boxes are provided here which can be used to specify whether or not it is possible to connect to the LoadMaster WUI using the following protocols; SSLv3, TLS1.0, TLS1.1, TLS1.2, or TLS1.3. TLS1.1, TLS1.2, and TLS1.3 are enabled by default. It is not recommended to only have SSLv3 selected because SSLv3 is only supported by some old browsers. When connecting to the WUI using a web browser, the highest security protocol which is mutually supported by both the browser and the WUI will be used.

Note: If FIPS mode is enabled, the only available options are TLS1.1 and TLS1.2.

WUI Cipher set

Select the relevant cipher set to use for WUI access. For information on each of the cipher sets available, refer to the Cipher Sets section.

TLS1.3 Cipher sets

Select the cipher sets for the TLS1.3 protocol to be allowed using any combination of supported ciphers over an Admin WUI access protocols. By default, the following three cipher sets are enabled:

- TLS_AES_256_GCM_SHA384

- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_128_GCM_SHA256

Note: The TLS1.3 cipher sets only become available when the **TLS1.3** protocol is enabled.

Note: TLS1.3 ciphers are always available regardless of whether the selected cipher set contains any TLS1.3 ciphers, which is how the OpenSSL libraries behave. TLS1.3 ciphers can only be removed from offered ciphers by disabling the TLS1.3 protocol (which is not recommended).

CAUTION: Disabling all three default cipher sets may cause loss of connectivity to the LoadMaster.

Intermediate and CA Certificates

The default behavior for UI access authentication is to accept a client certificate validated by any of those in the trusted LoadMaster store. In LoadMaster firmware version 7.2.55, users can now specify the specific Intermediate Certificate to be used to validate the Client Certificate for UI access authentication.

Available certificates are listed in the **Available Certificates** select list on the left. To assign or unassign a certificate, select it and click the right or left arrow button. Then click **Set WUI Intermediate Certificates**. Multiple certificates can be selected by holding **Ctrl** on your keyboard and clicking each required certificate.

Note: In LoadMaster firmware version 7.2.55, support was added for **Intermediate and CA Certificates**. By default, all intermediate certificates are listed under **All Certificates** list used for WUI access. If any certificate is moved from All Certificates list to **Available Certificates** list, then the All Certificates list is renamed to **Assigned Certificates** list. In this scenario the certificates listed in Assigned Certificates list will be used for WUI access.

WUI Session Management

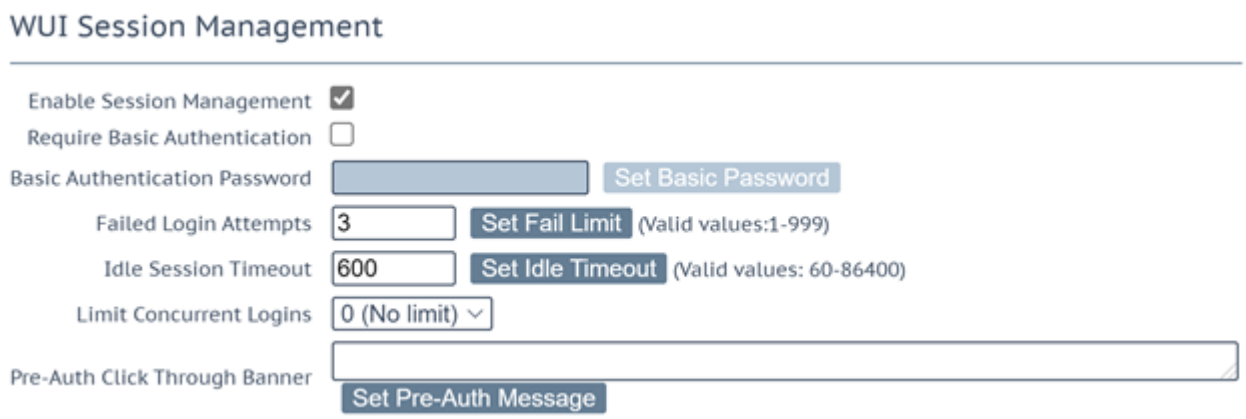

Session management is enabled by default on all LoadMasters initially deployed with LTS firmware versions or above.

Note: If you perform a factory reset of the LoadMaster, you must enable the **Enable Session Management** check box and disable the **Require Basic Authentication** check box to successfully run APIv2 (JSON-format) requests.

The level of user permissions determine what WUI Session Management fields can be seen and modified. Refer to the table below for a breakdown of permissions.

Control	Bal user	User with 'All Permissions'	User with 'User Administration' permissions	All other users
Session Management	Modify	View	View	None
Require Basic Authentication	Modify	View	View	None
Basic Authentication Password	Modify	View	View	None
Failed Login Attempts	Modify	Modify	View	None
Idle Session Timeout	Modify	Modify	View	None
Limit Concurrent Logins	Modify	Modify	View	
Pre-Auth Click Through Banner	Modify	Modify	View	None
Currently Active Users	Modify	Modify	View	None
Currently Blocked Users	Modify	Modify	View	None

When using WUI Session Management, it is possible to use one or two steps of authentication.

If **Enable Session Management** check box is ticked and **Require Basic Authentication** is disabled, the user only needs to log in using their local username and password. Users are not prompted to log in using the **bal** or **user** logins.

If the **Enable Session Management** and **Require Basic Authentication** check boxes are both selected, there are two levels of authentication enforced to access the LoadMaster WUI. The initial level is Basic Authentication where users log in using the **bal** or **user** logins, which are default usernames defined by the system.

The purpose of the **user** user is so that administrators can provide credentials of the **user** user to people, instead of providing the **bal** credentials. The password for the **user** user, can be set by configuring the **Basic Authentication Password** text box. Only the **bal** user is permitted to set the **Basic Authentication Password**.

Once logged in using Basic Authentication, the user then must log in using their local username and password to begin the session.

Enable Session Management

Selecting the **Enable Session Management** check box enables the WUI Session Management functionality. This will force all users to log in to the session using their normal credentials.

When this check box is checked, the user is required to login to continue to use the LoadMaster.

Note: LDAP users need to login using the full domain name. For example; an LDAP username should be **test@progress.com** and not just **test**.

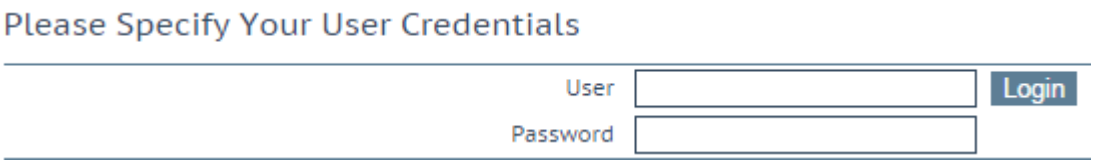

After a user has logged in, they may log out by clicking the **Logout** button, , in the top right-hand corner of the screen.

Once the WUI Session Management functionality is enabled, all the WUI Session Management options appear.

Require Basic Authentication

If WUI Session Management and Basic Authentication are both enabled, there are two levels of authentication enforced to access the LoadMaster WUI. The initial level is Basic Authentication where users log in using the **bal** or **user** logins, which are default usernames defined by the system.

Once logged in with Basic Authentication, the user then must log in using their local username and password to begin the session.

Basic Authentication Password

The Basic Authentication password for the **user** login can be set by typing the password into the **Basic Authentication Password** text box and clicking the **Set Basic Password** button.

The password needs to be at least 8 characters long and should be a mix of alpha and numeric characters. If the password is considered to be too weak, a message appears asking you to enter a new password.

Only the **bal** user is permitted to set the Basic Authentication password.

Failed Login Attempts

The number of times that a user can fail to login correctly before they are blocked can be specified within this text box. The valid values that may be entered are numbers between **1** and **999**.

If a user is blocked, only the **bal** user or other users with **All Permissions** set can unblock a blocked user.

If the **bal** user is blocked, there is a 'cool-down' period of 10 minutes before the **bal** user can login again. The **bal** user is unblocked if there are no login attempts for 10 minutes.

Idle Session Timeout

The length of time (in seconds) a user can be idle (no activity recorded) before they are logged out of the session. The valid values that may be entered are numbers between **60** and **86400** (between one minute and 24 hours).

Note: Any page that refreshes automatically will not time out from the WUI **Idle Session Timeout** setting. For example, the **Real Time Statistics** page, **GSLB Statistics** page, **WAF False Positive Analysis** page, and so on.

Limit Concurrent Logins

This option enables LoadMaster administrators to limit the maximum number of concurrent login sessions logins a single user can have to the LoadMaster WUI at any one time.

The values that can be selected range from 0 to 9.

A value of 0 allows an unlimited number of logins.

The value entered represents the total number and is inclusive of any **bal** user logins.

Pre-Auth Click Through Banner

Set the pre-authentication click through banner that is displayed before the LoadMaster WUI login page. This field can contain plain text or HTML code but not JavaScript. For security purposes, you cannot use the **‘** (single quote) and **“** (double-quote) characters. This field accepts up to 5,000 characters.

Note: Line breaks in the **Pre-Auth Click Through Banner** field in the LoadMaster UI must be in HTML code (such as <br/> or <p> tags).

Active and Blocked Users

Only the **bal** user or users with ‘All Permissions’ set can use this functionality. Users with ‘User Administration’ permissions set can view the screen but all buttons and input fields are greyed out. All other users cannot view this portion of the screen.

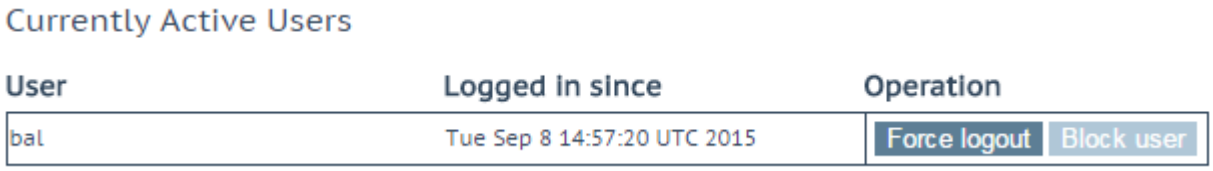

Currently Active Users

The user name and login time of all users logged into the LoadMaster are listed within this section.

To immediately log out a user and force them to log back into the system, click the **Force logout** button.

To block a user from being able to log in to the system, click the **Block user** button**.** The user will not be able to log back in to the system until they are unblocked or until the LoadMaster reboots. Clicking the **Block user** button does not force the user to log off, to do this, click the **Force logout** button.

If a user exits the browser without logging off, that session will remain open in the currently active users list until the timeout has reached. If the same user logs in again, before the timeout is reached, it would be within a separate session.

Currently Blocked Users

The user name and login time of when the user was blocked are listed within this section.

To unblock a user to allow them to login to the system, click the **Unblock** button.

OCSP Configuration

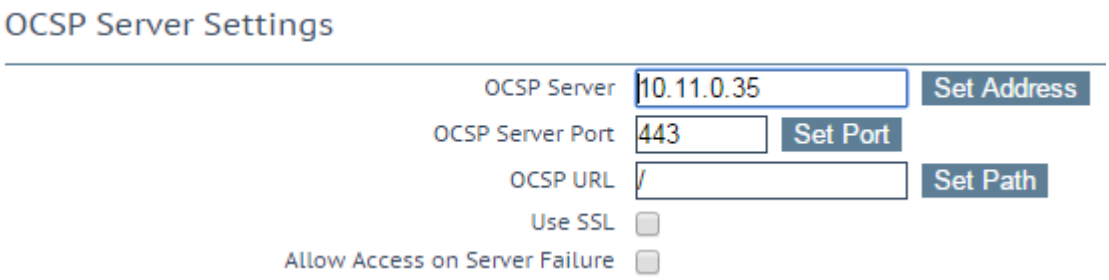

OCSP Server

The address of the OCSP server. This can either be in IP address or Fully Qualified Domain Name (FQDN) format.

OCSP Server Port

The port of the OCSP server.

OCSP URL

The URL to access on the OCSP server.

Use SSL

Select this to use SSL to connect to the OCSP server.

Allow Access on Server Failure

Treat an OCSP server connection failure or timeout as if the OCSP server had returned a valid response, that is, treat the client certificate as valid.

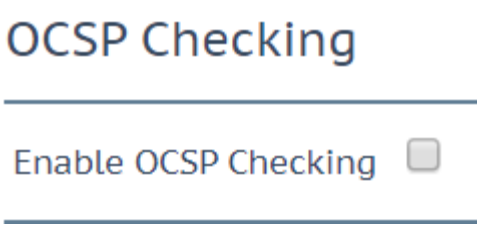

OCSP Checking

The **Enable OCSP Checking** UI control (and associated API) are all that is required in order to enable OCSP checking for outbound management connections that use certificate authentication (e.g., LDAP and remote logging).

- If the **OCSP Server/Port/URL** options are not set, all OCSP checking depends on the OCSP server setting in the AIA information from the certificate to be validated and this information is optional. If this information is not present or is invalid, no checking will be performed.
- If the **OCSP Server/Port/URL** options are set, then any certificate that does not have an OCSP server set in the AIA section will be checked using the provided OCSP server details.
- If both the certificate and the **OCSP Server/Port/URL** are set with the OCSP server address details, then only the information available in the certificate will be used to validate the certificate.If the details provided for OCSP server in the certificate are invalid, the OCSP checking will not switch to the LoadMaster OCSP server settings to validate the certificate.

The above behavior with respect to the OCSP Server/Port/URL settings also applies to OCSP checking of server certificate chains.

It should also be noted that OCSP checking for real server connections is not enabled by the above control. Real server OCSP certificate checks are enabled by the **Force Real Server Certificate** Checking option.

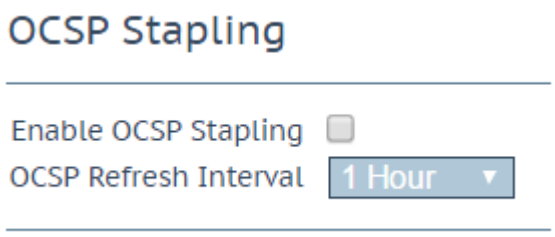

Enable OCSP Stapling

If the **Enable OCSP Stapling** check box is enabled, the LoadMaster verifies certificates for all external connections originated by the LoadMaster (except for re-encrypted connections to the Real Servers). Select this check box to enable the LoadMaster to respond to OCSP stapling requests. If a client connects using SSL and asks for an OCSP response, this is returned. Only Virtual Service certificates are validated. The system holds a cache of OCSP responses that are sent back to the client. This cache is maintained by the OCSP daemon. When the OCSP daemon sends a request to the server, it uses the name specified in the certificate (in the **Authority Information Access** field). If it cannot resolve this name, then it uses the default OCSP server specified in the **OCSP Server** text box.

OCSP Refresh Interval

Specify how often the LoadMaster should refresh the OCSP stapling information. The OCSP daemon caches the entry for up to the amount of time specified here, after which it is refreshed. Valid values range from 1 hour (default) to 24 hours.

LDAP Configuration

To get to the **LDAP Configuration** screen, expand **Certificates & Security** and click **LDAP Configuration**. This screen provides a management interface for LDAP endpoints. These LDAP endpoints may be used in three different areas:

- Health checks
- SSO domains
- WUI authentication



Any existing **LDAP Endpoints** are listed here, with an option to **Modify** and **Delete**. If an LDAP endpoint is in use, it cannot be deleted.

There is also an option to add a new LDAP endpoint. Type a name for the endpoint and click **Add**. Spaces and special characters are not permitted in the LDAP endpoint name.

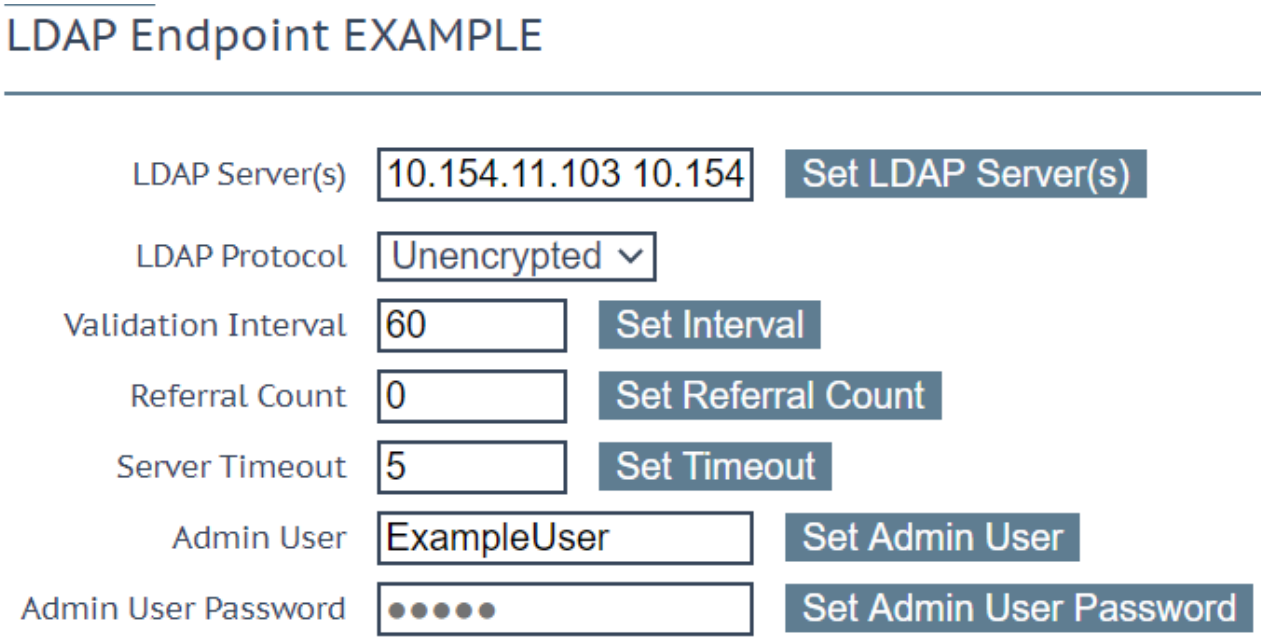

LDAP Server(s)

Specify a space-separated list of LDAP servers to be used. For windows Admin Controller (AC)/Domain Controller (DC), the scope of access for multiple domains and Permitted Groups is set to universal. Port numbers can also be specified if required. If you have multiple domains and are using **Permitted Groups**, sometimes it is necessary to include the Global Catalog port number, otherwise the **Permitted Groups** will fail. The default port is **3268**. For example, **10.110.20.23:3268**.

The LoadMaster uses OCSP to check the validity of the server certificates supplied by configured LDAPS servers. If these checks fail, connections to the server are not permitted.

LDAP Protocol

Select the transport protocol to use when communicating with the LDAP server.

Note: If you create an SSO domain with the **Authentication Protocol** set to **Certificates**, ensure to set the **LDAP Protocol** to **LDAPS** in the LDAP endpoint.

Validation Interval

Specify how often you should revalidate the user with the LDAP server.

Referral Count

The LoadMaster offers beta functionality to support LDAP referral replies from Active Directory Domain Controllers. If this is set to **0**, referral support is not enabled. Set this field to a value between **1** and **10** to enable referral chasing. The number specified will limit the number of hops (referrals chased).

Note: Multiple hops may increase authentication latency. There is a performance impact that depends on the number and depth of referrals required in your configuration.

Note: You must have intimate knowledge of your Active Directory structure to set the referral limit appropriately. The same credentials are used for all lookups, and so on.

Note: The use of Active Directory Global Catalog (GC) is the preferred configuration as the primary means of resolution instead of enabling LDAP referral chasing. A GC query can be used to query the GC cache instead of relying on LDAP and the referral process. Using Active Directory GC has little or no performance drag on the LoadMaster. For steps on how to add/remove the GC, refer to the following TechNet article: https://technet.microsoft.com/en-us/library/cc755257(v=ws.11).aspx

Server Timeout

Specify the LDAP server timeout in seconds. The default value is **5**. Valid values range from **5** to **60**.

Admin User

Type the username of an administrator user.

Admin User Password

Type the password for the specified administrator user.

Intrusion Detection Options (IPS/IDS)

SNORT is an Intrusion Prevention System (IPS) and an Intrusion Detection System (IDS). You can import SNORT rules into the LoadMaster and apply them to HTTP/HTTPS connections. You can also create your own rules using the SNORT 2.8 and 2.9 rule sets.

Note: The LoadMaster supports SNORT rules version 2.9 and below.

You can enable the rules for a Virtual Service by selecting the **Detect Malicious Requests** check box in **Virtual Services > View/Modify Services > Modify > Advanced Properties**.

Download the SNORT Rules

You can download the SNORT rule set from the SNORT website. In the **Rules** section, under **Community** - click **community-rules.tar.gz** to start the download.

Install the SNORT Rules

To install the SNORT rules on the LoadMaster, follow the steps below:

1. In the main menu of the LoadMaster WUI, go to **Certificates & Security > IPS / IDS**.
2. Click **Choose File** next to **Detection Rules**.
3. Browse to and select the previously downloaded **community-rules.tar.gz** file.
4. Click **Install new Rules**.
5. Select your chosen **Detection Level**.

Deactivate/Activate the SNORT Rules

You can modify the **community-rules.tar.gz** file by commenting out or un-commenting out rules. You can do this by opening the file as an archive using a file archive tool such as 7-Zip:

1. Open 7-Zip.
2. Click **File** and select **Open**.
3. Browse to the **community-rules.tar.gz** file.
4. Double-click the file to open the archive.
5. Continue double-clicking until you can see the following files:

- community.rules
- AUTHORS
- LICENSE
- sid-msg.map
- VRT-License.txt

1. Right-click **community.rules**.
2. Select **Edit** to open the file in a text editor (the edit shortcut key is **F4**).
3. Search for the desired rule by Signature ID (SID), for example, **sid:2067**.
4. To deactivate a rule, comment out the rule by adding a hash symbol (**#**) at the beginning of the line.
5. To activate a rule, un-comment the rule by deleting the **#** at the beginning of the line.
6. After your modifications are complete, click **File > Exit** to close the text editor.
7. When prompted to save the file, click **Yes**.

For further details, refer to the following Knowledge Base article: How to configure Intrusion Protection on LoadMaster (IPS+SNORT).

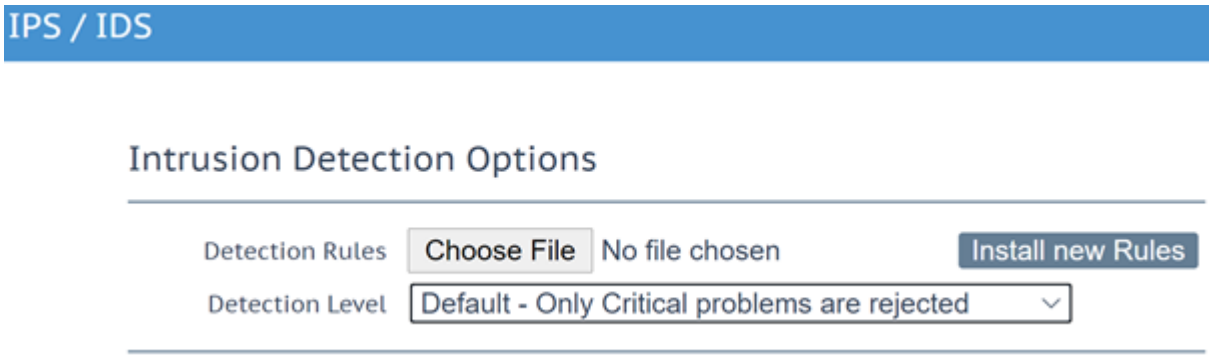

Detection Rules

Select the relevant detection rules and click the **Install New Rules** button to install them.

If you are implementing SNORT rules, please remember the following:

- The destination port must be $HTTP_PORTS
- A ‘msg’ may be optionally set
- The flow must be set to ‘to_server,established’

- The actual filter may be either ‘content’ or ‘pcre’
- Additional ‘http_’ parameters may be set
- The classtype must be set to a valid value

Note: To get updated or customized SNORT rules, please refer to the SNORT website: https://www.snort.org/.

Detection Level

Supports four levels of what to do when problems are encountered:

- **Low** – only logging with no rejection
- **Default** – only critical problems rejected
- **High** – Serious and critical problems rejected
- **Paranoid** – All detected problems rejected

The four levels of severity are classified according to the classtype value in the SNORT rules configuration file. If the severity is less than the configured value, a diagnostic is generated and the connection is dropped. The levels of severity correspond to the following values - **Low** = 1, **Default** = 2, and **High** = 3. You can view the SNORT rule classtypes and their corresponding values in the table below.

Classtype	Value
not-suspicious	3
unknown	3
bad-unknown	2
attempted-recon	2
successful-recon-limited	2
successful-recon-largescale	2
attempted-dos	2
successful-dos	2
attempted-user	1
unsuccessful-user	1
successful-user	1
attempted-admin	1
successful-admin	1

Classtype	Value
rpc-portmap-decode	2
shellcode-detect	1
string-detect	3
suspicious-filename-detect	2
suspicious-login	2
system-call-detect	2
trojan-activity	1
unusual-client-port-connection	2
network-scan	3
denial-of-service	2
non-standard-protocol	2
protocol-command-decode	3
web-application-activity	2
web-application-attack	1
misc-activity	3
misc-attack	2
icmp-event	3
kickass-porn	1
inappropriate-content	1
policy-violation	1
default-login-attempt	2
sdf	2

SSL Options

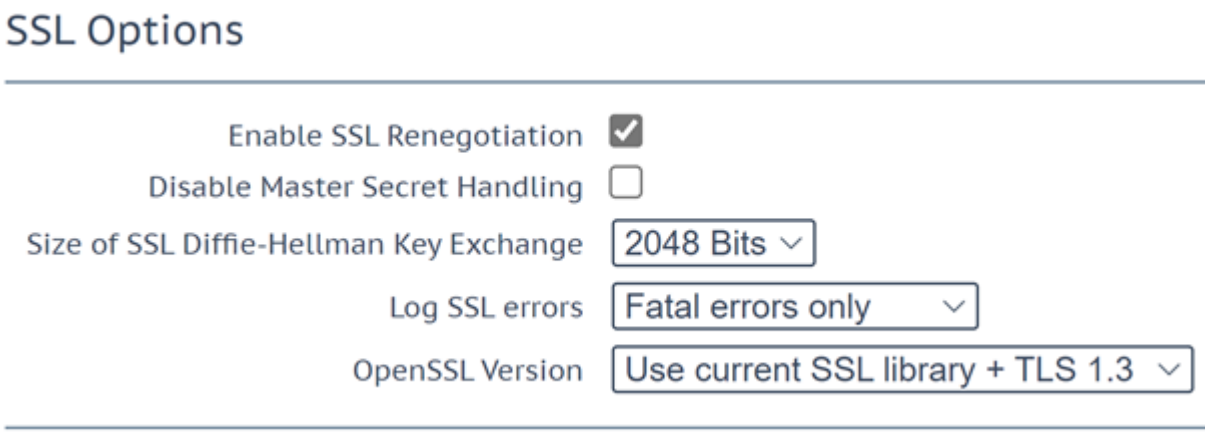

Enable SSL Renegotiation

When SSL renegotiation is enabled on the LoadMaster, either the client or server can initiate renegotiation within an established SSL session. If disabled, any attempt by either party to renegotiate will result in termination of the connection. In FIPS mode, this setting has no effect because SSL renegotiation is not supported under FIPS compliance.

Disable Master Secret Handling

In LoadMaster firmware version 7.2.52, the **Disable Master Secret Handling** check box was added. By default, the LoadMaster processes the Master Secret SSL Extension. This can cause problems for some legacy clients so it is possible to disable this processing.

Size of SSL Diffie-Hellman Key Exchange

Select the strength of the key used in the Diffie-Hellman key exchanges. If this value is changed, a reboot is required to use the new value. The default value is **2048 Bits**.

As of LoadMaster firmware version 7.2.53, you can select **4096** as a value in the **Size of SSL Diffie-Hellman Key Exchange** drop-down list.

After upgrading from a version prior to 7.2.53, it can take up to 30 minutes (on smaller models) to generate the 4k key. If you cannot see the **4096** option in the drop-down list 30 minutes after upgrading, try restart the login process.

CAUTION: During the upgrade from a version prior to 7.2.53, a new 4096-bit DHE key is generated. On smaller LoadMasters, this can lead to significant CPU and memory consumption that could impact regular Virtual Service traffic. So, Progress Kemp strongly recommends that this update be performed in a maintenance interval.

The use of 4K key will result in a significant degradation of performance compared to when using the 2K key.

Log SSL errors

Set the level of SSL error reporting in the logs. By default, the LoadMaster will not log common SSL alerts.

You can increase the verbosity of SSL error logging by setting this value to either of the following:

- **Fatal errors only** - only fatal errors are logged
- **Include Client errors** – this setting logs all client errors reported to the LoadMaster.
- **All errors** – this setting logs all SSL errors observed in the LoadMaster, including all common alerts and warnings that may or may not indicate an actual issue.

OpenSSL Version

By default, the LoadMaster uses the latest version of OpenSSL. This may cause performance problems on heavily-loaded sites. It is possible using the **OpenSSL Version** field to switch back to the old library which should alleviate some of these problems. Using the old library means that there is no support for TLS 1.3. Therefore, the **TLS1.3** check box is no longer available in the **SSL Properties** section of the Virtual Service modify screen.

If you switch from using the old library to using the current library in the **OpenSSL Version** field, **TLS1.3** is automatically re-enabled on all Virtual Services.

Note: This option is not applicable for following LoadMaster/ECS Connection Manager models:

- LM-X25

- LM-X40 Rev 05

- LM-X40M

- LM XHC 25G/40G/100G

- ECS Connection Manager H3 Rev 02

- ECS Connection Manager H3M

- ECS Connection Manager H3 25G/40G/100G

For these LoadMaster models, the **OpenSSL Version** field is available but the LoadMaster will continue to use the current OpenSSL implementation even if the **OpenSSL Version** field is set to **Use older SSL library - no TLS 1.3**.

CAUTION: Switching the OpenSSL version causes a total SSL outage during the switch. This operation should not be performed during working hours.

10

Web Application Firewall

Note: Utilizing WAF can have a significant performance impact on the LoadMaster deployment. Please ensure that the appropriate resources are allocated.

For virtual and bare metal LoadMaster instances, a minimum of 2GB of allocated RAM is required for the operation of WAF. The default memory allocation for Virtual LoadMasters and LoadMaster Bare Metal instances before LoadMaster Operating System version 7.1-22 is 1 GB of RAM. If this default allocation has not been changed, modify the memory settings before proceeding with the WAF configuration. If the check box to enable WAF is grayed out, it could mean that the LoadMaster does not have enough memory to run WAF.

There is a WAF engine open connection limit of 64000 per Virtual Service.

Note: To reduce the risk of the WAF connection limit being reached, the open connections will be closed after 20 seconds when a remote log server is too slow to respond. There is no impact on connections in use.

CAUTION: Legacy WAF (**WAF Options (Legacy)**) was fully deprecated as part of the v7.2.61 release. Deprecated means that Progress Kemp fully removed **WAF Options (Legacy)** from the LoadMaster. If you are running **WAF Options (Legacy)** and upgrade to the v7.2.61 release, a warning will be provided and the new WAF engine will be enabled with default values. We recommend updating your LoadMaster to use the latest WAF feature prior to upgrading to 7.2.61 so that you can configure the WAF engine to suit your configuration at the earliest possible convenience. For more details, refer to the following section: Converting from Legacy WAF to OWASP WAF of WAF Feature description document.

Contents

- Access Settings

- Export Logs
- Custom Rules
- False Positives

Access Settings

Note: The IP/FQDN and port of the Progress Kemp licensing server must be open on your firewall. The FQDN and IP address is **licensing.kemp.ax 52.166.52.190**, and the port is **443**. The old FQDNs for the Progress Kemp licensing server are **alsi.kemptechnologies.com** and **alsi2.kemptechnologies.com**, and the IP address for the old FQDNs is **52.136.251.129**. These also may need to be open depending on your LoadMaster version.

Note: Progress Kemp-provided daily updates are only available when you have an Enterprise Plus subscription. For further details on the subscription tiers, go to LoadMaster Support Subscriptions.

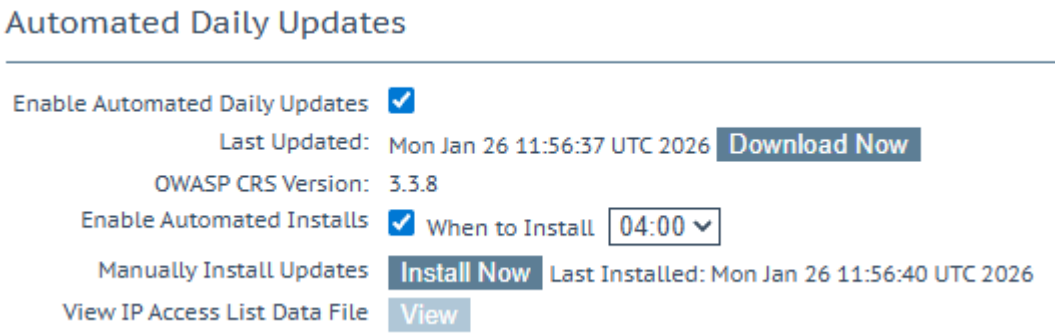

Enable Automated Daily Updates

Enables the automatic download of daily updates

Note: The **Download Now** and **Install Now** options are grayed out if WAF support license has expired. If this is the case, contact Progress Kemp to renew your subscription.

OWASP CRS version

Displays the latest OWASP Core Rule Set (CRS) version running on the LoadMaster.

To check the OWASP CRS current version, refer to the link https://coreruleset.org.

Enable Automated Installs

Enables automatic installation of the daily updates

By default, the **Enable Automated Installs** and **Manually Install Updates**options are grayed out. The rules must be downloaded for the first time before these options become available.

When to install

Select the time (hour of the day) to install the daily updates automatically.

Note: The daily updates must be assigned to a Virtual Service to take effect.

Download Now

Downloads the daily updates immediately.

A warning message appears if the rules have not been updated in the last seven days or if they have not been downloaded at all.

Install Now

Installs the daily updates manually.

Note: The daily updates must be assigned to a Virtual Service in order to take effect.

View IP Access List Data File

Click **View** to see the IP access list data file. This file contains a list of blocked IP addresses for known bad actors. If the **Enable IP Reputation Blocking** check box is enabled for a Virtual Service (**Virtual Services > View/Modify Services > Modify > WAF**), client addresses are checked against the IP access list file and are blocked if a match is found.

Export Logs

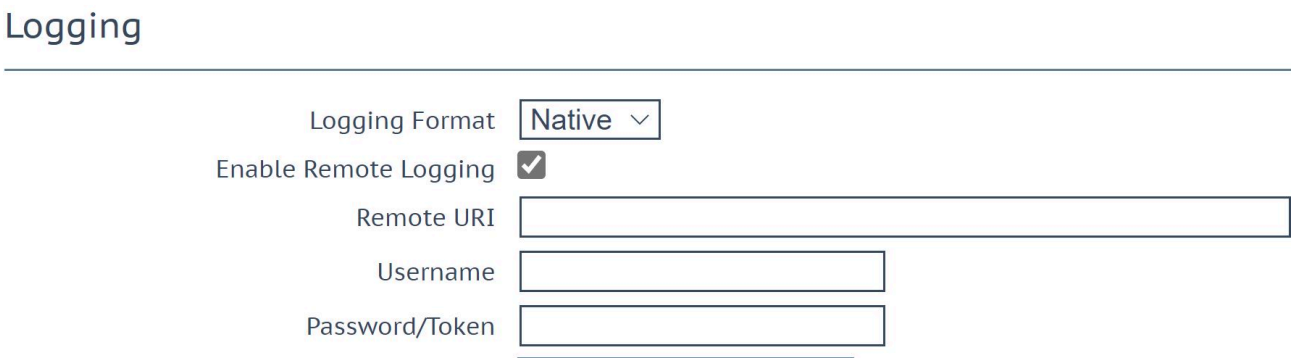

Logging Format

Depending on what format you want the audit logs to appear in, you can select the logging format option from drop-down list:

- Native
- JSON
- Splunk

Note: The **Splunk** option is only available if **Enable Remote Logging** is selected.

Note: Do not use **SPLUNK** as the **Username** if you are using a **Logging Format** other than **Splunk**.

Enable Remote Logging

This check box enables you to enable or disable remote logging for WAF.

Remote URI

Specify the Uniform Resource Identifier (URI) for the remote logging server.

Username

Specify the username for the remote logging server.

Password/Token

Specify the password or token for the remote logging server.

Custom Rules

Third-party rules can be uploaded to the LoadMaster. You can also write your own custom rules, which can be uploaded. These rules must be in the ModSecurity rule format in order to upload correctly. The **Custom Rules** screen enables you to upload **WAF Custom Rules** (.conf) and associated **WAF Custom Rule Data** (.data or .txt) files. The first character in the filename must be an alpha character or an underscore (**_**). The other characters in the filename can include full stops (**.**) or dashes (**-**). You can also upload gzip-compressed Tarball files (.tar.gz), which contain multiple rule and data files.

Note: We do not recommend using the WAF rule "redirect" action in custom rules because of the impact this has on system performance. You should use content rules instead for that purpose.

Note: Based on OWASP recommended best practices, the order of WAF rule processing was changed in LoadMaster firmware version 7.2.57. As of 7.2.57, custom rules are processed before the OWASP CRS rules.

To check the order that the rules are processed in:

1. Navigate to **System Configuration > Logging Options > System Log Files > Debug Options** and click the **Enable Logging** button next to **Enable WAF Debug Logging**.
2. When WAF debug logging is enabled, if there are WAF debug logs - the option **WAF Debug Log File** becomes available in the **System Log Files** screen. Click **View** to view the WAF debug log file.
3. You can see what order the rules are processed in. Lines that say **Invoking rule** specify when the rule was processed.

