



# Appliance Administration Manual

Version 8.8.1 Revision 1.0.3



# Table of Contents

Chapter 1 - Introduction	15
About this Manual	15
About the Appliance	15
Latest Version	15
Appliance Configuration Overview	16
Appliance Security	16
Security Mode	16
Passwords	17
The 'loadbalancer' WebUI account	17
The 'root' Linux account	17
Ports Used by the Appliance	17
Additional Information	18
Deployment Guides	18
Quick Start & Configuration Guides	18
Contacting Support	18
Chapter 2 - Load Balancing Concepts	19
Load Balancing - the Basics	19
Supported Protocols	19
Layer 4 & Layer 7	19
Load Balancing Algorithms	19
Round Robin / Weighted Round Robin	19
Least Connection / Weighted Least Connection	19
Destination Hashing	19
Real Server Agent	19
Layer 4 vs Layer 7	20
The Basics	20
Performance	20
Persistence	20
Real Server Changes	20
Transparency	20
Other Considerations	20
Does Your Application Cluster Correctly Handle its own State?	20
Replication Solutions for Shared Data	21
Solutions for Session Data	21
Persistence	21
What do You do if Your Application is not Stateless?	21
Loadbalancer.org Supported Persistence Options	21
What are Your Objectives?	22
What are VIPs and RIPs?	22
What is a Floating IP Address?	23
Loadbalancer.org Terminology	23
Chapter 3 - Topologies & Load Balancing Methods	25
One-Arm and Two-Arm Topologies	25
Supported Load Balancing Methods	25
Layer 4 DR Mode	26
Layer 4 NAT Mode	27
Layer 4 SNAT Mode	30
Layer 7 SNAT Mode	31
Which Load Balancing Method Should I Use?	31

Mode Summary .....	31
Layer 4 DR Mode .....	32
Layer 4 NAT Mode .....	32
Layer 4 SNAT Mode .....	32
Layer 7 SNAT Mode .....	32
Our Recommendation .....	32
Chapter 4 - Appliance Fundamentals .....	33
Hardware Appliance Installation .....	33
Virtual Appliance Installation .....	33
Supported Hypervisors .....	33
Host Requirements .....	34
Virtual Hardware Resource Allocations .....	34
Download & Extract the Appliance .....	34
Hypervisor Deployment .....	35
VMware Host Client .....	35
VMware vSphere Client .....	35
VMware Workstation Player .....	36
VMware Tools / Open VM Tools .....	36
To Verify that Open VM Tools is running .....	36
Microsoft Hyper-V .....	36
Linux Integration Services .....	37
KVM .....	37
Nutanix AHV .....	37
XEN .....	37
Cloud Appliance Installation .....	37
Configuring Initial Network Settings .....	37
Appliance Access & Configuration Methods .....	41
Accessing the WebUI .....	41
Main Menu Options .....	43
Root Access .....	44
Keyboard Layout .....	44
Chapter 5 - Appliance Management .....	45
Network Configuration .....	45
Physical/Virtual Adapters .....	45
Configuring IP Addresses .....	45
Configuring Bonding .....	46
Configuring VLANs .....	48
NIC Offloading .....	49
Configuring MTU Settings .....	49
Configuring Default Gateway & Static Routes .....	50
Management Gateway .....	50
Configuration Example .....	51
Policy Based Routing (PBR) .....	51
Configuring Hostname & DNS Configuration .....	52
Portal Management .....	53
System Date & Time Configuration .....	53
Auto Configuration using NTP Servers .....	53
Manual Configuration .....	54
Appliance Internet Access via Proxy .....	54
SMTP Relay Configuration .....	55
Syslog Server Configuration .....	55

SNMP Configuration .....	57
Installing the License Key .....	58
Running OS Level Commands .....	58
Restoring Manufacturer's Settings .....	59
Using the WebUI .....	59
Using the Console / SSH Session .....	59
Restarting & Reloading Services .....	60
Appliance Restart & Shutdown .....	62
Appliance Software Update .....	62
Determining the Current Software Version .....	62
Creating a Backup .....	63
Online Update .....	63
Auto-Check for Updates .....	63
Configuring Online Update .....	63
Manual Check for Updates .....	63
Offline Update .....	64
Updating a Clustered Pair .....	64
Appliance Security Features .....	65
Security Mode .....	65
Users & Passwords .....	66
Linux 'root' User Account .....	66
WebUI User Accounts .....	66
External Authentication .....	69
Firewall Configuration .....	72
Manual Firewall Configuration .....	72
Firewall Lock-down Wizard .....	73
Conntrack Table Size .....	75
Appliance Security Lockdown Script .....	75
SSH Keys .....	77
Appliance Configuration Files & Locations .....	78
Chapter 6 - Configuring Load Balanced Services .....	79
Introduction .....	79
Layer 4 Services .....	79
The Basics .....	79
Creating Layer 4 Virtual Services .....	79
Defining a New Layer 4 VIP .....	79
Duplicating an Existing Layer 4 VIP .....	81
Modifying a Layer 4 VIP .....	81
Creating Layer 4 Real Servers (RIPs) .....	83
DR Mode Considerations .....	84
The ARP Problem .....	84
Detecting the ARP Problem .....	85
Solving the ARP Problem for Linux .....	85
Solving the ARP Problem for Solaris .....	90
Solving the ARP Problem for Mac OS X/BSD .....	91
Solving the ARP Problem for Windows Servers .....	91
Solving the ARP Problem - Possible Side Effect for Windows 2012 & Later .....	97
Other Windows Settings that May Cause Issues .....	97
Configuring Your Application to Respond to Both the RIP and VIP .....	98
Windows Firewall Settings .....	99
NAT Mode Considerations .....	99



NAT Mode Potential Issues	99
One-Arm (Single Subnet) NAT Mode	100
Firewall Marks	101
Firewall Marks - Auto Configuration	101
Firewall Marks - Manual Configuration	102
Layer 4 - Advanced Configuration	107
Layer 7 Services	109
The Basics	109
Creating Layer 7 Virtual Services (Using the Wizard)	109
Creating Layer 7 Virtual Services	111
Defining a New Layer 7 VIP	111
Duplicating an Existing Layer 7 VIP	112
Modifying a Layer 7 VIP	112
ACLs (aka Content Switching) and URL Rewriting	121
Modifying HTTP Header Fields	130
HTTP Header Field Modification Examples	133
Creating Layer 7 Real Servers (RIPs)	134
Layer 7 - Custom Configurations	135
Configuring Manually Defined Virtual Services	135
HAProxy Error Codes	138
Transparency at Layer 7	138
Enabling Transparency	138
Inserting Headers	139
Using TProxy to modify the Source IP Address	140
Configuration Examples	141
Layer 7 - Advanced Configuration	145
Floating IPs	147
SSL Termination	148
Concepts	148
SSL Termination on the Real Servers (SSL Pass-through)	148
SSL Termination on the Load Balancer (SSL Offloading)	149
Using STunnel or Pound to Terminate SSL	149
Using HAProxy to Terminate SSL	150
Certificates	150
Let's Encrypt	153
Creating a SSL Termination	154
Server Name Indication (SNI)	158
SSL Termination on the Load Balancer with Re-encryption (SSL Bridging)	159
Using STunnel or Pound to Terminate SSL	159
Using HAProxy to Terminate SSL	160
SSL - Advanced Configuration	161
Pound Global Settings	161
STunnel Global Settings	162
HTTP to HTTPS Redirection	162
When Terminating SSL on the Real Servers	162
When Terminating SSL on the Load Balancer	163
Using STunnel or Pound	163
Using HAProxy	164
Server Feedback Agent	165
Windows Agent	165
Linux/Unix Agent	168

Custom HTTP Agent .....	169
Configuring VIPs To Use The Agent .....	169
Global Server Load Balancing (GSLB) .....	169
Key Concepts .....	169
Key features .....	170
GSLB Configuration .....	170
External Health Check Scripts (GSLB) .....	173
GSLB Multi-site Example .....	175
Conceptual Overview .....	175
Appliance Configuration .....	177
DNS Server Configuration .....	180
GSLB Diagnostics .....	181
Configuring the Appliance via CLI, API & Direct Service Calls .....	182
Command Line Interface (CLI) .....	183
LBCLI Command Reference .....	183
Add an ACL rule .....	183
Create a Certificate Signing Request .....	183
Add a floating IP address .....	184
Add a global name .....	184
Add a member .....	184
Add a pool .....	184
Add a topology .....	185
Add a header .....	185
Add a health check .....	186
Add a layer 4 real server .....	186
Add a layer 4 virtual service .....	186
Add a layer 7 real server .....	187
Add a layer 7 virtual service .....	187
Add a PBR rule .....	189
Add a static IP address .....	189
Add a static route .....	189
Add an SSL Termination .....	189
Add a WAF .....	190
Clone a health check .....	190
Clone a virtual service .....	190
Clone a virtual service .....	191
Delete an ACL rule .....	192
Delete a floating IP .....	193
Delete a global name .....	193
Delete a member .....	193
Delete a pool .....	193
Delete a member .....	194
Delete a header rule .....	194
Delete a health check .....	194
Delete a header rule .....	194
Delete a real server .....	195
Delete a virtual service .....	195
Delete a static IP address .....	195
Delete a static route .....	195
Delete an SSL termination .....	196
Delete a WAF .....	196