CAUTION: Please be aware that enabling WAF debug logs generate logs that may include Personally Identifiable Information as defined under the General Data Protection Agreement (EU GDPR). You should follow your organization's best practice to protect this information which may include anonymizing, deleting, or encrypting the data within the logs.

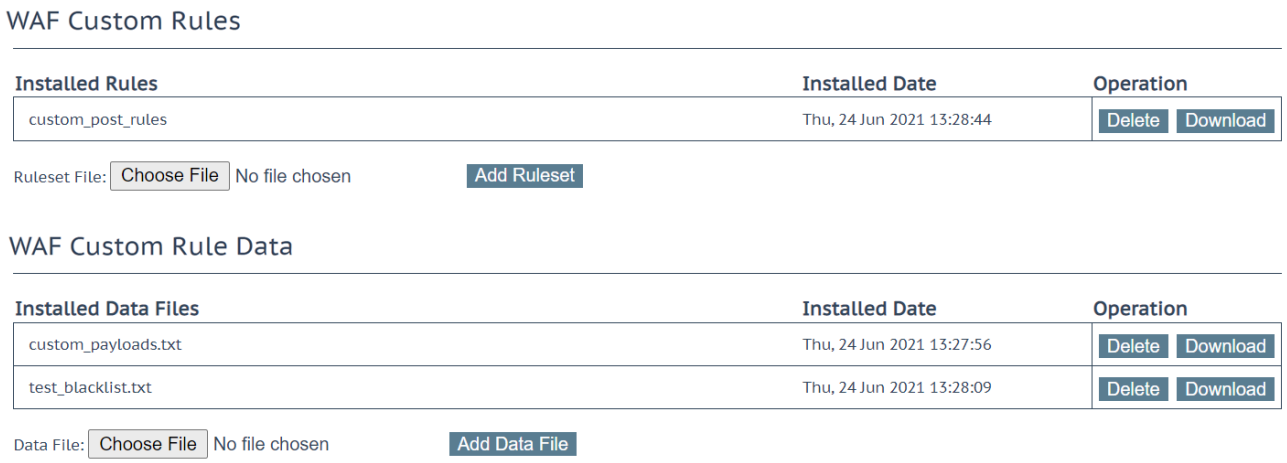

Installed Rules: **Choose File**.

Individual custom rules can be uploaded as .conf files. Alternatively, you can load a package of rules in a .tar.gz file. After selecting the rule files to be uploaded, click the **Add Ruleset** button.

WAF Custom Rule Data : **Choose File**.

Upload any additional data files with associated data for the rules.

The additional files are for the rules' associated data files. If you uploaded a Tarball when uploading the rules, the rules and data files can be packaged together. After selecting the data files to be uploaded, click the **Add Data File** button.

Note: Custom rules and data files can be deleted or downloaded by clicking the relevant buttons.

Note: If a rule is assigned to a Virtual Service, it will not be available for deletion.

False Positives



To perform False Positive analysis, at least one Virtual Service should be running WAF with OWASP rules and anomaly scoring. Select the appropriate Virtual Service from the drop down list and the rules being triggered will be shown, along with the WAF log information.

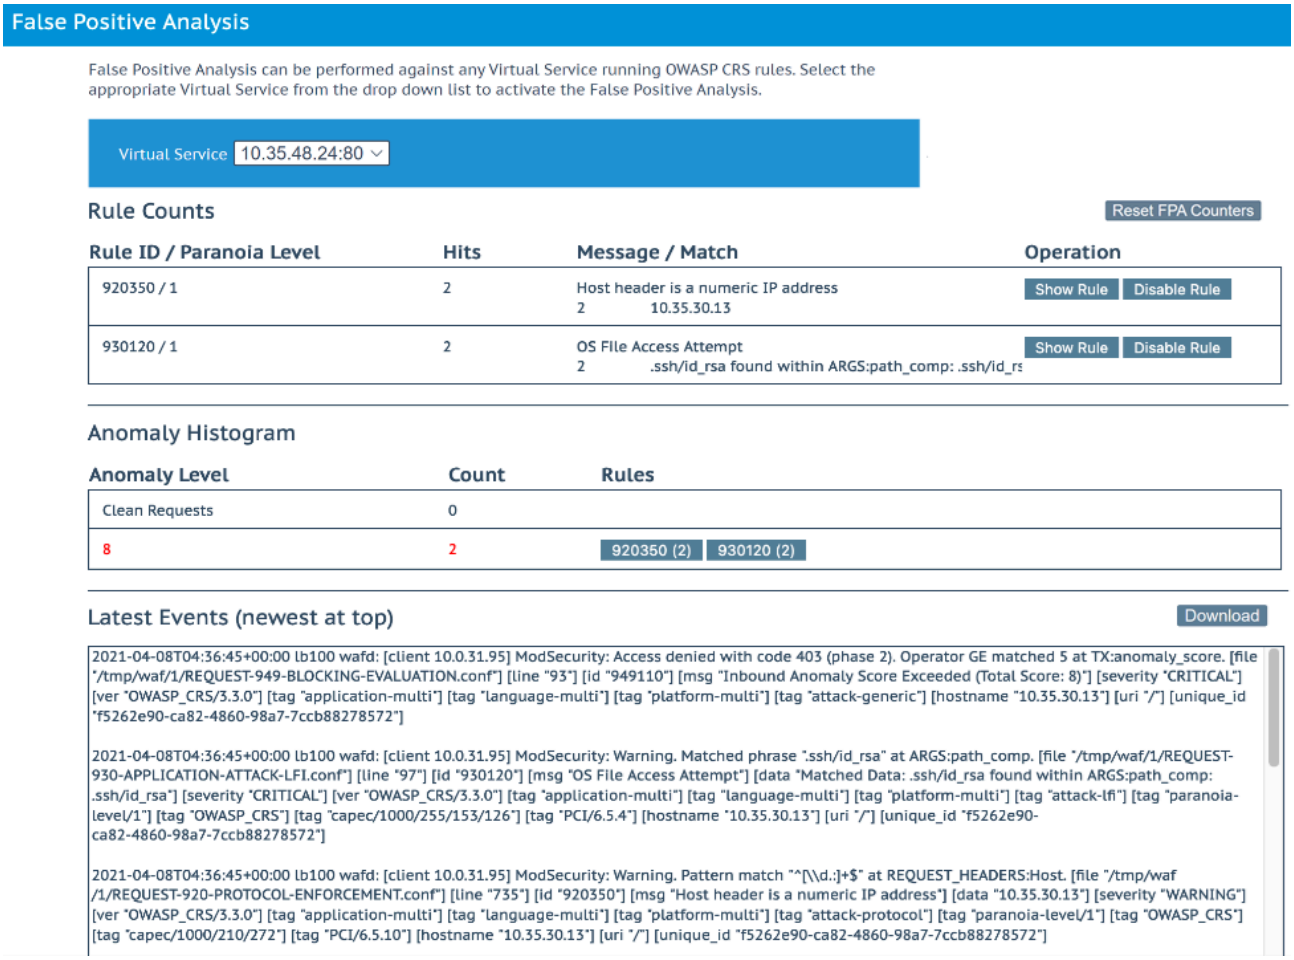

Rule Counts

The Rule Counts section displays information on any rules that are being triggered by requests. It displays the Rule ID, the paranoia level the rule is running under, the number of hits per requests that have triggered the rule and the message or match for the request are displayed for each rule that is triggered.

Clicking the **Show Rule** button in the **Operation** column displays the contents of the rule file associated with the triggered rule. This opens in a separate tab and the URL contains the triggered rule id..

The rule can be disabled by clicking the **Disable Rule** button.

Reset FPA Counter

Reset all False Positive Analysis Counters (Anomaly Histogram and Latest Events) for the virtual service. Clearing the Latest Events does not remove the logs from the LoadMaster, they are still available under **System Configuration > Logging Options > System Log Files > WAF Event Log File**.

Anomaly Histogram

The first row of the **Anomaly Histogram** section displays how many requests have been run without triggering a rule.

Each subsequent row gives details of rules that have been triggered and which are affecting the Anomaly Score. In each row the cumulative Anomaly Score, the number of requests which have triggered the rule and

the rule details are provided

Latest Events (newest at top)

Displays the event details for each rule that is triggered. These messages are in the standard ModSecurity log format and contains the anomaly score, the warning message, the attack state, and the paranoia level.

Download

Click the **Download** button to download the displayed WAF event logs details.

11

System Configuration

Refer to the following sections for details about the **System Configuration** WUI options, such as **Network Setup**, **HA and Clustering**, and **System Administration**.

Contents

- Network Setup
- HA and Clustering
- QoS/Limiting
- System Administration
- Logging Options
- Troubleshooting
- Miscellaneous Options

Network Setup

Refer to the following sections for details on the LoadMaster **Network Setup** options, including **Interfaces**, **Default Gateway**, and **Packet Routing Filter**.

Contents

- Interfaces
- Host & DNS Configuration

- Default Gateway
- Additional Routes
- Packet Routing Filter
- VPN Management
- Route Based VPN

Interfaces

This screen provides settings related to the external network and internal network interfaces. Some important points to note about interfaces are listed below:

- A Virtual Service can only be added to a known network or subnet. A network or subnet is considered known if there is an interface that has an IP address in that range.
- You cannot have two interfaces on the same subnet because this can cause routing problems.
- Do not use the same IP address on an interface and a Virtual Service. There is one exception to this rule - when using a LoadMaster on Azure.

The screen has the same information for the **eth0** and **eth1** Ethernet ports. The example below is for **eth0** on a non–HA (High Availability) unit.

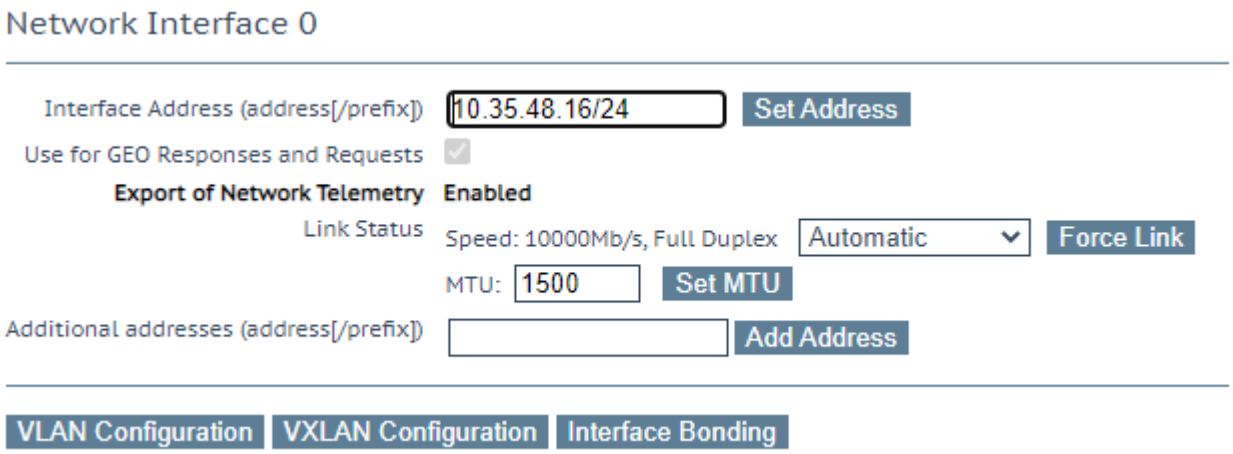

Interface Address

Within the **Interface Address (address[/prefix])** text box you can specify the Internet address of this interface. When **Allow Multi Interface Access** is enabled in **Certificates & Security > Remote Access**, and you set an IP address on another interface (other than eth0), a pop-up appears asking "Would you like to enable admin WUI access for this interface?". Click **OK** to enable the **Allow Administrative WUI Access** check box. If you click **Cancel**, the **Allow Administrative WUI Access** check box remains disabled and WUI access is not granted on that interface until that option gets enabled.

Note: If **Allow Multi Interface Access** option is enabled and the Admin WUI interface has been changed, a restart of web access is required to detect the change of interface. To restart the web server, login to the admin console (lsetup) as bal. Navigate to the **Local Admin > Web Address** and select the **Immediately**

Stop Web Server Access option and click **Ok**. When the page refreshes, select the **Immediately Start Web Server Access** option. To exit, select the **return to previous menu** and then select the **Exit Configuration**.

Cluster Shared IP address

Specify the shared IP address which can be used to access the cluster. This is also used as the default source address when using Server NAT.

Note: The clustering options are only available on LoadMasters which have a clustering license and where clustering has been configured. To add the clustering feature to your license, please contact a Progress Kemp representative. For further information on clustering, refer to the LoadMaster Clustering, Feature Description.

Use for Cluster checks

Use this option to enable cluster health checking between the nodes. At least one interface must be enabled.

Use for Cluster Updates

Select this option to use this interface to synchronize the configuration with all of the other nodes in the cluster.

Speed

By default, the **Speed** of the link is automatically detected. In certain configurations, this speed is incorrect and must be forced to a specific value.

Note: On AWS Elastic Network Adapter (ENA) interfaces, no speed is displayed when it cannot be read from the interface.

Note: For ENA driver interfaces, the speed displayed behaves similar to bonded interfaces in that the speed of the overall interface is the aggregate of the number of separate logical interfaces that have been defined on that interface is displayed.

Use for Default Gateway

The **Use for Default Gateway** check box is only available if the **Enable Alternate GW support** is selected in the **Network Options** screen. If the settings being viewed are for the default interface this option is grayed out and selected. To enable this option on another interface, go to the other interface by clicking it in the main menu on the left. Then this option is available to select. Once you select this option, the **Default Gateway** screen appears. Set the new default gateway. A notification displays to inform you that the default gateway has changed.

Note: If you reboot before setting the new default gateway, the gateway is removed from the configuration.

Allow Administrative WUI Access

This option is only available when the **Allow Multi Interface Access** check box is enabled in **Certificates & Security > Remote Access**.

When both of these options are enabled, the WUI can be accessed from the IP address of the relevant interface, and any **Additional addresses** set up for that interface.

Note: There is only one interface attached to all of these addresses so there may be issues with this unless the certificate used is a wildcard certificate. For more information on certificates, refer to the SSL Accelerated Services Feature Description.

Note: There is a maximum of 64 network interfaces that can be tracked and a maximum of 1024 total addresses where the system will listen on.

Use for GEO Responses and Requests

By default, only the default gateway interface is used to listen for and respond to DNS requests. This field gives you the option to listen on additional interfaces.

Note: This option cannot be disabled on the interface containing the default gateway. By default, this is eth0.

When this option is enabled, GEO also listens on any **Additional addresses** that are configured for the interface.

Export of Network Telemetry

This indicates if network telemetry monitoring is enabled or disabled for this interface (depending on what interfaces are selected on the **Network Telemetry** screen). For further details, refer to the Network Telemetry section.

MTU

Within the **MTU** field you can specify the maximum size of Ethernet frames that will be sent from this interface. The valid range is **512** - **9216**.

Note: The valid range of **512** - **9216** may not apply to VLMs as the range will be dependent on the hardware the VLM is running on. It is advised to check your hardware restrictions.

Additional addresses

Using the **Additional addresses** field allows the LoadMaster to give multiple addresses to each interface, as aliases. This is sometimes referred to as a “router on a stick”. It allows both IPv4 and IPv6 addresses in standard IP+CIDR format, so this can also be used to do a mixed mode of IPv4 and IPv6 addresses on the same interface. Any of the subnets that are added here will be available for both virtual IPs and real server IPs.

HA

If the unit is part of a HA configuration, the following screen will be displayed when one of the interfaces is clicked.

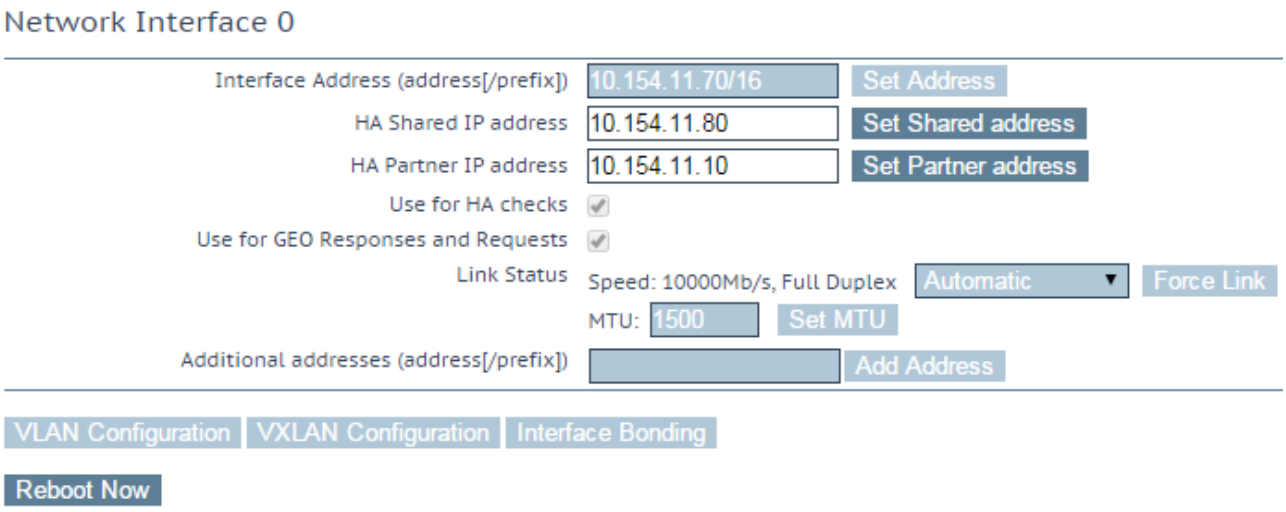

This screen tells the user:

- This is the **Master** machine of the pair (top-rightof the screen)
- This LoadMaster is up and the paired machine is down (green and red icons)
- The IP address of this LoadMaster
- The **HA Shared IP address**. This is the IP address used to configure the pair.
- The IP address of the paired machine
- This interface is enabled for HA health checking
- This interface is used as the Default Gateway
- The speed of the link is automatically detected
- Any alternate addresses on this interface

Creating a Bond/Team

Before creating a bonded interface please note the following:

- You can only bond interfaces higher than the parent, so if you choose to start with eth1, you can then bond eth2, eth3 and above, but you cannot bond eth0 (unless you start with eth0)
- Bond links first if you need VLAN tagging then add VLANs after the bond has been configured
- To add a link to a bonded interface, any IP addressing must first be removed from the link to be added
- Enabling the Active-Backup mode generally does not require switch intervention
- Bonding **eth0** with **eth1** can lead to serious issues and is not allowed to occur

Click **Interface Bonding** to request the bond.

Confirm the bond creation by clicking **Create a bonded interface**.

Acknowledge the warning dialogs.

Using the Web User Interface (WUI) select the **System Configuration > Interfaces > bndx** menu option.

If you do not see the **bndX** interface, refresh your browser, then select the bonded interface and click the **Bonded Devices** button.

Select the desired bonding mode.

Add the additional interfaces to this bond.

Configure the IP and Subnet Mask on the bonded interface.

Removing a Bond/Team

Remove all VLANs on the bonded interface first; if you do not remove them they will automatically be assigned to the physical port at which the bond started.

Select the **System Configuration > Interfaces > bndx** menu option. If you do not see the **bndX** interface refresh your browser, then select the bonded interface, then click the **Bonded Devices** button.

Unbind each port by clicking **Unbind Port**, repeat until all ports have been removed from bond.

Once all child ports have been unbounded, you can unbond the parent port by clicking **Unbond this interface** button.

Adding a VLAN

Select the interface and then select the **VLAN Configuration** button.



Add the **VLAN Id** value and select the **Add New VLAN** menu option.

Repeat as needed. To view the VLANs, select the **System Configuration > Network Setup** menu option and expand the drop-down list.

Removing a VLAN

Note: Before removing a VLAN, please ensure that the interface is not being used for other purposes, for example as a multicast interface, WUI interface, SSH interface or a GEO interface.

To remove a VLAN select the **System Configuration > Network Setup** menu option and select the appropriate VLAN ID from the drop-down list.

Once selected, delete the IP and then click **Set Address**. Once the IP has been removed you will have the option to delete the VLAN, by clicking the **Delete this VLAN** button.

Repeat as needed. To view the VLANs select the **System Configuration > Interfaces** menu option and select the appropriate VLAN ID from the drop-down list.

Adding a VXLAN

Select the relevant interface and then click the **VXLAN Configuration** button.

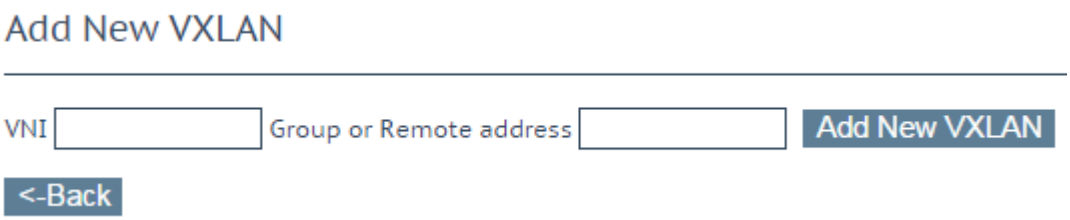

Enter a new VXLAN Network Identifier (VNI) in the **VNI** text box. Enter the multicast group or remote address in the **Group or Remote address** text box. Click **Add New VXLAN**.

To modify the VXLAN, go to **System Configuration > Interfaces** and select the VXLAN from the drop-down list.



On this screen, the interface address of the VXLAN can be specified. The VXLAN can also be deleted from this screen.

If HA is enabled, HA parameters can be set in the VXLAN:

- The **HA Shared IP address**. This is the IP address used to configure the HA pair.
- The IP address of the partner machine
- Specify whether or not this interface is used for HA health checking

Host & DNS Configuration

Set Hostname

Hostname lb100 Set Hostname

DNS NameServer (IP Address)	Operation
8.8.8.8	Delete
8.8.4.4	Delete

Add Nameserver

IP Address Add

DNS Search Domains	Operation
Kemp.LAB.INTRA	Delete

Add Search Domain

Domain Add

DNS Resolver Options

Enable DNSSEC Resolver
Automatically Update DNS Entries
DNS Update Interval 60 Set Update Interval
Reload DNS Entries for RS Errors
Resolve DNS Names now Run Resolver Now

Host IP Address	Host FQDN	Operation
10.154.33.233	example.com	Delete

Add/Modify Hosts for Local Resolution

IP Address Host FQDN Add/Modify

Set Hostname

Set the hostname of the local machine by entering the hostname in the **Hostname** text box and clicking **Set Hostname**. Only alphanumeric characters are allowed.

Add NameServer (IP Address)

Enter the IP address of a DNS server to resolve names locally on the LoadMaster in this field and click **Add**. A maximum of three DNS servers are allowed.

Note: It is not possible to delete the last remaining NameServer if the DNSSEC client is enabled. You can disable the DNSSEC client on the **Host & DNS Configuration** screen.

Add Search Domain

Specify the domain name to prepend to requests to the DNS NameServer in this field and click **Add**. A maximum of six Search Domains are allowed.

Add/Modify Hosts for Local Resolution

These fields provide the ability to manipulate the host file from the LoadMaster WUI. Specify the IP address and the host FQDN for the entry.

Enable DNSSEC Resolver

By default, the LoadMaster DNSSEC client is disabled. Only enable this option if needed. In some circumstances, the DNSSEC validation takes a significant amount of time to fail. This can cause the LoadMaster to appear to freeze or hang.

Selecting this option enables DNSSEC capabilities on the LoadMaster. You must add at least one **Nameserver** before DNSSEC can be enabled. The LoadMaster must be rebooted after changing the DNSSEC option to activate/deactivate the feature. When the setting is changed, it cannot be changed again until the LoadMaster has been rebooted.

Note: When using HA – the DNSSEC option must be configured on both devices separately.

DNSSEC works with the following utilities in the LoadMaster:

- Vipdump
- Ping and ping6
- Syslog
- SNMP
- Wget
- NTP
- SMTP
- Real Servers

Automatically Update DNS Entries

When this option is enabled, the LoadMaster attempts to automatically update any changed DNS names

(based on the **DNS Update Interval**):

- If the address is not found, or if it is the same as before – nothing is done (except a log entry is generated).
- If the address is different, the Real Server entry is updated with the new address, if possible.
- If the new address is invalid for some reason, for example if it is a non-local address and the **Enable Non-Local Real Servers** option is disabled, no changes are made and a log is generated.

DNS Update Interval

Set the update interval for DNS entries. Valid values range from 1 to 60 (minutes). The default value is 60.

Reload DNS Entries for RS Errors

When this option is enabled, DNS entries are reloaded when health checks have errors and an FQDN is associated with the Real Server IP address.

Resolve DNS Names now

Clicking the **Run Resolver Now** button forces a new resolution of DNS names. The behavior is the same as the **Automatically Update DNS Entries** option, except this is a manual (not an automatic) check.

Default Gateway

The network interface addresses must be set before configuring the default gateway.

The LoadMaster requires a default gateway through which it can communicate with the Internet.

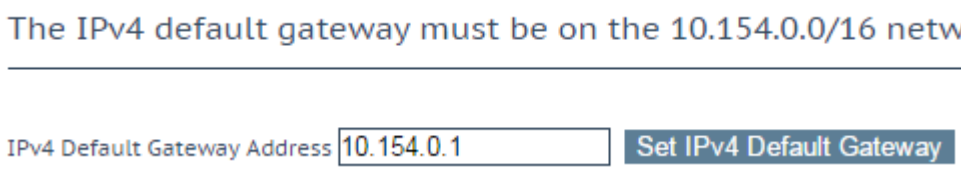

If both IPv4 and IPv6 addresses are being used on the LoadMaster, then both an IPv4 and IPv6 Default Gateway Address are required.

Note: IPv4 and IPv6 default gateways must be on the same interface.

Additional Routes



Further routes can be added. These routes are static and the gateways must be on the same network as the LoadMaster. To segment traffic you can also leverage the Virtual Service level default gateway.

Packet Routing Filter

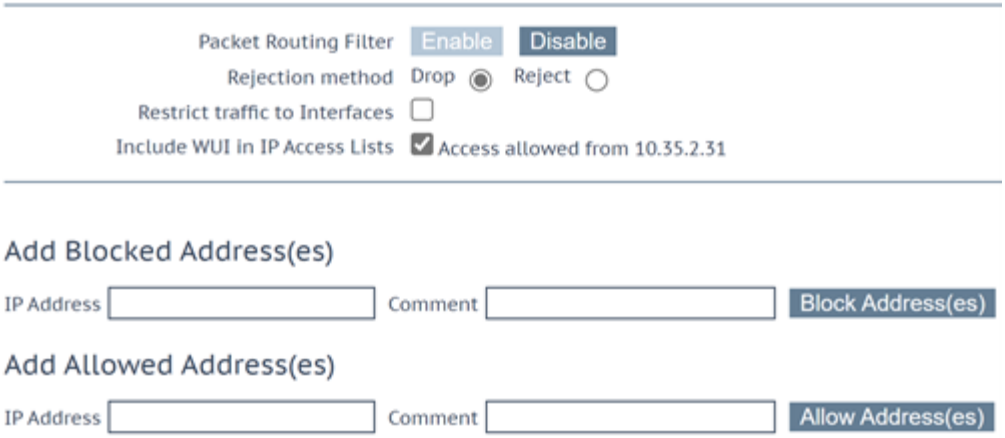

Packet Routing Filter

When the packet routing filter is activated, it restricts traffic to the LoadMaster but client access to services running on the interface addresses (SSH 22, HTTPS 443, SNMP 161, and DNS 53) is unaffected. Enabling SNAT prevents you from blocking traffic to a Virtual Service that has the same IP address as the LoadMaster's default gateway interface. This can affect Azure or any cloud platforms that use a single IP address.

If the packet routing filter is not activated, the LoadMaster also acts as a simple IP-forwarder.

Note: The **Reject/Drop blocked packets** and **Restrict traffic to Interfaces** fields will not be displayed if the **Packet Routing Filter** is disabled.

Reject/Drop blocked packets

When an IP packet is received from a host, which is blocked using the Access Control Lists (ACLs), the request is normally ignored (dropped). The LoadMaster may be configured to return an ICMP reject packet, but for security reasons it is recommended to drop any blocked packets silently.

Restrict traffic to Interfaces

This setting enforces restrictions upon routing between attached subnets. This option disabled by default.

Include WUI in IP Access lists

By default, the access control lists on the **Packet Routing Filter** page control access to the Virtual Services on the LoadMaster but not WUI access. If the **Include WUI in IP Access Lists** option is enabled, access to the WUI is also controlled by the access control lists. Enabling the **Include WUI in IP Access Lists** option allows the WUI to be accessed only from the IP address that enabled the check box - a message appears next to the check box saying **Access allowed from <IPAddress>**. This protects you from locking yourself out of the WUI. Attempts to log in from other IP addresses will be denied. If access to the WUI is needed from other IP addresses, you must add them to the allowed list.

The **Include WUI in IP Access Lists** option is only designed to work with one IP address.

Enabling the **Include WUI in IP Access lists** option does not affect any current access to Virtual Services. For example:

- If you enable the **Include WUI in IP Access lists** option from an IP address that is not in either the allowed or blocked lists, that IP address has access to the WUI and all the Virtual Services.
- Other IP addresses not in either the allowed or blocked list do not have access to the WUI but they do have access to all Virtual Services.

If you add IP addresses using the **Add Allowed Address(es)** fields, the connectivity for all other IP addresses will be blocked to the Virtual Services.

With the **Include WUI in IP Access lists** option disabled, access to the WUI is not affected by the packet filter.

If you need to, you can disable the access control lists using the console interface.

Add Blocked Address(es)

The LoadMaster supports an Access Control List (ACL) system. Any host or network entered into the ACL will be blocked from accessing any service provided by the LoadMaster.

The ACL is only enabled when the Packet Filter is enabled. The allowlist allows a specific IP address or address range access. If the address or range is part of a larger range in the block list, the allow list will take precedence for the specified addresses.

If a user does not have any addresses listed in their block list and only has addresses listed in their allow list, then only connections from addresses listed on the allow list are allowed and connections from all other addresses are blocked.

This option allows a user to add or delete a host or network IP address to the Access Control List. In addition to IPv4 addresses - IPv6 addresses are allowed in the lists if the system is configured with an IPv6 address family. Using a network specifier specifies a network.

For example, specifying the address **192.168.200.0/24** in the block list will block all hosts on the 192.168.200 network.

Note: A static port Virtual Service, with an access list defined to block particular traffic, will not work correctly if you also have a wildcard Virtual Service on the same IP address. The wildcard Virtual Service will accept the traffic after the static port Virtual Service denies it.

Note: It is recommended to use a separate IP address in this case to avoid unexpected behavior resulting from this interaction.

VPN Management

Note: The **VPN Management** link/screen will only be available if the LoadMaster is licensed for IPsec tunneling.

Note: This is document covers the policy-based VPN functionality that is available in the latest LoadMaster Long Term Support (LTS) release. As of LoadMaster firmware version 7.2.53, the **VPN Management** menu option changed to **Policy Based VPN** and support was added for route-based VPNs.

For further information on IPsec tunneling, including step-by-step instructions on how to set it up, refer to the IPSec Tunnelling Feature Description.



Connection Name

Specify a unique name to identify the connection.

Create

Create a uniquely identifiable connection with the specified name.

View/Modify

View or modify the configuration parameters for this connection.

Delete

Delete this connection.

Note: All associated configuration will be permanently deleted. A connection can be deleted at any time, even if it is running.

View/Modify VPN Connection

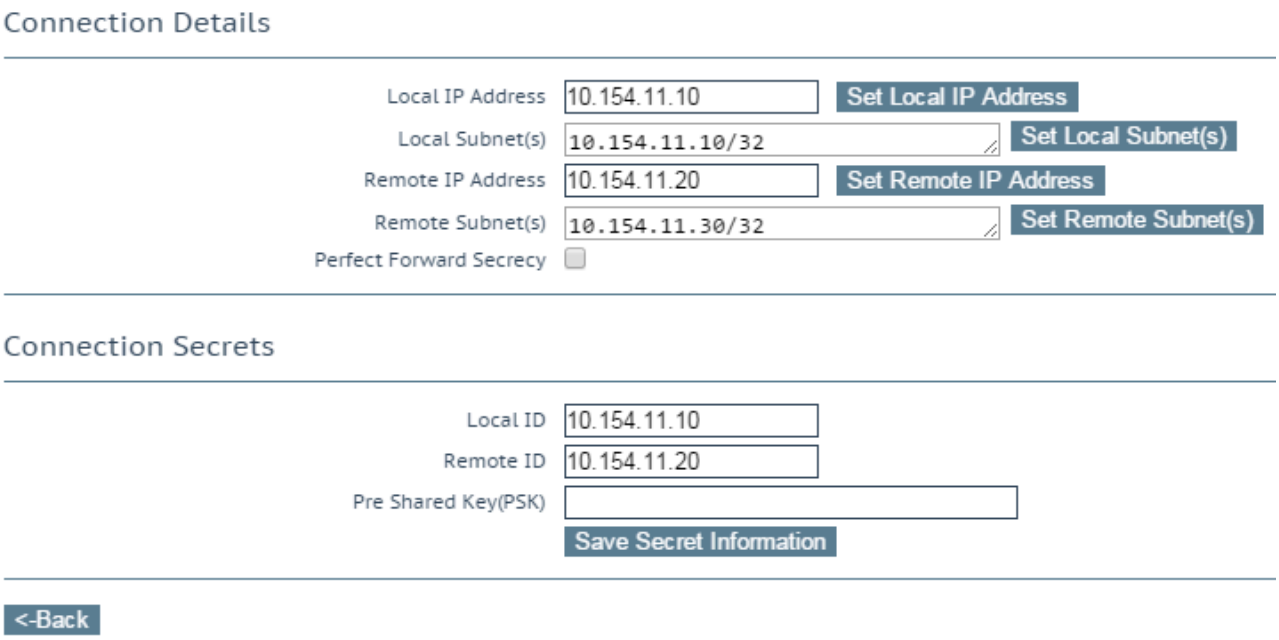

When initially creating a connection, or when modifying a connection, the **View/Modify VPN Connection** screen appears.

Local IP Address

Set the IP address for the local side of the connection.

In non-HA mode, the **Local IP Address** should be the LoadMaster IP address, that is, the IP address of the default gateway interface.

In HA-mode, the **Local IP Address** should be the shared IP address. This will be automatically populated if HA has already been configured. For more information on setting up tunneling in a HA configuration, refer to the next section.

Local Subnet Address

When the **Local IP Address** is set the **Local Subnet Address** text box is automatically populated. The local IP can be the only participant if applicable, given the /32 CIDR. Review the **Local Subnet Address** and update it if needed. Ensure to click **Set Local Subnet Address** to apply the setting, whether the address has been changed or not. Multiple local subnets can be specified using a comma-separated list. Up to 10 IP addresses can be specified.

Remote IP Address

Set the IP address for the remote side of the connection. In the context of an Azure endpoint, this IP address is expected to be the public-facing IP address for the Virtual Private Network (VPN) Gateway device.

Remote Subnet Address

Set the subnet for the remote side of the connection. Multiple remote subnets can be specified using a

comma-separated list. Up to 10 IP addresses can be specified.

Perfect Forward Secrecy

Activate or deactivate the Perfect Forward Secrecy option.

Note: The cloud platform being used will determine what the **Perfect Forward Secrecy** option should be set to. **Perfect Forward Secrecy** is needed for some platforms but is unsupported on others. To find out what will work with your cloud platform, refer to the IPsec Tunneling Feature Description document.

Local ID

Identification for the local side of the connection. This may be the local IP address. This field is automatically populated with the same address as the **Local IP Address** if the LoadMaster is not in HA mode.

If the LoadMaster is in HA mode, the **Local ID** field will be automatically set to **%any**. This value cannot be updated when the LoadMaster is in HA mode.

Remote ID

Identification for the remote side of the connection. This may be the remote IP address.

Pre Shared Key (PSK)

Enter the pre-shared key string.

Save Secret Information

Generate and save the connection identification and secret information.

Route Based VPN

Prior to LoadMaster firmware version 7.2.53, the LoadMaster only supported policy-based Virtual Private Networks (VPNs). As of LoadMaster version 7.2.53, route-based VPN support has been introduced. The route-based VPN functionality is available as an add-on pack which you can download from the LoadMaster Add-on Packages page. When you have the add-on file, install it in **System Configuration > System Administration > Update Software**. Reboot the LoadMaster after installing the add-on to activate it (**System Configuration > System Administration > System Reboot**).

Note: The route-based VPN add-on is only supported in firmware version 7.2.53 and above.

Note: The LoadMaster route-based VPN functionality is based on the strongSwan IPsec VPN solution. For more details on strongSwan documentation, refer to the strongSwan website link https://www.strongswan.org/.

After installing the add-on and rebooting the LoadMaster, a new main menu option becomes available: **System Configuration > Network Setup > Route Based VPN**. Specify a unique name to identify the connection and click **Create**.

After configuring the details of the VPN connection, you can also see the **Connection Debug** options:

- Stop and start the IPsec daemon on the LoadMaster.
- Display the connection status.
- Display routes.
- Show logs.

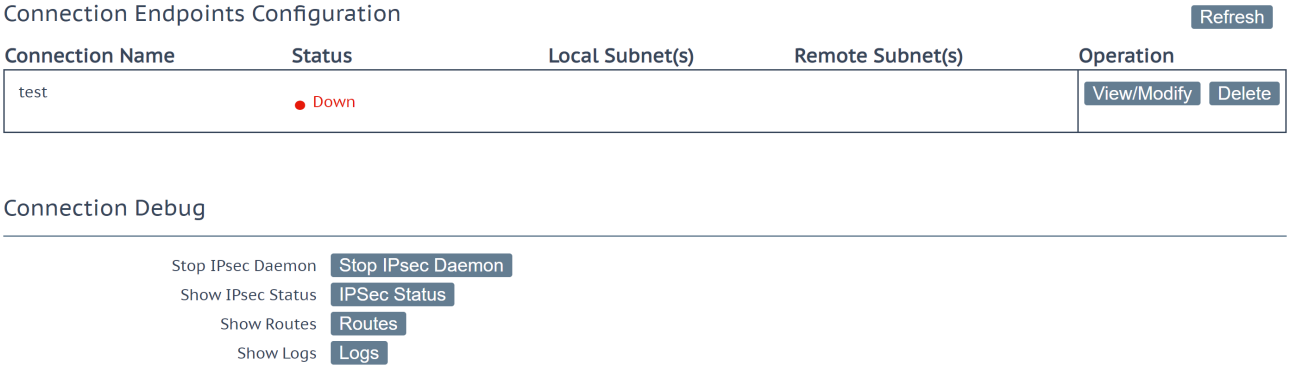

Connection Name

Specify a unique name to identify the connection.

Create

Create a uniquely identifiable connection with the specified name.

View/Modify

View or modify the configuration parameters for this connection.

Delete

Delete this connection.

Note: Ensure the **Connection Name** is the same as the connection name in the **ipsec.conf** file you will be uploading.

Note: The **Connection Name** must be at least three and at most 20 characters in length. Valid characters are **a-z, A-Z, 0-9, _,** and **-**.

After creating the connection initially, you must upload the connection configuration, route configuration, and secrets files based on the strongSwan standard format. These files provide support for multiple configuration option which the user can use to configure strongswan IPsec.

Note: Ensure to set the left IP address in the ipsec.conf file to your LoadMaster IP address (eth0).

Once the files have been uploaded and validated, the **Connection Details** are shown in the fields on the **View/Modify** VPN Connection screen.

View/Modify VPN Connection

Connection Details

Local IP Address	10.35.30.109	
Local Subnet(s)	0.0.0.0/0	
Remote IP Address	10.35.44.42	
Remote Subnet(s)	0.0.0.0/0	
Connection Config file	Choose File No file chosen	Update Config
Route Config file	Choose File No file chosen	Update Route

Connection Secrets

Secrets file	Choose File No file chosen	Update Secrets

Connection Debug

Start Connection	Start Connection
Show IPsec Status	IPSec Status
Show Logs	Logs

<-Back

Local IP Address

Display the IP address for the local side of the connection.

In non-HA mode, the **Local IP Address** should be the LoadMaster IP address, that is, the IP address of the default gateway interface.

In HA-mode, the **Local IP Address** should be the shared IP address. This will be automatically populated if HA has already been configured. For more information on setting up tunneling in a HA configuration, refer to the next section.

Local Subnet(s)

Display the **Local Subnet(s)** for the local side of the connection.

Remote IP Address

Display the IP address for the remote side of the connection. In the context of an Azure endpoint, this IP address is expected to be the public-facing IP address for the Virtual Private Network (VPN) Gateway device.

Remote Subnet(s)

Display the subnet for the remote side of the connection.

Connection Config file

Click the **Choose File**, browse to and select the file and click **Update Config**. The configuration file contains the default value **conn %deafult** for all the IPsec connections configured. This configuration is given directly to the strongswan IPsec engine. The format of Connection Config file contains the following parameters:

conn %default

dpddelay=30

dpdtimeout=120

dpdaction=restart

ikelifetime=28800s

#keylife=20m

keylife=1d

rekeymargin=3m

keyingtries=1

authby=secret

keyexchange=ikev2

mobike=yes

ike=aes256-sha384-ecp384

esp=aes256-sha384-ecp384,aes256gcm16

conn routevpn

left=10.35.45.170

leftsubnet=10.35.99.170/32

right=65.51.241.146

rightsubnet=10.0.70.1/32

leftid=78.56.45.56

rightid=%any

mark=1

auto=start

ikelifetime=1h

lifetime = 1h

margintime = 9m

rekeyfuzz = 100%

ike=aes256gcm16-sha384-modp2048

esp=aes256gcm16-sha384-modp2048

Route Config file

Click the **Choose File**, browse to and select the file and click **Update Route**. The Route Config file contains routing information to route packets to a specific remote network using the IPsec tunnel created. The format of the Route Config file contains the following parameters:

<LEFT IP> <RIGHT IP> <REMOTE NETWORK> <SOURCE IP>

Secrets file

Click the **Choose File**, browse to and select the file and click **Update Secrets**. The Secrets file contains the remote and local endpoint address and the secrets to communicate with the remote endpoint. The format example of a Secret file is:

10.35.45.170 65.51.241.146 %any: PSK "fE31$I#%w&"

Show IPsec Status

You can display the connection status by clicking **IPSec Status**.

Show Logs

You can display the connection logs by clicking **Logs**.

HA and Clustering

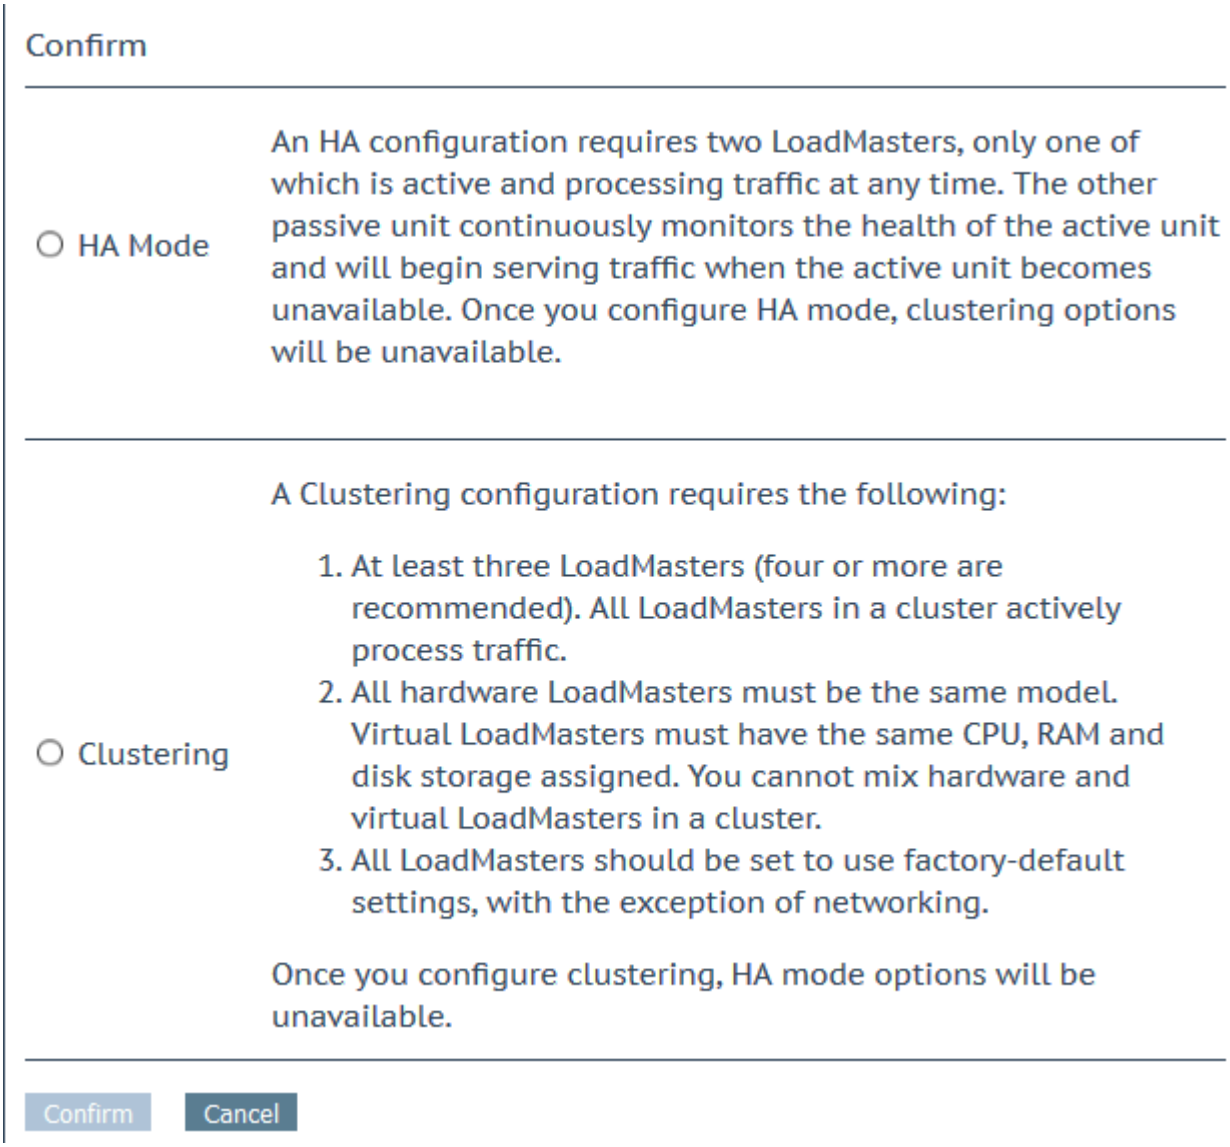

This section in the WUI is only called **HA and Clustering** if you have a LoadMaster license with clustering enabled. If you do not have clustering, this section will be called **HA Parameters** and you will not see the screen shown above. If clustering has been configured, this section will be called **Cluster Control**.

This screen describes both **HA Mode** and **Clustering**. Select the relevant option and click **Confirm** to continue.

Note: Once clustering is configured, the HA mode options will be unavailable.

Contents

- HA Mode
- Cluster Control

HA Mode

Note: If you are using the LoadMaster in a cloud environment, refer to the Azure HA Parameters or AWS HA Parameters section.

Note: If you are using the LoadMaster for AWS product, refer to the **AWS HA Parameters** section.

The role of the appliance can be changed by setting the HA Mode. If **HA (First) Mode** or **HA (Second) Mode** is selected as the **HA Mode**, a prompt will appear reminding to add a shared IP. Changing the HA Mode will require a reboot, so after the details are set, click the **Reboot** button provided. Once the LoadMaster has rebooted, the HA menu option will be available in the **System Configuration** section provided the role is not “Non HA Mode”. HA will NOT work if both machines are specified the same.

When logged into the HA cluster, use the shared IP address to view and set full functionality to the pair. If you log into the direct IP address of either one of the devices the menu options are quite different (see menus below). Logging into one of the LoadMaster directly is usually reserved for maintenance.



When a LoadMaster is in HA mode, the following screen appears when you select the **HA Parameters** menu option.

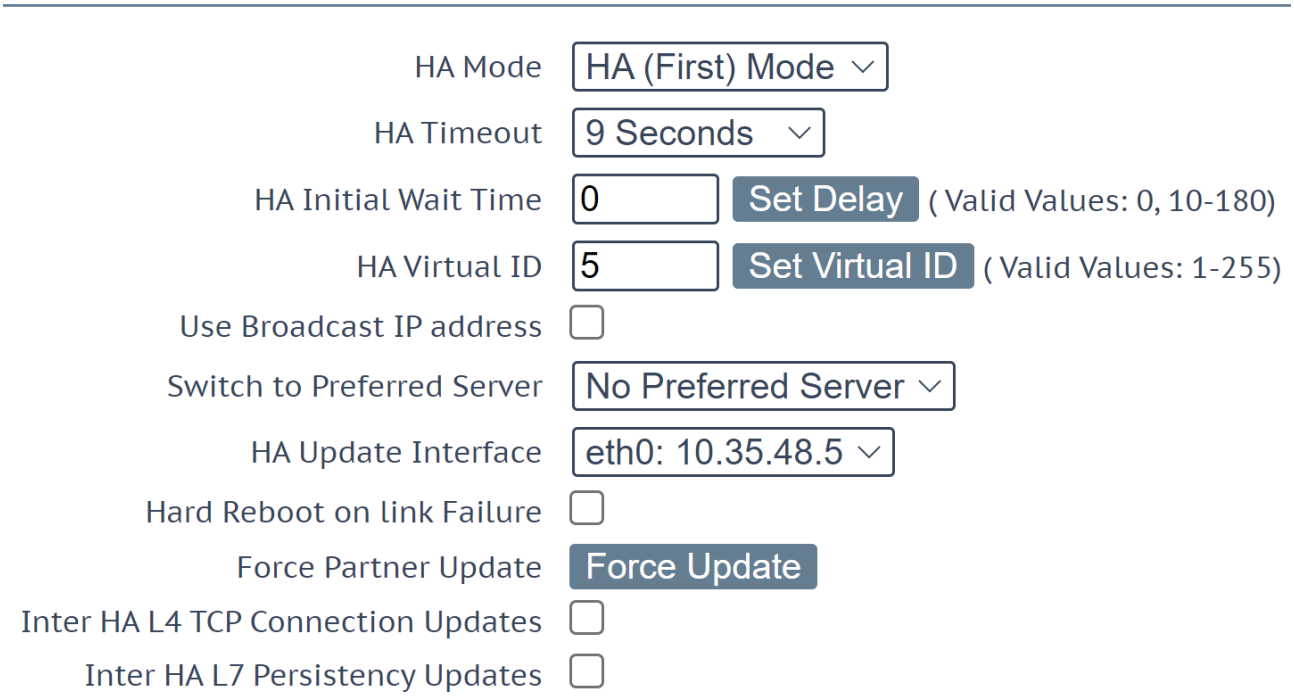

HA Status

At the top of the screen, next to the time, icons are shown to denote the real-time status of the LoadMaster units in the cluster. There will be an icon for each unit in the cluster. You can open the WUI for the first or second HA unit by clicking the relevant status icon.

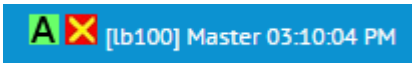

The possible icons are:

Green (with 'A')		The unit is online and operational and the HA units are correctly paired. The A in the middle of the square indicates that this is the active unit.
Green (without 'A')		The unit is online and operational and the HA units are correctly paired. The absence of an 'A' in the middle of the square indicates that this is the standby unit.
Red/Yellow		The unit is not operational. It may be

		offline or misconfigured. The unit is not ready to take over. It may be offline or incorrectly paired.
Blue		When the unit reboots more than 3 times in 5 minutes it moves into a pacified state. In this state the machine is only accessible using the direct machine WUI (not the shared WUI), and, it is not participating in any HA activity, that is, no changes from the active unit are received and it will not take over if the active unit fails. To clear the pacified state, the unit must remain stable (without rebooting) for at least five minutes. After this stability period, a subsequent reboot clears the pacified state and returns the unit to normal operation.
Grey		The machine is in an indeterminate state and may require a reboot to return to operation. In some cases, this may mean both machines are active, that is, both are set to active, and something has gone seriously wrong. Call Progress Kemp Support for assistance with this issue, if rebooting does not solve it.
No HA icons		If the HA status squares are not appearing in the WUI, it probably means that HA is not enabled. Go to **System Administration** and select the HA option. Ensure the **HA Mode** is set to either **First** or **Second**.

In HA mode each LoadMaster will have its own IP address used only for diagnostic purposes directly on the unit. The HA pair have a shared IP address over which the WUI is used to configure and manage the pair as a single entity.

Note: Both HA1 and HA2 must be on the same subnet with the same default gateway and be in the same physical site. They must not be separated by an intra-site link and must use the same gateway to return traffic.

HA Mode

If using a single LoadMaster, select Non-HA Mode. When setting up HA mode, one LoadMaster must be set to HA (First) and the other HA (Second). If they are both set to the same option, HA will not operate.

HA Timeout

The time that the Active machine must be unavailable before a switchover occurs. With this option, the time it takes an HA cluster to detect a failure can be adjusted from 3 seconds to 15 seconds in 3 second increments. The default value is 9 seconds. A lower value will detect failures sooner, whereas a higher value gives better protection against a DOS attack.

HA Initial Wait Time

How long after the initial boot of a LoadMaster, before the machine decides that it should become active. If the partner machine is running, then this value is ignored. This value can be changed to mitigate the time taken for some intelligent switches to detect that the LoadMaster has started and to bring up the link.

HA Virtual ID

When using multiple HA LoadMaster clusters on the same network, this value uniquely identifies each cluster so that there are no potential unwanted interactions.

Note: All LoadMasters on the network that are or will be configured into HA pairs must be assigned unique **HA Virtual ID** numbers.

As of the 7.2.36 release, the LoadMaster selects a virtual ID based on the shared IP address of the first configured interface (the last 8 bits). It is selected and displayed once both the shared address and the partner address have been set. You can change the value to whatever you want (in the range 1 – 255) or you can keep it at the value it already selected. Please ensure the virtual ID is unique on each LoadMaster on the network.

Use Broadcast IP address

By default, the LoadMaster uses an IP multicast address when sending CARP packets. Enabling this option forces the use of the IP broadcast address instead.

Switch to Preferred Server

By default, neither partner in a HA cluster has priority. So that when a machine restarts after a switchover, the machine becomes the standby and stays in that state until forced to Active. Specifying a preferred host means that when this machine restarts, it will always try to become active and the partner will revert to standby mode. If a preferred server is specified, this will result in a double failover event because if the active unit fails, the standby unit will take over as active and when the preferred unit comes back up it will take over as active.

HA Update Interface

The interface used to synchronize the HA information within the HA cluster.

Hard Reboot on link Failure

In LoadMaster firmware version 7.2.53, a new option, **Hard Reboot on link Failure**, was introduced. When the **Hard Reboot on link Failure** check box is enabled, the LoadMaster configured in HA reboots if any

configured interface loses connectivity with the network (that is, experiences a link failure). The reboot occurs regardless of the LoadMaster's HA status (Primary or Standby).

The **Hard Reboot on link Failure** check box is available in the **System Configuration > HA Parameters** screen when both of these are true:

- High Availability (HA) is configured
- The **Switch to Preferred Server** option is set to **No Preferred Server**.

Note: The **Hard Reboot on link Failure** check box will be unavailable, if you select a preferred server from the Switch to Preferred Server drop-down list.

You cannot have a preferred server if **Hard Reboot on link Failure** is enabled - if you did, it could lead to circular swapping between the active and standby LoadMaster units.

Force Partner Update

Immediately forces the configuration from the active to standby unit without waiting for a normal update.

Inter HA L4 TCP Connection Updates

When using L4 services, enabling updates will allow L4 connections to be maintained across a HA switchover by sharing the connection table. This option is ignored for L7 services.

L4 Sync Threshold

The minimum number of incoming packets that a connection needs to receive before the connection will be synchronized.

Default value: 3 seconds

Valid values range from the 1 to the <*SyncPeriod-1*>. For example, if the **L4 Sync Period** is set to 50, the maximum **L4 Sync Threshold** would be 49.

L4 Sync Period

A connection gets synchronized every time the number of its incoming packets modulus sync_period equals the threshold.

Default value: 50 seconds

Valid values range from the <*SyncThreshold+1*> to 255.

L4 Sync Refresh Period

This is the allowed difference time (in seconds) in the reported connection timer that triggers a new sync message.

Default value: 0

Valid values range from **0** to **10**.

Inter HA L7 Persistence Updates

When using L7 services, enabling this option will allow persistence information to be shared between the HA partners. If an HA failover occurs, the persistence information will not be lost. Enabling this option can have a significant performance impact.

HA Multicast Interface

The network interface used for multicast traffic which is used to synchronize Layer 4 and Layer 7 traffic when Inter-HA Updates are enabled.

Use Virtual MAC Addresses

Enabling this option forces the MAC address to switch between a HA pair during a switchover which is useful when gratuitous ARPs (used in communicating changes in HA IP addresses to switches) are not allowed.

Note: This option is only available for hardware LoadMasters.

Contents

- Azure HA Parameters
- AWS HA Parameters

Azure HA Parameters

This screen is only available in LoadMaster for Azure products.

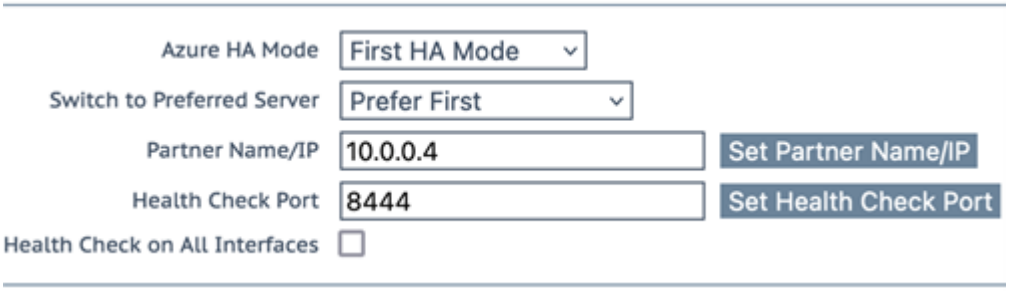

Azure HA Mode

Select the required HA mode for this unit. There are three options:

- First HA Mode
- Second HA Mode
- Non HA Mode

If you are only using a single LoadMaster, select **Non HA Mode**.

When using HA mode, one machine must be specified as the **First (Active)** and the second machine must be specified as the **Second (Stand-by)**.

Note: HA will not work if both units have the same value selected for the **Azure HA Mode**.

Note: Synchronization of Virtual Service settings only occurs from the active to the standby. Changes made

---

to the active will be replicated to the standby. However, changes made to the standby are never replicated to the active unit.

Note: If the active unit fails, connections will be directed to the standby unit. The active unit is active and will never become the standby, even if it fails. Similarly, the standby unit will never become the active unit. When the active unit comes back up, connections will automatically be directed to the active unit again.



You can tell, at a glance, which unit is active, and which is the standby, by checking the mode in the top bar of the LoadMaster.

Switch to Preferred Server

There are two possible values to select:

- **No Preferred Host:** Each unit takes over when the other unit fails. No switchover is performed when the partner is restarted.
- **Prefer First:** The HA1 (active) unit always takes over. This is the default option.

Partner Name/IP

Specify the host name or IP address of the HA partner unit.

Health Check Port

Set the port over which the health check will be run. The port must be the same on both the active and standby unit for HA to function correctly.

Health Check on All Interfaces

When this option is enabled, the health check listens on all interfaces. This is required when using a multi-arm configuration. If this is disabled, the health check listens on the primary eth0 address (this is the default behavior).

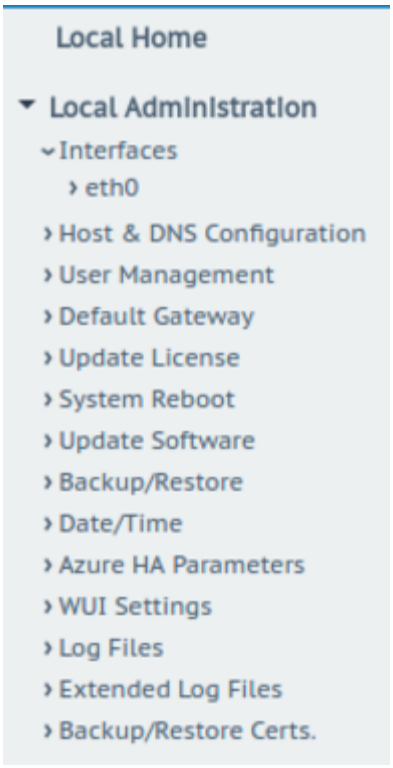

If a unit is in standby mode, WUI access is restricted to **Local Administration** only. Full WUI access is available if the unit is in an active or unchecked state.

AWS HA Parameters

This screen is only available in LoadMaster for Amazon Web Services (AWS) products.

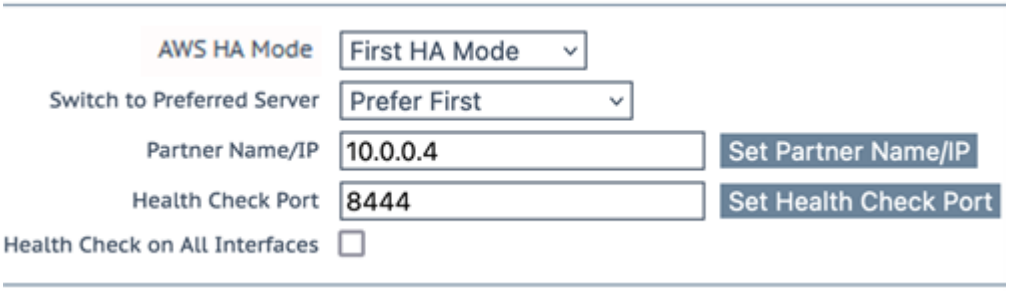

AWS HA Mode

Select the required HA mode for this unit. There are three options:

- First HA Mode
- Standby HA Mode
- Non HA Mode

If you are only using a single LoadMaster, select **Non HA Mode**.

When using HA mode, one machine must be specified as the **First (Active)** and the second machine must

be specified as the **Second (Stand-by)** .

Note: HA will not work if both units have the same value selected for the **AWS HA Mode**.

Note: Synchronization of Virtual Service settings only occurs from the active unit to the standby unit. Changes made to the active unit will be replicated to the standby unit. However, changes made to the standby are never replicated to the active unit.

Note: If the active unit fails, connections will be directed to the standby unit. The active unit is active and will never become the standby, even if it fails. Similarly, the standby unit will never become the active unit. When the active unit comes back up, connections will automatically be directed to the active unit again.



You can tell, at a glance, which unit is active, and which is the standby, by checking the mode in the top bar of the LoadMaster.

Switch to Preferred Server

There are two possible values to select:

- **No Preferred Host:** Each unit takes over when the other unit fails. No switchover is performed when the partner is restarted.
- **Prefer First:**The HA1 (active) unit always takes over. This is the default option.

Partner Name/IP

Specify the host name or IP address of the HA partner unit.

Health Check Port

Set the port over which the health check will be run. The port must be the same on both the active and standby unit for HA to function correctly.

Health Check on All Interfaces

When this option is enabled, the health check listens on all interfaces. This is required when using a multi-arm configuration. If this is disabled, the health check listens on the primary eth0 address (this is the default behavior).



If a unit is in standby mode, WUI access is restricted to **Local Administration** only. Full WUI access is available if the unit is in an active or unchecked state.

Cluster Control

The **Cluster Control** option will only be available on LoadMasters which have a clustering license. To add the clustering feature to your license, please contact a Progress Kemp representative. For further information on clustering, refer to the LoadMaster Clustering Feature Description.

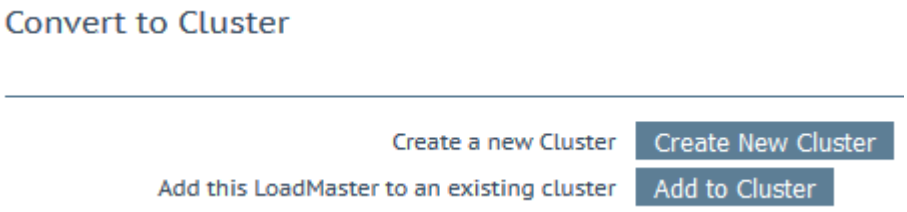

Create New Cluster: If setting up a new cluster, click this button.

Add to Cluster: Add this LoadMaster to an already existing cluster.

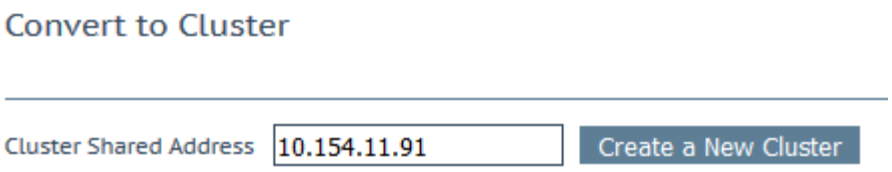

When the **Create New Cluster** button is clicked, the screen above will appear which prompts to set the shared IP address of the cluster. The shared IP address is the address which will be used to administer the cluster.

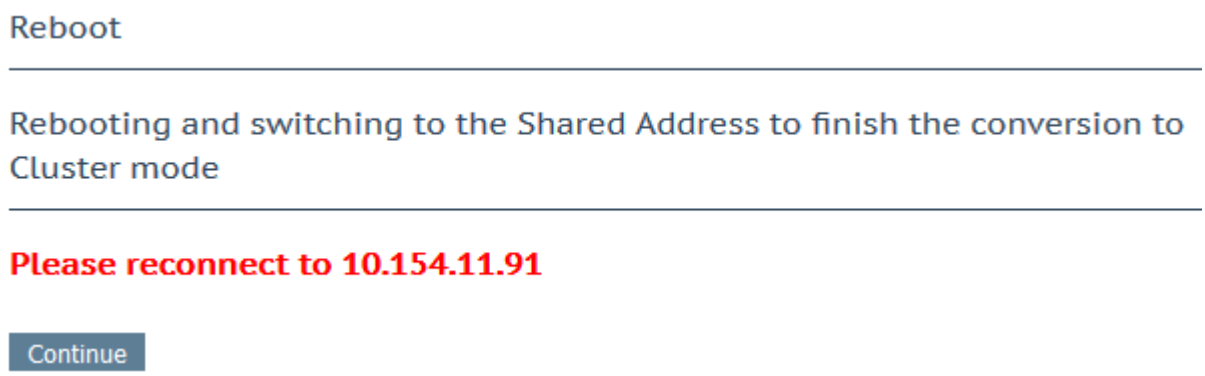

When the **Create a New Cluster** button is clicked, the LoadMaster reboots. A message will appear asking to reconnect to the shared IP address that was just set.

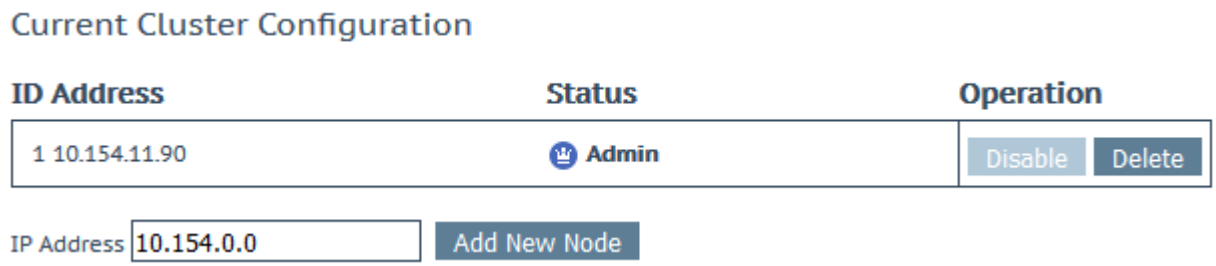

After creating a cluster, the **Cluster Control** screen in the WUI of the shared IP address will allow the addition of LoadMaster nodes into the cluster.

Note: A LoadMaster can only be added to a cluster when the cluster is available and the LoadMaster is waiting to join the cluster. Refer to the LoadMaster Clustering Feature Description for further information and steps.

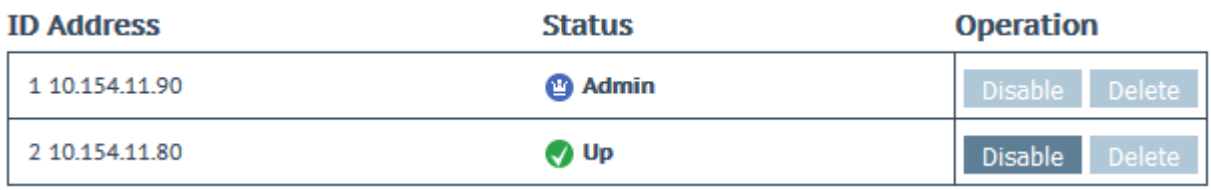

The **Cluster Control** screen, in the shared IP address WUI, displays details for each of the nodes in the cluster.

Show Options: Clicking the **Show Options** button will display the **Cluster Parameters** section which contains two additional fields which can be used to set the **Cluster Virtual ID** and **Node Drain Time**. For further information, refer to the LoadMaster Clustering Feature Description.

ID: The cluster ID.

Address: The IP address of the LoadMaster node. If a second IP address appears in brackets after the first one - the second IP address is the IP address of the interface port. Icons are displayed depending on the status:

Icon	Status	Description
	Admin	The node is the primary control node.
	Disabled	The node is disabled - connections will not be sent to that node.
	Starting	The node is starting (enabling).
	Up	The node is up.
	Down	The node is down.
	Draining	The node has been disabled and the connections are being shut down in an orderly fashion. Drain stopping lasts for 10 seconds by default. This can be updated by changing the **Node Drain Time** value on the **Cluster Control** screen. For more information, refer to the LoadMaster Clustering Feature Description.

Operation: The different operations that can be performed in relation to the notes:

- **Disable:** Disable the node. Nodes that are disabled will first go through drain stopping. During the drain stopping time, the connections are shut down in an orderly fashion. After the drain, the node will be disabled and no traffic will be directed to that node.
- **Enable:** Enable the node. When a node comes up, it will not be immediately be brought into rotation. It will only come online after it has been up for 30 seconds.

- **Delete:** Delete a node from the cluster. When a node is deleted it becomes a regular single LoadMaster instance. If the LoadMaster is later added back in to the cluster, any configuration changes that have been made in the shared IP address will propagate to the node LoadMaster.
- **Reboot:** When performing a cluster-wide firmware update, a **Reboot** button will appear on this screen after uploading the firmware update patch. For step-by-step instructions on how to perform a cluster-wide firmware update, refer to the LoadMaster Clustering Feature Description.

Add New Node: Add a new node with the specified IP address to the cluster.

Cluster Parameters

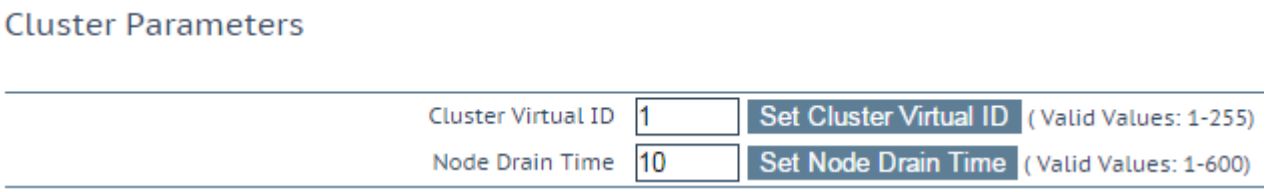

When the **Show Options** button is clicked, the **Cluster Parameters** section appears. This section contains two additional WUI options - **Cluster Virtual ID** and **Node Drain Time**.

Cluster Virtual ID

When using multiple clusters or LoadMaster HA systems on the same network, the virtual ID identifies each cluster so that there are no potential unwanted interactions. The cluster virtual ID is set to **1** by default, but it can be changed if required. Valid IDs range from 1 to 255. Changes made to an admin Loadmaster propagate across all nodes in the cluster.

Node Drain Time

When a node is disabled, the connections that are still being served by the node are allowed to continue for the amount of seconds specified in the **Node Drain Time** text box. No new connections will be handled by the node during this time. The **Node Drain Time** is set to **10** seconds by default, but it can be changed if required. Valid values range from 1 to 600 (seconds).

During the drain time the status changes to Draining until the specified drain time elapses.

When the drain time has elapsed the status changes to disabled.

QoS/Limiting

Refer to the following sections for details on the LoadMaster **QoS/Limiting** WUI options.

Contents

- Global Limits
- Limiter Options
- Client Limiting
- URL Based Limiting

Global Limits



In the **Global Limits** section, you can configure the following options:

- **Maximum Concurrent Connections:** Limit the maximum number of simultaneous connections (combined total of TCP and UDP connections) allowed to the LoadMaster. Setting the limit to 0 disables this option. Valid values are 0 - 100000000.

Note: The maximum values are based on the hardware or Virtual LoadMaster that is in use and may vary per model.

- **Global Connections/s Limit:** Limit the maximum number of connection attempts (per second). Setting the limit to 0 disables this option. Valid values are 0 - 1000000.
- **Global HTTP Requests/s Limit:** Limit the maximum number of HTTP request attempts (per second). This has no effect on non-HTTP traffic. Setting the limit to 0 disables this option. Valid values are 0 - 1000000.

Note: The **Global Limits** take precedence over the other limits configured. For example, if you set the **Client Concurrent Connection Limit** to **5000** but the global **Maximum Concurrent Connections** limit is set to **50**, then 50 is the limit that is enforced.

Note: If the total number of connections from all clients exceed the global limit, they will be dropped.

- **Global Bandwidth Limit:** The global bandwidth limit. Setting the limit to **0** disables bandwidth limiting. Units are in kilobits/second. The minimum value is 16 kilobits/second (2 kilobytes/second). The maximum is 99999999 (which is just under 100 Gbit) but most LoadMasters have a bandwidth limit set in the license and the license bandwidth limit will be enforced if the value specified in the **Global Bandwidth Limit** field is greater than that. When calculating bandwidth, data in both directions are tracked and used in the calculation.

Note: You can also configure per-Virtual Service bandwidth limiting. For further details, refer to the Rate Limiting Feature Description. The lowest of the three possible limits (global, client, and Virtual Service) that is reached first is enforced. Note that the global limit is for all Virtual Services, the Virtual Service limit is for the current Virtual Service with multiple clients, and the client limit is for a single client.