Disable the API endpoint	196
Disable Floating IP	196
Drain a real server	197
Edit a global name	197
Edit a member	197
Edit a pool	197
Edit a topology	198
Modify a health check	198
Edit a layer 4 real server	199
Edit a layer 4 virtual service	199
Edit a layer 7 real server	200
Edit a layer 7 virtual service	200
Edit a SSL termination	201
Edit a WAF	202
Enable the API endpoint	202
Enable Floating IP	203
Remove and add a floating IP address	203
Remove all ACLs for a virtual service	203
Remove all static IP addresses	203
Remove all static routes	204
Show the properties of the API endpoint	204
Show the default route	204
Show the DNS servers	204
Show the firewall settings	204
Show the graphing settings	204
Show the GSLB reports	204
Show the hostname	205
Show the interface offload settings	205
Show the management gateway	205
Show the interface MTU settings	205
Show the NTP servers	205
Show the HTTP proxy server settings	205
Show the security settings	206
Show the email smarthost	206
Show the email smarthost	206
Show the syslog settings	206
Show the current date and time	206
Show the URL of the updates server	206
Create a cluster	206
Halt a real server	207
Clear all HAProxy stick tables	207
Install a licence	207
List ACL rules	207
List SSL certificates	207
List floating IP addresses	208
List global names	208
List members	208
List pools	208
List topologies	208
List header rules	208
List health checks	209

List PBR rules .....	209
List static IP addresses .....	209
List static routes .....	209
List WAFs .....	209
Lockdown access to the WebUI .....	209
Report the node statuses .....	210
Bring a real server online .....	210
Reboot appliance .....	210
Regenerate default certificate .....	210
Reload WebUI .....	210
Reload GSLB .....	210
Reload HAProxy .....	211
Reload Heartbeat .....	211
Reload LDirectord .....	211
Reload STunnel .....	211
Reload Syslog .....	211
Reload WAF .....	211
Break a cluster .....	211
Restart Autoscaled .....	211
Restart AZHA .....	211
Restart Collectd .....	212
Restart Firewall .....	212
Restart HAProxy .....	212
Restart Heartbeat .....	212
Restart LDirectord .....	212
Restart Pound .....	212
Restart SNMP .....	212
Restart STunnel .....	212
Restart Syslog .....	213
Restart WAF .....	213
Restore factory defaults .....	213
Set default route .....	213
Set DNS servers .....	213
Set firewall properties .....	213
Set graphing properties .....	214
Set hostname .....	214
Set interface offload .....	214
Set the management gateway .....	214
Set the MTU .....	215
Set NTP server .....	215
Set HTTP proxy .....	215
Set security settings .....	215
Set the email smarthost .....	216
Set SNMP properties .....	216
Set syslog properties .....	216
Set time and date .....	217
Enable or disable update check .....	217
Show configuration .....	217
Show configuration for LDirectord .....	217
Show configuration for Layer 4 virtual services .....	217
Show configuration for HAProxy .....	217

Show configuration for Layer 7 virtual services. . . . .	218
Shutdown appliance. . . . .	218
Show GSLB status. . . . .	218
Create a technical support download. . . . .	218
Get uptime. . . . .	218
Running Ibccli from a remote Linux Host . . . . .	218
Running Ibccli from a remote Windows Host. . . . .	218
Application Programming Interface (API) . . . . .	219
Enabling the API . . . . .	219
API Endpoint . . . . .	219
Creating the JSON Request . . . . .	219
JSON Syntax Validation . . . . .	220
Sending API Requests. . . . .	220
Using ipvsadm to configure Layer 4 Services . . . . .	222
Using Linux socket commands to configure Layer 7 Services. . . . .	223
Chapter 7 - Web Application Firewall (WAF). . . . .	225
Introduction . . . . .	225
Implementation Concept . . . . .	225
Creating a New WAF Gateway . . . . .	227
Step 1 - Create the Layer 7 VIP . . . . .	227
Step 2 - Define the associated Real Servers (RIPs) . . . . .	227
Step 3 - Define the WAF Gateway . . . . .	228
Step 4 - Reload Services to Apply the new Settings . . . . .	228
Step 5 - View Configured Services. . . . .	228
Fail Open and Fail Closed Designs . . . . .	228
Configuring a WAF Gateway for Fail Open Operation (Default) . . . . .	229
Configuring a WAF Gateway for Fail Closed Operation . . . . .	230
WAF Gateway Settings . . . . .	231
Disable Web Application Firewall . . . . .	231
Ruleset . . . . .	231
Paranoia Level. . . . .	231
Rule Engine Traffic Blocking . . . . .	232
Process Request Data. . . . .	232
Process Response Data . . . . .	232
Inbound Anomaly Score . . . . .	232
Outbound Anomaly Score . . . . .	232
Audit Mode . . . . .	233
WAF Proxy Timeout. . . . .	233
Enable Cache Acceleration . . . . .	233
Double Login Enable . . . . .	233
WAF - Advanced Configuration. . . . .	234
PCRE Match Limit. . . . .	234
PCRE Match Limit Recursion . . . . .	234
Working With the Core Rule Set . . . . .	235
What is the Core Rule Set?. . . . .	235
Core Rule Set Map. . . . .	235
Anomaly Scoring . . . . .	235
Overview of Anomaly Scoring . . . . .	235
How Anomaly Scoring Mode Works. . . . .	235
Summary of Anomaly Scoring Mode . . . . .	236
Anomaly Score Thresholds . . . . .	236

Severity Levels .....	237
Paranoia Levels .....	237
Introduction to Paranoia Levels .....	237
Description of the Four Paranoia Levels .....	238
Choosing an Appropriate Paranoia Level .....	239
Setting the Paranoia Level .....	239
How Paranoia Levels Relate to Anomaly Scoring .....	239
False Positives and Tuning .....	240
What are False Positives? .....	240
Example False Positive .....	240
Why are False Positives a Problem? .....	241
Tuning Away False Positives .....	241
Directly Modifying CRS Rules .....	241
Rule Exclusions .....	242
Adding Custom WAF Configuration .....	247
Per-Transaction Resource Limits .....	248
Default Limits .....	248
Example 1 .....	249
Example 2 .....	249
Handling Large HTTP Requests .....	250
Large File Uploads .....	250
Requests Containing Large Amounts of Non-File Upload Data .....	251
Requests with Massive Header Fields .....	251
WAF Gateway Error Logs .....	252
Logging Mechanism Overview .....	252
Viewing the Error Logs .....	252
Default View .....	252
Simple View .....	252
Breakdown View .....	253
Fixes View .....	253
Breakdown of a Log Entry .....	254
Chapter 8 - Real Server Health Monitoring & Control .....	257
Configuring Health Checks .....	257
Health Checks for Layer 4 Services .....	257
Health Checks for Layer 7 Services .....	262
External Health Check Scripts .....	265
Default Scripts .....	265
Adding Additional Health Check Scripts .....	266
Using Script Templates .....	266
Uploading External Files .....	267
Testing External Health Check Scripts at the Command Line .....	268
Simulating Health Check Failures .....	269
Disabling Health Checks .....	269
Fallback Server .....	269
Local Fallback Server .....	270
Using a Separate Dedicated Server .....	270
Using a Layer 7 VIP .....	270
Configuring A Real Server as the Fallback Server .....	271
Configuring Primary/Secondary Real Servers .....	271
Configuring Email Alerts for Virtual Services .....	271
Layer 4 .....	271