Limiter Options

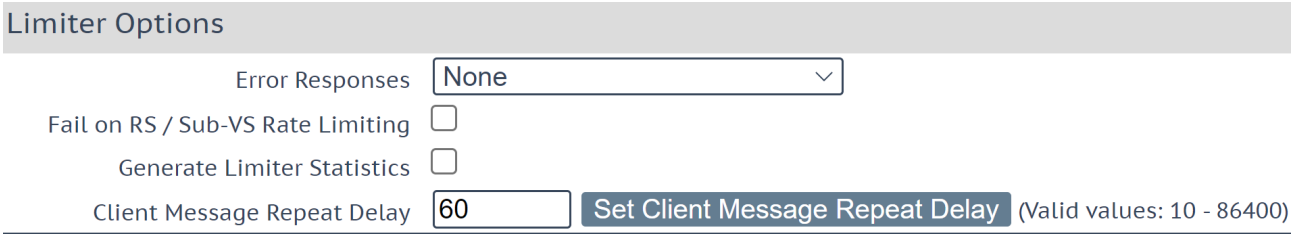

In the **Limiter Options** section, you can configure the following options:

- **Error Responses:** By default, the LoadMaster simply drops any connections when the RPS limit is reached. The system can send a 429 or 503 HTTP error response instead (followed by a close) if you select the appropriate option in this drop-down list.
- **Fail on RS/Sub-VS Rate Limiting:** If rate limiting is activated for a Real Server (RS) or a SubVS, the LoadMaster normally tries to select a different RS/SubVS to use for the connection. Enabling this check box forces the request to fail if the RS that was selected (for example, by persistence) was rate limited. An error response is sent back if one is selected in the **Error Responses** drop-down list.
- **Generate Limiter Statistics:** Enabling this option generates a global summary syslog message every five seconds containing the current state of the limiting QoS subsystem.

Note: This option is disabled by default. Depending upon your client limiting configuration, this can generate a lot of log messages which could be resource intensive.

- **Client Message Repeat Delay:** Set the minimum time after a client is no longer limited before a new message is generated. If a client generates a message and continues to be blocked for continuously hitting the limit, no new message is generated. Only if the client goes quiet for the delay period will a new message be generated. Valid values range from 10 - 86400 seconds. The default value is 60 seconds.

Client Limiting

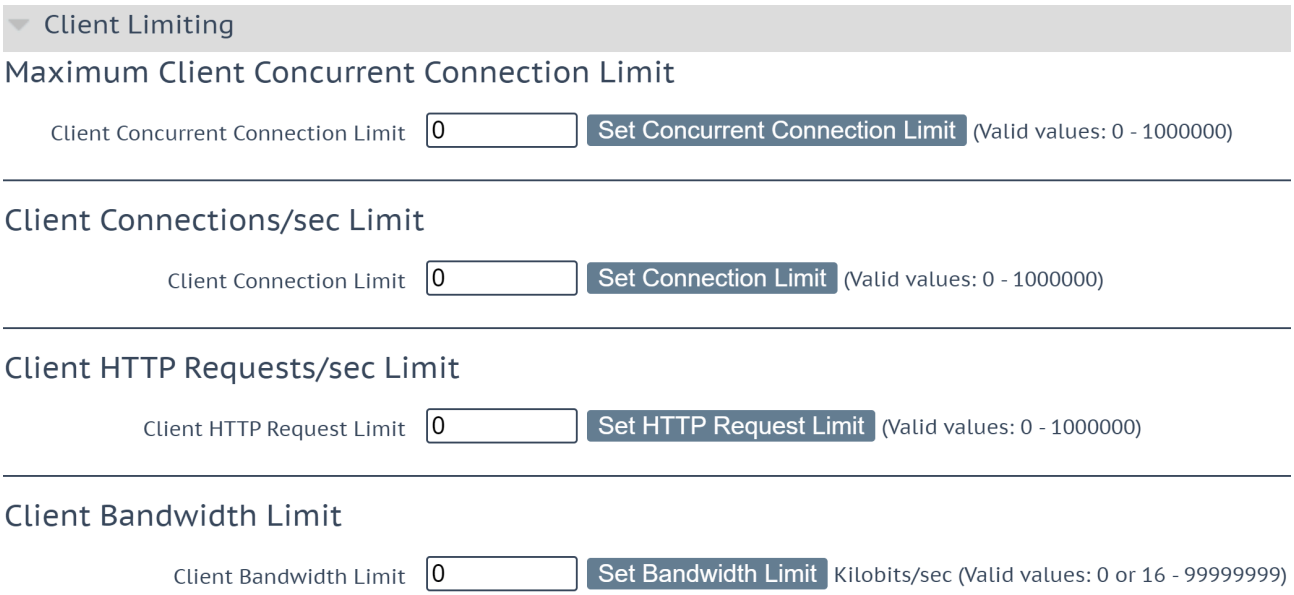

Maximum Client Concurrent Connection Limit

In this section, you must configure the global **Client Concurrent Connection Limit** before you get options to configure concurrent connection limits for particular addresses or networks. The **Client Concurrent Connection Limit** limits the default maximum number of concurrent connection attempts (per second) from a specific host. Setting the limit to 0 disables this option. Valid values range from 0 - 1000000.

Client Connections/sec Limit

In this section, you must configure the global **Client Connection Limit** before you get options to configure the CPS limits for particular addresses or networks. The **Client Connection Limit** limits the default maximum number of connection attempts (per second) from a specific host. Setting the limit to 0 disables this option. Valid values range from 0 - 1000000.

Client HTTP Requests/sec Limit

In this section, you must configure the global **Client HTTP Request Limit** before you get options to configure the RPS limits for particular addresses or networks. The **Client HTTP Request Limit** limits the default maximum number of HTTP request attempts (per second) from a specific host. This has no effect on non-HTTP traffic. Setting the limit to 0 disables this option. Valid values range from 0 - 1000000.

Client Bandwidth Limit

In this section, you must configure the global **Client Bandwidth Limit** before you get options to configure the bandwidth limits for particular addresses or networks. The **Client Bandwidth Limit** limits the default maximum number of bandwidth attempts (per second) from a specific host. Setting the limit to **0** disables this option. Units are in kilobits/second. The minimum value is 16 kilobits/second (2 kilobytes/second). The maximum is the value configured in the **Global Bandwidth Limit**. When calculating bandwidth, data in both directions is tracked and used. This means both the client and server-side data is tracked and used as part of the calculation.

For more information and instruction, refer to the Rate Limiting Feature Description.

URL Based Limiting

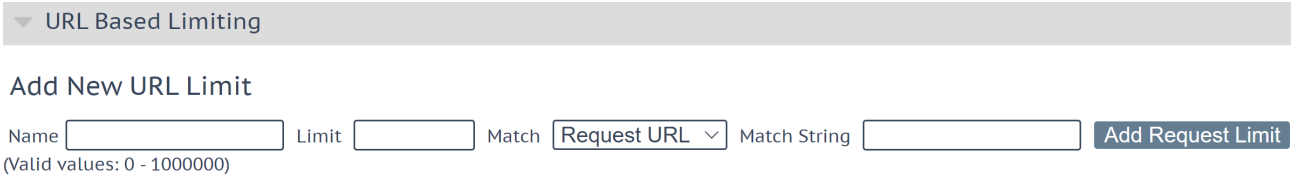

The **URL Based Limiting** is based on options in a HTTP request. A request consists of a URL, Method, and request headers. **Host** and **User-Agent** are request headers. The LoadMaster URL-based limiting rules inspect based on what is selected in the **Match** drop-down list (**Request URL**, **Host**, **User Agent**, **Method**, **!Request URL**, **!Host**, **!User Agent**, or **!Method**. If the limit is hit the LoadMaster sends a response code (as set in the **Error Responses** drop-down list in the **Limiter Options** section).

For more information and instruction, refer to the Rate Limiting Feature Description.

System Administration

These options control the base-level operation of the LoadMaster. It is important to know that applying changes to these parameters in a HA pair must be done using the floating management IP. Many of these options will require a system reboot. When configuring these parameters, only the active system in a pair is affected.

Contents

- User Management
- License Management
- System Reboot
- Update Software
- Backup/Restore
- Date/Time

User Management

The content below describes the different user management WUI fields. For further information on user management and WUI authentication, refer to the User Management Feature Description.

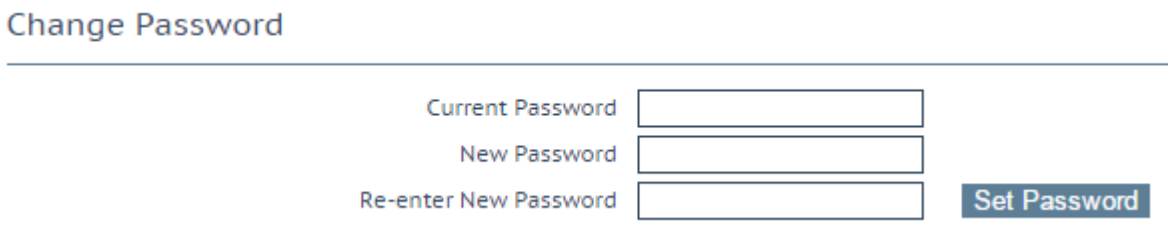

The **Change Password** section can be used to change the appliance password. This is a local change only

and does not affect the password of the partner appliance in a HA deployment.



Minimum Password length

Set the minimum password length for all local user passwords. After selecting a different value for this field, you must refresh the page for the new value to be enforced.

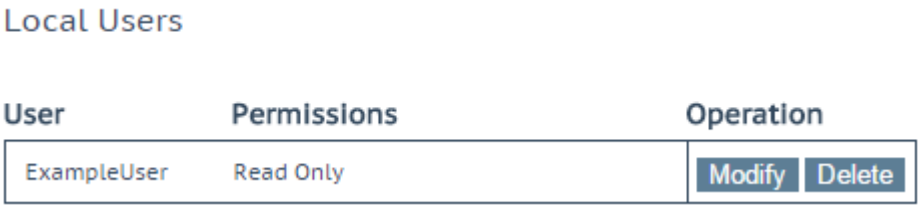

The **Local Users** section lists any existing local users. Two options are available for existing users:

- **Modify:** Change details for an existing local user, such as their permissions and password. For further information, refer to the Modify User section.
- **Delete:** Delete the relevant user.



New users can be added in the **Add User** section.

Usernames can be a maximum of 64 characters long. Usernames can start with a digit and can contain alphanumeric characters, in addition to the following special characters:

=~^._+#@/-

The minimum password length is defined by what is set in the **Minimum password length** field. All characters are allowed.

The **Use RADIUS Server** option allows you to determine if RADIUS server authentication will be used when the user is logging in to the LoadMaster. The RADIUS Server details must be setup before this option can be used.

When RADIUS authentication is in use, the LoadMaster passes the user's details to the RADIUS server and the RADIUS server informs the LoadMaster whether the user is authenticated or not. For further details on how to configure the RADIUS Server details please refer to the WUI Authentication and Authorization section and the RADIUS Authentication and Authorization, Technical Note.

Note: When Session Management is enabled, the Use RADIUS Server option is not available within this screen. Please refer to the WUI Authentication and Authorization section for further information on how to configure RADIUS Server when Session Management is enabled.

When Session Management is enabled, a check box called **No Local Password** will be displayed in the **Add User** section. This option can be enabled if client certificate authentication or SAML authentication will be used to authenticate this user when they are accessing the LoadMaster. To enable client certificate authentication, set the **Admin Login Method** in the **Remote Access** screen. For further information, refer to the Remote Access section or the User Management Feature Description.

Note: Certificate-based authentication will be deprecated at some point in the future.

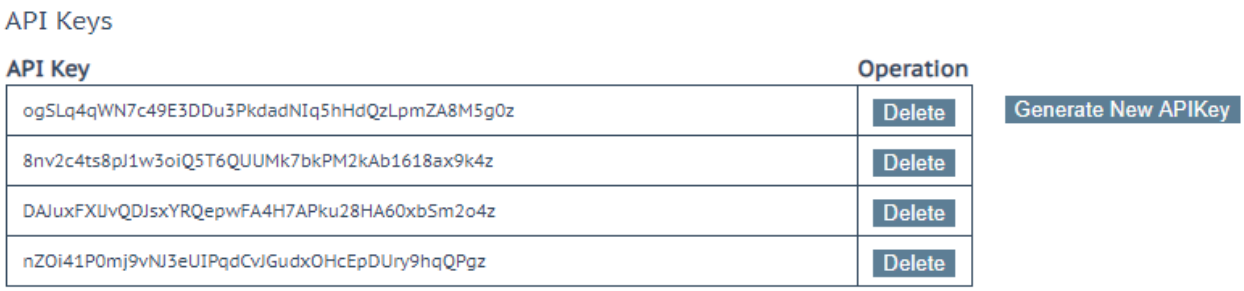

When running API commands, you can authenticate using an API key. An API key is a unique identifier used to authenticate a user.

The **API Keys** section on the **User Management** screen displays any API keys currently generated for the logged in user. You can have up to 16 API keys per user - if you try to create more, the oldest is silently deleted. The oldest API key is listed at the top.

To generate an API key for a specific user, go to the **Modify** screen for that particular user.

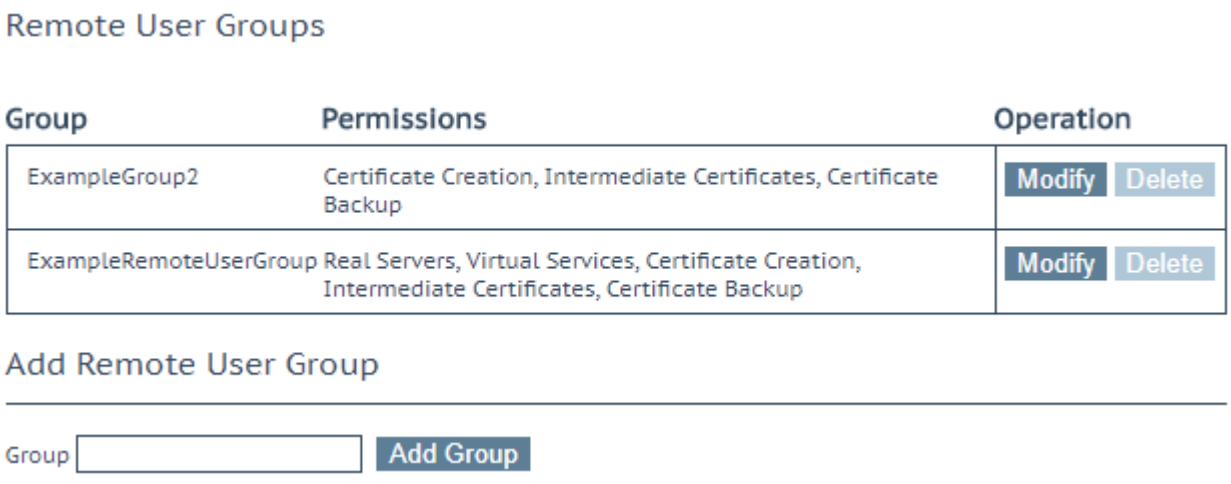

The **Remote User Groups** section displays any remote user groups that have been created. The group name and associated permissions are displayed. These groups can be selected for LDAP WUI authentication in the following location: **Certificates & Security > Remote Access > WUI Authorization Options**. For further information, refer to the WUI Authentication and Authorization section.

Note: It is important to select and apply the group, or groups. If there are no groups selected, no group checking is performed and remote users can log in without a group check.

You can add a new remote user group by entering a name and clicking **Add Group**.

Note: The following characters are permitted in the group name: alphanumeric characters, spaces, or the following special symbols: **=~^._+#,@/-**.

Click **Modify** to edit the group permissions.

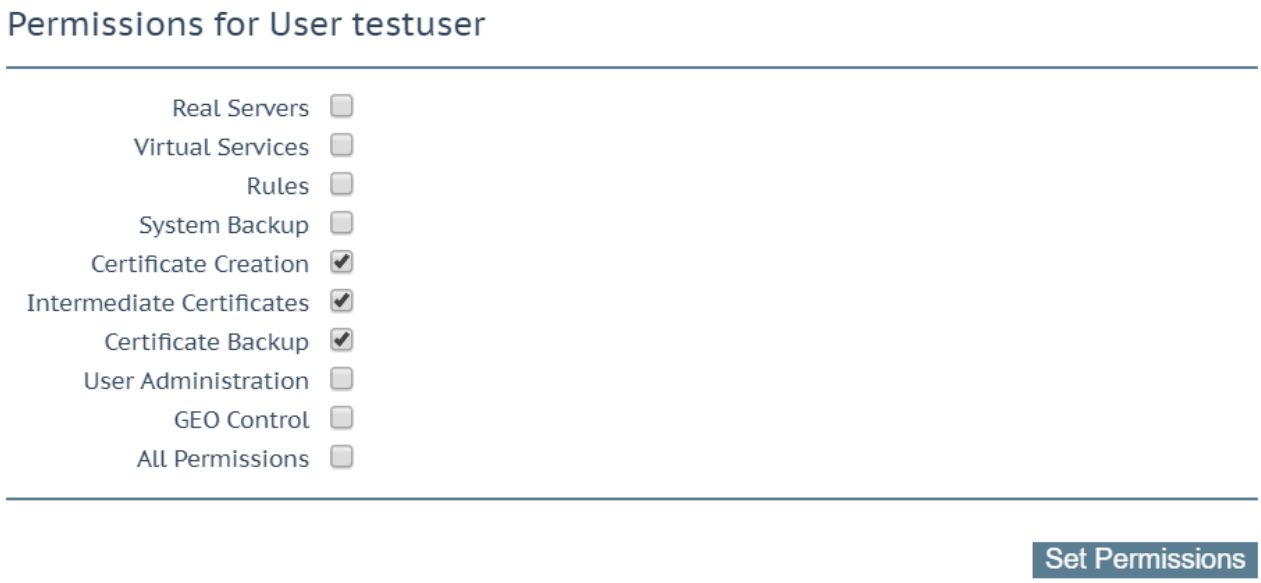

For further information regarding group permissions, refer to the User Management Feature Description.



There are two permissions relating to Virtual Services - **Virtual Services** and **Add Virtual Services**.

The **Add Virtual Services** permission is only visible when the **Allow Extended Permissions** check box is selected on the **User Management** screen.

The Virtual Service operations allowed differ based on what combination of options you have selected. For a summary of these connotations, refer to the table below:

Allow Extended Permissions	Virtual Services	Add Virtual Service	Operations Allowed	Operations not Allowed
Enabled	Enabled	Disabled	• View existing Virtual Services	• Add Virtual Service

Allow Extended Permissions	Virtual Services	Add Virtual Service	Operations Allowed	Operations not Allowed
			• Modify existing Virtual Services • Change Virtual Service port	• Duplicate Virtual Service • Change Address • Export template
Enabled	Disabled	Enabled	• View existing Virtual Services	• Add Virtual Service • Duplicate Virtual Service • Change Address • Export template • Modify existing Virtual Services • Change Virtual Service port
Enabled	Enabled	Enabled	• Add Virtual Service • Duplicate Virtual Service • Change address • Export template • View existing Virtual Services • Modify existing Virtual Services • Change Virtual Service port	Not applicable
Enabled	Disabled	Disabled	View existing Virtual Services	Not applicable
Disabled	Enabled	Disabled	• Add Virtual Service • Duplicate Virtual Service • Change address	Not applicable

Allow Extended Permissions	Virtual Services	Add Virtual Service	Operations Allowed	Operations not Allowed
			• Export template • View existing Virtual Services • Modify existing Virtual Services • Change Virtual Service port	
Disabled	Disabled	Disabled	View existing Virtual Services	• Add Virtual Service • Duplicate Virtual Service • Change address • Export template • Modify existing Virtual Service • Change Virtual Service port

Modify User

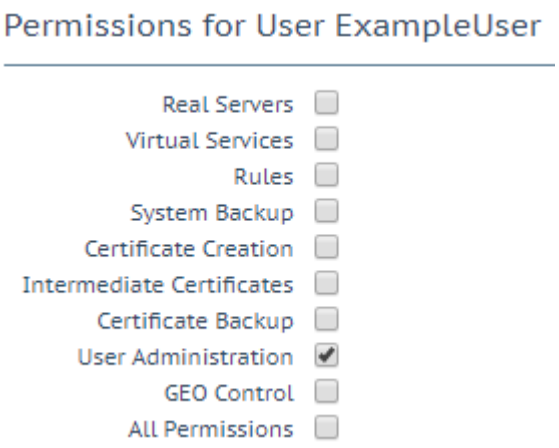

In this screen you may set the level of user permissions. This determines what configuration changes the user is allowed to perform. The primary user (**bal**) always has full permissions. Secondary users may be restricted to certain functions.

For further information regarding user permissions, please refer to the User Management Feature Description.



The **Change Password** section can be used to modify a user's password. It is also possible to enable and disable RADUIS server authentication for a user.

Note: When Session Management is enabled, the Use RADIUS Server option is not available within this screen. Please refer to the WUI Authentication and Authorization section for further information on how to configure RADIUS Server when Session Management is enabled.

When Session Management is enabled, a check box called **No Local Password** will be displayed in the **Change Password** section. This option can be enabled if client certificate authentication will be used to authenticate this user when they are accessing the LoadMaster. To enable client certificate authentication, set the **Admin Login Method** in the **Remote Access** screen. For further information, refer to the Remote Access or the User Management Feature Description.

Named users, even those without User Administration privileges, can change their own passwords. When a named user clicks the **System Administration > User Management** menu option the **Change Password** screen appears.

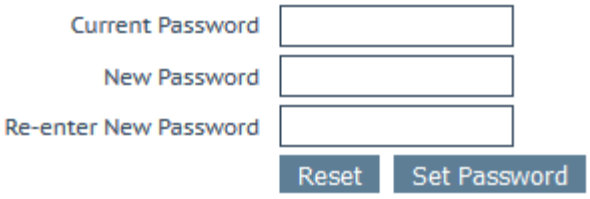

From within this screen, users can change their own password. The minimum password length is defined by what is set in the **Minimum password length** field. All characters are allowed, except **\"'`'**. Once changed, a confirmation screen appears after which the users will be forced to log back in to the LoadMaster using their new password.



When running API commands, you can authenticate using an API key. An API key is a unique identifier used to authenticate a user.

The **API Keys** section on the **Modify** user screen displays any API keys currently generated for that specific user. You can have up to 16 API keys per user - if you try to create more, the oldest is silently deleted. The oldest API key is listed at the top.

To generate an API key for a specific user, click **Generate New APIKey**.

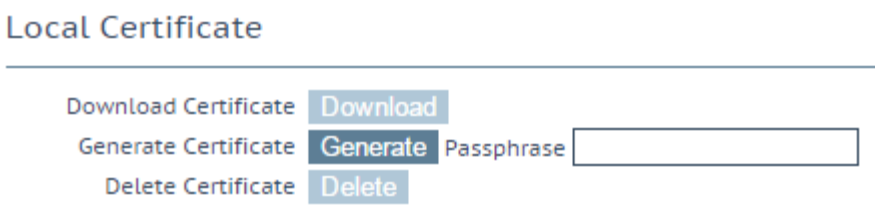

In the **Local Certificate** section, a certificate can be generated for the user. A **Passphrase** can be optionally set which is used to encrypt the private key. Once that certificate has been downloaded, it can be used as a client certificate to allow password-less access to the LoadMaster API. Users with 'User Administration' permissions are able to manage local certificates for themselves and other users.

To enable client certificate authentication to the LoadMaster, set the **Admin Login Method** in the **Remote Access** screen. For further information, please refer to the Remote Access section or the User Management Feature Description.

License Management

Note: In LoadMaster firmware version 7.2.53, the **Update License/Owner** button is used to update the LoadMaster license to change the ownership of the LoadMaster license. For further details, refer to the Licensing Feature Description.

This screen displays the activation date and the expiration date of the current license. You would use the **License Management** function if your license has changed, for example if:

- You have renewed support
- You have renewed your license
- You have changed your license type

Before updating the license in the LoadMaster, you must either contact your Progress Kemp representative, or use the **Upgrade** option displayed on the home page of UI. After you have contacted Progress Kemp or used the upgrade option, there are two main ways to update a license – using the **Online** method and using the **Offline** method. Refer to the sections below to find out details about the screens for each method.

You can also license a LoadMaster using LoadMaster 360. At present, LoadMaster 360 only supports perpetual LoadMaster licenses and pooled LoadMaster licensing. Service Provider License Agreements (SPLA) are not currently supported by LoadMaster 360. For details on how to update a LoadMaster license using LoadMaster 360, refer to the Technical Note - LoadMaster Licensing in LoadMaster 360 document.

For more information and instructions, refer to the Licensing Feature Description.

A **Kill License** button may appear on the **License Management** screen:

- If your LoadMaster license is either a Permanent (PERM), Temporary (TEMP), or Service Provider License Agreement (SPLA) license that was obtained from the Progress Kemp Licensing Server (either online or offline), then clicking the **Kill License** button will invalidate your license on the Progress Kemp licensing server and in the inventory of the Progress ID that was used to license the LoadMaster.
- If your LoadMaster license is a Service Provider License Agreement (SPLA) license that was obtained from a local Kemp 360 Central instance in your configuration, then clicking the **Kill License** button will return the LoadMaster license to the available license pool on Kemp 360 Central, and the LoadMaster will

return to an unlicensed state.

Note: Do not click **Kill License** unless instructed to by Progress Kemp Support.

Online Method

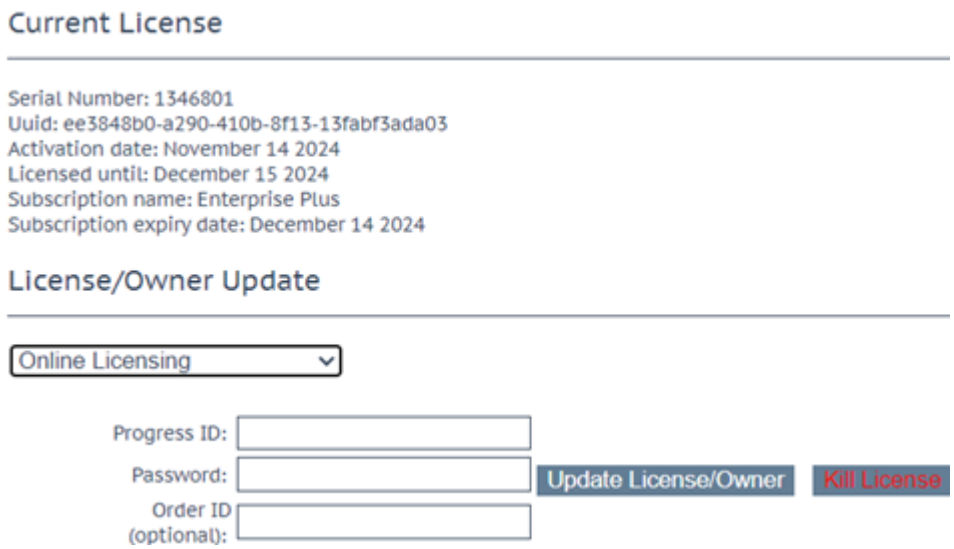

To upgrade the license using the online method, the LoadMaster must be connected to the internet. You will need to enter your Progress ID (formerly Kemp ID) and **Password** to license using the online method.

A reboot is recommended after updating the license.

Note: Do not click **Kill License** unless instructed to by Progress Kemp Support.

Offline Method

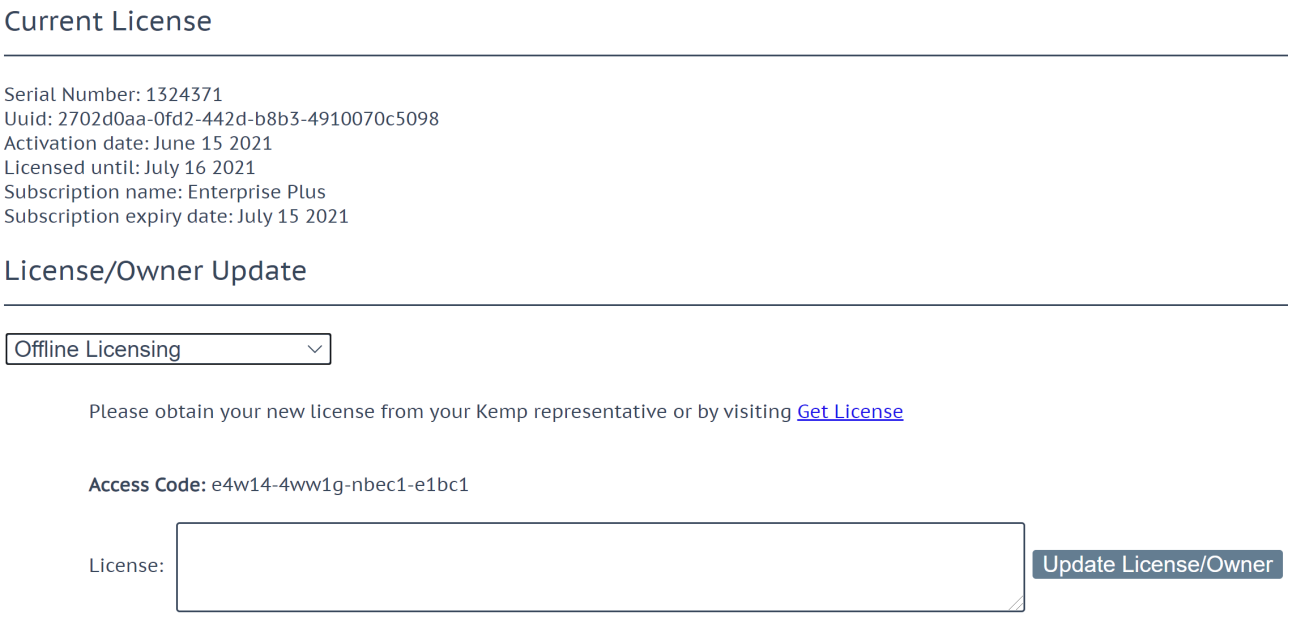

To upgrade the license using the offline method, you need to enter license text in the LoadMaster. You can either get this from Progress Kemp or using the **Get License** link. For more information and instructions, refer to the Licensing Feature Description.

A reboot is recommended after updating the license. If upgrading to an ESP license, a reboot is required after the update.

Debug Checks

If you try to license and there are any issues, a number of checks are performed automatically and the results and associated error message are displayed.



These checks perform the following tasks:

- Ping Default Gateway
- Ping DNS Servers
- Ping Licensing Server

System Reboot

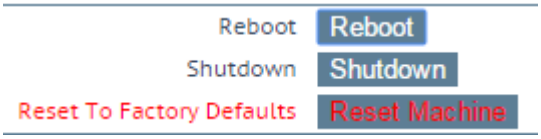

Reboot

Reboot the appliance.

Shutdown

Clicking this button attempts to power down the LoadMaster. If, for some reason, the power down fails, it will at a minimum halt the CPU.

Reset Machine

Reset the configuration of the appliance with exception of the license and username and password information. This only applies to the active appliance in a HA pair.

Note: If you perform a factory reset of the LoadMaster, you must enable the **Enable Session Management** check box and disable the **Require Basic Authentication** check box to successfully run APIv2 (JSON-format) requests.

Update Software

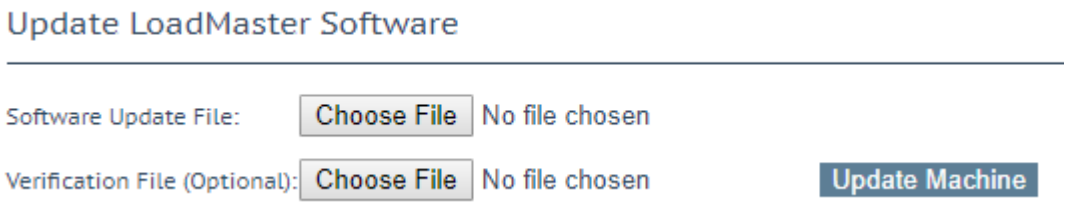

You can download firmware patches from the Downloads page.

The **Verification File** field appears if the **Update Verification Options** field is set to **Required** or **Optional** in **System Configuration > Miscellaneous Options > WUI Settings** (this is set to **Required** by default).

By default, the integrity of a patch or add-on is verified at install time. This is done without the use of a secondary XML verification file.

Setting the **Update Verification Options** to **Required** or **Optional** allows the use of a secondary XML verification file to be uploaded at the same time as you are uploading a patch or add-on, which allows you to verify the digital signatures used to sign the update file archive. The patch/add-on will not be installed if the patch cannot be verified against the XML file.

If you are upgrading a LoadMaster to version 7.2.51 or above and it is set to require validation, you must supply one of the two XML verification files supplied with this release:

- 7.2.<release/build number>.RELEASE.PATCH-64-MULTICORE-preV7.2.51.0.checksum.xml: Use this file when upgrading to a LoadMaster running a release that is prior to LMOS 7.2.51.
- 7.2.<release/build number>.RELEASE.PATCH-64-MULTICORE.checksum.xml: Use this file when repeating an upgrade to a LoadMaster that is already running 7.2.51.0 or above and you want to repeat the upgrade process.

If the **Update Verification Options** drop-down list is set to **Required** it is not possible to upgrade to LoadMaster firmware 7.2.51 or above. It is also not possible to downgrade from 7.2.51 to a version previous to that. To workaround this, set the **Update Verification Options** field to **Optional** or **No verification file - deprecated** before downgrading or upgrading between the specified versions.

Note:

When updating LoadMaster hardware, starting with 7.2.54.5, you can only update a hardware LoadMaster with an image that supports the boot method used by the hardware. The newer NG models (introduced in 2023) use the UEFI boot method, while earlier models use the BIOS boot method. If you attempt to update a hardware appliance with an updated image that doesn't support the boot method used by the appliance, the update process stops and the following message is displayed:

The system cannot be updated with the provided image, which is not supported on the hardware platform on which you are attempting to install it.

All LoadMaster -NG hardware models (introduced in 2023) require UEFI-capable patches, and so cannot be updated with any patch earlier than 7.2.54.6, or between 7.2.55.0 and 7.2.59.1 (inclusive).

For more information, refer to the knowledge base article: LoadMaster Firmware Upgrade Path.

Upgrading to 7.2.61 or above (legacy WAF removal)

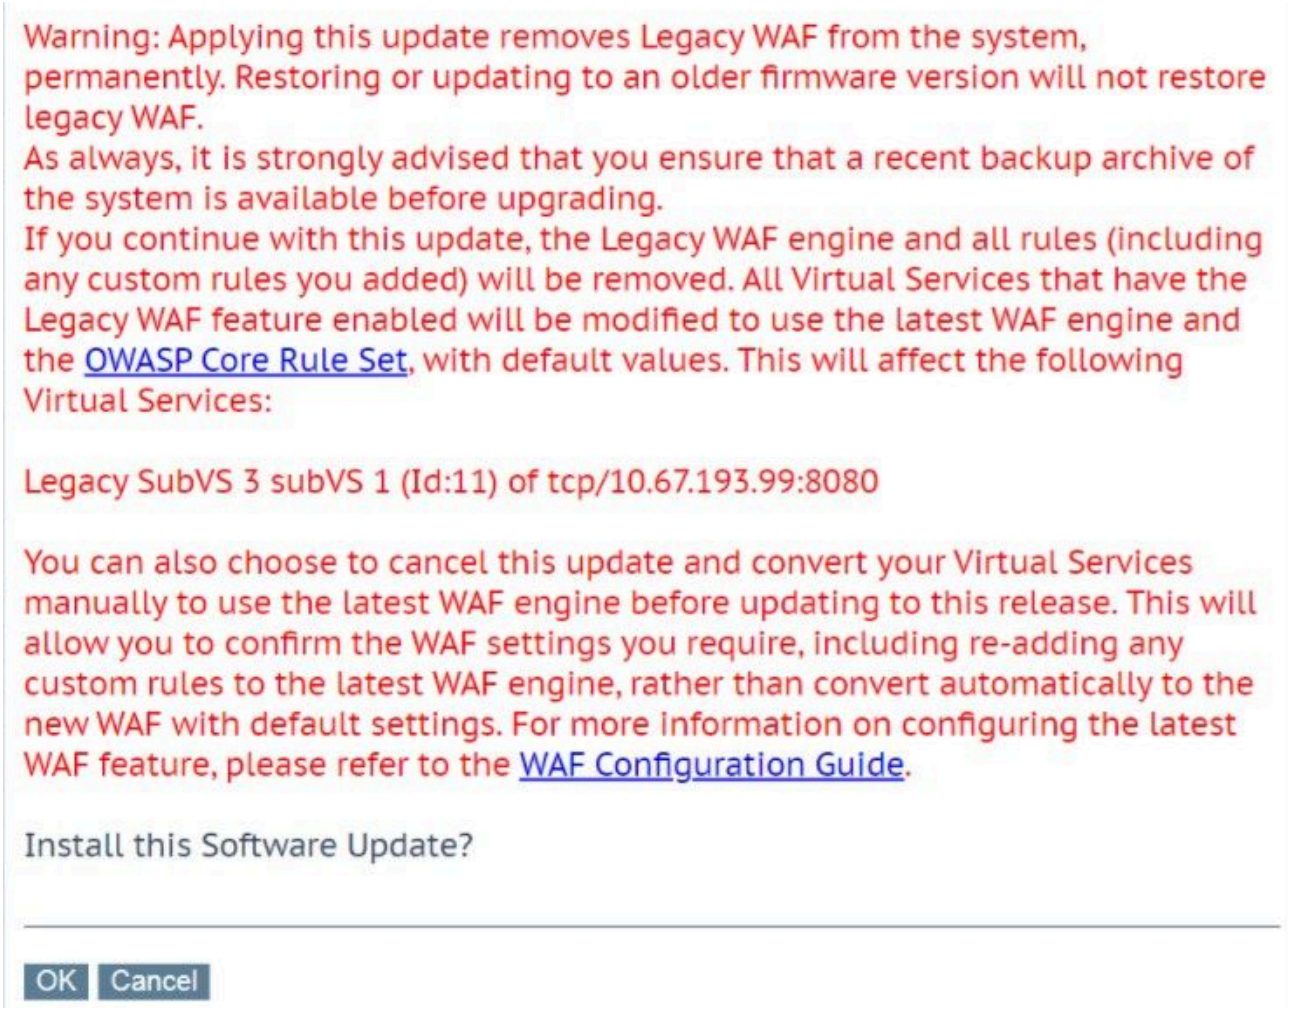

If you upgrade from a LoadMaster version prior to 7.2.61, during the upgrade any legacy WAF-enabled Virtual Services get converted to use the OWASP WAF with the default values configured. During upgrade, you get notified with a warning message that indicates the list of Virtual Services and SubVSs that have legacy WAF enabled. To continue to use the legacy WAF functionality, click **Cancel** (to remain on the previous version) or click **OK** to proceed with the LoadMaster firmware version upgrade. For more details, refer to the following section: Converting from Legacy WAF to OWASP WAF of WAF Feature description document.

When you upgrade to 7.2.61, legacy WAF functionality is removed from the LoadMaster. If you downgrade from 7.2.61 or above to a version below 7.2.61, legacy WAF will not be restored.

Update Machine

Note: You cannot directly downgrade from version 7.2.62 (or above) to any release version earlier than 7.2.61. If you try to, an error message appears that says "Invalid software update". To downgrade to a version earlier than 7.2.61, first downgrade to 7.2.61 and then downgrade again.

After you have downloaded the firmware you can unzip it. You can then browse to the file and upload the firmware directly onto the LoadMaster. The firmware will be unpacked and validated on LoadMaster. If the patch is validated successfully you will be asked to confirm the release information. To complete the update you will need to reboot the appliance. This reboot can be deferred if needed.

Note: A checksum file is provided inside the zipped firmware download. Note that this checksum file matches the firmware file itself (contained within the zip file). It does not match the zipped download. The checksum available for download on the Progress Kemp website matches the zipped firmware download.

Update Cluster

Note: The **Update Cluster** option will only be available on LoadMasters which have a clustering license. To add the clustering feature to your license, please contact a Progress Kemp representative. For further information on clustering, refer to the LoadMaster Clustering Feature Description.

The firmware on all LoadMasters in a cluster can be updated using the shared IP address by clicking the **Update Cluster** button. For step-by-step instructions on how to perform a cluster-wide software update, refer to the LoadMaster Clustering Feature Description.

Restore Software

If you have completed an update of LoadMasters firmware you can use this option to revert to the previous build.

Installed Addon Packages

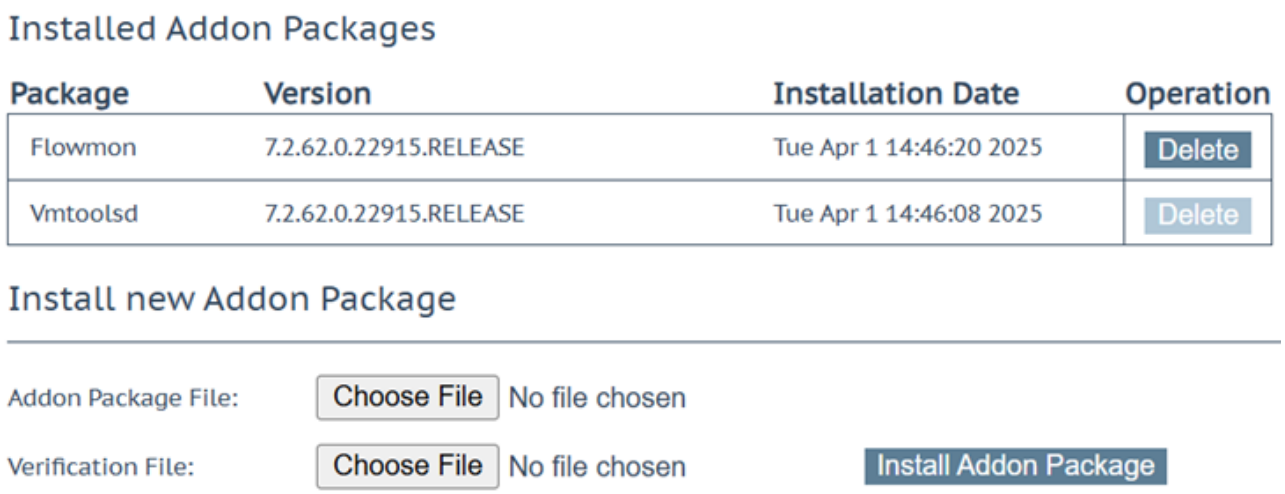

Add-on packages can be installed in the LoadMaster. Add-on packages provide features that are additional to those already included in the LoadMaster. We plan on creating further add-on packages in the future.

Note: Progress Kemp only support add-on packages that can be downloaded from the LoadMaster Add-on Packages page.

Note: If an add-on package of the same name is uploaded, the existing one will be overwritten/updated.

Note: The **Verification File** field appears if the **Display Verify Update Option** is selected in **System Configuration > Miscellaneous Options > WUI Settings** (this is enabled by default). By default, the integrity of a patch or add-on is verified at install time. This is done without the use of a secondary XML verification file.

Note: If an installed add-on package cannot be started, the text will display in red and the hover text will show that the package could not be started.

How to Install an Add-On Package

Follow these steps to install an add-on package on a standalone unit or on two units configured into an High Availability (HA) pair. If you are using clustering, perform this procedure on each unit in the cluster as if it was a standalone unit:

1. Download the add-on package and the associated XML verification file from the LoadMaster Add-on Packages page. The XML verification file checksum refers to the add-on package itself, not the zip archive. Unzip the add-on package before proceeding.
2. If the LoadMaster to which you are applying the add-on is *not* in HA mode, skip this step. Otherwise, if you have a pair of LoadMasters configured into HA mode, apply the add-on first to the standby unit:
   1. Log into the current **standby** unit using that unit's IP administrative IP address.
   2. Go to **Local Administration > Update Software**.
   3. Under **Install new Addon Package** click the **Choose File** button next to **Addon Package File**.

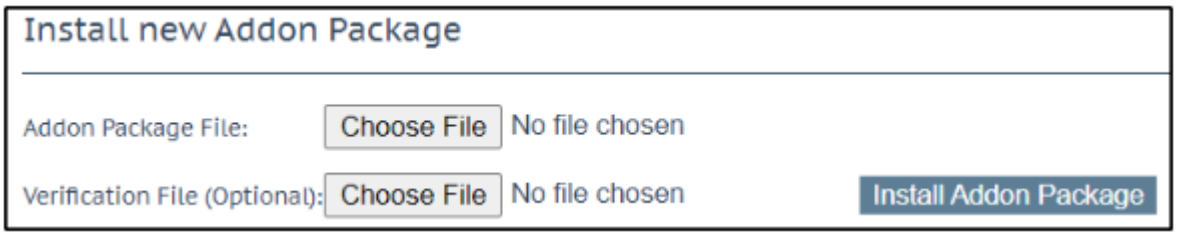

   4. Browse to the add-on package file downloaded in step 1, select the file, and click **Open**.
   5. Click the **Choose File** button next to **Verification File**.
   6. Browse to the add-on package verification file downloaded in step 1, select the file, and click **Open**.
   7. Click **Install Addon Package**.
   8. After the installation is finished, go to **System Configuration > System Administration > System Reboot** and click the **Reboot** button. The unit will return to standby mode after the reboot completes.
3. For a standalone/clustered unit, or the active unit in an HA pair, log into the administrative interface of the unit; or, for an HA pair, log into the shared IP address (which connects to the unit currently in active mode).
4. Go to **System Configuration > System Administration > Update Software**.
5. Under **Install new Addon Package** click the **Choose File** button next to **Addon Package File**.

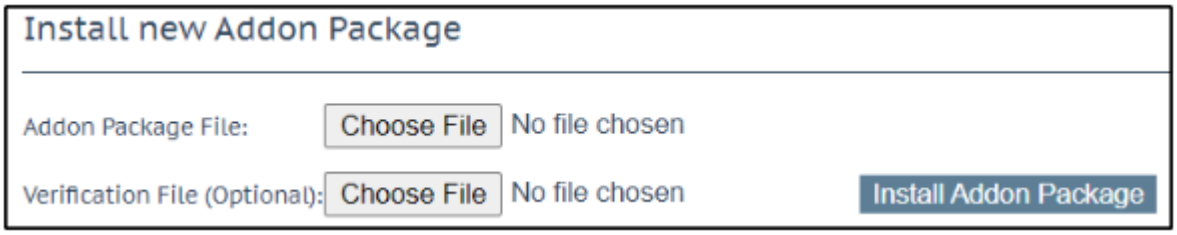

6. Browse to the add-on package file downloaded in step 1, select the file, and click **Open**.
7. Click the **Choose File** button next to **Verification File**.

8. Browse to the add-on package verification file downloaded in step 1, select the file, and click **Open**.
9. Click **Install Addon Package**.
10. After the installation is finished, go to **System Configuration > System Administration > System Reboot** and click the **Reboot** button. If this is the active unit in an HA pair, the system will either assume the standby role after the reboot or will attempt to assume the active role, depending on the configuration of the unit.

Backup/Restore

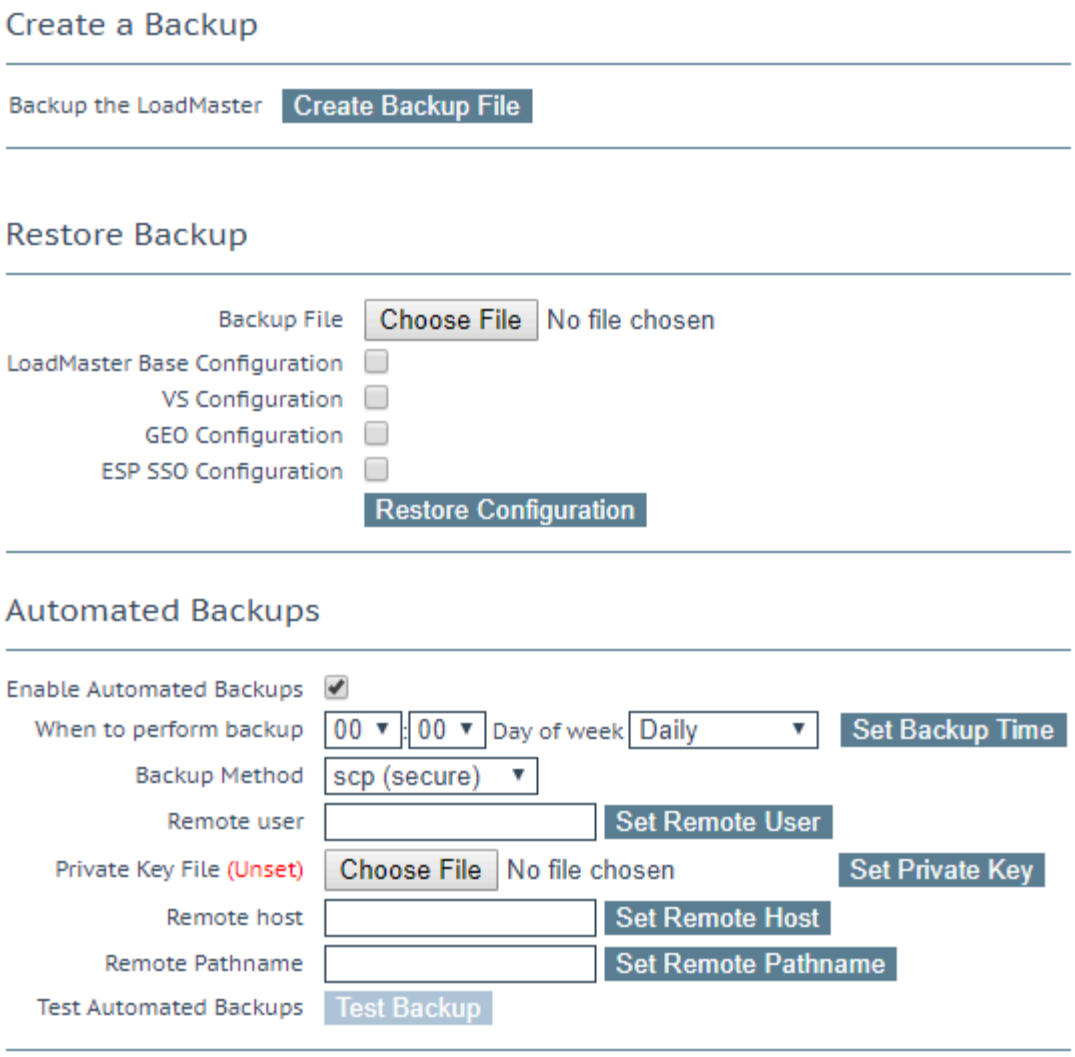

Create Backup File

Generate a backup that contains the Virtual Service configuration, the local appliance information and statistics data. License information and SSL Certificate information is not contained in the backup.

For further details on what is not included in LoadMaster backups, refer to the following page: Important Notes.

For ease of identification, the Backup file name includes the LoadMaster's hostname.

By default, the LoadMaster includes a Netstat output in backups taken. When this is included, backups take longer to complete. You can stop the Netstat output from being included by disabling the **Include Netstat in Backups** option in the **Troubleshooting** screen (**System Configuration > Troubleshooting**).

Restore Backup

When performing a restore (from a remote machine), the user may select what information should be restored:

- **VS Configuration**
- **LoadMaster Base Configuration**
- **GEO Configuration**
- **ESP SSO Configuration** (This restores the SSO domains, LDAP endpoints and SSO custom image sets. This does not restore the Virtual Service settings - use the **VS Configuration** option to restore those.)
- A combination of the options

Note: It is not possible to restore a single machine configuration onto a HA machine or restore a HA configuration onto a single machine.

Note: It is not possible to restore a configuration with ESP-enabled Virtual Services onto a machine which is not enabled for ESP.

Note: A WAF configuration can only be restored onto a LoadMaster with a WAF license.

Note: If you restore a backup taken on LoadMaster version 7.2.58 or below (with Legacy Web Application Firewall (WAF) enabled on some Virtual Services) to a fresh install of 7.2.59 or above (where Legacy WAF is no longer available) - the LoadMaster will enable OWASP WAF on the affected Virtual Services with default WAF settings. If this happens, a message is displayed listing the Virtual Services and SubVSs that were affected. You should review these Virtual Services and make any modifications needed.

Automated Backups

If the **Enable Automated Backups** check box is selected, the system may be configured to perform automated backups on a daily or weekly basis.

For ease of identification, the Backup file name includes the LoadMaster's hostname.

If the automated backups are not being performed at the correct time, ensure the NTP settings are configured correctly. For further information, refer to the Date/Time section.

When to perform backup

Specify the time (24 hour clock) of backup. Also select whether to backup daily or on a specific day of the week. When ready, click the **Set Backup Time** button.

Note: In some situations, spurious error messages may be displayed in the system logs, such as:Dec 8 12:27:01 Kemp_1 /usr/sbin/cron[2065]: (system) RELOAD (/etc/crontab)Dec 8 12:27:01 Kemp_1 /usr/sbin/cron[2065]: (CRON) bad minute (/etc/crontab)

Note: These can be safely ignored and the automated backup will likely still complete successfully.

Backup Method

Select the file transfer method for automated backups:

- Ftp (insecure)
- scp (secure)
- sftp (secure)

If using scp or sftp, the **Private Key File** must be supplied.

Remote user

Set the username required to access remote host.

Private Key File

If using scp as the backup method, the **Private Key File** must be provided. This is the SSH private key generated using ssh-keygen on the remote scp server.

Remote password

The **Remote password** is used when the **Backup Method** is set to **Ftp (insecure)**. Set the password required to access remote host. This field accepts alphanumeric characters and most non-alphanumeric characters. Disallowed characters are as follows:

- Control characters
- ‘ (apostrophe)
- ` (grave)
- The delete character

Remote host

Set the IP address or hostname of the remote host to which you want the backup archives sent, optionally followed by a colon and the port number. If no port is specified, the default port for the selected protocol is used.

Remote Pathname

Set the location on the remote host to store the file.

Test Automated Backups

Clicking the **Test Backup** button performs a test to check if the automated backup configuration is working correctly. The results of the test can be viewed within the System Message File.

Date/Time

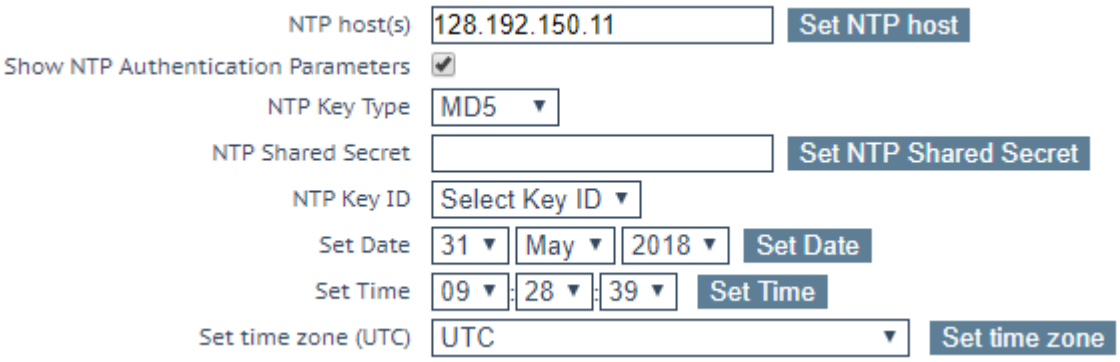

You can manually configure the date and time of LoadMaster or leverage an NTP server.

NTP host(s)

Specify the host which is to be used as the NTP server. NTP is a strongly preferred option for a HA cluster. For a single unit it is at the user's discretion. Clicking the **Set NTP host** button will refresh the time based on the details configured.

If you do not have a local NTP server, refer to www.pool.ntp.org for a list of public NTP server pools which can be used.

Note: The time zone must always be set manually.

Show NTP Authentication Parameters/Disable NTP Authentication

The LoadMaster supports NTPv4 which uses cryptographic signing to query a secure NTP server. This uses a simple authorization scheme which uses a shared secret and key to validate that the response from the server is actually valid. Enable the **Show NTP Authentication Parameters** check box to display the parameters that are needed to support NTP authenticated requests. If you select the **Show NTP Authentication Parameters** checkbox and change any of the parameters, the name of the check box changes to **Disable NTP Authentication**.

Note: For the NTPv4 feature to work, a file must be created on the server (/etc/ntp.keys), which has the following format:

<keyid> M <secret string> ... <keyid> M <secret string>

To enable the use of the key, specify the keyid in the trustedkey line of /etc/ntp.conf, for example, if the keyid is 5 then you have to specify "trustedkey5". The trustedkey value can take multiple values, for example trustedkey 1 2 3 4 5 9 10).

NTP Key Type

Select the **Show NTP Authentication Parameters** check box to display the **NTP Key Type** drop-down list. Select either the MD5, SHA-1, or legacy SHA NTP key type. **MD5** is the default value.

NTP Shared Secret

The NTP shared secret string. The NTP secret can be a maximum of 20 ASCII characters long or 40 hexadecimal characters long.

NTP Key ID

Select the NTP key ID. The values range from 1 to 99. Different key IDs can be used for different servers.

Set Date

Set the date on the LoadMaster.

Set Time

Set the time on the LoadMaster.

Set time zone

Set the time zone where the LoadMaster is located.

Logging Options

Logging of LoadMaster events can be both pushed and also pulled from the appliance. It is important to note that log files on LoadMaster are not historical, if the appliance reboots the logs are reset. It is important to keep a record of events generated on LoadMaster on a remote facility.

Contents

- System Log Files
- Extended Log Files
- Syslog Options
- SNMP Options
- Email Options
- SDN Log Files

System Log Files

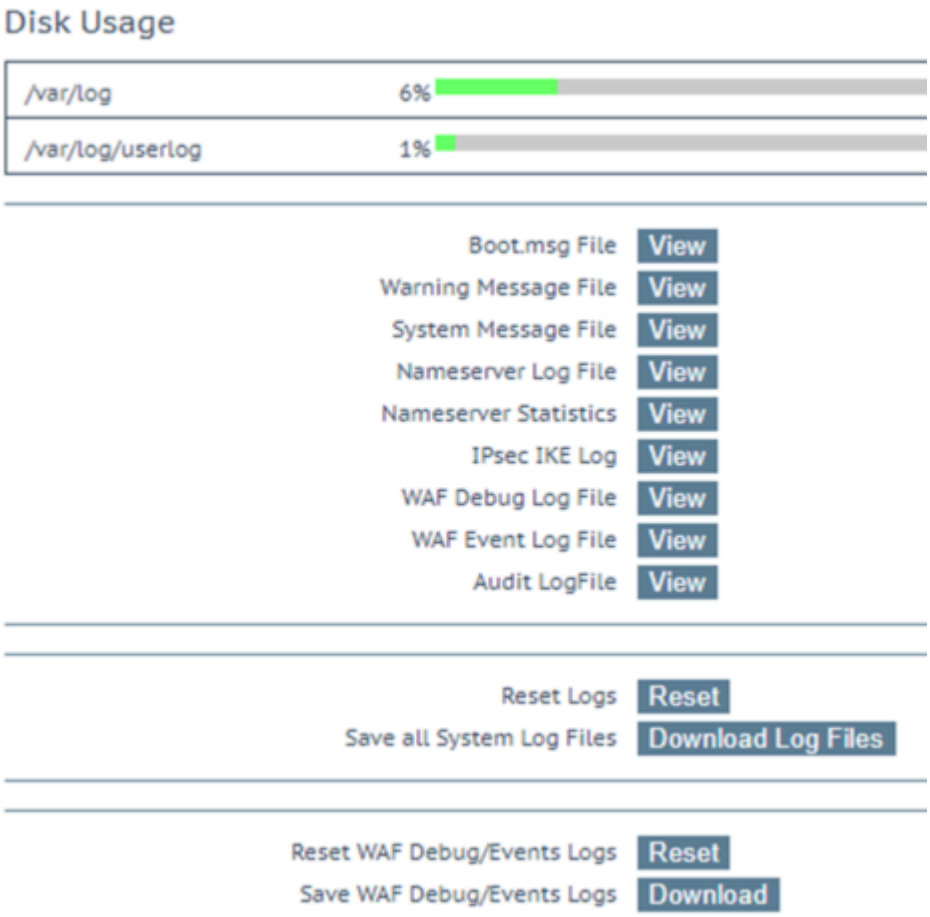

The LoadMaster generates syslog messages using the following format:

```
YYYY-MM-DDThh:mm:ss+hh:mm Hostname Subsystem: Message Data
```

An explanation of the format is as follows:

- **Date/Time Stamp**: `YYYY-MM-DDThh:mm:ss+hh:mm`

  This includes the full date and time along with the timezone offset.

- **Hostname**: The name of the LoadMaster host that generated the message.

- **Subsystem**: The component or module that originated the message (for example, `kernel`, `wafd`, `ssomgr`, and so on).

- **Message Data**: The actual content of the message.

All syslogs are rotated every seven days (older logs are removed).

Disk Usage - This section provides an indication of the percentage used/free of the log partition. Color-

coding is used to highlight different usage levels:

- 0% to 50%: green
- 50% to 90%: orange
- 90% to 100%: red

Boot.msg File - contains information, including the current version, during the initial starting of LoadMaster.

Warning Message File - contains warnings logged during the operation of LoadMaster.

System Message File - contains system events logged during the operation of LoadMaster. This includes both operating system-level and LoadMaster internal events.

Nameserver Log File - show the DNS name server log.

Nameserver Statistics - show the latest name server statistics.

IPsec IKE Log - show the IPsec IKE log.

Note: The **IPsec IKE Log** button does not appear if **VPN Management** is not configured under **System Configuration > Network Setup > VPN Management**.

WAF Debug Log File - contains debug traces useful for debugging WAF issues. Only enable this option when requested to do so by Progress Kemp Technical Support.

Note: The **WAF Debug Log File** button does not appear if there are no WAF debug logs.

WAF Event Log File- contains logs for most recently triggered WAF rules.

Note: The **WAF Event Log File** button does not appear if there are no WAF event logs.

Audit LogFile - contains a log for each action which is performed by a user; either using the API or the WUI. This will only function if session management is enabled. For further information on session management, refer to the Admin WUI Access section.

Clear Logs - clears the warning and system message logfiles. You can either click **Clear All** to clear all the system log files or click the arrow to select specific log files to clear.

Save Logs - this option allows you to send saved system log files to Progress Kemp Support. Click **Save All** to save all system log files to your computer. The logs are compressed. You can view/read them using a text editor after you uncompress the file.

The LoadMaster automatically rotates logs so the hard disk does not fill up with log files.

Debug Options

The LoadMaster has a range of features that will help you and Progress Kemp Support staff with diagnosing connectivity issues. Clicking **Debug Options** brings up the screen shown below.

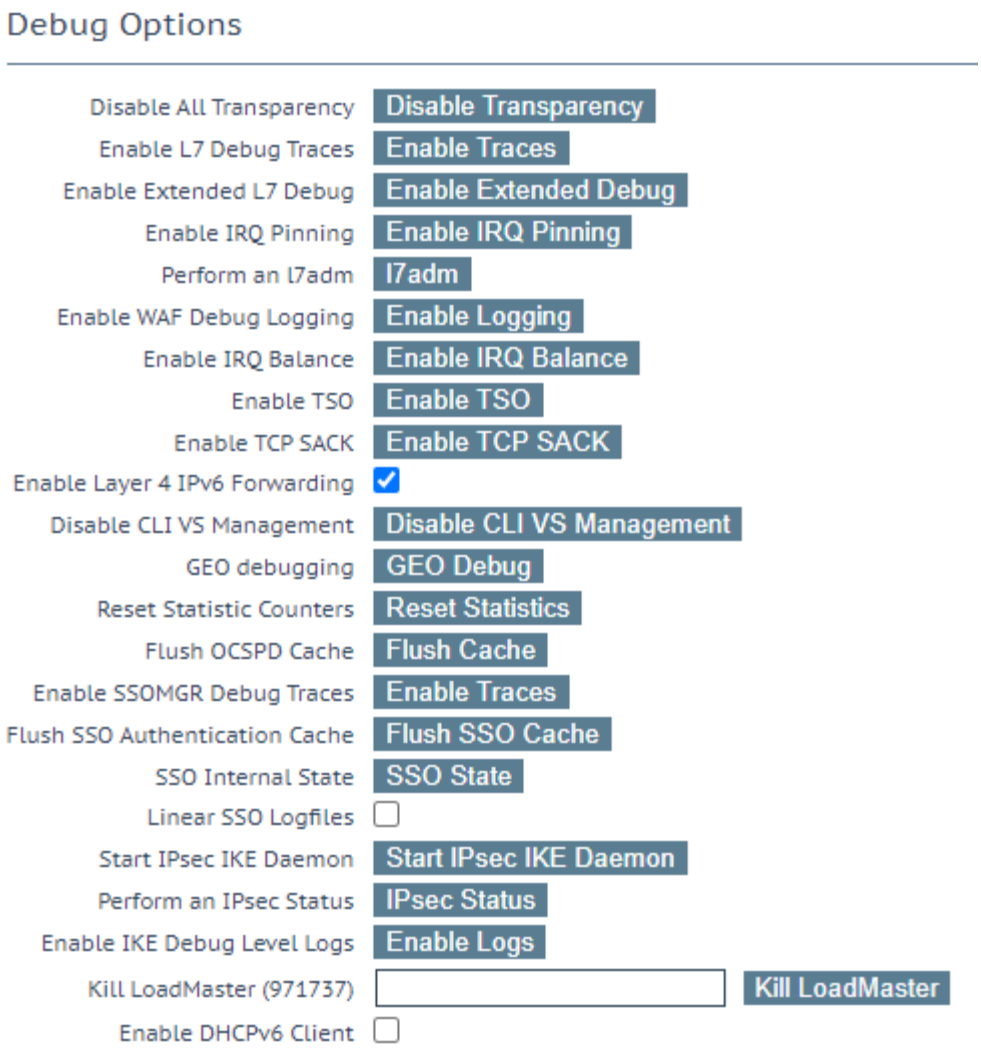

Note: WARNING – Progress Kemp does not recommend using debug commands during normal operation. They should ideally only be used in conjunction with a Progress Kemp Support Technician's recommendations.

Note: Debug commands have performance impacts on the LoadMaster and may expose your system to additional security vulnerabilities during the time they are running.

Disable All Transparency

Disables transparency on every Virtual Service and forces them to use Layer 7. Use with caution.

Note: This option is only for debugging and does not replace the normal controls to enable and disable transparency on a per-Virtual Service basis.

Note: Using this option to disable transparency saves a copy of the configuration file before disabling transparency. When transparency is turned back on (not all Virtual Services may have had transparency turned on before the change), the original configuration is restored. Therefore, any changes to the

configuration during this time are lost. This includes creating new Virtual Services.

Enable L7 Debug Traces

This option enables debugging on all Layer 7 (L7) connections.

Generates log traffic in the message files. Due to the large number of files being logged it slows down L7 processing.

CAUTION: Enabling this option can consume more resources and it is possible that some authorization parameters may be exposed. Only enable this option if recommended by Progress Kemp Support.

Enable Extended L7 Debug

Click **Enable Extended Debug** to enable L7 extended debug options.

In LoadMaster firmware version 7.2.53, the **Enable Extended L7 Debug** option was enhanced. Enabling this option may be needed when performing extensive testing.

Once the **Enable Extended L7 Debug** option is enabled, a **Process Debug** button appears. Clicking **Process Debug** displays the list of processes and the debug level.

The user has to sanitize the logs before providing them externally. Logs should only be enabled (debug level is set to 1) for debug purposes and disabled (debug level is set to 0) immediately after. Logs should be removed from the LoadMaster as soon as possible when the debugging is completed.

CAUTION: Enabling the Extended L7 Debug option can consume more resources and it is possible that some authorization parameters may be exposed. Only enable this option if recommended by Progress Kemp Support.

When extended debugging is enabled, an additional **Extended Debug** configuration section becomes available in the Virtual Service modify screen (**Virtual Services > View/Modify Services > Modify**) for all VSs. When using Sub-Virtual Services (SubVSs), the **Extended Debug** settings are also inherited by the SubVS, so that a single call can be logged in its entirety. It is also possible to enable debug on a single SubVS if required and a further option to limit the logging by specifying the client IP address. All logs associated with this feature are recorded in the system messages file **messages.txt**.



There are two fields that you can configure in this section:

- **L7 Debug Level:** Set the level of Layer7 debugging for this Virtual Service. Possible values are as follows:
  - No Debug
  - Call Tracing

- Full Debug
- Full Debug + HTTP Headers

Call Tracing is a basic level log that displays most relevant operations, while **Full Debug** displays all available debug logs, which is the same as the global setting of **Enable L7 Debug Traces** but on a per-VS level.

Note: By default, the **L7 Debug Level** is set to **No Debug** for all Virtual Services and SubVSs. To enable logging for a particular Virtual Service or SubVS, you must set the **L7 Debug Level** to **Call Tracing** or **Full Debug** in the **Extended Debug** section of the Virtual Service or SubVS modify screen.

CAUTION: Setting the L7 Debug Level to Full Debug + HTTP Headers may expose sensitive information.

- **Client to Trace:** It is also possible to limit the debug information even further by specifying a client IP address (you can specify an IPv4 or IPv6 address). If an address is specified, only connections coming from that specific client IP are logged/traced. This allows debugging capability from a single address.

Enable IRQ Pinning

Click the button to enable Interrupt Request Line (IRQ) pinning. This is disabled by default.

CAUTION: Only enable this option in consultation with Progress Kemp Support.

Note: When you change the IRQ pinning option from off to on, IRQ pinning is enabled on all network interfaces that are assigned an IP address. When IRQ pinning is enabled and you add an IP address to an unconfigured interface, that interface will not have IRQ pinning enabled until you either toggle the IRQ pinning off and back on again, or the system is rebooted.

Perform an l7adm

Note: Displays raw statistics about the L7 subsystem.

Enable WAF Debug Logging

You can enable logging to get detailed information about traffic that is analyzed by Web Application Firewall. Information that is contained in the logs includes the time that LoadMaster WAF received the request from your LoadMaster resource, detailed information about the request, and the action for the rule that each request matched.

Enable WAF debug traces.

CAUTION: Please be aware that enabling this option will generate logs that may include Personally Identifiable Information as defined under the General Data Protection Agreement (EU GDPR). You should follow your organization's best practice to protect this information which may include anonymizing, deleting, or encrypting the data within the logs.

Note: This generates a lot of log traffic. It also slows down WAF processing. Only enable this option when

---

requested to do so by Progress Kemp Technical Support. Kemp does not recommend enabling this option in a production environment.

The WAF debug logs are never closed and they are rotated if they get too large. WAF needs to be disabled and re-enabled in all WAF-enabled Virtual Service settings to re-enable the debug logs. Alternatively, perform a rule update, with rules that are relevant for the Virtual Service(s).

Enable IRQ Balance

Enable this option only after consulting with Progress Kemp support staff.

Enable TSO

Enable TCP Segmentation Offload (TSO).

Note: Only modify this option after consultation with Progress Kemp Technical Support. Changes to this option will only take affect after a reboot.

Enable TCP SACK

Click this button to enable TCP SACK (Selective ACKnowledgement) processing. This is a global setting that affects all Layer 7 Virtual Services. It only works if TCP SACK is enabled on a Virtual Service client and the LoadMaster.

Enable Layer 4 IPv6 Forwarding

This option, when enabled (the default), supports IPv6 forwarding behaviour as supported on LoadMaster versions previous to LoadMaster Operating System (LMOS) 7.2.50. When this option is disabled, it supports the IPv6 forwarding behavior required by the IPv6 standards. If IPv6 forwarding behavior that conforms to the IPv6 standards is required in your deployment, you should disable this option.

Enable/Disable CLI VS Management

Enable or disable the Command Line Interface (CLI) Service Management function.

GEO Debugging

Clicking **GEO Debug** displays a GEO Debug page. The GEO Debug page displays the GEO configuration related information and various GEO Debug Options. For more details on GEO Debugging, refer to the GEO Debug section.

Display RAID Information

Note: The **Display RAID Information** and **Display RAID Disks Information** buttons only appear if a RAID controller is installed on the LoadMaster.

Display the Redundant Array of Independent Disks (RAID) controller details. Some example information is below:

```
---------------------------------------------------------------

Controller details
```

```
----------------------------------------------------------------

- Chip ID.............: 10

- Parent Controller Index: 255

- OS Physical Name.....: /dev/sda

- Serial Number.......: 427491329

- AES Power on State.....: 0

- Sata Ports.............: 2

----------------------------------------------------------------

Raid Port 0 details

----------------------------------------------------------------

- Raid Model Name.............: H/W RAID1

- Raid Serial Number..........: OUEYEXCXTQ53GE1BSOSN

- EZBackup Disk Support.......: 0

- Port Multiplier port........: 0

- Raid Capacity...............: 953 (29 GB)

- Raid Capacity low word.....: 0

- Raid State.................: 1 (Active)

- Raid Status................: 3 (Normal)

- Raid Level.................: 1 (Raid 1 (mirror))

- Mark Type..................: 0

- Active Member..............: 15

- Active Level...............: 0

- Rebuild Priority...........: 3

- Standby Timer..............: 0

- Total members in the RAID....: 2
```

```
Member disk 0

- Ready.....................: 1

- Lba 48 Bit Support.......: 1

- SATA Page.................: 0

- SATA Port.................: 0

- SATA Base.................: 0

- SATA Size.................: 953

------------------------------------------------------------

Member disk 1

- Ready.....................: 1

- Lba 48 Bit Support.......: 1

- SATA Page.................: 0

- SATA Port.................: 1

- SATA Base.................: 0

- SATA Size.................: 953

------------------------------------------------------------
```

Display RAID Disks Information

Display details about the RAID disks. Some example information is below:

```
------------------------------------------------------------------------

Sata Port 0 details

------------------------------------------------------------------------

- Disk Model Name..............: 32GB SATA Flash Drive

- Disk Serial Number...........: C0122916B01000000074

- Disk Firmware Version........: SFDC001D

- EZBackup Disk Support........: 1
```

```
- Port Multiplier port.........: 15

- Disk Capacity................: 954 (29 GB)

- Port Type....................: 2 (RAID)

- Port Speed...................: 2 (GB)

- Page 0 State.................: 2

- Page 0 Raid Index............: 0

- Page 0 Member Index..........: 0

- Page 0 Raid Name.............:

- Page 0 Raid Serial Number....:

- Page 0 Raid Segment Base.....: 0

- Page 0 Raid Size.............: 953

- Page 0 Raid EZ Backup Support: 0

- Page 1 State.................: 0

- Page 1 Raid Index............: 0

- Page 1 MemberIndex...........: 0

- Page 1 Raid Name.............:

- Page 1 Raid Serial Number....:

- Page 1 Raid Segment Base.....: 0

- Page 1 Raid Size.............: 0

- Page 1 Raid EZ Backup Support: 0

- PortErrorStatus..............: 0

-----------------------------------------------------------------

Sata Port 1 details

-----------------------------------------------------------------

- Disk Model Name..............: 32GB SATA Flash Drive
```

```
- Disk Serial Number...........: E011321290100000005A

- Disk Firmware Version........: SFDC001D

- EZBackup Disk Support........: 1

- Port Multiplier port.........: 15

- Disk Capacity................: 954 (29 GB)

- Port Type....................: 2 (RAID)

- Port Speed...................: 2 (GB)

- Page 0 State.................: 2

- Page 0 Raid Index............: 0

- Page 0 Member Index..........: 1

- Page 0 Raid Name.............:

- Page 0 Raid Serial Number....:

- Page 0 Raid Segment Base.....: 0

- Page 0 Raid Size.............: 953

- Page 0 Raid EZ Backup Support: 0

- Page 1 State.................: 0

- Page 1 Raid Index............: 0

- Page 1 MemberIndex...........: 0

- Page 1 Raid Name.............:

- Page 1 Raid Serial Number....:

- Page 1 Raid Segment Base.....: 0

- Page 1 Raid Size.............: 0

- Page 1 Raid EZ Backup Support: 0

- PortErrorStatus..............: 0
```

Reset Statistic Counters

Reset all statistics counters to zero and delete any old graphs. This also deletes the Round Robin Database (RRD) files but these files are automatically recreated when needed.