Global Layer 4 Email Settings .....	271
VIP Level Settings .....	272
Configuring a Smart Host (SMTP relay) .....	272
Layer 7 .....	272
Real Server Monitoring & Control using the System Overview .....	273
Real Server Monitoring .....	273
Real Server Control .....	274
Ordering of VIPs .....	275
Sort by Column .....	275
Drag & Drop .....	276
Real Server Monitoring & Control using the HAProxy Statistics Report .....	276
Real Server Monitoring .....	276
Real Server Control .....	277
Chapter 9 - Appliance Clustering for HA .....	279
Introduction .....	279
Clustered Pair Concepts .....	279
Primary/Secondary Operation .....	279
Pair Communication .....	279
Heartbeat .....	279
Primary Secondary Replication .....	279
Settings that are NOT Replicated to the Secondary Appliance .....	280
Manually Forcing Appliance Synchronisation .....	280
To Create an HA Pair (Add a Secondary) .....	280
To Break an HA Pair (Remove a Secondary) .....	282
Promoting a Secondary to Primary .....	283
Configuring Heartbeat .....	284
Communication Method .....	285
Peer Failure Detection .....	286
Routing Failure Detection .....	286
Email Alerts .....	286
Fail-back Settings .....	286
Configuring a Smart Host (SMTP relay) .....	286
Connection State & Persistence Table Replication .....	287
Layer 4 VIPs .....	287
Layer 7 VIPs .....	288
Clustered Pair Diagnostics .....	288
Heartbeat State Diagnostics .....	289
Split Brain Scenarios .....	290
Forcing Primary/Secondary Failover & Failback .....	290
Testing & Verifying Primary/Secondary Replication & Failover .....	291
Heartbeat Log .....	293
Late Heartbeats .....	294
Chapter 10 - Configuration Examples .....	295
Introduction .....	295
Initial Network Settings .....	295
Example 1 - One-Arm DR Mode (Single Appliance) .....	295
Configuration Overview .....	295
Network Settings .....	295
Virtual Service (VIP) .....	296
Real Servers (RIPs) .....	297
Physical Real Server Changes - Solve the ARP Problem .....	297



Basic Testing & Verification .....	297
Example 2 - One-Arm Layer 4 SNAT Mode (Single Appliance) .....	298
Configuration Overview .....	298
Network Settings .....	298
Virtual Service (VIP) .....	299
Real Servers (RIPs) .....	299
Basic Testing & Verification .....	300
Example 3 - Two-Arm NAT Mode (Clustered Pair) .....	300
Configuration Overview .....	301
Primary Unit - Network Settings .....	301
Secondary Unit - Network Settings .....	302
Virtual Service (VIP) .....	302
Real Servers (RIPs) .....	303
Physical Real Server Changes - Set the Default Gateway .....	304
Create the HA Clustered Pair .....	304
Checking the Status .....	305
Verify Heartbeat Settings .....	305
Verify the Secondary Configuration .....	306
Basic Testing & Verification .....	306
Example 4 - One-Arm SNAT Mode & SSL Termination (Single Appliance) .....	306
Configuration Overview .....	306
Network Settings .....	307
Virtual Service (VIP) .....	307
Real Servers (RIPs) .....	308
SSL Termination .....	309
Basic Testing & Verification .....	309
Example 5 - Load Balancing FTP .....	310
Layer 4 .....	310
Layer 7 .....	310
Active Mode .....	310
Passive Mode .....	312
Chapter 11 - Diagnostics & Troubleshooting .....	316
The Appliance .....	316
Resource Utilisation .....	316
Administration Log .....	316
Apache Logs .....	316
Apache Access Log .....	316
Apache Error Log .....	316
Layer 4 & Layer 7 Virtual Services .....	316
Checking Service State using the System Overview .....	316
VIP(s) Fail to appear in the System Overview .....	317
VIPs & RIPs are Green but Users still Can't Connect .....	318
Layer 7 VIPs .....	318
Layer 4 VIPs .....	319
Diagnosing Real Server Issues .....	319
Requests are Not being Load Balanced as Expected .....	321
Log File - Layer 4 Virtual Services .....	321
Interpreting the log file .....	321
Log File - layer 7 Virtual Services .....	322
Interpreting the log file .....	322
SSL Termination - Pound .....	323

SSL Termination - STunnel	323
WAF	324
GSLB	324
Heartbeat	324
Chapter 12 - Monitoring & Reporting	325
Reports	325
Layer 4 Status	325
Layer 4 Traffic Rate	325
Layer 4 traffic Counters	326
Layer 4 Current Connections	327
Layer 4 Current Connections (Resolve Hostnames)	328
Layer 7 Status	328
Layer 7 Stick Table	328
GSLB Reports	329
Graphs	329
Resource Utilisation Graphs	329
Layer 4 and Layer 7 Virtual Services/Real Servers	330
Graphing Options	334
Layer 7 (HAProxy) Prometheus Exporter	335
Grafana	335
SNMP Reporting	335
MIB Files	335
SNMP for Layer 4 Services	336
Monitoring Layer 4 VIPs & RIPs using SNMP	336
SNMP for Layer 7 Services	337
Monitoring Layer 7 VIPs & RIPs using SNMP	338
SNMPv3	339
Chapter 13 - Useful Tools & Utilities	340
Useful Diagnostics Tools	340
Netstat	340
Telnet	340
Tcpdump	341
Ethtool	341
NMAP	341
Wireshark	342
Windows Specific Tools	342
Microsoft Network Monitor	342
WinSCP	342
PuTTY	342
Chapter 14 - Backup & Restore and Disaster Recovery	343
Backup & Restore	343
Backup	343
Local XML Backup	343
Restore Options	344
Restore	344
Local XML Restore	345
Restore System Defaults	345
Template Deployment	346
Disaster Recovery	346
Being Prepared - Creating Backups	346
Firmware Recovery using a USB Memory Stick	347

Disaster Recovery After Node (Primary or Secondary) Failure .....	349
Chapter 15 - Technical Support .....	353
Introduction .....	353
WebUI Support Options .....	353
Contact Us .....	353
Technical Support Download .....	353
Useful Links .....	354
Remote Support .....	354
Live Chat .....	354
Appendix .....	355
Front & Rear Panel Layouts .....	355
Enterprise 1G .....	355
Enterprise 10G/25G/40G/50G .....	355
Enterprise 100G .....	355
IPMI (Remote Management) Configuration .....	355
iDRAC (Remote Management) Configuration .....	359
iDRAC IP Address .....	359
Logging in to iDRAC .....	359
iDRAC Password Reset .....	359
iDRAC Licensing .....	359
More Information .....	359
Appliance IPv4 Address Format (CIDR notation) .....	359

# Chapter 1 - Introduction

## About this Manual

This document covers all required administration information for v8.8.x Loadbalancer.org appliances.

## About the Appliance

The core software is based on LBOS-7 which is a customized Linux build maintained by Loadbalancer.org, LVS, Ldirectord, Linux-HA, HAProxy & STunnel. Full root access is provided which enables complete control of all settings.

The appliance is available in the following formats: hardware, virtual (VMware, HyperV, KVM, Nutanix & XEN) and cloud based (Amazon, Azure & GCP).

Appliances can be deployed as single units or as a clustered pair.

### Note

Loadbalancer.org always recommend that clustered pairs should be used where possible for high availability and resilience, this avoids introducing a single point of failure to your network. For more information on configuring an HA pair please refer to [Chapter 9 - Appliance Clustering for HA](#).

## Latest Version

The latest version of the appliance (v8.8.1) includes the following branding changes, new features, improvements, bug fixes and security updates:

### Branding

- As we move into 2023, Loadbalancer.org will be rolling out a new look & feel for our brand.
- We believe the new logo and color scheme is more reflective of our true company personality.
- The new branding will be introduced in this update of the appliance's WebUI.
- The complete transition to the new brand-look will happen gradually over the next few months.

### New Features

- Rolling online updates. Starting with this release you will now be able to update to the latest version without having to perform incremental updates.
- Added the facility to manage Loadbalancer Portal configuration via the WebUI.
- Added "First" Scheduler and "Last Successful" persistence method.
- Added new ACL type "Reported SNI" to act on the value of the SNI field of an unterminated TLS connection.
- Added Template Importer/Exporter to help with deployment automation.
- Virtual servers can now be bound to traffic on ANY interface by using 0.0.0.0.

### Improvements

- Improved gateway address validation to prevent entering an invalid ip.
- Open-vm tools for VMware based deployments on new installations.

- Free-type ACLs may now be multiline.
- Improved HAProxy config error reporting to add more detail.
- Improvements to LBCLI including input and output validation.
- The STunnel service is now supervised to improve high availability.VIPs
- Updated OpenSSH to 9.0p1.
- HAProxy Stats page will now auto-refresh by default.