Flush OCSPD Cache

When using OCSP to verify client certificates, OCSPD caches the responses it gets from the OCSP server. This cache can be flushed by pressing this button. Flushing the OCSPD cache can be useful when testing, or when the Certificate Revocation List (CRL) has been updated.

Stop IPsec IKE Daemon

Stop the IPsec IKE daemon on the LoadMaster.

Note: If this button is clicked, the connection for all tunnels will go down.

Perform an IPsec Status

Display the raw IPsec status output.

Enable IKE Debug Level Logs

Control the IPsec IKE log level.

Flush SSO Authentication Cache

Clicking the **Flush SSO Cache** button flushes all Single Sign-On (SSO) records, resets all authentication server statuses, resets the KCD domain (if relevant) and re-reads the configuration. This has the effect of logging off all clients using Single Sign-On to connect to the LoadMaster.

SSO Internal State

Click **SSO State** to show the internal state for the SSO Manager.

SSO LDAP server timeout

Set the SSO LDAP server timeout value in seconds (default value is 5 seconds).

Linear SSO Logfiles

By default, older log files are deleted to make room for newer log files, so that the filesystem does not become full. Selecting the **Linear SSO Logfiles** check box prevents older files from being deleted.

Note: When using Linear SSO Logging, if the log files are not periodically removed and the file system becomes full, access to ESP-enabled Virtual Services will be blocked, preventing unlogged access to the virtual service. Access to non-ESP enabled Virtual Services are unaffected by the Linear SSO Logfile feature.

Kill LoadMaster

Permanently disables all LoadMaster functions. The LoadMaster can be re-enabled by being relicensed.

Note: Please do not kill your LoadMaster without consulting Progress Kemp Technical Support.

Note: The **Kill LoadMaster** option will not be available in LoadMasters which are tenants of the Progress Kemp Multi-Tenant LoadMaster.

Enable DHCPv6 Client

When this option is enabled, the DHCPv6 client will run on the primary interface. This provides the capability to obtain an IPv6 address on boot. If you want DHCPv6 to be run on every boot, keep this option enabled. However, this is a long running process and it keeps running in the background when it is enabled so if you only need an IPv6 address to be assigned and you do not need to renew and release the IPv6 address you should disable this option after the IPv6 address is assigned.

Extended Log Files

The **Extended Log Files** screen provides options for logs relating to the ESP and WAF features.

The ESP and WAF audit logs are rotated every 30 days (older logs are removed). WAF remote logs are rotated every seven days.

Note: If debug logging is enabled, it is possible that sensitive information may appear in the logs. If you are concerned by this, clear all the logs immediately after disabling debug logging.

To get to the **Extended Log Files** screen – in the LoadMaster WUI, go to **System Configuration > Logging Options > Extended Log Files**.

Note: The WAF logs are not generated in real time – they can be up to two minutes behind what the WAF engine is actually processing.

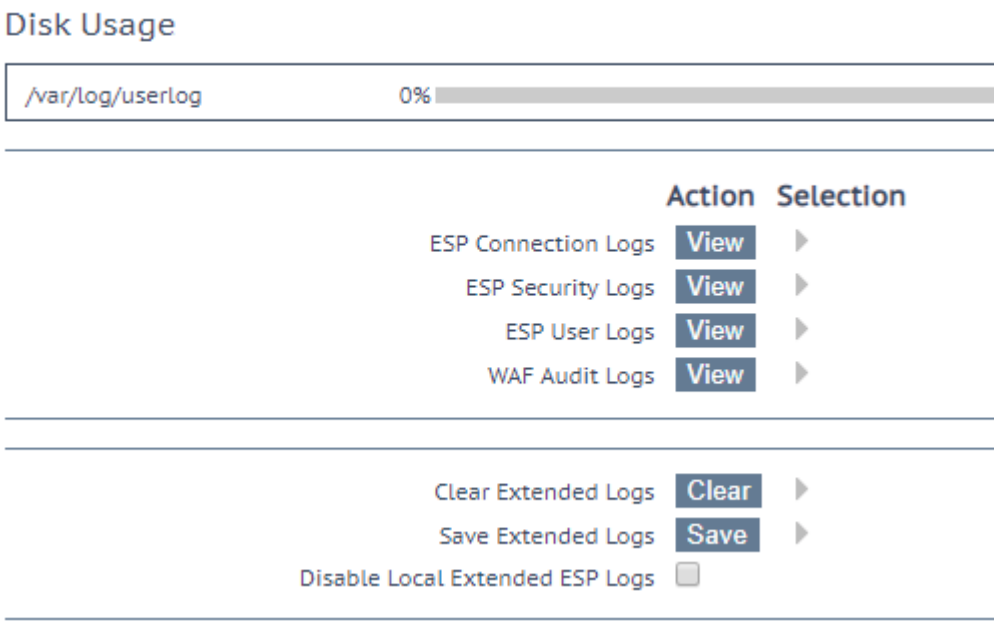

Disk Usage - This section provides an indication of the percentage used/free of the log partition. Color-coding is used to highlight different usage levels:

- 0% to 50%: green
- 50% to 90%: orange
- 90% to 100%: red

There are multiple log files relating to ESP stored on the LoadMaster. These are listed below the **Disk Usage** section. These logs are persistent across LoadMaster reboots.

You can select one of the **View** or **Save Action** buttons with the default filter options to apply the action to the various log files (Connection Logs, Security Logs, and so on). For the **Clear** button, you must first select which logs to clear using the **Selection** controls.

To access the **Selection Controls**, click one of the right caret icons ▸ at the right of the buttons. For example, clicking on the icon to the right of the **Clear** and **Save** buttons, displays these controls.

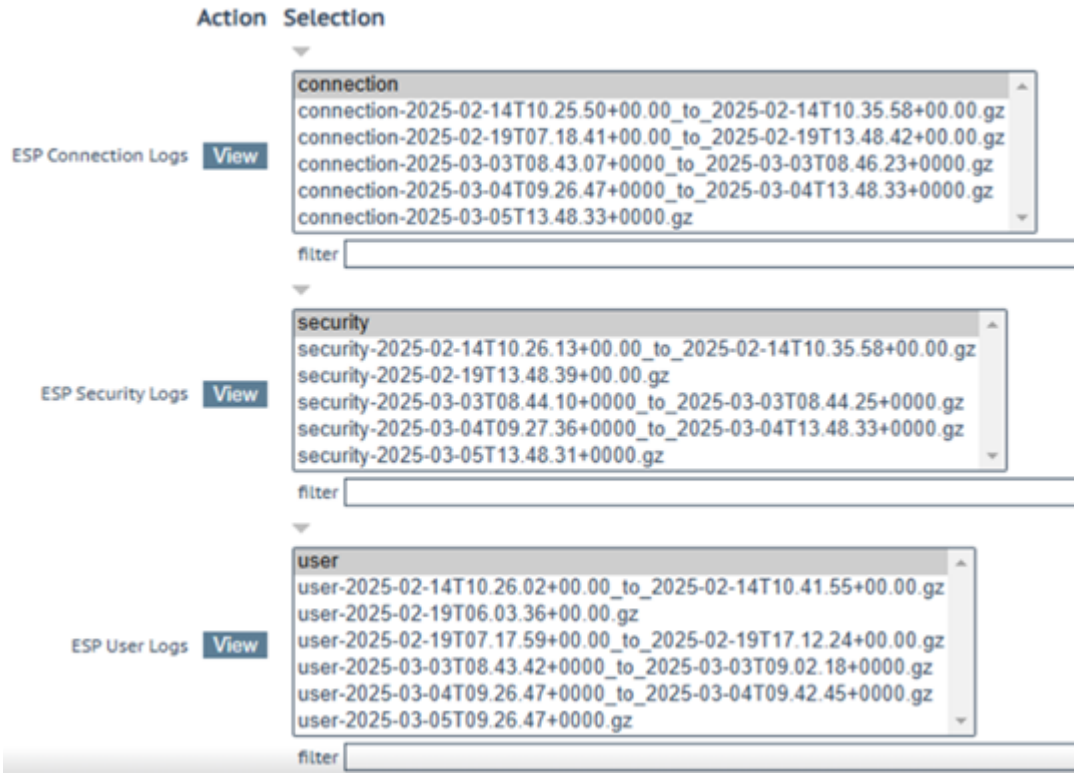

You can filter the logs to clear or save by selecting a subset of log files from the multiple pick list on the right.

- **ESP Connection Logs:** logs recording each connection.
- **ESP Security Logs:** logs recording all security alerts.
- **ESP User Logs:** logs recording all user logins.

Note: In LoadMaster firmware version 7.2.51, ESP user logs were expanded to be more useful and applicable to enterprise customers with extensive logging infrastructure. User Authentication, Authorization, and Accounting (AAA) information is included in the logs, including the time of request, username, domain, AAA server, AAA protocol type, AAA result, and error message. For further details, refer to the ESP Logs Technical Note.

Note: In LoadMaster firmware version 7.2.53, the ESP client session logging was further enhanced. The LoadMaster logs: - The initially created ESP session - The time when the LoadMaster cleared the session from the cache. Note that if the entire cache is cleared, a single log message is recorded at the time of clearing, which notes that all existing sessions at that time were cleared form the cache. - If an ESP session is deleted (when the user logs out from the application, when the session expires, or the user enters invalid credentials). The time when the LoadMaster cleared the session is also logged.

- **WAF Audit Logs:** recording WAF logs based on what has been selected for the **Audit mode** drop-down list in the **WAF Options** section of the Virtual Service modify screen. The number listed in each log entry corresponds to the ID of the Virtual Service. To get the Virtual Service ID, first ensure that the API interface is enabled (**Certificates & Security > Remote Access > Enable API Interface**). Then, in a web browser address bar, enter **https://<LoadMasterIPAddress>/access/listvs**. Check the **index** of the Virtual Service. This is the number that corresponds to the number on the audit log entry.

To view the logs, please select the relevant options and click the relevant **View** button.

One or more archived log files can be viewed by selecting the relevant file(s) from the list of file names and clicking **View**. You can filter the log files by entering a word(s) or regular expression in the **filter** field and clicking **View**.

Note: If you use quotes in regular expressions in the LoadMaster WUI, there are limitations. For more information, refer to the section Limitations of Using Regular Expressions in the LoadMaster WUI.

Clear Extended Logs

Log files can be deleted by filtering one or more individual log files in the log file list or selecting a specific log type (for example, connection, security, or user) in the log file list and clicking **Clear**. Click **OK** on any warning messages.

Save Extended Logs

Click the arrow to expand the options. Select a file type (for example, **connection**). All extended logs can be saved to a file by clicking **Save**. This saves a file to your machine.

Specific log files can be saved by filtering one or more individual log files in the log file list or selecting a specific log type (for example connection, security or user) in the log file list and clicking **Save**.

Disable Local Extended ESP Logs

If **Disable Local Extended ESP Logs** is disabled (the default option), messages are written to the extended ESP logs expediently and are not sent to any remote syslog servers that are defined.

If **Disable Local Extended ESP Logs** is enabled, no messages are written to the extended ESP logs and messages are only sent to the remote logger (if one is defined). If a remote logger is not defined, no logs are recorded.

You can no longer configure the system to both populate the local extended ESP logs and send the same messages to remote syslog servers, as it was in previous releases.

Clear Temporary WAF Remote Log Data

Clear the temporary WAF remote log data.

Save Temporary WAF Remote Logs Data

Save the temporary WAF remote log data.

Syslog Options

The LoadMaster can produce various warning and error messages using the syslog protocol. Syslog messages include a full timestamp and the LoadMaster host name in every message. Syslog messages comply with RFC5424. These messages are normally stored locally.

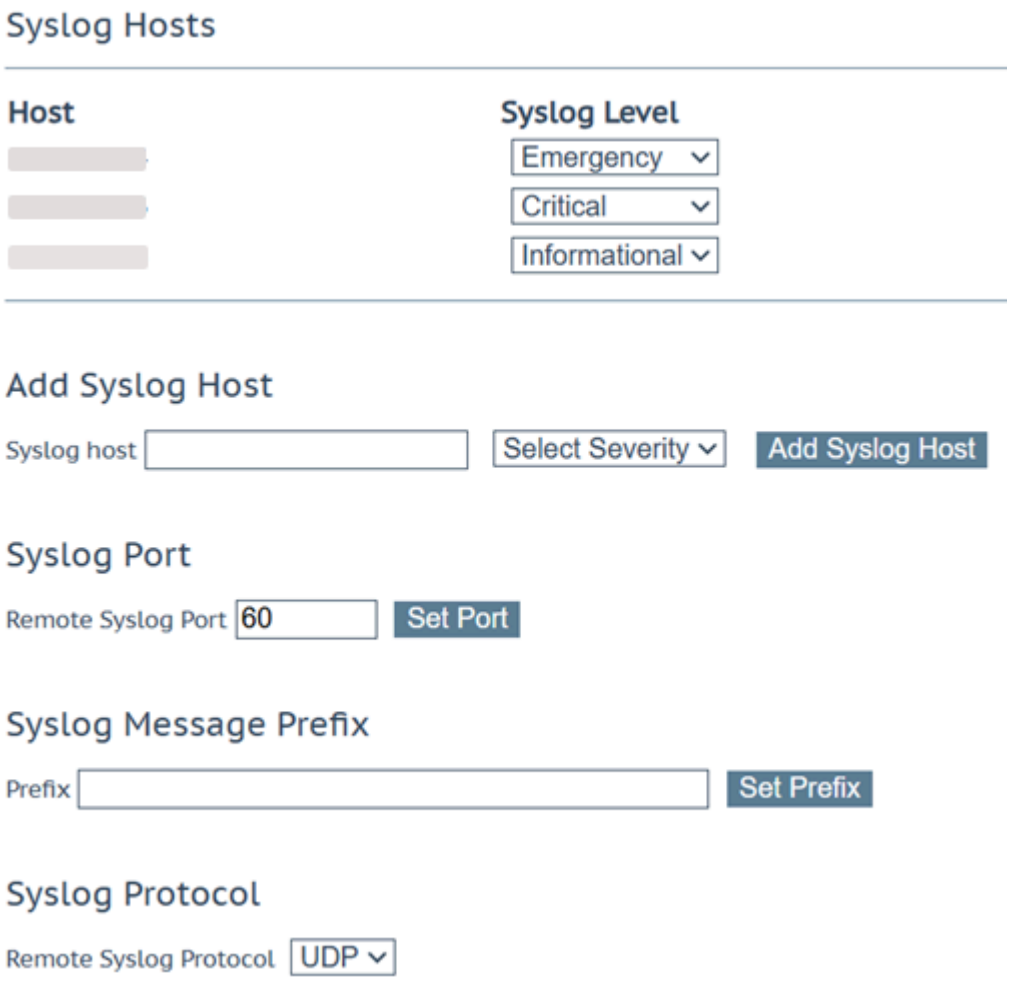

It is also possible to configure the LoadMaster to transmit these error messages to a remote syslog server by entering the relevant IP address in the **Syslog host** text box, selecting the severity and clicking **Add Syslog Host**. Syslog messages are transmitted securely using TLS to remote servers.

The LoadMaster uses OCSP to check the validity of the server certificates supplied by configured syslog servers. If these checks fail, connections to the server are not permitted.

To delete a hosts entry, set the severity level to **None**.

Six different error message levels are defined and each message level may be sent to a different server. Notice messages are sent for information only; Emergency messages normally require immediate user

action.

Note: Up to 10 individual IP addresses can be specified. If there were more than 10 hosts configured in a previous LoadMaster version, after upgrading - all entries are displayed but no more can be added.

Examples of the type of message that may be seen are shown below:

- **Emergency**: Kernel-critical error messages
- **Critical**: Unit 1 has failed and unit 2 is taking over as active (in a HA setup)
- **Error**: Authentication failure for root from 192.168.1.1
- **Warn**: Interface is up/down
- **Notice**: Time has been synced
- **Info**: Local advertised Ethernet address

Note: One point to note about syslog messages is they cascade in an upwards direction. Thus, if a host is set to receive WARN messages, the message file includes messages from all levels above WARN but none for levels below.

If you enter the same host address again, the old entry for the same host is replaced. There is no need to have multiple entries for the same host because a single entry covers the syslog level that is defined, plus all other levels that are of higher priority. So, you only need to include one entry with the lowest level priority required.

You can also specify a non-standard port for syslog transfer by entering it into the **Remote Syslog Port** text box and clicking **Set Port**:

- If the **Remote Syslog Port** is not configured, logging is done on UDP on port 514.
- If the **Remote Syslog Port** is configured as 601, logging is done on TCP on port 601.
- If the **Remote Syslog Port** is configured as any port other than 601, logging is done on secured TCP, that is, over SSL on the configured port.

To enable a syslog process on a remote Linux server to receive syslog messages from the LoadMaster, the syslog must be started with the “-r” flag.

Syslog Message Prefix

Enter the fixed string in the **Prefix** text box and click **Set Prefix** to add it to the beginning of all syslog messages. You can use this prefix to include identifiers such as a license key, which may be required by remote syslog consumers for authentication or tracking purposes. There is a limit of 80 characters for the **Prefix** field. Most characters are supported.

You can specify the protocol to use when connecting to a remote syslog server by selecting the appropriate option in the **Remote Syslog Protocol** drop-down list.

Server Certificate Validation

This check box only appears when **TLS** is selected as the **Remote Syslog Protocol**.

When **Server Certificate Validation** is enabled, it ensures that the host name or IP address that was used to initiate the secure connection resides in the Certificate Subject or Subject Alternative Names (SAN) of the

certificate.

Server Certificate Validation is disabled by default.

SNMP Options

The fields on the **SNMP Options** screen vary depending on whether you enable SNMP V3. Refer to the relevant section below for details on the fields depending on your configuration.

SNMP Options

With this menu, the SNMP configuration can be modified.

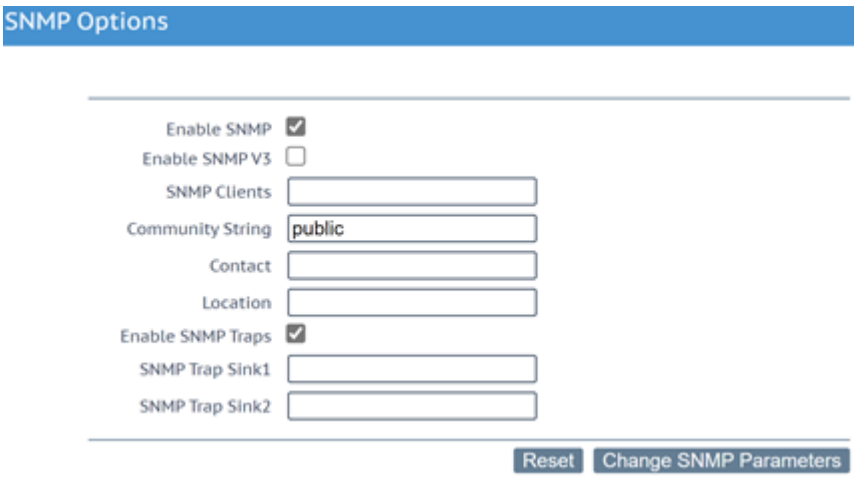

Enable SNMP

This check box enables or disables SNMP. This option allows the LoadMaster to respond to SNMP requests.

By default, SNMP is disabled.

CAUTION: There is a significant performance hit when SNMP is enabled.

When the feature is enabled, the following traps are generated:

- **ColdStart**: generic (start/stop of SNMP sub-system)
- **VsStateChange**: (Virtual Service state change)
- **RsStateChange:** (Real Server state change)
- **HaStateChange**: (HA configuration only: LoadMaster failover)

Note: When using SNMP monitoring of ESP-enabled Virtual Services that were created using a template, ensure to monitor each SubVS directly rather than relying on the master service. This is because the Authentication Proxy sub-service will always be marked as up and, as a consequence, so will the master service.

The information regarding all LoadMaster-specific data objects is stored in the below enterprise-specific MIBs (Management Information Base).

MIB file	Related Data
IPVS-MIB.txt	Virtual Server stats
B100-MIB.txt	L7 LoadMaster configuration and status info
ONE4NET-MIB.txt	Enterprise ID
CERTS-MIB.txt	SSL certificate information

These MIBs (located on the Other Downloads page) need to be installed on the SNMP manager machine to be able to request the performance/config-data of the LoadMaster using SNMP.

To retrieve the performance/config-data using SNMP for the Dell-specific LoadMaster hardware, you must upload the ENTITY-SENSOR-MIB (.1.3.6.1.2.1.99) and the ENTITY-MIB (.1.3.6.1.2.1.47).

The description of the counters can be taken from the LoadMaster MIBs (the description clause). Apart from just reading the MIB this can be done for Linux with the command:

snmptranslate -Td -OS <oid>

where <oid> is the object identifier in question.

For legacy reasons the Progress Kemp object identifier is specified as **one4net**.

Example: <oid> = .1.3.6.1.4.1.one4net.ipvs.ipvsRSTable.rsEntry.rSConns

snmptranslate -Td –Ov
.1.3.6.1.4.1.one4net.ipvs.ipvsRSTable.rsEntry.rSConns.1.3.6.1.4.1.12196.12.2.1.12

rSConns OBJECT-TYPE

-- FROM IPVS-MIB

SYNTAXCounter32

MAX-ACCESSread-only

STATUScurrent

DESCRIPTION"the total number of connections for this RS"

::= { iso(1) org(3) dod(6) internet(1) private(4) enterprises(1) one4net(12196) ipvs(12) ipvsRSTable(2) rsEntry(1) 12 }

The data object defined in the LoadMaster MIBS is a superset to the counters displayed by the WUI.

Note: The data objects on the LoadMaster are not writable, so only GET requests (GET, GET-NEXT, GET-BULK, and so on) should be used.

You can retrieve SSL certificate information including the file name, certificate Subject Name, certificate serial number, certificate start date, certificate end date, and certificate issuer information using the **CERTS-MIB.txt** file. SNMP can display this information for up to 256 SSL certificates.

You can also retrieve disk space usage details using SNMP. The /var/log and /var/log/userlog partition information is available.

Regarding the data partitions, these are not under the Progress Kemp OID, they are under a standard OID. The details are as follows:

- Path where the disk is mounted: .1.3.6.1.4.1.2021.9.1.2.1
- Path of the device for the partition: .1.3.6.1.4.1.2021.9.1.3.1
- Total size of the disk/partition (kBytes): .1.3.6.1.4.1.2021.9.1.6.1
- Available space on the disk: .1.3.6.1.4.1.2021.9.1.7.1
- Used space on the disk: .1.3.6.1.4.1.2021.9.1.8.1
- Percentage of space used on disk: .1.3.6.1.4.1.2021.9.1.9.1
- Percentage of inodes used on disk: .1.3.6.1.4.1.2021.9.1.10.1

SNMP Clients

With this option, the user can specify from which SNMP management hosts the LoadMaster will respond to.

Note: If no client has been specified, the LoadMaster will respond to SNMP management requests from any host.

When the **Enable SNMP V3** option is enabled, this option is not available.

SNMP Community String

This option allows the SNMP community string to be changed. The default value is “public”.

Allowed characters in the **Community String** are as follows: **a**-**z**, **A**-**Z**, **0**-**9**, **_.-@()?#%^+~!**.

When the **Enable SNMP V3** option is enabled, this option is not available.

Contact

This option allows the SNMP Contact string to be changed. For example, this could be e-mail address of the administrator of the LoadMaster.

Note: The **Email Options** in the LoadMaster (**System Configuration > Logging Options > Email Options**) must be configured for email to be sent successfully.

SNMP Location

This option allows the SNMP location string to be changed.

Note: This field accepts the following characters: **a-z A-Z 0-9 _ . - ; , = : { } @ ( ) ? # % ^ + ~ !**

Note: Do not enter a hashtag symbol (**#**) as the first character in the Location.

SNMP traps

When **Enable SNMP Traps** is selected and an IP address is specified in an SNMP trap sink text box, a trap is generated when an important event happens to a LoadMaster, Virtual Service, or Real Server. These are sent to the SNMP trap sinks. If a change is made, the LoadMaster waits for all changes to finish and then waits five seconds before reading it. At that point, all changes will have stabilized and SNMP traps can then be sent. If there are any state changes within the five second wait, the state changes are handled and then the wait is restarted.

Enable/Disable SNMP Traps

This toggle option enables and disables the sending of SNMP traps.

Note: SNMP traps are disabled by default.

Send SNMP traps from the shared address

Note: This check box is only visible when the LoadMaster is in HA mode.

By default, SNMP traps are sent using the IP address of the active HA unit as the source IP address. Enabling this option will send SNMP traps from the active HA unit using the shared IP address.

SNMP Trap Sink1

This option allows the user to specify a space-separated list of hosts to which a SNMPv1 trap will be sent when a trap is generated.

SNMP Trap Sink2

This option allows the user to specify a space-separated list of hosts to which a SNMPv2 trap will be sent when a trap is generated.

Contents

- SNMP V3 Options

SNMP V3 Options

With this menu, the SNMP V3 configuration can be modified.

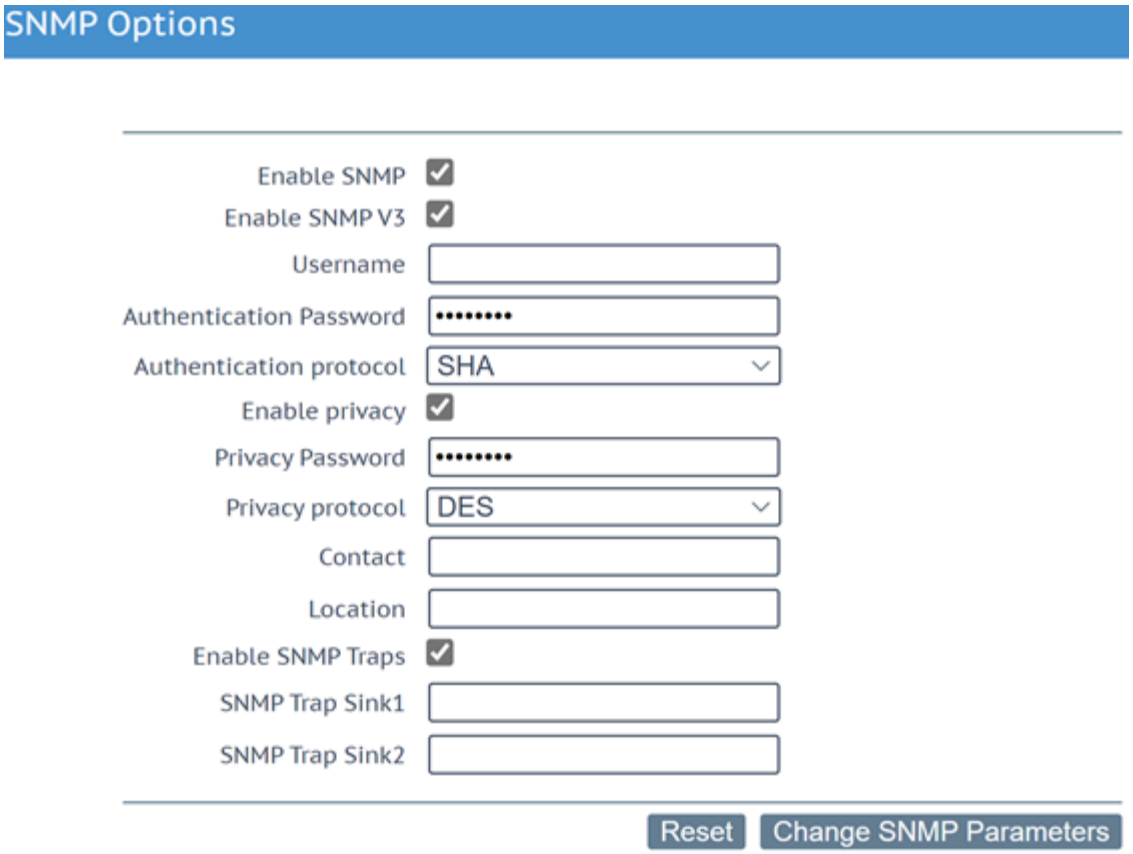

Enable SNMP V3

This check box enables SNMPv3. SNMPv3 primarily adds security and remote configuration enhancements to SNMP.

When this option is enabled, the following additional fields become available:

- Username

- Authentication Password

- Authentication protocol

- Enable Privacy checkbox

Note: The **Username** and **Authentication Password** must be set for SNMPv3 to work, while the **Private Password** field is optional.

Note: The password must be at least 8 characters long.

Authentication protocol

Select the relevant **Authentication protocol** - **MD5** or **SHA**. **SHA** is recommended because **MD5** is unsecure and deprecated.

Enable Privacy

This checkbox can be used to set the **Privacy Password** and **Privacy Protocol** for SNMPv3.

Privacy protocol

Select the relevant **Privacy protocol** - **AES** or **DES**. **AES** is recommended because **DES** is deprecated.

Contact

This option allows the SNMP Contact string to be changed. For example, this could be e-mail address of the administrator of the LoadMaster.

SNMP Location

This option allows the SNMP location string to be changed.

Note: This field accepts the following characters:

---

a-z A-Z 0-9 _ . - ; , = : { } @ ( ) ? # % ^ + ~ !

Note: Do not enter a hashtag symbol (**#**) as the first character in the Location.

SNMP traps

When **Enable SNMP Traps** is selected and an IP address is specified in an SNMP trap sink text box, a trap is generated when an important event happens to a LoadMaster, Virtual Service, or Real Server. These are sent to the SNMP trap sinks. If a change is made, the LoadMaster waits for all changes to finish and then waits five seconds before reading it. At that point, all changes will have stabilized and SNMP traps can then be sent. If there are any state changes within the five second wait, the state changes are handled and then the wait is restarted.

Enable/Disable SNMP Traps

This toggle option enables and disables the sending of SNMP traps.

Note: SNMP traps are disabled by default.

Send SNMP traps from the shared address

Note: This check box is only visible when the LoadMaster is in HA mode.

By default, SNMP traps are sent using the IP address of the active HA unit as the source IP address. Enabling this option will send SNMP traps from the active HA unit using the shared IP address.

SNMP Trap Sink1

This option allows the user to specify a list of hosts to which a SNMPv1 trap will be sent when a trap is generated.

SNMP Trap Sink2

This option allows the user to specify a list of hosts to which a SNMPv2 trap will be sent when a trap is generated.

Email Options

This screen permits the configuration of email alerting for LoadMaster events. Email notification can be delivered for six predefined informational levels. Each level can have a distinct email address and each level supports multiple email recipients. Email alerting depends on a mail server, support for both an open relay mail server and a secure mail server is provided.



A sample email alert is above. This is from the **Info** level. Syslog emails contain one or more lines of syslogs (they are grouped together if possible).

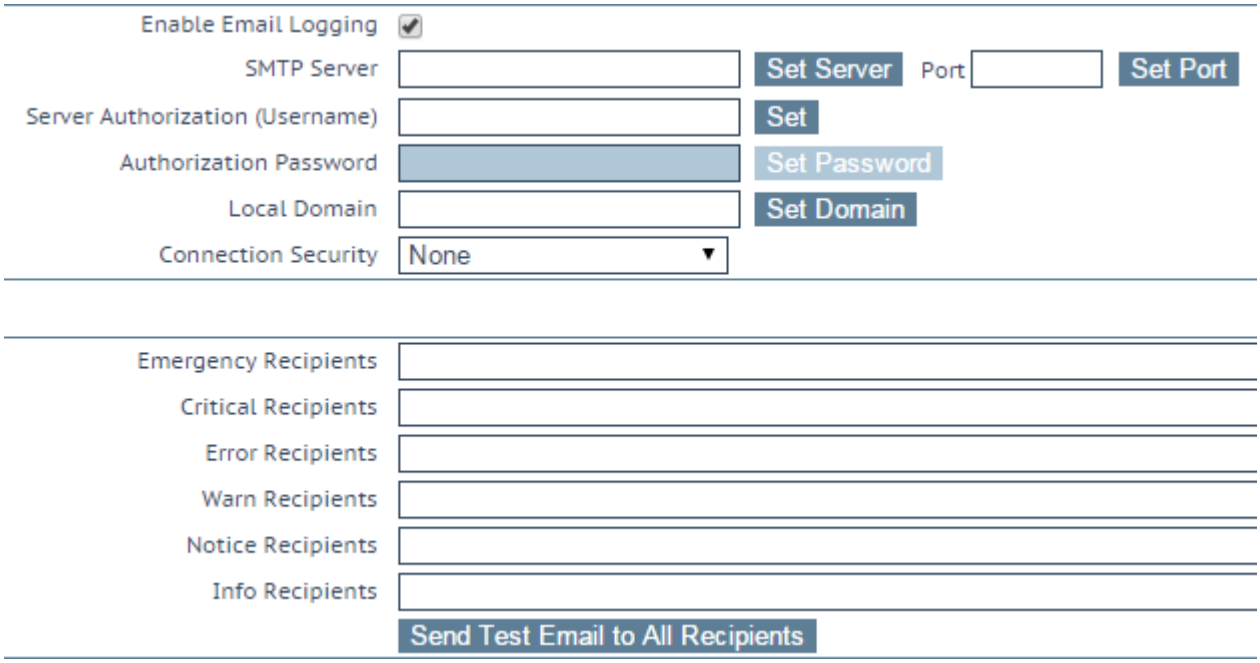

SMTP Server

Enter the FQDN or IP address of the mail server. If you are using an FQDN ensure to set the DNS Server.

Port

Specify the port of the SMTP server which will handle the email events.

Server Authorization (Username)

Enter the email address to use as the 'From' address in the email headers.

Authorization Password

Enter the password if your mail server requires authorization for mail delivery. This is not a required if you

mail server does not require authorization.

Local Domain

Enter the top-level domain, if your mail server is part of a domain. This is not a required parameter.

Connection Security

Select the type of security for the connection;

- None
- STARTTLS, if available
- STARTTLS
- SSL/TLS

Set Email Recipient

In the various **Recipients** text boxes, enter the email address that corresponds with the level of notification desired. Notifications will be sent for the level of severity, plus anything with a higher severity – so there is no need to enter the email address in multiple text boxes as that will lead to duplicate notifications being sent. For example, any email address entered into the **Critical Recipients** text box will get critical emails, but will also get emergency emails.

Multiple email addresses are supported by a comma-separated list, such as:

Info Recipients: **info@progress.com, sales@progress.com**

Error Recipients: **support@progress.com**

Clicking the **Send Test Email to All Recipients** button sends a test email to all the listed email recipients.



An example email alert is shown above. The **Subject** of the email contains the relevant highest alert level. There can be multiple alerts in a single email - they are collated together for a period of 30 seconds to avoid flooding inboxes.

SDN Log Files

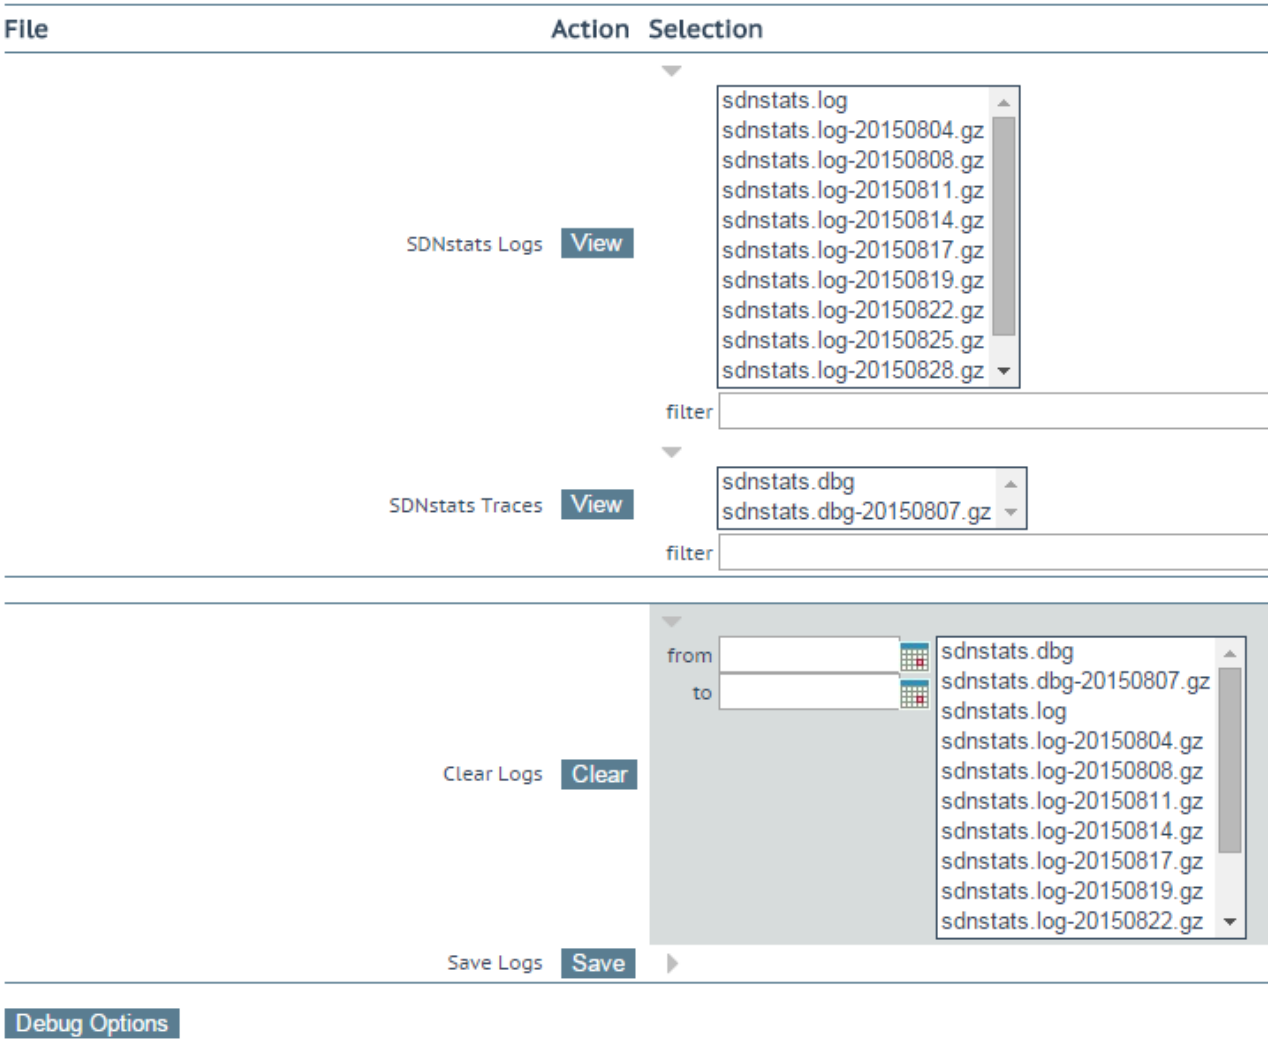

The **SDN Log Files** screen provides options for logs relating to the SDN feature. To view all of the options click the ▼ icons.