## Bug Fixes

- Fixed having to specify a Floating IP and/or a Gateway when creating a new PBR rule.
- Fixed issue where L4 services would restart when autoscaling in AWS.
- Fixed setup wizard to validate passwords correctly.
- Backend-only virtual services no longer cause a conflict with layer 7 logging.

## Security Updates

None in this release.

## Appliance Configuration Overview

Initial network configuration is carried out at the console using the Network Setup Wizard. Once the wizard has been run, load balanced services can be configured using the WebUI; either using the Setup Wizard (for Layer 7 services) or by manually defining the Virtual Services (VIPs) and associated Real Servers (RIPs).

By default, the WebUI is accessible on HTTPS port **9443**, this can be changed if required. For more information please refer to the "Appliance Security" section below.

We always recommend that where possible two appliances are deployed as a clustered pair for high availability and resilience, this avoids introducing a single point of failure to your network.

We recommend that the Primary appliance should be fully configured first, then the Secondary appliance should be added to create an HA pair. Once the HA pair is configured, load balanced services will be automatically synchronized from the Primary to Secondary appliance. Load balanced services should then be configured & modified on the Primary appliance and the Secondary will be automatically kept in sync. For more information on configuring an HA pair, please refer to [Chapter 9 - Appliance Clustering for HA](#).

### Note

For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

## Appliance Security

### Note

For full details of all security related features, please refer to [Appliance Security Features](#).

## Security Mode

To control how the appliance is accessed and which features are enabled, 3 security modes are provided:

- **Secure - (default)** - in this mode:
  - the WebUI is accessible on HTTPS port **9443**. If you attempt to access the WebUI on HTTP port **9080** you will be redirected to HTTPS port **9443**
  - access to the *Local Configuration > Execute shell command* menu option is disabled
  - the ability to edit the firewall script & the firewall lockdown wizard script is disabled
  - 'root' user console & SSH password access are disabled
- **Custom** - in this mode, the security options can be configured to suit your requirements
- **Secure - Permanent** - this mode is the same as **Secure** but the change is *irreversible*

**Important** | Only set the security mode to **Secure - Permanent** if you are 100% sure this is what you want!

To configure the Security Mode:

1. Using the WebUI, navigate to: *Local Configuration > Security*.
2. Select the required *Appliance Security Mode* - if **Custom** is selected, configure the additional options according to your requirements.
3. Configure the *HTTPS Port for Web User Interface*, *Web Interface SSL Certificate* and *Ciphers to use* according to your requirements.
4. Click **Update**.

## Passwords

The password for the 'root' user Linux account and the 'loadbalancer' WebUI user account are set during the Network Setup Wizard. These can be changed at any time.

**Note** | The passwords for the cloud products are either set to a default value or are configured during instance deployment. Also, for Enterprise AWS and Enterprise Azure it's not possible to directly log in as root. For more details, please refer to the relevant [Quick Start Configuration Guide](#).

### The 'loadbalancer' WebUI account

This can be changed using the WebUI menu option: *Maintenance > Passwords*.

### The 'root' Linux account

As explained in [Security Mode](#) above, 'root' user console & SSH password access are disabled by default. Once enabled, the 'root' password can be changed at the console, or via an SSH session using the following command:

```
# passwd
```

## Ports Used by the Appliance

By default, the appliance uses the following TCP & UDP ports:

Protocol	Port	Purpose
TCP	22	SSH
TCP & UDP	53	DNS

Protocol	Port	Purpose
TCP & UDP	123	NTP
TCP & UDP	161	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback page
TCP	9443	WebUI - HTTPS

## Additional Information

### Deployment Guides

Comprehensive deployment guides are available that focus on load balancing specific applications. They cover the configuration of the load balancer and also any application specific configuration changes that are required to enable load balancing. All guides are available on our website at the following URL: <https://www.loadbalancer.org/support/deployment-guides/>.

### Quick Start & Configuration Guides

The following related documentation may also be useful:

- [Quick Start Guide - Hardware & Virtual](#)
- [Quick Start Configuration Guide - Amazon AWS](#)
- [Quick Start Configuration Guide - Microsoft Azure](#)
- [Quick Start Configuration Guide - Google Cloud Platform](#)

### Contacting Support

If you have any questions regarding the appliance or need assistance with load balancing your application, please don't hesitate to contact [support@loadbalancer.org](mailto:support@loadbalancer.org).



# Chapter 2 - Load Balancing Concepts

## Load Balancing - the Basics

Loadbalancer.org appliances enable two or more servers to be combined into a cluster. This enables inbound requests to be distributed across multiple servers which provides improved performance, reliability and resilience. Appliances can also be deployed as a clustered pair (our recommended solution) which creates a highly-available configuration.

## Supported Protocols

Loadbalancer.org appliances support virtually any TCP or UDP based protocol including HTTP, HTTPS, FTP, SMTP, RDP, SIP, IMAP, POP, DNS etc. etc.

## Layer 4 & Layer 7

Load balancing at layer 4 and layer 7 is supported. LVS (*Linux Virtual Server*) is utilized at layer 4 whilst HAProxy is used at layer 7.

## Load Balancing Algorithms

The Loadbalancer.org appliance supports several different load balancing algorithms. Each one has its advantages and disadvantages and it depends on the specific application which is the most appropriate to use. Usually the default method *Weighted Least Connection* is a good solution which works well in most situations. The following sections summarize each method supported.

### Round Robin / Weighted Round Robin

With this method, incoming requests are distributed to Real Servers in a sequential manner relative to each Real Server's weight. Servers with a higher weight receive more requests. A server with a weight of 200 will receive 4 times the number of requests than a server with a weight of 50. Weightings are relative, so it makes no difference if Real Server #1 and #2 have weightings of 50 and 100 respectively or 5 and 10.

### Least Connection / Weighted Least Connection

With this method, incoming requests are distributed to Real Servers with the fewest connections relative to each Real Server's weight. Servers with a higher weight receive more requests. A server with a weight of 200 will receive 4 times the number of requests than a server with a weight of 50. Again, weightings are relative, so it makes no difference if Real Server #1 and #2 have weightings of 50 and 100 respectively or 5 and 10. *This is the default method for new VIPs.*

### Destination Hashing

With the method, requests are distributed to Real Servers by looking up the destination IP in a static hash table. This algorithm is designed for use with web proxies and is supported with Layer 4 DR mode Virtual Services only. For more information on this method please refer to [Modifying a Layer 4 VIP](#).

### Real Server Agent

To compliment the methods above, Loadbalancer.org appliances also support Real Server (i.e backend server) agents. This permits the load balancing algorithm to be dynamically modified based on each Real Server's running characteristics. For example, one Real Server could have a run-away process that is consuming excessive CPU resources or RAM. Without the agent, the load balancer has no way of knowing this and would continue to send requests to the overloaded server based on the algorithm selected. With the agent installed on the Real Server, feedback is provided to the load balancer and the algorithm is then adjusted to reduce requests that are sent to that server. For more information please refer to [Server feedback Agent](#).

## Layer 4 vs Layer 7

A fundamental choice when setting up the load balancer is whether to configure the services at layer 4 or layer 7.

### The Basics

At layer 4 the primary protocols used are TCP and UDP. These protocols are not aware of upper level protocols such as FTP, HTTP, HTTPS, DNS, RDP etc. Therefore the load balancer can only make load balancing decisions based on details available at layers 4 and below such as port numbers and IP addresses. At layer 7, the load balancer has more information to make load balancing related decisions since more information about upper level protocols is available.

Layer 7 load balancing uses a proxy at the application layer (HAProxy). Requests are terminated on the load balancer, and the proxy generates a new request which is passed to the chosen Real Server.

### Performance

Due to the increased amount of information at layer 7 and the fact that a proxy is in use, performance is not as fast as at layer 4. If raw throughput is a primary concern, then layer 4 is probably the better choice.

### Persistence

Persistence (aka affinity or sticky connections) is the ability to ensure that a specific client connects back to the same server within a specific time limit. It is normally required when the session state is stored locally on the Real Server rather than in a separate database. At layer 4, Source IP persistence is the only option. At layer 7, additional methods are available such as HTTP cookie persistence where the load balancer sets a cookie to identify the session and Microsoft Connection Broker where the load balancer is able to utilize the redirection token for reconnecting users to existing sessions.

### Real Server Changes

For layer 4 DR mode, the **ARP Problem** (more information is available [here](#)) has to be solved, for layer 4 NAT mode, the default gateway on the Real Servers must be the load balancer. For layer 4 SNAT mode and layer 7 SNAT mode the Real Servers do not need to be changed in any way.

### Transparency

Transparency refers to the ability to see the originating IP address of the client. For layer 4 DR mode and NAT mode, connections are transparent. For layer 4 SNAT mode and layer 7 SNAT mode, the IP address of the load balancer is recorded as the source address (for Layer 7 SNAT mode, this can also be set to a user defined address). For layer 7 SNAT mode, additional configuration steps can be taken to force the client IP to be logged. Options include using TProxy to re-write the source address or by enabling support for X-Forwarded-For or Proxy Protocol headers. For more information, please refer to [Transparency at Layer 7](#).

## Other Considerations

### Does Your Application Cluster Correctly Handle its own State?

#### Note

Load balancers work most effectively if the application servers are completely stateless. This means that if an application server (i.e. Real Server) fails and is automatically taken out of the cluster, then all the current user sessions will be transferred to other servers in the cluster without the users needing to re login to the application again. **If your application doesn't have a persistent data store then you can't have seamless fail over for your backend servers.**

Do your web servers store persistent information on local drives?

- Images (jpeg, png, gif etc.)
- Files (html, php, asp etc.)

If so, these files either need to be on shared storage, or they need to be replicated to all of the nodes in the cluster.

### Replication Solutions for Shared Data

On Linux systems you can use the Rsync command to replicate files. On Windows systems you also have Rsync, although you may prefer the Robocopy command. Usually, you'll upload your content to one primary server and then replicate it to the other servers in the cluster.

### Solutions for Session Data

Standard ASP and PHP session data is stored locally by default, leaving your session data in a local store will prevent you from implementing seamless application server fail-over in your cluster. If an application server fails, all of the local session data will be lost and your user will need to re-log in and possibly lose shopping baskets etc.

This problem is easily resolvable by implementing a shared persistent data store for the cluster. This is usually either done with a shared backend database or a shared memory solution.

### Persistence

Persistence is required by many web applications. Once a user has interacted with a particular server all subsequent requests are sent to the same server. It is normally required when the session state is stored locally to the web server as opposed to a shared database.

### What do You do if Your Application is not Stateless?

Some applications require state to be maintained such as:

- Terminal Services/Remote Desktop Services
- SSH
- FTP (upload)
- SMTP (incoming)

You may also find that you are unable to modify your HTTP/HTTPS based application to handle shared session data.

For these cases, you can use persistence based on source IP address. You lose the ability to have transparent fail-over, but you do still get increased capacity and manageability. This persistence problem occurs with all load balancers and all vendors use standard methods and technologies to mitigate the issue.

### Loadbalancer.org Supported Persistence Options

For layer 4, source IP address persistence is available and is enabled by default for all new Virtual Services.

For layer 7, the following persistence options are supported:

- Source IP Address
- HTTP Cookie (the default for new Virtual Services)
- Application Cookie

- SSL Session ID
- MS Session Broker
- RDP Client Cookie
- HTTP Cookie with Fallback to Source IP
- X-Forwarded-For with Fallback to Source IP
- Stick On Fallback
- Last Successful

Cookies can offer more even traffic distribution and also handle any clients where the source IP address may change during the session (e.g. mega proxies).

SSL session ID based persistence is useful in certain circumstances, although due to the way some browsers operate - notably older versions of Internet Explorer, the session ID can be renegotiated frequently (every few seconds) which effectively breaks the persistence.

**Note** | For details on all layer 7 persistence options, please refer to [Modifying a Layer 7 VIP](#).

**Note** | It's also possible to define other custom persistence types if required using the custom configuration option available for layer 7 Virtual Services. For more information, please refer to [Layer 7 - Custom Configurations](#).

## What are Your Objectives?

It's important to have a clear focus on your objectives and the required outcome for the successful implementation of your load balancing solution. If the objective is clear and measurable, you know when you have achieved your goal.

Load balancers have a number of flexible features and benefits for your technical infrastructure and applications.

The primary question to consider is: **Are you looking for increased performance, reliability, ease of maintenance or all three?**

Performance	A load balancer can increase performance by allowing you to utilize several commodity servers to handle the workload of one application.
Reliability	Running an application on one server gives you a single point of failure. Utilizing a load balancer to present multiple servers improves application availability but moves the point of failure to the load balancer. We always advise that you deploy load balancers as clustered pairs to remove this single point of failure. For more information, please refer to <a href="#">Chapter 9 - Appliance Clustering for HA</a> .
Maintenance	Using the appliance, you can easily bring servers on and off line to perform maintenance tasks, without disrupting your users.

**Note** | In order to achieve all three objectives, your application must handle persistence correctly. For more information, please refer to [Does Your Application Cluster Correctly Handle its own State?](#).

## What are VIPs and RIPs?

Load balancer vendors typically use the term Virtual IP address (VIP) to describe the address that the load

balanced cluster is accessed from. It's important to understand that 'VIP' refers to both the physical IP address and also to the logical load balancer configuration. Likewise, 'RIP' refers to both the Real Server's physical IP address and its representation in the logical load balancer configuration.

#### Note

It's not possible to configure a VIP on the same IP address as any of the network interfaces. This ensures that services are able to 'float' (move) between Primary and Secondary appliances when using an HA Pair.

### What is a Floating IP Address?

A floating IP address (FIP) is automatically created whenever a VIP is configured. The FIP address is the same as the VIP address. Since the FIP must be able to move between the Primary and Secondary appliance, it's not possible to configure a VIP/FIP on the same IP address as an interface as mentioned in the note above. FIPs can also be manually defined to provide a 'floating default gateway' for layer 4 NAT mode configurations. This allows the default gateway for the NAT mode Real Servers to be brought up on the Secondary should the Primary fail.

### Loadbalancer.org Terminology

Load Balancer	An IP based traffic manager for server clusters.
Primary	The normally active appliance in a HA Clustered Pair of load balancers.
Secondary	The normally passive appliance in a HA Clustered Pair of load balancers.
Virtual Service	The main building block used to define load balanced services. It defines the IP address clients connect to, which Real Servers are load balanced and other settings such as health check options, persistence options and timeout settings.
Real Server	The actual backend server being load balanced. Multiple Real Servers are associated with a Virtual Service.
VIP	Virtual IP address - the IP address of the load balanced cluster of RIPs, the address presented to connecting clients. Also used as an acronym for Virtual Service.
RIP	The Real IP address - the IP address of a backend server in the cluster. Also used as an acronym for Real Server.
Floating IP	The Floating IP Address is automatically created whenever a Virtual Service is configured, the floating IP address is the same as the VIP address. It enables services to be moved between the Primary and Secondary appliance.
WebUI / WUI	Web User Interface. Used to configure and manage the appliance.
Layer 4	Part of the seven layer OSI model. Also a descriptive term for load balancing methods that routes packets based on TCP/IP header information.
Layer 7	Part of the seven layer OSI model. Also a descriptive term for a proxy based load balancing method that distributes packets based on the entire TCP/IP header and also the payload information at the application layer.
DR Mode	Direct Routing (aka DSR/Direct Server Return) is a standard layer 4 load balancing technique that distributes packets by altering only the destination MAC address of the packet.
NAT Mode	Network Address Translation is a standard layer 4 load balancing technique that changes the destination of packets to and from the VIP (external subnet to internal cluster subnet).
Layer 4 SNAT Mode	Source Network Address Translation - similar to NAT mode but also modifies the source address of all outgoing traffic to be the load balancer.

Layer 7 SNAT Mode	Source Network Address Translation - the load balancer acts as a proxy for all incoming & outgoing traffic.
SSL Termination	The SSL certificate is installed on the load balancer in order to decrypt HTTPS traffic on behalf of the cluster.
MASQUERADE	Descriptive term for standard firewall technique where internal servers are represented as an external public IP address. Sometimes referred to as a combination of SNAT & DNAT rules.
One-Arm	The load balancer has one physical network card connected to one subnet.
Two-Arm	The load balancer has two interfaces connected to two subnets - this can be achieved by using two network adapters, or by creating VLANs on a single adapter.
GW	The Default Gateway for a backend server in the cluster.
Eth0	Usually the internal Ethernet interface, although this is optional. Also known as Gb0 on the Enterprise 10G, 50G and 100G.
Eth1	Usually the external Ethernet interface, although this is optional. Also known as Gb1 on the Enterprise 10G, 50G and 100G.
Eth2	Third Ethernet interface.
Eth3	Fourth Ethernet interface.
Eth4	Fifth Ethernet interface (Enterprise 100G only, also depends on choice of adapter cards).
Eth5	Sixth Ethernet interface (Enterprise 100G only, also depends on choice of adapter cards).