View SDNstats Logs

To view the SDNstats logs please select the relevant log files and click **View**.

The **sdnstats.log** file is the main, rolling log file. The .gz files are backups of logs for a particular day.

One or more archived log files can be viewed by selecting the relevant file(s) from the list of file names and clicking the **View** button. The log files can be filtered by entering a word(s) or regular expression in the **filter** field and clicking the **View** button.

View SDNstats Traces

Note: This option is only available if SDNstats debug logging is enabled (**System Configuration > Logging**

Options > SDN Log Files > Debug Options > Enable Debug Log).

To view the SDNstats logs please select the relevant log files and click **View**.

One or more archived log files can be viewed by selecting the relevant file(s) from the list of file names and clicking **View**. The log files can be filtered by entering a word(s) or regular expression in the **filter** field and clicking **View**.

```
Apr 19 16:26:32 gstatsv2.py:iter:491 One minute timer
Apr 19 16:26:37 gstatsv2.py:run:506 Calling iter
Probing(10.35.7.10,8443,https=True):
        [HP VAN] SUCCESS [Version] 2.5.20.1227
```

The traces show probing results – this indicates if the LoadMaster can successfully communicate with the SDN controller.

Clear Logs

All SDN logs can be deleted by clicking the **Clear** button.

A specific range of log files can be filtered by specifying a date range using the **from** and **to** fields. Specifying a date range will simply select the relevant log files that apply in the right-hand box. Individual log files can still be selected/deselected as needed on the right.

Note: Important: If the **sdnstats.log** file is selected, all logs in that file will be cleared, regardless of what dates are selected in the date range fields.

Save Extended Logs

All SDN logs can be saved to a file by clicking the **Save** button.

Specific log files can be saved by filtering on a specific date range and/or selecting one or more individual log files in the log file list in the log file list and clicking the **Save** button.

Debug Options

To get to the SDN Debug Options screen, click the **Debug Options** button on the **SDN Log Files** screen.

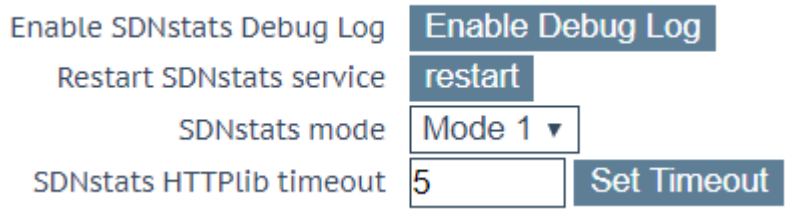

Enable Debug Log

Enable SDNstats debug logging.

To view the SDN Statistics logs, open **System Configuration > Logging Options > SDN Log Files,** select

the log file you wish to view and click the **View** button.

Note: Debug logging should only be enabled when troubleshooting because it will impact performance of the LoadMaster.

Restart SDNstats service

When troubleshooting issues with SDN, the entire SDN service can be restarted. Restarting the connection will not affect any traffic connections - it just restarts the connection between the LoadMaster and the SDN controller.

If successful the **Process ID** will change to a new id.

The **Process ID** can be found by clicking the **Debug** button in **System Configuration > Logging Options > System LogFiles** and clicking the **ps** button.

Note: This will restart the connection to all attached SDN controllers.

SDNstats mode

There are two modes that can be used to gather the SDN statistics.

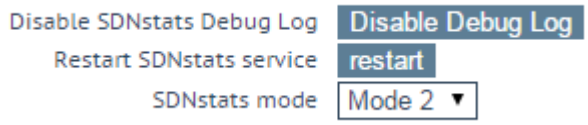

The mode can be set by going to **System Configuration > Logging Options > SDN Log Files > Debug Options** and setting the **SDNstats mode**.

The modes are described below:

- **Mode 1:** When set to mode 1, the statistics are taken from the switch port that is connected to the server and the statistics are relayed back to the LoadMaster.
- **Mode 2:** When set to mode 2, the information is taken from all of the switch ports along the path.

SDNstats HTTPlib timeout

This field enables you to increase the amount of time to wait for the SDN controller to respond. This can reduce the possibility of time outs caused by latency in the environment. Valid values for this field range from 5 to 60.

Troubleshooting

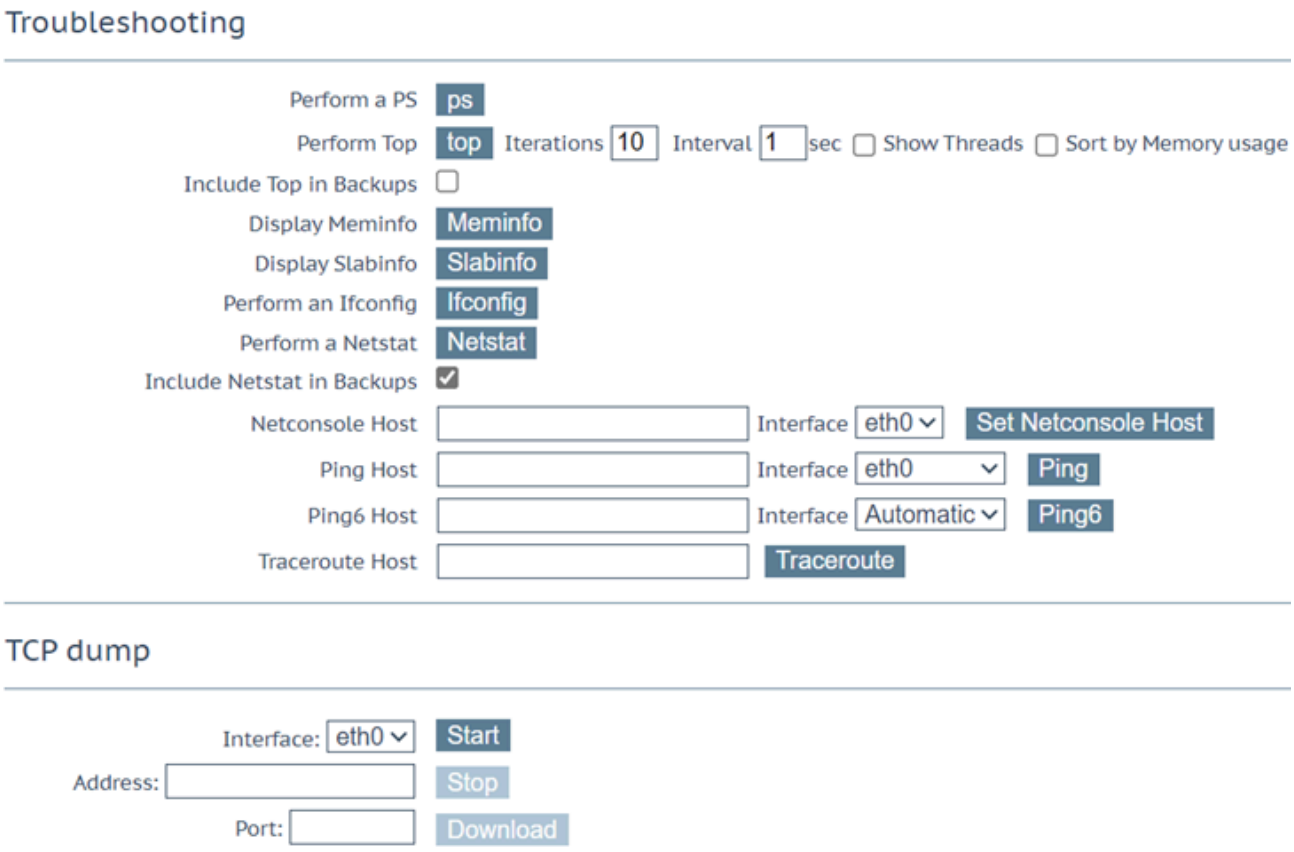

Perform a PS

Performs a **ps** on the system.

Perform a Top

Performing a **top** command displays memory, CPU, and I/O usage for the LoadMaster. You can specify the number of samples and an interval between them (the default is 10 samples and a 1 second interval). You can also show threads and/or sort by memory usage by selecting the appropriate check boxes. By default the results are sorted by CPU usage.

Include Top in Backups

By default, the LoadMaster does not include a top output in backups. This can be enabled by selecting this check box. When included in backups, **top** is run using the default parameters (regardless of what is configured in the WUI) and is sorted by memory usage.

Display Meminfo

Displays raw memory statistics.

Display Slabinfo

Displays raw slab statistics.

Perform an Ifconfig

Displays raw Ifconfig output.

Perform a Netstat

Displays Netstat output.

Include Netstat in Backups

By default, the LoadMaster includes a Netstat output in backups taken. When this is included, backups take longer to complete. You can stop the Netstat output from being included by disabling this option.

Netconsole Host

If directed by a Progress Kemp support engineer, you can use this feature to send critical kernel logs to a syslog server in the event of the LoadMaster failing or crashing. A syslog server configured as the Netconsole Host receives all critical kernel messages.

You can select which interface the **Netconsole Host** is set to using the **Interface** dropdown.

Note: Please ensure that the **Netconsole Host** IP specified is on the selected interface as errors may occur if it is not.

Note: The **Netconsole Host** parameter can be set to either an IPv4 or IPv6 address, the only restriction being that the Netconsole IP address must be in the same address family as the LoadMaster's IP address on the selected **Interface**. In other words, if the LoadMaster's IP address on the selected **Interface** is an IPv4 address, then the **Netconsole Host** IP address must be an IPv4 address as well. If you want to use an IPv6 address for **Netconsole Host**, then you must select an **Interface** on which the LoadMaster has an IPv6 address.

Note: Netconsole is not configurable on a bonded interface.

Ping Host

Performs a ping on the specified host. The interface which the ping should be sent from can be specified in the **Interface** drop-down list. The **Automatic** option selects the correct interface to ping an address on a particular network.

The interface tries to determine if the address to ping is an IPv4 or IPv6 address and selects the correct command to perform the ping. For an address in numeric form this is simple, however this is not possible for non-numeric addresses so they will always be treated as an IPv4 address.

Ping6 Host

Perform a ping6 of a specific IPv6 host.

Traceroute Host

Perform a traceroute of a specific host.

TCP dump

A TCP dump can be captured either by one or all Ethernet ports. You can specify the address and port parameters in the WUI. To specify additional, optional parameters - use the Application Programming Interface (API) to perform a TCP dump.

You can stop and start the dump. You can also download it to a particular location. The results of the TCP dump can then be analysed in a packet trace analyser tool such as Wireshark.

For more information, refer to the Packet Trace Guide Technical Note.

Miscellaneous Options

Refer to the following sections for details on the LoadMaster **Miscellaneous Options**, including **WUI Settings**, **L7 Configuration**, and **Network Options**.

Contents

- WUI Settings
- L7 Configuration
- Network Options
- SDN Configuration
- Kemp 360 Central Activation Settings

WUI Settings

Only the **bal** user or users with 'All Permissions' set can use this functionality. Users with different permissions can view the screen but all buttons and input fields are grayed out.

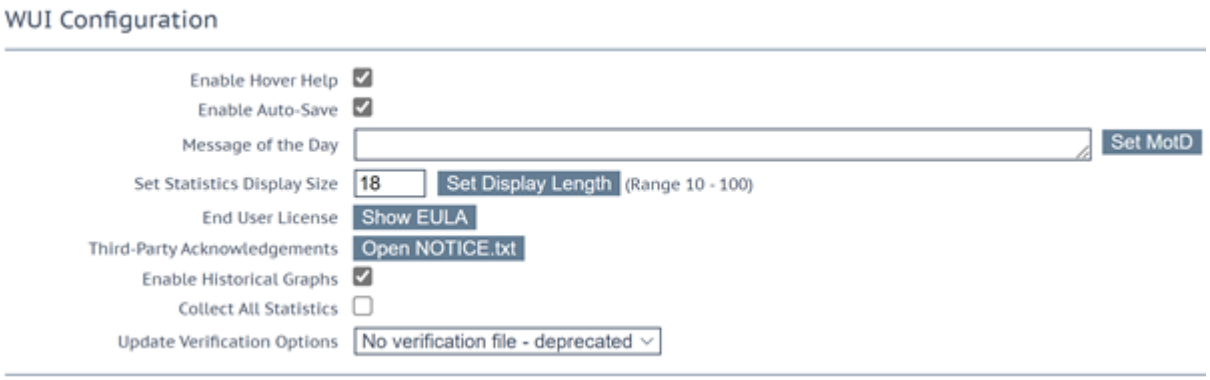

Enable Hover Help

Enables blue hover notes shown when the pointer is held over a field.

Enable Auto-Save

Auto-save is enabled by default. When auto-save is enabled, changes you make in the LoadMaster UI are applied instantly. For example, if you select an option in a drop-down list, the change is applied instantly.

If auto-save is disabled, several default UI behaviors are modified. For example, as in snapshot below the unsaved changes for Persistence Options are indicated by a background color change.

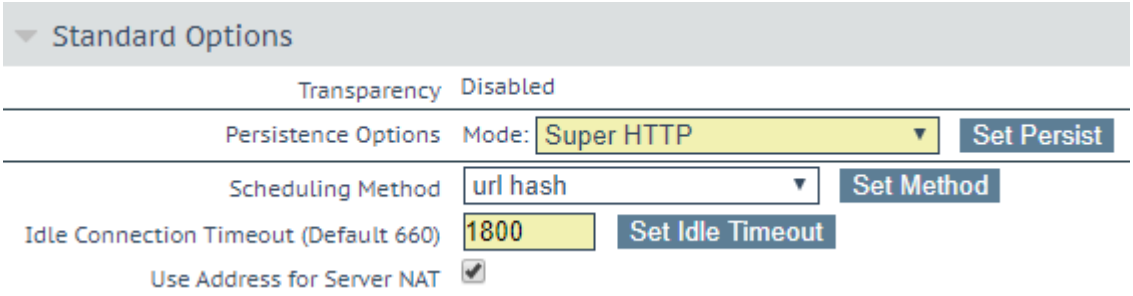

You must confirm drop-down selections by clicking the associated button (when auto-save is enabled, there is no button for some fields - the change is automatically applied once selected).

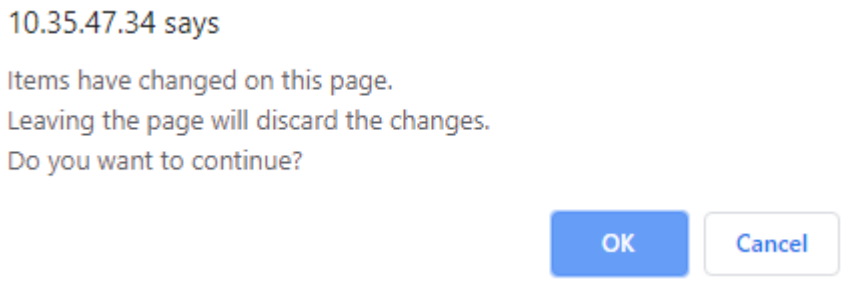

You are warned if you attempt to navigate away from a page with unsaved changes.

CAUTION: Disabling auto-save is currently a beta feature.

Note: If auto-save is disabled and you downgrade the LoadMaster firmware to a version older than 7.2.49, you cannot enable the auto-save feature because the check box will be unavailable.

The Progress Kemp documentation is written with auto-save enabled.

Message of the Day (MOTD)

Type in text into the field and click the **Set MotD** button. This message will be displayed within the LoadMaster Home screen.

If WUI Session Management is enabled, the MOTD is displayed on the login screen rather than the Home screen.

Note: The maximum allowed message length is 5,000 characters. HTML is supported, but not required. Single quotes (‘) and double quotes (“) are not allowed, though you can use the equivalent HTML character codes. For example, entering &#34it&#39s allowed&#34 would result in a MOTD of “it’s allowed”.

Note: Line breaks in the **Message of the Day** field in the LoadMaster UI must be in HTML code (such as <br/> or <p> tags).

Set Statistics Display Size

This sets the maximum number of rows that can be displayed in the Statistics page. The allowable range is between 10 and 100 rows being displayed on the page.

End User License

Click the **Show EULA** button to display the LoadMaster End User License Agreement.

Third-Party Acknowledgements

Click **Open NOTICE.txt** to view the third-party acknowledgments applicable to the LoadMaster release you are running. The text file opens in a new tab or window.

Enable Historical Graphs

Enable the gathering of historical statistics for the Virtual Services and Real Servers.

Collect All Statistics

By default, this option is disabled. This means that only the statistics for the Virtual Services and Real Servers that are configured to be displayed on the home page are collected. Enabling this option will force the LoadMaster to collect statistics for all Virtual Services and Real Servers.

Note: If there are a large number of Virtual Services and Real Servers this option can cause CPU utilization to become very high.

Update Verification Options

This field is set to **Required** by default. The other options are **Optional** and **No verification file - deprecated**. The different options are described below:

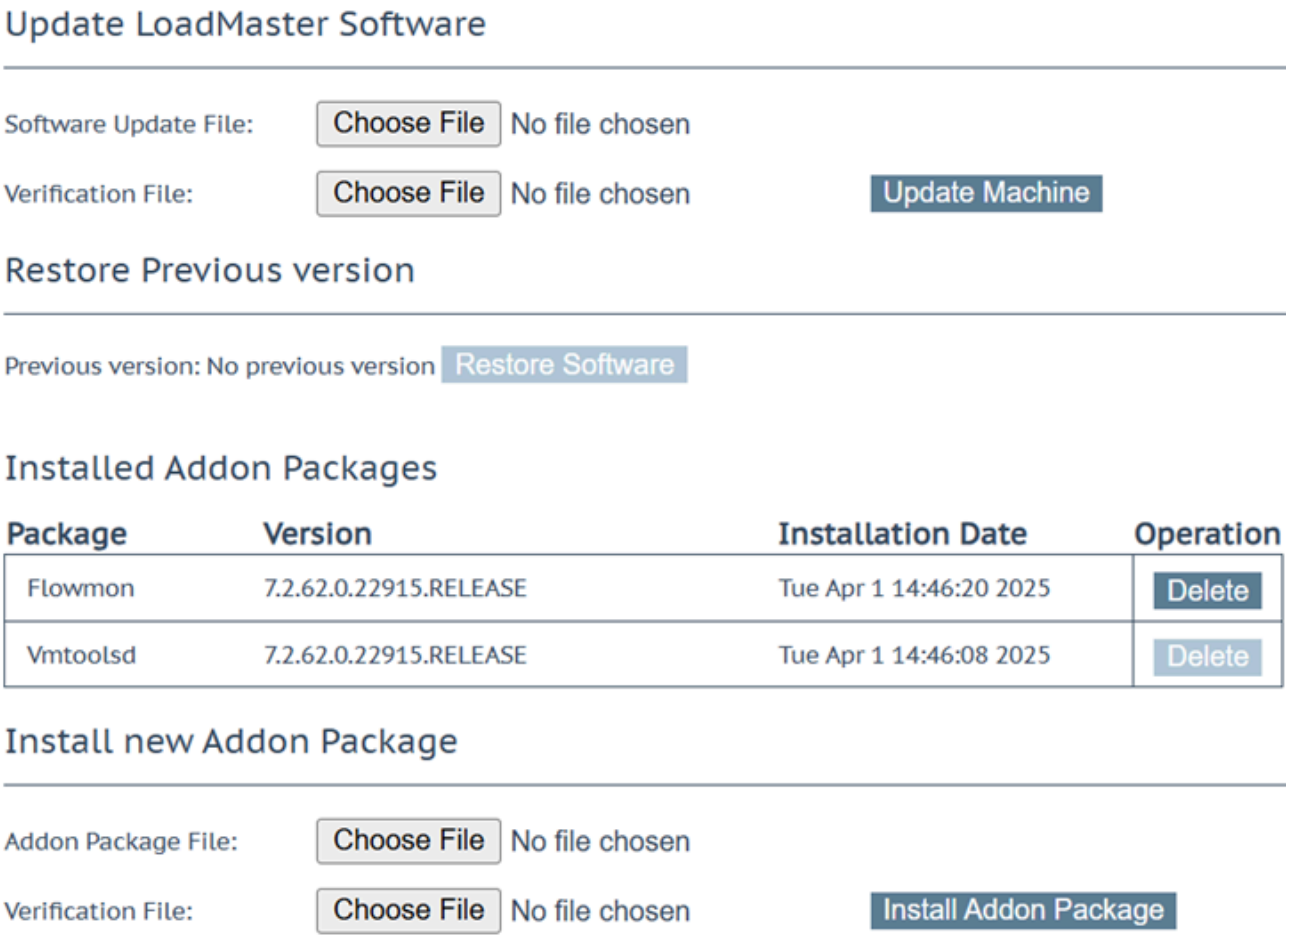

- **Required:** A **Verification File** option is displayed in the **Update Software** screen for updating the LoadMaster software and installing add-on packages. In this case, the verification file is mandatory and you cannot update the LoadMaster software or install an add-on without providing the verification file.

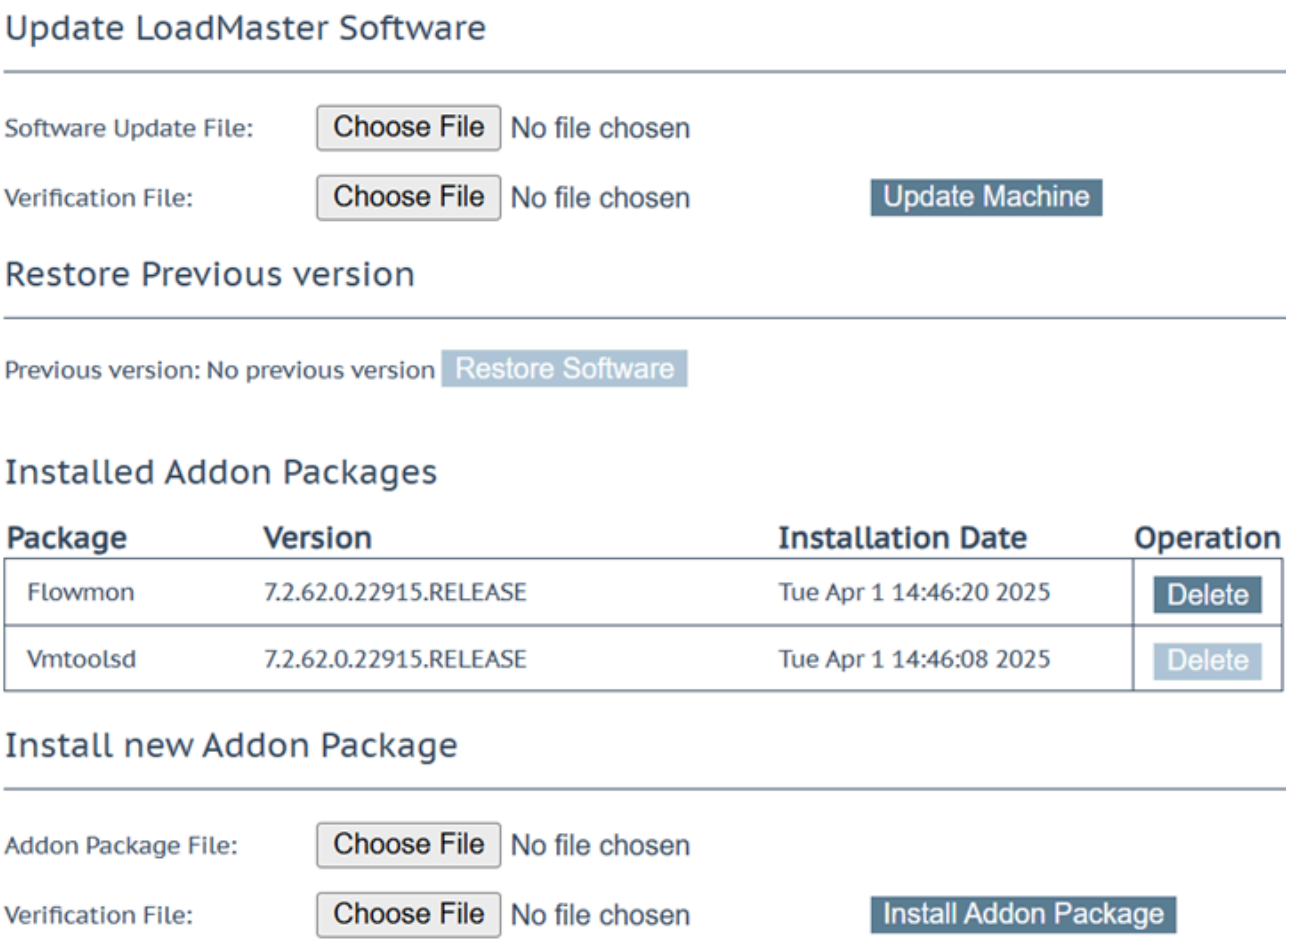

- **Optional:** A **Verification File** option is displayed in the **Update Software** screen for updating the LoadMaster software and installing add-on packages. In this case, the verification file is optional and you can update the LoadMaster software or install an add-on without providing the verification file.

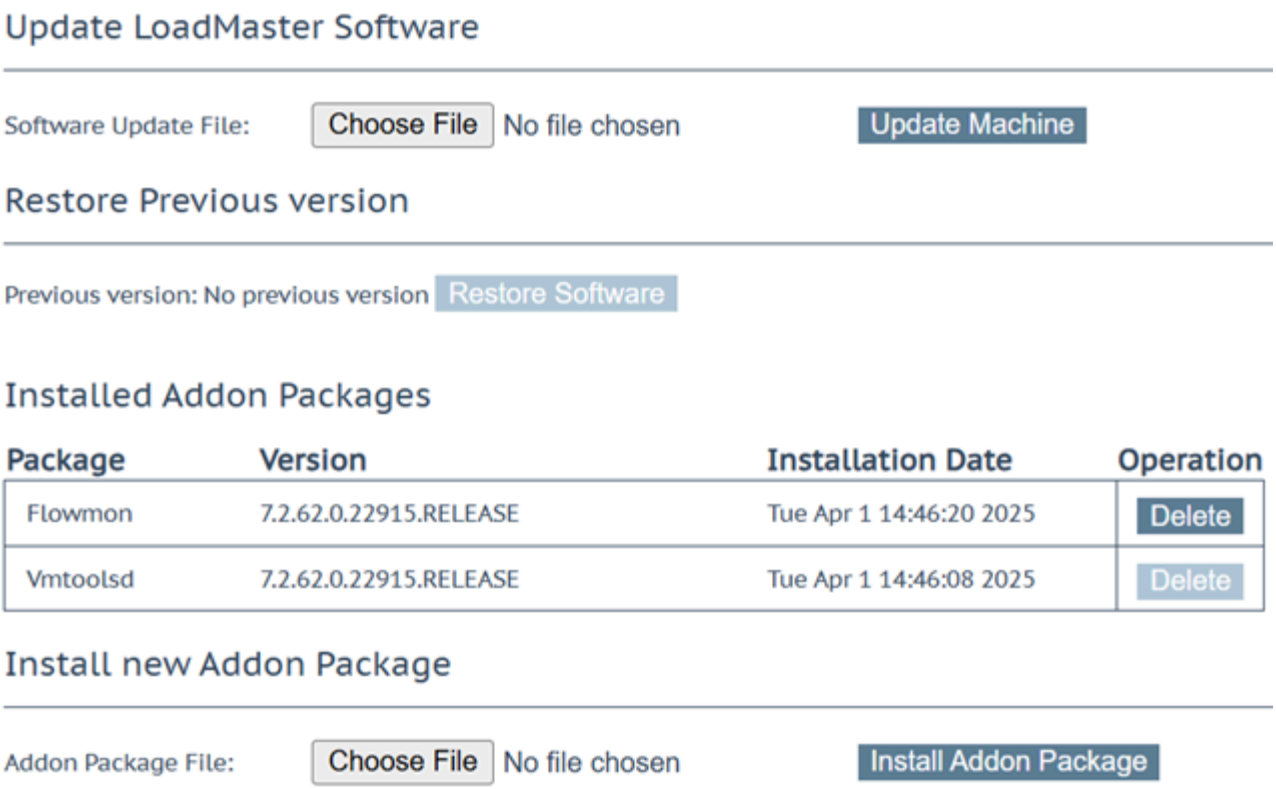

- **No verification file - deprecated:** The **Verification File** option is not displayed in the **Update Software** screen. In this case, you can update the LoadMaster software or install an add-on without providing the verification file. This option is provided for compatibility with legacy releases only and is deprecated.

The **Verification File** option provides an extra level of verification above the integrity checks that are always performed on update and add-on images, by allowing a secondary XML verification file to be uploaded at the same time that you are uploading a patch or add-on. This XML file is used to verify the digital signatures associated with the update or add-on which will not be installed if the patch cannot be verified against the uploaded XML file.

For FIPS machines - the **Update Verification Options** field is not shown because the uploading of a secondary XML verification file is mandatory when uploading a patch or add-on.

L7 Configuration

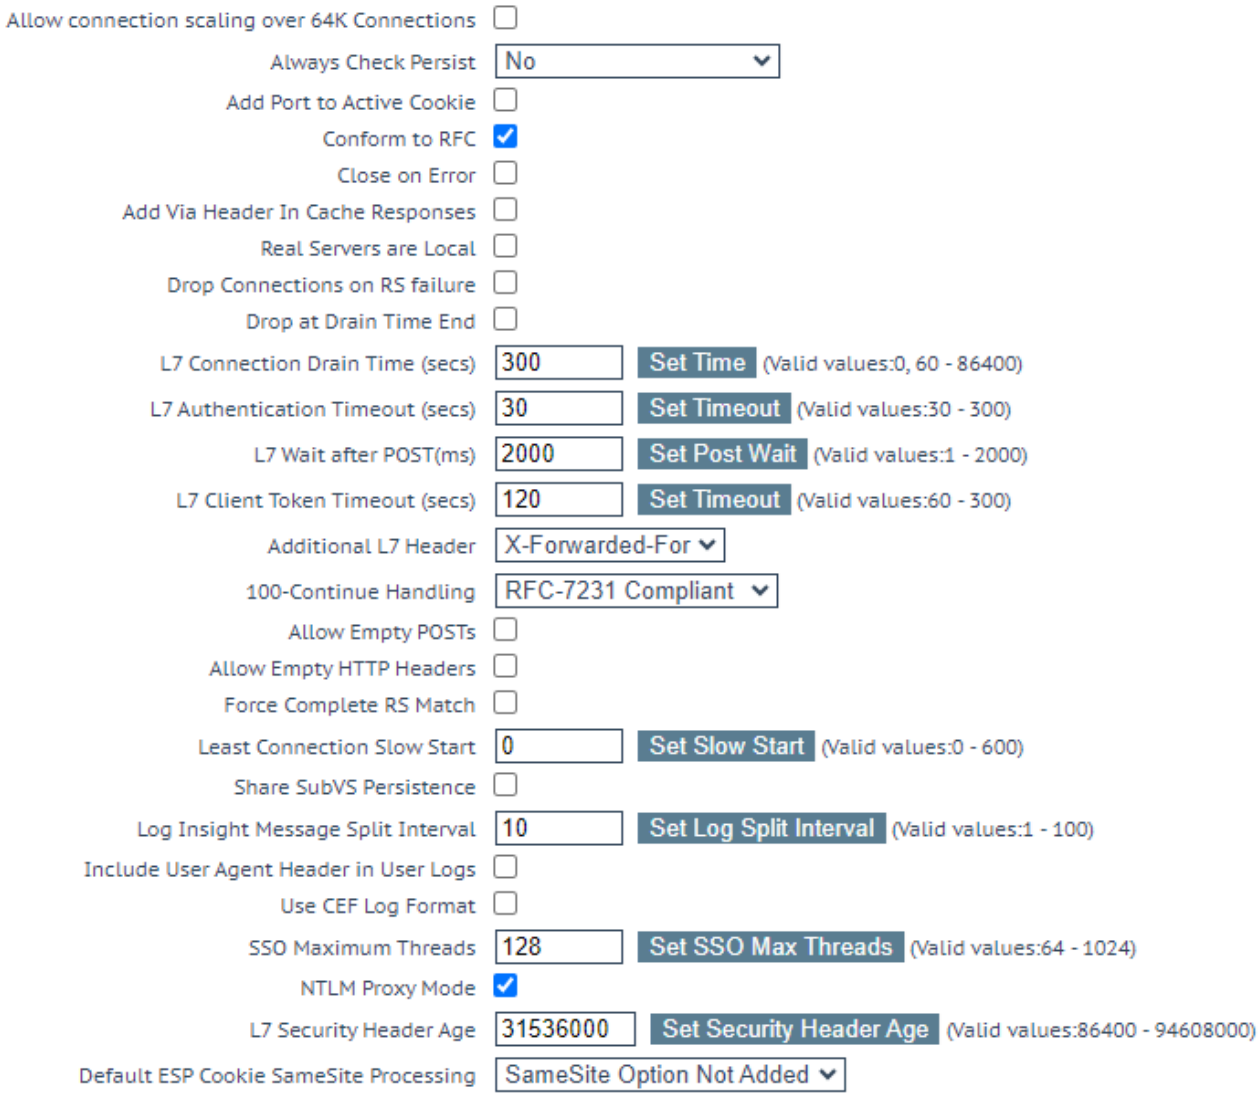

Allow Connection Scaling over 64K Connections

Under very high load situations, Port Exhaustion can occur. Enabling this option will allow the setting of Alternate Source Addresses which can be used to expand the number of local ports available.

Note: If more than 64K concurrent connections are required, enable the **Allow Connection Scaling over 64K Connections** option and set the Virtual Service IP as the alternate address in the **Alternate Source Addresses** input field. This allows each Virtual Service to have its own pool of source ports.

Note: Transparent Virtual Services are capped at 64K concurrent connections. This limit is on a per Virtual Service basis.

Note: If, after selecting this option, you set some Alternate Source Addresses, you will not be able to deselect the **Allow connection scaling over 64K Connections** option.

Always Check Persist

By default, the L7 module will only check persist on the first request of a HTTP/1.1 connection. Selecting **Yes** for this option will check the persistence on every request. Selecting **Yes – Accept Changes** means that all persistence changes will be saved, even in the middle of a connection.

Add Port to Active Cookie

When using active cookies, the LoadMaster creates the cookie from (among other things) the IP address of the client. However, if many clients are behind a proxy server, all of those clients come from the same IP address. Turning this on adds the clients source port to the string as well, making it more random.

Conform to RFC

This option addresses parsing the header of a HTTP request in conformance with RFC 1738.

The request consists of 3 parts: GET /pathname HTTP/1.1 and when "conform" is on, the LoadMaster scans through the pathname until it finds a space. It then presumes that the next thing is HTTP/1.x. If the pathname contains spaces and the browser is conformant to the RFC, the pathname will have the spaces escaped to "%20" so the scan for a space will function correctly.

However, on some non-conformant browsers, spaces are not escaped and the wrong pathname is processed. And since the system cannot find the HTTP/1.x, the LoadMaster will reject the request.

Turning off this feature forces the LoadMaster to assume that the pathname extends to the last space on the line. It is then assumed that what follows is HTTP/1.x. So making pathnames with spaces in them useable – however, it is non-conformant to the RFC 1738.

Close on Error

If the LoadMaster has to send back a failure report to the client, for example if a file is newer in the cache; this forces the LoadMaster to close the connection after sending the response. You can continue using the connection after sending a failure report, but some systems could become confused. This option forces the close instead of continuing.

Add Via Header In Cache Responses

The relevant HTTP RFC states that proxies should add a Via header to indicate that something came from the cache. Unfortunately, older LoadMaster versions did not do this. This check box is used to enable backward compatibility with older versions (if needed).

Real Servers are Local

The LoadMaster has an automatic detection of local/non-local clients for the purpose of transparency (selective transparency). This works well in most cases, but it does not work well if the client is actually a Real Server. Turning this option on helps the LoadMaster to determine that a Real Server is actually local, therefore making selective transparency work.

When this option is enabled in a two-armed environment (with clients and Real Servers on the second interface) the Real Servers are treated as if they are local to the clients, that is, non-transparent. If the Real Servers are on a completely different network, then they cannot be local and will always be treated as not local. Local is defined as being on the same network.

Enabling this option requires careful network topology planning and should not be attempted before contacting the Progress Kemp Support team.

Drop Connections on RS Failure

This only applies when a Real Server fails a health check. This is unrelated to the **Drop at Drain Time End** option. When this option is enabled, if the Real Server fails the health check, the connection is dropped. The client must re-establish a connection. If this option is disabled, the connection is not dropped.