## Chapter 3 - Topologies & Load Balancing Methods

### One-Arm and Two-Arm Topologies

The number of 'arms' is a descriptive term for how many interfaces are used to connect a device to a network. It's common for a load balancer that uses a routing method (NAT) to have a two-arm configuration although one-arm is also supported. Proxy based load balancers (SNAT) commonly use a one-arm configuration although two-arm is also supported.

#### One-Arm


The VIP and the load balanced servers are located in a single subnet. The load balancer requires a single network interface adapter - eth0 in the diagram below.



#### Two-Arm

Here, 2 subnets are used. The VIP is located in one subnet and the load balanced Real Servers are located in the other. The load balancer requires 2 interfaces, one in each subnet as shown in the diagram below.

**Note** | This can be achieved by using two network adapters, or by creating VLANs on a single adapter.



**Note** | Typically *eth0* is used as the internal interface and *eth1* is used as the external interface. This is not a requirement - each interface can be used for any purpose.

### Supported Load Balancing Methods

The Loadbalancer.org appliance is one of the most flexible load balancers available. The design allows different load balancing modules to utilize the core high availability framework of the appliance. Multiple load balancing methods can be used at the same time or in combination with each other. The table below describes the methods supported by the appliance.



Layer	Method	Comments	Topology	Note
Layer 4	DR (Direct Routing)	Ultra-fast local server based load balancing <ul style="list-style-type: none"> <li>Requires the <b>ARP Problem</b> to be solved on each Real Server - for more details please refer to <b>DR Mode Considerations</b></li> </ul>	One-Arm (*)	1
Layer 4	NAT (Network Address Translation)	Fast Layer 4 load balancing <ul style="list-style-type: none"> <li>The appliance must be the default gateway for the Real Servers</li> </ul>	One or Two-Arm	1
Layer 4	TUN	Similar to DR but works across IP encapsulated tunnels	One-Arm	2
Layer 4	SNAT (Source Network Address Translation)	Fast layer 4 SNAT supporting both TCP & UDP <ul style="list-style-type: none"> <li>Very simple to implement</li> <li>Requires no Real Server configuration changes</li> </ul>	One or Two-Arm	3
Layer 7	SSL Termination (STunnel & Pound)	Usually required in order to process cookie persistence in HTTPS streams on the load balancer <ul style="list-style-type: none"> <li>SSL Termination is processor intensive</li> </ul>	One or Two-Arm	4
Layer 7	SNAT (Source Network Address Translation using HAProxy)	Layer 7 allows greater flexibility including full SNAT and remote server load balancing, cookie insertion and URL switching <ul style="list-style-type: none"> <li>Very simple to implement</li> <li>Requires no Real Server configuration changes</li> <li>Not as fast as Layer 4 methods</li> </ul>	One or Two-Arm	4

(\*) DR mode can also be used in a multi-homed configuration where real servers are located in different subnets. In this case, the load balancer must have an interface in the same subnet to enable layer 2 connectivity which is required for DR mode to operate.


## Notes

1. Recommended for high performance fully transparent and scalable solutions.
2. Only required for Direct Routing implementation across routed networks (rarely used).
3. Recommended when you want to load balance both TCP and UDP but you're unable to use DR mode or NAT mode due to network topology or Real Server related reasons.
4. Recommended if HTTP cookie persistence is required, also used for several Microsoft applications such as Exchange, SharePoint & Remote Desktop Services and for overall deployment simplicity since Real Servers can be on any accessible subnet and no Real Server changes are required.

## Layer 4 DR Mode

One-arm direct routing (DR) mode is a very high performance solution that requires little change to your existing infrastructure.


**Note** | Kemp, Brocade, Barracuda & A10 Networks call this *Direct Server Return* and F5 call it *N-Path*.



- DR mode works by changing the destination MAC address of the incoming packet to match the selected Real Server on the fly which is very fast.
- When the packet reaches the Real Server it expects the Real Server to own the Virtual Services IP address (VIP). This means that you need to ensure that the Real Server (and the load balanced application) respond to both the Real Server's own IP address and the VIP.
- The Real Servers should not respond to ARP requests for the VIP. Only the load balancer should do this. Configuring the Real Servers in this way is referred to as ***Solving the ARP Problem***. For more information please refer to **DR Mode Considerations**.
- On average, DR mode is 8 times quicker than NAT for HTTP, 50 times quicker for Terminal Services and much, much faster for streaming media or FTP.
- The load balancer must have an Interface in the same subnet as the Real Servers to ensure layer 2 connectivity required for DR mode to work.
- The VIP can be brought up on the same subnet as the Real Servers, or on a different subnet provided that the load balancer has an interface in that subnet.
- Port translation is not possible with DR mode, e.g. VIP:80 → RIP:8080 is not supported.
- DR mode is transparent, i.e. the Real Server will see the source IP address of the client.

## Layer 4 NAT Mode

Layer 4 NAT mode is a high performance solution, although not as fast as layer 4 DR mode. This is because real server responses must flow back to the client via the load balancer rather than directly as with DR mode.



- The load balancer translates all requests from the Virtual Service to the Real Servers.
- NAT mode can be deployed in the following ways:
  - **Two-arm (using 2 Interfaces)** (as shown above) - Here, 2 subnets are used. The VIP is located in one subnet and the load balanced Real Servers are located in the other. The load balancer requires 2 interfaces, one in each subnet.

#### Note


This can be achieved by using two network adapters, or by creating VLANs on a single adapter.

- Normally eth0 is used for the internal network and eth1 is used for the external network although this is optional. Any interface can be used for any purpose.
- If the Real Servers require Internet access, *Autonat* should be enabled using the WebUI menu option: *Cluster Configuration > Layer 4 - Advanced Configuration*, the external interface should be selected.
- The default gateway on the Real Servers must be set to be an IP address on the load balancer.

#### Note

For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can 'float' (move) between Primary and Secondary appliances.

- Clients can be located in the same subnet as the VIP or any remote subnet provided they can route to the VIP.
- **One-arm (using 1 Interface)** - Here, the VIP is brought up in the same subnet as the Real Servers.



- To support remote clients, the default gateway on the Real Servers must be an IP address on the load balancer and routing on the load balancer must be configured so that return traffic is routed back via the router.

#### Note

For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can 'float' (move) between Primary and Secondary appliances.

- To support local clients, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer. For more information please refer to [One-Arm \(Single Subnet\) NAT Mode](#).
- If you want Real Servers to be accessible on their own IP address for non-load balanced services, e.g. RDP, you will need to setup individual SNAT and DNAT firewall script rules for each Real Server or add additional VIPs for this.
- Port translation is possible with Layer 4 NAT mode, e.g. VIP:80 → RIP:8080 is supported.
- NAT mode is transparent, i.e. the Real Servers will see the source IP address of the client.

## NAT Mode Packet re-Writing

In NAT mode, the inbound destination IP address is changed by the load balancer from the Virtual Service IP address (VIP) to the Real Server. For outbound replies the load balancer changes the source IP address of the Real Server to be the Virtual Services IP address.

The following table shows an example NAT mode setup:

Protocol	VIP	Port	RIP	Port
TCP	10.0.0.20	80	192.168.1.50	80

In this simple example all traffic destined for IP address 10.0.0.20 on port 80 is load-balanced to the real IP address 192.168.1.50 on port 80.

Packet rewriting works as follows:

1) The incoming packet for the web server has source and destination addresses as:

Source	x.x.x.x:34567	Destination	10.0.0.20:80
--------	---------------	-------------	--------------

2) The packet is rewritten and forwarded to the backend server as:

Source	x.x.x.x:34567	Destination	192.168.1.50:80
--------	---------------	-------------	-----------------

3) Replies return to the load balancer as:

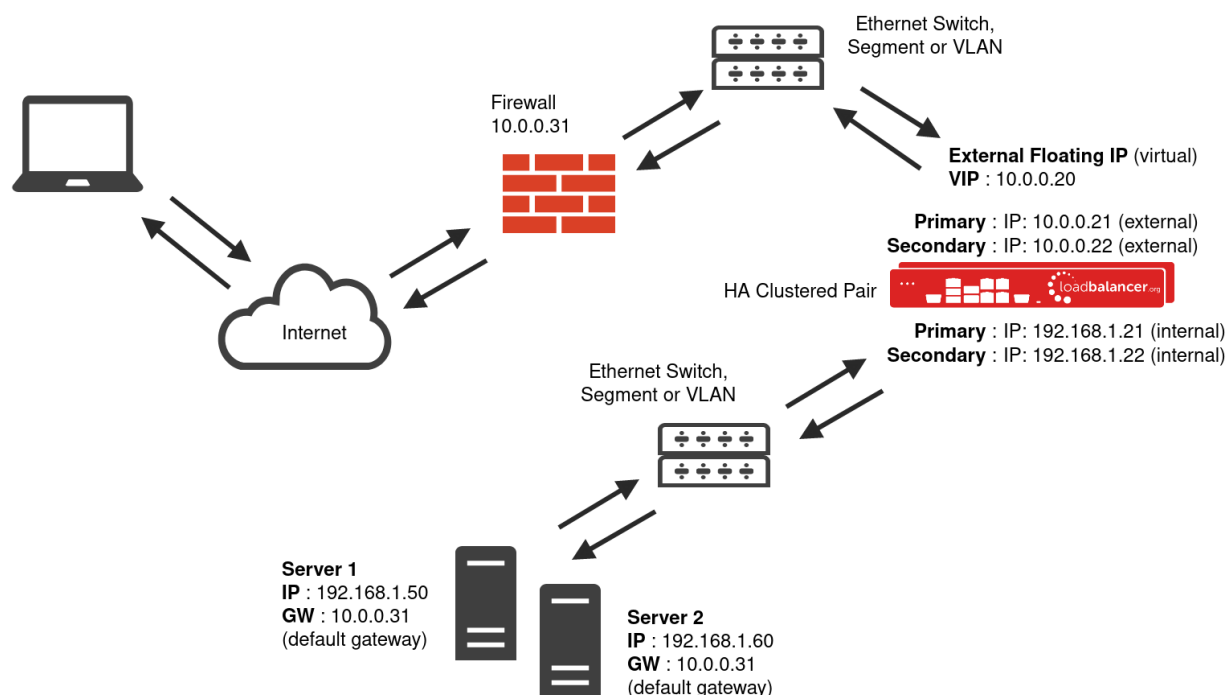
Source	192.168.1.50:80	Destination	x.x.x.x:34567
--------	-----------------	-------------	---------------

4) The packet is written back to the VIP address and returned to the client as:

Source	10.0.0.20:80	Destination	x.x.x.x:34567
--------	--------------	-------------	---------------

## Layer 4 SNAT Mode

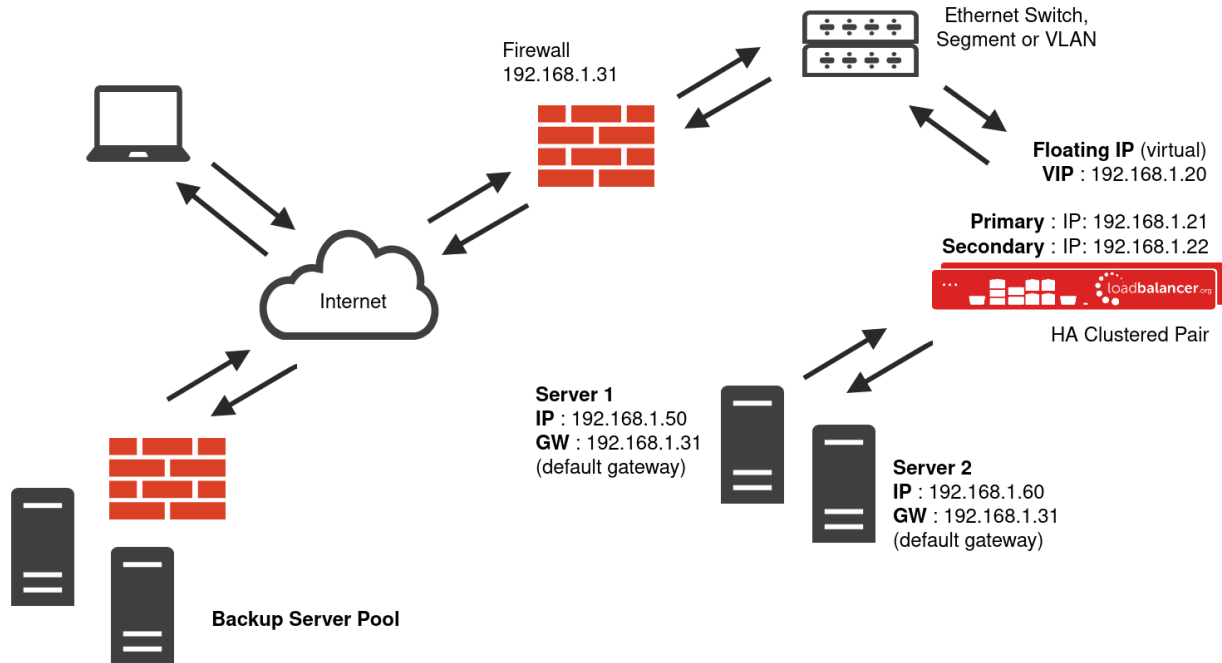
Layer 4 SNAT mode is a high performance solution, although not as fast as Layer 4 NAT mode or Layer 4 DR mode.



- The load balancer translates all requests from the external Virtual Service to the internal Real Servers in the same way as NAT mode - please refer to [Layer 4 NAT Mode](#) for more information.
- Layer 4 SNAT mode is not transparent, an iptables SNAT rule translates the source IP address to be the load balancer rather than the original client IP address.
- Layer 4 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, eth0 is normally used for the internal network and eth1 is used for the external network although this is not mandatory.
- If the Real Servers require Internet access, Autonat should be enabled using the WebUI option: *Cluster Configuration > Layer 4 - Advanced Configuration*, the external interface should be selected.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 4 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 4 SNAT mode VIPs and layer 7 SNAT mode VIPs because the required firewall rules conflict.

## Layer 7 SNAT Mode

Layer 7 SNAT mode uses a proxy (HAProxy) at the application layer. Inbound requests are terminated on the load balancer and HAProxy generates a new corresponding request to the chosen Real Server. As a result, Layer 7 is typically not as fast as the Layer 4 methods. Layer 7 is typically chosen when either enhanced options such as SSL termination, cookie based persistence, URL rewriting, header insertion/deletion etc. are required, or when the network topology prohibits the use of the layer 4 methods.



- Because layer 7 SNAT mode is a full proxy, any server in the cluster can be on any accessible subnet including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's own IP address by default, or any other local appliance IP address if preferred (e.g. the VIP address). This can be configured per layer 7 VIP. If required, the load balancer can be configured to provide the actual client IP address to the Real Servers in 2 ways. Either by inserting a header that contains the client's source IP address, or by modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. For more information on these methods please refer to [Transparency at Layer 7](#).
- Layer 7 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, eth0 is normally used for the internal network and eth1 is used for the external network although this is not mandatory.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

### Note

For detailed configuration examples using various modes, please refer to [Chapter 10 - Configuration Examples](#).

## Which Load Balancing Method Should I Use?

### Mode Summary

## Layer 4 DR Mode

This mode offers the best performance and requires limited physical Real Server changes. The load balanced application must be able to bind to the Real Server's own IP address and the VIP at the same time. This mode requires the **ARP Problem** to be solved as described [here](#). Layer 4 DR mode is transparent, i.e. the Real Servers will see the source IP address of the client.

## Layer 4 NAT Mode

This mode is also a high performance solution but not as fast as DR mode. It requires the default gateway of each Real Server to be the load balancer and supports both one-arm and two-arm configurations. Layer 4 NAT mode is transparent, i.e. the Real Servers will see the source IP address of the client.

## Layer 4 SNAT Mode

This mode is also a high performance solution but not as fast as the other layer 4 modes. It does not require any changes to the Real Servers and can be deployed in one-arm or two-arm mode. This mode is ideal for example when you want to load balance both TCP and UDP but you're unable to use DR mode or NAT mode due to network topology or Real Server related reasons. Layer 4 SNAT mode is non-transparent, i.e. the Real Servers will see the source IP address of the load balancer.

## Layer 7 SNAT Mode

This mode offers greater flexibility but at lower performance levels. It supports HTTP cookie insertion, RDP cookies, Connection Broker integration and works very well with either Pound or STunnel when SSL termination is required. It also enables content switching and header manipulation rules to be implemented. It does not require any changes to the Real Servers and can be deployed in one-arm or two-arm mode. HAProxy is a high performance solution, but since it operates as a full proxy it cannot perform as fast as the layer 4 solutions. Layer 7 SNAT mode is non-transparent by default, i.e. the Real Servers will see the source IP address of the load balancer. This mode can be made transparent through the use of TProxy.

## Our Recommendation

Where possible we recommend that Layer 4 Direct Routing (DR) mode is used. This mode offers the best possible performance since replies go directly from the Real Servers to the client, not via the load balancer. It's also relatively simple to implement. Ultimately, the final choice does depend on your specific requirements and infrastructure.