Drop at Drain Time End and L7 Connection Drain Time (secs)

L7 Connection Drain Time (secs): When a Real Server is disabled, the sessions that are still persistent to that Real Server continue to be served until the Drain Time has expired or until there are no more active sessions being handled by the Real Server. No new sessions are handled by the Real Server.

If **Drop at Drain Time End** is disabled the LoadMaster, in theory, could still forward traffic to the Real Server indefinitely if the session never times out or closes. For this reason, in most scenarios, the **Drop at Drain Time End** option should be enabled.

If **Drop at Drain Time End** is enabled, regardless of the persistence and the active sessions, once the **L7 Connection Drain Time** hits zero, all connections to this Real Server (relating to this Virtual Service) will be dropped.

These two options control how existing connections are handled when a Real Server in a Layer 7 Virtual Service is disabled. They do not apply to Virtual Services operating at Layer 4.

Drop at Drain Time End is *disabled* by default, which means that the **L7 Connection Drain Time** setting is not enforced. In this case, when a Real Server is disabled, no new connections will be sent to the Real Server and all existing connections are immediately closed.

When **Drop at Drain Time End** is *enabled*:

- No new connections will be sent to the Real Server unless there is a valid persistence record for the client present on the LoadMaster.
- Existing *non-persistent* connections (that is, connections from a client with no matching persistence record on LoadMaster) are closed immediately.
- Existing *persistent* connections to the Real Server remain open until the **L7 Connection Drain Time** expires. Setting the drain time to **0** forces all the connections to be dropped immediately when a Real Server is disabled. The default setting is **300** seconds.

Note the following for Virtual Services with SubVSs:

- When a Real Server in a SubVS is disabled, the **L7 Connection Drain Time** setting is effective only when the **Persistence Options** value is set *in the SubVS configuration*. If persistence is set on the parent Virtual Service level instead, then a default 10-second drain time is enforced for the Real Server. So, in the following scenario:
  - If the **Persistence Options** are set in the parent Virtual Service, and;
  - There are no **Persistence Options** set in the SubVS, and;
  - The global **L7 Connection Drain Time** setting is not set to **0**;

  The 10-second drain time overrides the **L7 Connection Drain Time** in this specific case.
- The **L7 Connection Drain Time** does not apply when a SubVS is disabled.

L7 Authentication Timeout (secs)

This option supports the integration with 3rd party, multi-factor, authentication solutions which may have

secondary processes such as SMS or telephone verification. This setting determines how long (in seconds) the SSO form waits for authentication verification to complete before timing out.

L7 Client Token Timeout (secs)

The duration of time (in seconds) to wait for the client token while the process of authentication is ongoing (used for RSA SecurID and RADIUS authentication). The range of valid values is 60 to 300. The default value is 120.

L7 Wait after POST(ms)

In LoadMaster firmware version 7.2.51, a new option was introduced that is applicable when performing Kerberos Constrained Delegation (KCD) back-end authentication. The field is called **L7 Wait after POST**. This option is configurable in the LoadMaster User Interface (UI). The **L7 Wait after POST** option allows you to change the length of time to wait for a 401 response from a POST before sending the remainder of the POST body. Valid values for the wait period range from 1 to 2000 milliseconds (ms). The default value is 2000. If KCD is not being used, this option has no effect.

Additional L7 Header

This enables Layer 7 header injection for HTTP/HTTPS Virtual Services. Header injection can be set to **X-ClientSide** (LoadMaster specific), **X-Forwarded-For**, or **None**. The default value is **X-Forwarded-For**.

100-Continue Handling

Determines how **100-Continue Handling** messages are handled. The available options are:

- **RFC-2616 Compliant**: conforms with the behavior as outlined in RFC-2616
- **Require 100-Continue:** forces the LoadMaster to wait for the 100-Continue message
- **RFC-7231 Compliant:** ensures the LoadMaster does not wait for 100-Continue messages. This is the default value.

Note: Modifying how 100 Continue messages are handled by the system requires an understanding of the relevant technologies as described in the RFCs listed above. It is recommended that you speak with a Progress Kemp Technical Support engineer before making changes to these settings.

Allow Empty POSTs

By default the LoadMaster blocks POSTs that do not contain a Content-Length or Transfer-Encoding header to indicate the length of the requests payload. When the **Allow Empty POSTs** option is enabled, such requests are assumed to have no payload data and are therefore not rejected.

Note: In version 7.1-24 and later releases, the supported Content-Length limit has been increased to 2TB (from 2GB).

Allow Empty HTTP Headers

By default, this option is disabled. If this option is disabled, the LoadMaster drops (does not process) any HTTP header that has an empty value. The LoadMaster only drops the header with an empty value, not the whole request. If this option is enabled, the LoadMaster processes headers with empty values and forwards them to the Real Server.

Force Complete RS Match

By default, when the LoadMaster is trying to locate a Real Server for use with content switching, it tries to use the same Real Server as currently selected, even if the port is not the same. Enabling this option forces the port to also be compared.

Least Connection Slow Start

When using the **Least Connection** or **Weighted Least Connection** scheduling methods, a period can be specified globally using the **Least Connection Slow Start** field during which the number of new connections are throttled and gradually increased to a Real Server which has come online and has returned to the scheduling process. When any Real Server is brought back into service and the **Least Connection Slow Start** is set to a non-zero value, the LoadMaster throttles new traffic to the Real Server so that is it not potentially overwhelmed by a sudden stream of traffic. Under testing, the observed Connections Per Second (CPS) rate limit is observed to grow slowly over the specified time period until the full Real Server connection capacity is permitted. The slow start applies regardless of the reason the Real Server was removed from the scheduling process (for example, manually disabled, rate-limited, and so on).

The value of this **Slow Start** period can be between **0** (disabled - this is the default) and **600** seconds.

The **Least Connection Slow Start** feature can be used in conjunction with the **Connection Rate Limit** feature that was introduced in LoadMaster firmware version 7.2.51.

Share SubVS Persistence

By default, each SubVS of a Virtual Service has an independent persistence table. Enabling this option will allow the SubVS to share this information. For this to work, the persistence mode must be the same on all SubVSs within that Virtual Service. A reboot is required to activate this option.

Note: The only **Persistence Mode** that cannot be shared is **SSL Session ID**.

When setting up shared SubVS persistence, there are some requirements to get this feature fully functional:

- All Real Servers in the SubVS need to be the same
- The **Persistence Mode** needs to be the same across all SubVSs
- The timeouts need to be set with the same timeout value

If the above requirements are not correct, the persistence may not work correctly either within the SubVS or across the SubVSs.

Log Insight Message Split Interval

The **Log Insight Split Interval** value controls how many syslog messages should be sent to each server in the pool before moving to the next server. For example, if there are three Log Insight nodes and the **Log Insight Message Split Interval** is set to **1** - a single message is sent to server A, and then to server B and then server C before again distributing a message to server A.

Include User Agent Header in User Logs

When enabled, the User Agent header field gets added to the User Logs.

Use CEF Log Format

When enabled, the ESP logs are generated in Common Event Format (CEF). CEF log format is easily consumable for Security Information and Event Management (SIEM) tools, such as; Splunk, SolarWinds, LogRhythm, AlienVault, and so on.

SSO Maximum Threads

The maximum number of allowed threads for SSO authentication attempts. The range of valid values is 64 to 1024. The default value is 128.

NTLM Proxy Mode

In LoadMaster firmware version 7.2.48.4 Long Term Support (LTS) and 7.2.53, the **NTLM Proxy Mode** option was added to the LoadMaster. When upgrading from an older version of LoadMaster firmware to one of these versions (or above) the **NTLM Proxy Mode** option is not enabled by default. As a result, you must manually enable **NTLM Proxy Mode** after upgrading.

For all new deployments of LoadMasters after 7.2.48.4 LTS or 7.2.53, **NTLM Proxy Mode** is enabled by default.

When **NTLM Proxy Mode** is enabled, NTLM authorization works against the Real Servers. If **NTLM Proxy Mode** is disabled, the old insecure NTLM processing is performed.

Note: We highly recommend ensuring that **NTLM Proxy Mode** is enabled.

When **NTLM Proxy Mode** is enabled globally, the **Client Authentication Mode** in Virtual Services is called **NTLM-Proxy**. If **NTLM Proxy Mode** is disabled globally, the **Client Authentication Mode** in Virtual Services is called **NTLM**.

L7 Security Header Age

This option allows you to configure a custom timeout value to determines how long (in seconds) a security header (STS) send by LoadMaster is valid. The range of valid values is 86400 to 94608000. The default value is 31536000.

Default ESP Cookie SameSite Processing

This option allows you to set the default value of SameSite option for cookies sent by the LoadMaster during ESP processing. The SameSite attribute tells browsers when and how to process cookies in first-party or third-party situations. SameSite is used by a variety of browsers to identify whether or not to allow a cookie to be accessed.

The following are the SameSite options:

- **SameSite option not added**
- **SameSite=None:** Signals that the cookie data can be shared with third parties/external sites (for advertising, embedded content, etc).
- **SameSite=LAX:** Signals that the cookie may be used as a first party cookie but may also be used when accessing the site from an external site via a link clicked by the user.
- **SameSite=Strict:** It is a subset of lax and only enables the cookie be used in first party context and excludes its use when accessing via an incoming link from an external site.

Network Options

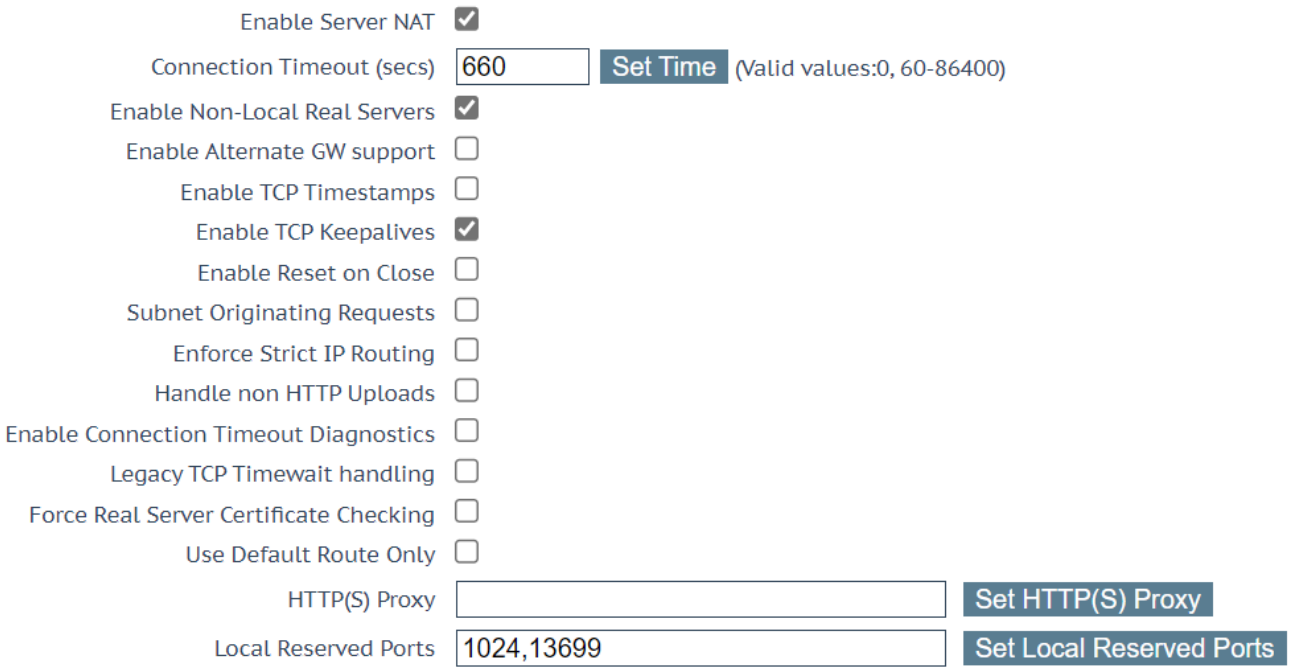

Enable Server NAT

This option enables Server Network Address Translation (SNAT). If this is disabled, the Real Server IP address is used when connecting.

If this is enabled, addresses that are of the same address family (IPv4/IPv6) as the primary address of the default gateway are NATed to the “primary address”. If the **Use Address for Server NAT** is enabled in the Virtual Service, the Virtual Service address will be used. For further information on the **Use Address for Server NAT** option, refer to the Standard Options section.

If the source address is not in the same family as the primary address, then the address will be SNATed to the first additional address which is on the same network as the default gateway for that address family.

For example, if the primary address of the default interface is an IPv6 address, then IPv6 addresses will be SNATed to that address. If the primary address is an IPv4 address, then IPv6 addresses will be SNATed to the first additional address (IPv6) which is on the same network as the IPv6 default gateway.

Similarly, if the primary address of the default interface is an IPv4 address, then IPv4 addresses will be SNATed to that address. If the primary address is an IPv6 address, then IPv4 addresses will be SNATed to the first additional address (IPv4) which is on the same network as the IPv4 default gateway.

Note: FTP and SNAT do not work together reliably all the time so this configuration is not supported.

Connection Timeout (secs)

The length of time (in seconds) that a connection may remain idle before it is closed. This value is independent of the Persistence Timeout value.

Setting a value of **0** will reset the value to the default setting of **660** seconds.

Enable Non-Local Real Servers

Allow non-local Real Servers to be assigned to Virtual Services. This may be needed if the LoadMaster can only have one interface and the Real Servers are on a different network to the interface. This option is enabled by default.

Enable Alternate GW support

If there is more than one interface enabled, this option provides the ability to move the default gateway to a different interface.

Enabling this option adds another option to the **Interfaces** screen – **Use for Default Gateway**.

Note: The **Enable Alternate GW support** option will appear on a different screen in GEO only LoadMasters.

Note: Alternate default gateway support is not permitted in a cloud environment.

Enable TCP Timestamps

The LoadMaster can include timestamps in the SYN on both connections from clients and connections to Real Servers.

Note: Note this may impact connections that are NATed and should only be enabled on consultation with Progress Kemp Customer Support.

Enable TCP Keepalives

By default the TCP keepalives are enabled which improves the reliability of TCP connections that are long lived (SSH sessions). Keepalives are not usually required for normal HTTP/HTTPS services, but may be required for FTP services, for example.

The keepalive messages are sent from the LoadMaster to the Real Server and to the client. Therefore, if the client is on a mobile network, there may be an issue with additional data traffic.

Enable Reset on Close

When this setting is disabled (the default), unencrypted and encrypted TCP connections to the LoadMaster on both the client and server sides are closed using the standard TCP exchange of FIN and ACK packets. In situations where a Virtual Service is under a high incoming connection load, the ability to establish new connections to the Virtual Service can be improved by turning on **Enable Reset on Close**; this tells the LoadMaster to close TCP connections with a single TCP RST (reset) packet, rather than the normal TCP closing exchange.

Subnet Originating Requests

With this option enabled, the source IP address of non-transparent requests will come from the LoadMaster's address on the relevant subnet, that is, the subnet where the Real Server is located or the subnet of the gateway that can route to the Real Server (if the Real Server is non-local and configured to use static route). For more information on configuring a static route, refer to the following knowledge base article: Creating a

Static Route.

This is the global option/setting.

Note: It is recommended that the **Subnet Originating Requests** option is enabled on a per-Virtual Service basis.

When the global option is disabled, the per Virtual Service **Subnet Originating Requests** option takes precedence, that is, it can be enabled or disabled per Virtual Service. This can be set in the **Standard Options** section of the Virtual Services properties screen (if **Transparency** is disabled). For more information on the per Virtual Service option, refer to the Standard Options section.

Note: If this option is switched on for a Virtual Service that has SSL re-encryption enabled, all connections currently using the Virtual Service will be terminated because the process that handles the connection must be killed and restarted.

Enable Strict IP Routing

When this option is selected, only packets which arrive at the machine over the same interface as the outbound interface are accepted.

Note: The **Use Default Route Only** option may be a better way to achieve this.

Handle non HTTP Uploads

Enabling this option ensures that non HTTP uploads (such as FTP uploads) function correctly.

Enable Connection Timeout Diagnostics

By default, connection timeout logs are not enabled. This is because they may cause too many unnecessary logs. If you wish to generate logs relating to connection timeouts, select the **Enable Connection Timeout** check box.

Legacy TCP Timewait Handling

Enable this option to revert to the legacy mode of reusing TCP timewait connections.

Note: Only enable the **Legacy TCP Timewait Handling** option after consulting with Progress Kemp Support.

Force Real Server Certificate Checking

By default, when re-encrypting traffic the LoadMaster does not check the certificate provided by the Real Server. This option forces the LoadMaster to verify that the certificate on the Real Server is valid, that is, the certificate authority and expiration are OK. This includes all intermediate certificates.

Use Default Route Only

Forces traffic from Virtual Services that have a Virtual Service gateway set, to only be routed to the interface where the Virtual Service gateway is located.

This setting can allow the LoadMaster to be directly connected to client networks without returning traffic directly, by using the Virtual Service gateway instead.

Note: Enabling this option affects all Virtual Services that have a Virtual Service gateway set.

Note: Other network options may affect routing such as **Subnet Originating Requests**, refer to the Routing Feature Description document for further details.

For further details on the **Use Default Route Only** option, refer to the Use Default Route Only section of the Routing Feature Description.

HTTP(S) Proxy

This option allows clients to specify the HTTP(S) proxy server and port the LoadMaster will use to access the internet. This must be an IP address and port (not an FQDN).

Local Reserved Ports

A single port, or list of comma-separated ports, that will not be used when initiating connections to Real Servers. Some things to note about this field are as follows:

- Ports specified in the list must be between 1024 and 63999 (inclusive).
- A port cannot appear more than once in the list.
- The total length of the list cannot exceed 128 characters.
- If you are making changes to an existing list, the entire list of ports must be specified.

This option would be used to remove specific ports from the list of ports used by the system to communicate with back-end Real Servers, typically because:

- those ports are dedicated for some other purpose than load balancing
- and, it is either desired that no connections are made to that server port for load balancing purposes, or that the server will simply drop load balancing connections received on that port.

Warning: To protect against port exhaustion during periods of heavy traffic, LoadMaster uses a wide set of ports (1024 to 63999). Significantly reducing the number of these ports available for connections to Real Servers could lead to port exhaustion (that is, dropped server-side connections), so best practice is to keep the number of ports removed low – a single-digit percentage of all ports.

SDN Configuration

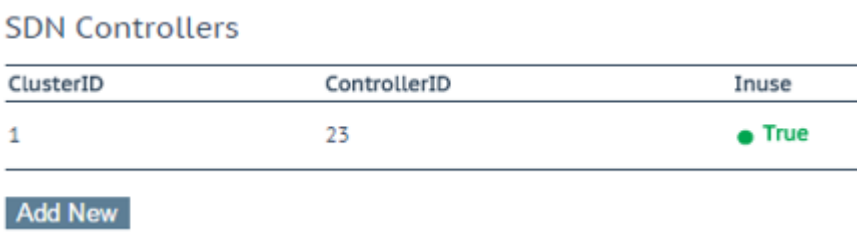

Add New

Add a new SDN controller connection.

Modify

Modify an existing SDN controller connection.

Delete

Delete an existing SDN controller connection.

SDN Controller Settings

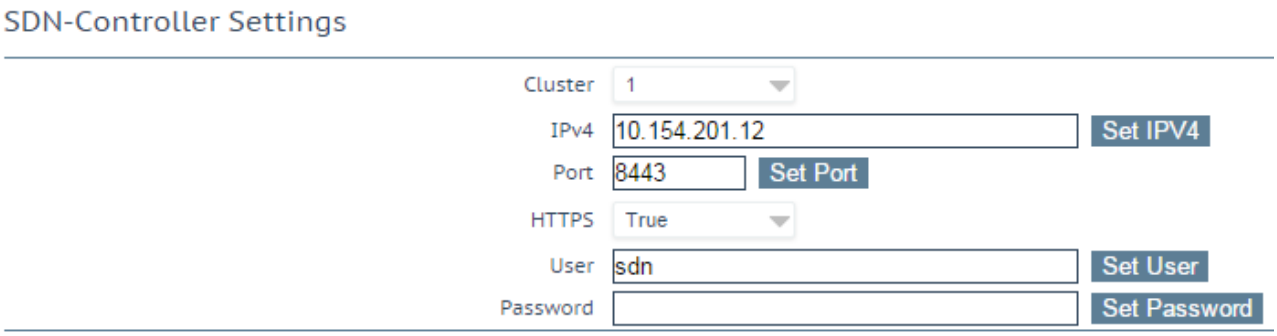

When adding a new SDN controller connection, initially a screen will appear asking for the **Cluster**, **IPv4** address and **Port**. After an SDN controller connection has been added, the settings can be updated by clicking **Modify** on the **SDN Statistics** screen.

Cluster

The cluster that the SDN controller will be a member of.

Note: Keep the **Cluster** field set to the default value.

IPv4

The IPv4 address of the SDN controller.

Port

The port of the SDN controller WUI.

Note: The default **Port** for the HP VAN Controller is **8443**.

Note: The default **Port** for the OpenDaylight SDN controller is **8181**.

HTTPS

Use HTTP/HTTPS to access the SDN controller.

User

The username to be used to access the SDN controller.

Password

The password of the user to be used to access the SDN controller.

Kemp 360 Central Activation Settings

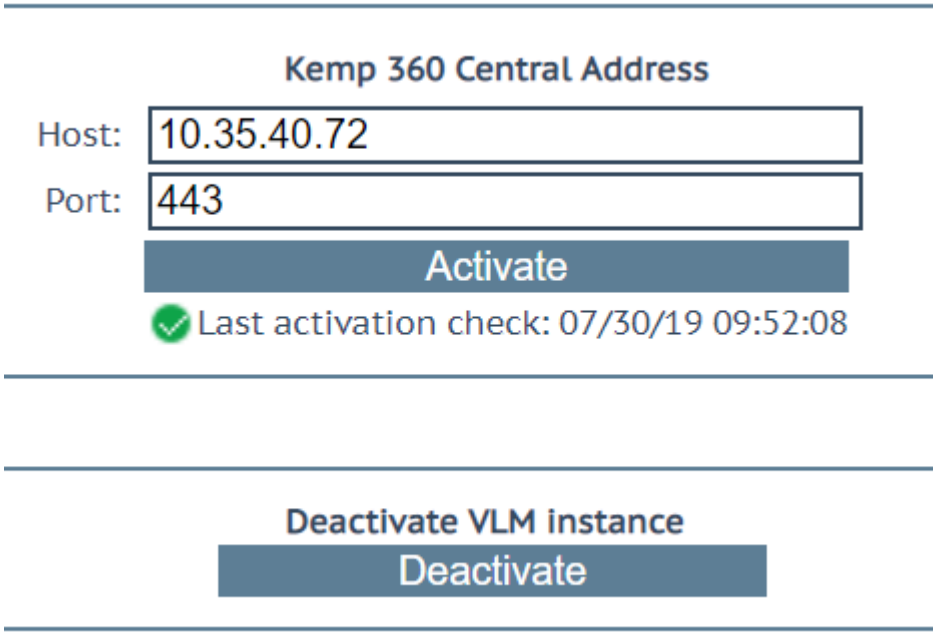

To view the Kemp 360 Central Activation Settings, click **System Configuration > Miscellaneous Options > Kemp 360 Central Activation Settings**. If the LoadMaster is already licensed from Kemp 360 Central, the Kemp 360 Central IP address and port appear here. If blank, you can use the controls on this page to provide the IP address and port of a Kemp 360 Central deployment from which you want to request a license. By activating your LoadMaster here, it is automatically added to Kemp 360 Central and the system can start obtaining statistics from it immediately.

The **Deactivation** button removes the current license from the LoadMaster, returning the LoadMaster to an unlicensed state. This stops any active traffic to the LoadMaster, and so should be invoked with great caution.

For more information, refer to the Kemp 360 Central Feature Description.

12

Network Telemetry

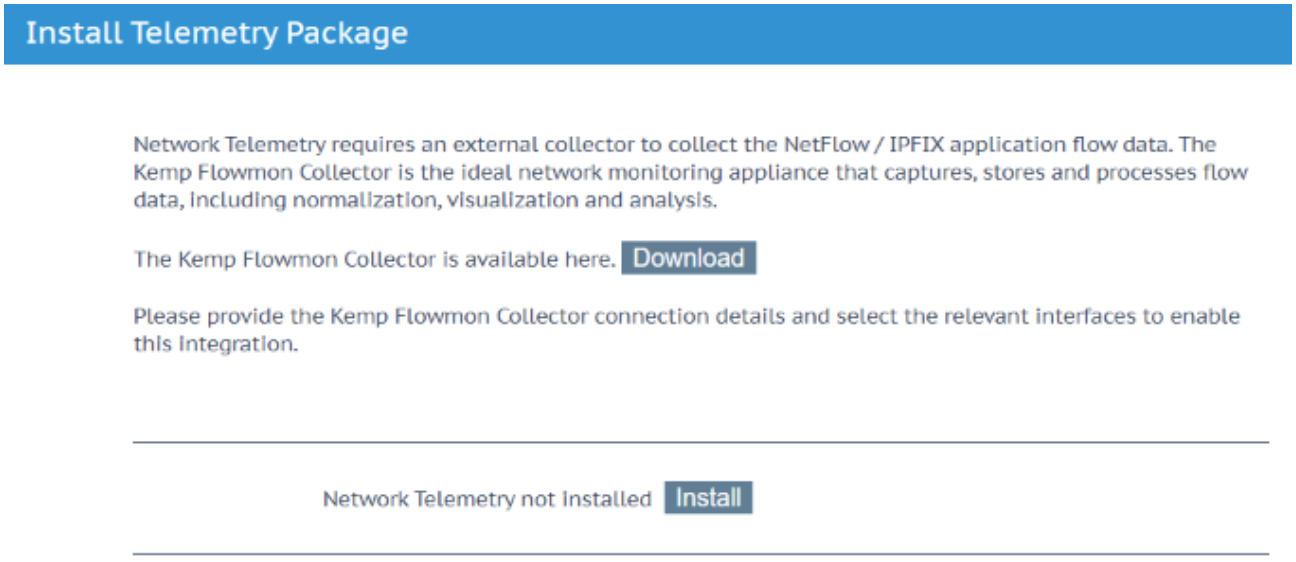

Network telemetry requires an external collector to collect the NetFlow/IPFIX application flow data. The Progress Kemp Flowmon Collector is the ideal network monitoring appliance that captures, stores, and processes flow data, including normalization, visualization, and analysis. You can download the Flowmon Collector by clicking **Download Flowmon Collector** on the Network Telemetry screen.

Network telemetry is available by default on all new LoadMaster deployments for firmware version 7.2.53 and above.

Note: Enabling Network Telemetry may impact performance throughput.

On Long Term Support (LTS) LoadMaster versions, or older versions of the LoadMaster that have been patched to a newer version, you may need to enable the network telemetry feature. To enable the network telemetry feature, click **Network Telemetry** in the main menu of the LoadMaster WUI and click **Install**.

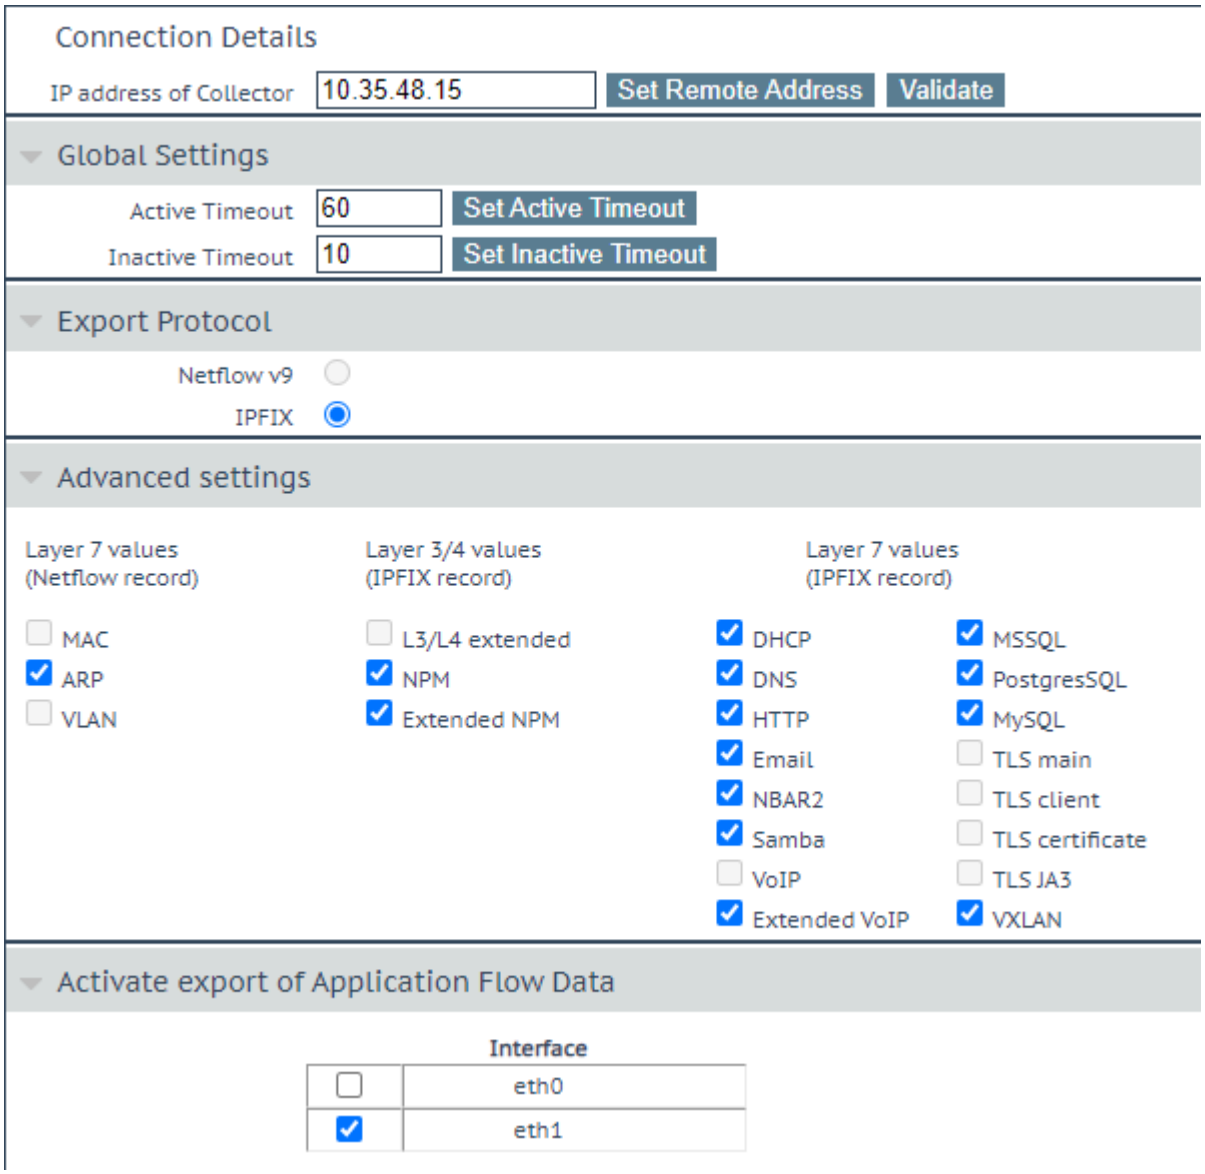

After you successfully install network telemetry on the LoadMaster, you should see a number of fields to configure on the **Network Telemetry** screen. Details on each of these options are below:

- **IP address of Collector:** Define the destination IP address or Fully Qualified Domain Name (FQDN) and port number of your IPFIX collector (for example, **1.1.1.1:2055** or **collector.local:3000**). The IPFIX export runs over the UDP protocol and you must ensure that the collector is reachable over the network from the LoadMaster. Once you configure the collector IP address or FQDN you can validate the network connectivity by clicking **Validate** and clicking **OK**. Validation is based on a plain ICMP ping message and it validates the IP or FQDN (not the port).

Note: Because the network telemetry is not encrypted it should only be exported over a secure network. This is not a LoadMaster limitation - this is because IPFIX is a plain text UDP packet stream.

- **Active Timeout:** Set the global active timeout value. The default value is **300**.
- **Inactive Timeout:** Set the global inactive timeout value. The default value is **30**.
- **Export Protocol:** The export protocol (**IPFIX** is currently the only selectable protocol).
- **Advanced Settings:** Enable/disable the check boxes here depending on what values you would like to collect.

Note: There are some check boxes in the **Advanced settings** section that are not possible to change at present. These will be configurable in a future release.

- **Activate export of Application Flow Data:** Select the relevant interface (or interfaces) to collect data for.

The network interface screens (for example, **System Configuration > Network Setup > Interfaces > eth0**) indicate if network telemetry monitoring is enabled or disabled for that interface (depending on what interfaces are selected on the **Network Telemetry** screen).

Note: In order to enable Network Telemetry on an interface, the interface must have an IP Address. Interfaces configured with Virtual LANs can not have Network Telemetry enabled unless an IP Address is assigned.

Note: When two LoadMasters are operating as a HA pair, the Network Telemetry traffic will present from the physical address of the LoadMaster and not the shared address. It is necessary to create a profile for both the LoadMasters in the HA pair on the Flowmon Collector.

For further details on network telemetry, refer to the Network Telemetry Feature Description.

13

Help

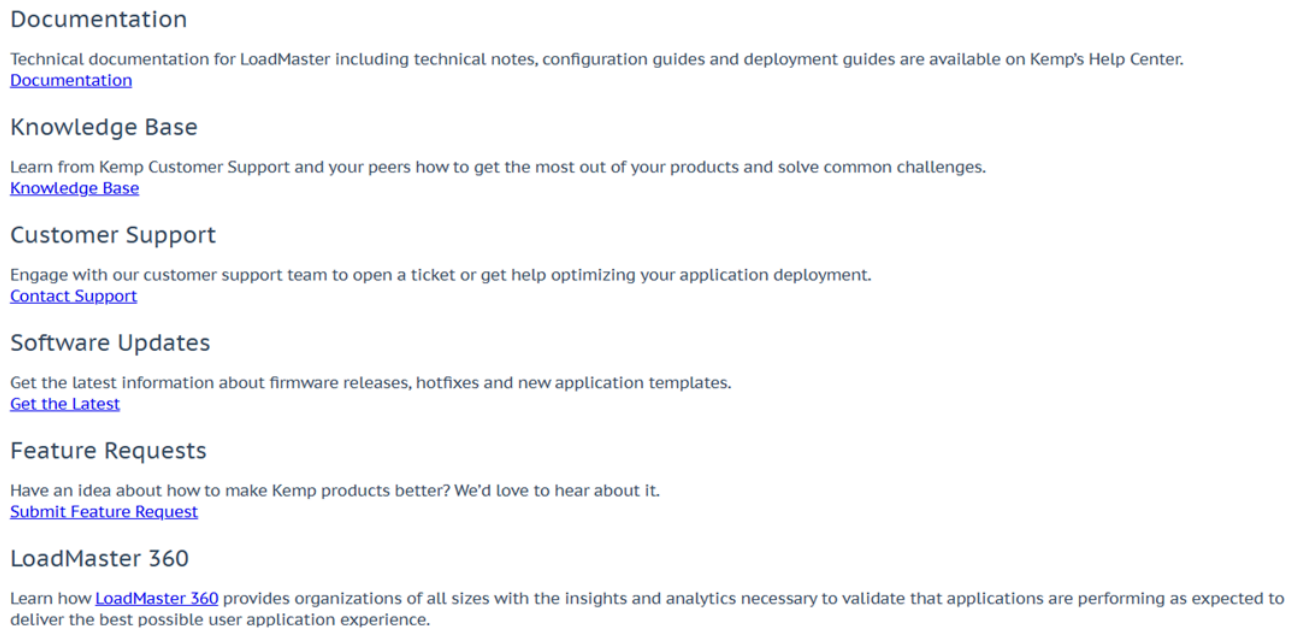

The **Help** screen provides a consolidated location for access to external Progress Kemp services.

Documentation

Access the Progress Kemp technical documentation, including Deployment Guides, Installation Guides, Feature Descriptions, Technical Notes, Overviews, Release Notes, and Interface Descriptions.

Network Telemetry Documentation

Access the documentation relating to the network telemetry feature.

Knowledge Base

Access Knowledge Base articles on a variety of subjects such as SSO/ESP, Fault Tolerance, Operational Maintenance, Applications, Security, Platforms, Routing/Switching, and Content Delivery.

Customer Support

Open a ticket with Progress Kemp's Customer Support team.

Software Updates

Get the latest information about firmware releases, hot fixes and new application templates.

Feature Requests

Take a look at existing feature requests submitted by other customers, and raise your own feature request.

LoadMaster 360

Take a look at LoadMaster 360, which is a powerful solution designed to empower organizations of all sizes with the tools they need to ensure optimal application performance. By providing comprehensive insights and analytics, LoadMaster 360 enables businesses to validate that their applications are functioning as intended, ensuring a seamless and exceptional user application experience.

14

References

Other useful documents are as follows:

Virtual Services and Templates Feature Description

RSA Two Factor Authentication Feature Description

Content Rules Feature Description

GEO Feature Description

GEO Sticky DNS Feature Description

Packet Trace Guide Technical Note

VMware Tools Add-On Package, Feature Description

Customizing SSO Authentication Forms Technical Note

Port Following Feature Description

SSL Accelerated Services Feature Description

Kerberos Constrained Delegation Feature Description

IPsec Tunneling Feature Description

Progress Kemp LoadMaster Product Overview

SDN Adaptive Load Balancing Feature Description

DoD Common Access Card (CAC) Authentication Feature Description

RESTful API Documentation

LoadMaster Licensing Feature Description

Feature Description LoadMaster Clustering

Microsoft Exchange 2010 Deployment Guide

RADIUS Authentication and Authorization Technical Note

User Management Feature Description