### Note

If you are using Microsoft Windows Real Servers (i.e. the backend servers) make sure that Windows **NLB** (Network Load Balancing) is **completely disabled** to ensure that this does not interfere with the operation of the load balancer.



## Chapter 4 - Appliance Fundamentals

### Hardware Appliance Installation


1. Remove all packaging and rack mount the appliance if required.
2. Connect the power lead.

**Note** | The power supply is an auto-sensing unit (100v to 240v).

3. Connect a network cable from your switch to one of the Ethernet ports, typically eth0 but this is not mandatory. If using a two-armed configuration connect another cable to a second Ethernet port, typically eth1 but again, this is not mandatory.
4. For a clustered hardware pair, the units must be able to communicate either via network (ucast), via serial cable or both. By default, ucast only is used. If serial is preferred or you want to use both methods, connect a serial cable between the two appliances.

**Note** | If a serial cable is used, Heartbeat must be configured for this using the WebUI option: *Cluster Configuration > Heartbeat Configuration* and enabling 'Serial'.

5. Attach a monitor & keyboard to the appliance.
6. Check mains power is on and press the power switch. The fans should start & front panel LEDs should light.



**Note** | The above image shows the Enterprise 1G. For network interface information for other models please refer to [Front & Rear Panel Layouts](#).

7. Now follow the on-screen instructions to configure the management IP address and other network settings as detailed in [Configuring Initial Network Settings](#).

### Virtual Appliance Installation

#### Supported Hypervisors

The VA is available for the following hypervisors:

- VMware ESXi : v4.0 & later
- Virtual Box : v4.0 & later

- Microsoft Hyper-V : v2012 & later
- KVM : Kernel version v2.6.20 & later
- XEN : v6.0 & later
- Nutanix AHV

## Host Requirements

To run the Loadbalancer.org Enterprise VA (irrespective of which Hypervisor is being used) the following basic server specifications must be met:

- 64bit CPU
- Virtual Technology hardware support - either Intel-VT-x or AMD-V compliant CPUs

## Virtual Hardware Resource Allocations

By default the appliance is allocated the following resources:

- 2 vCPUs
- 4GB RAM
- 20GB disk

The CPU and memory allocations are suitable for a PoC or for low to medium throughput production applications. For more demanding situations, they can be increased as needed. Resources required depend on multiple factors including the application being load balanced, the number of end-users, the anticipated throughput, the underlying physical hardware running the hypervisor and whether you'll be load balancing at layer 4 or layer 7. Therefore it's not realistic to make generic recommendations. If you need assistance in determining the resources required for your deployment, please contact support.

## Download & Extract the Appliance

1. Go to <https://www.loadbalancer.org/get-started/>.
2. Enter your details (all fields except email address are optional).
3. Click the **Download** button.
4. Now click the **Download** link next to the appropriate hypervisor.
5. Unzip the contents of the file to your chosen location.

### Note

All information provided is 100% confidential. If you're conducting a PoC (Proof of Concept) trial we may follow up with an email or phone call to see how you're getting on and offer assistance, but under no circumstances will Loadbalancer.org share your details with a third party.

### Note

The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.

### Note

The VA has 4 network adapters. For VMware only the first adapter (eth0) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

## Hypervisor Deployment

### VMware Host Client

1. Right-click **Host** in the VMware Host Client inventory and select **Create/Register VM**.
2. On the **Select creation type** page of the wizard, select **Deploy a virtual machine from an OVF or OVA file** and click **Next**.
3. On the **Select OVF and VMDK files** page, provide a unique name for the virtual machine.
4. Click the blue pane to open your local system storage, browse to the VA download location and select both the **.ovf** and **.vmdk** files.
5. Complete the remaining options according to your requirements and deploy the VA.

### VMware vSphere Client

The steps below apply to VMware ESX/ESXi & vSphere Client v6.7 and later.

#### Upgrading to the latest Hardware Version

When the appliance is deployed, the virtual hardware version is set to 7. This enables compatibility with ESX version 4.0 and later. You can upgrade to a later hardware version if required.

**Note** | Create a snapshot or backup of the virtual machine first before upgrading.

#### Installing the Appliance using vSphere Client

1. Right-click the inventory object where the appliance is to be located and select **Deploy OVF Template**.
2. In the **Select an OVF Template** screen, select the **Local File** option, click **Choose Files**, browse to the VA download location and select both the **.ovf** and **.vmdk** files and click **Next**.
3. In the **Select a name and folder** screen, type a suitable name for the appliance - this can be up to 80 characters in length.
4. Select the required location for the appliance - by default this will be the location of the inventory object from where the wizard was started and click **Next**.
5. In the **Select a compute resource** screen, select the required compute resource for the appliance - by default this will be the inventory object from where the wizard was started and click **Next**.
6. In the **Review details** screen, verify the template details and click **Next**.
7. In the **Select Storage** screen, first select the required storage location for the appliance.
8. Now select the required disk format and click **Next**.

**Note** | Loadbalancer.org recommends selecting a thick provision format. By default the appliance disk is 20GB.

9. In the **Select Networks** screen, select the required destination network using the drop-down next to **VM Network** and click **Next**.
10. In the **Ready to complete** screen, review the settings and click **Finish** to create the virtual appliance. To change a setting, use the **Back** button to navigate back through the screens as required.

#### Configure Additional Network Adapters

The appliance has 4 network adapters. By default only the first adapter is connected. This will be **eth0** when viewed in the appliance WebUI. If you require additional network adapters to be connected, follow the steps below:

1. Right-click the appliance, select **Edit Settings**.
2. Using the drop-down next to the relevant network adapter(s), select the required **Network** and tick (check) the **Connected** check-box and click **OK**.

Start the Appliance

Now power up the appliance.

### VMware Workstation Player

1. Select **Player > File > Open**.
2. Browse to the VA download location and select the **.ovf** file.
3. Modify the default name as required and click **Import**.

### VMware Tools / Open VM Tools

From v8.8.1 all new VA's use Open VM Tools rather than VMWare Tools.

When updating a v8.7.x appliance to v8.8.1 using online or offline update, VMWare tools is removed and Open VM Tools is installed.

### To Verify that Open VM Tools is running

Run the following command at the console or via an SSH session:

```
# service vmtoolsd status
vmtoolsd (pid 3514) is running...
```

#### Important

When using open-vm-tools, the VMware Tools status is shown as "Guest Managed" on the virtual machine Summary tab. The status Guest Managed means that you cannot use the vCenter Server to manage VMware Tools and you cannot use vSphere Update Manager to upgrade VMware Tools.

### Microsoft Hyper-V

The steps below apply to Windows Hyper-V 2012 & later.

Installing the Appliance using Hyper-V Manager

1. Start Hyper-V Manager, then using the right-click menu or the Actions pane select **Import Virtual Machine** and click **Next**.
2. In the **Locate folder** screen, browse to the location of the extracted download and select the **Loadbalancer.org Ent. VA** folder.
3. Click **Next** until you reach the **Choose Import Type** screen, select the option **Copy the virtual machine (create a new unique ID)** and click **Next**.
4. In the **Choose Folder For Virtual Machine Files** screen, tick (check) the checkbox **Store the Virtual Machine in different location**, then select a suitable location for the virtual machines files and click **Next**.
5. In the **Choose Folder to store Virtual Hard Disks** screen, select a suitable location for the virtual hard disk files and click **Next**.
6. In the **Completing Import Wizard** screen, verify that all settings are correct and click **Finish** to complete the import process. To change a setting, use the **Previous** button to navigate back through the screens as

required.

Once complete, the load balancer will appear in the **Virtual Machines** list.

**Note** | For a clustered pair, make sure that you select a different folder location in steps 4 & 5.

#### Configure Network Adapters

The appliance has 4 network adapters, these remain disconnected once deployment completes. To connect an adapter to a virtual switch:

1. Right-click the appliance, select **Settings**.
2. Select the first network adapter and set the required virtual switch that the adapter should be connected to. This will be **eth0** when viewed in the appliance WebUI.
3. Click **Apply**.

**Note** | If you need to connect additional network adapters repeat step 2 as required.

#### Start the Appliance

Now power up the appliance.

#### Linux Integration Services

Linux Integration Services are pre-installed by default. Therefore manual installation is not required.

#### KVM

The following steps should be followed on the KVM host:

1. Extract the archive to **/var/lib/libvirt/images/**.
2. `virsh define Loadbalancer*.xml`.
3. `virsh start Loadbalancer*`.

**Note** | Network cards are set to NAT by default so adjust as needed before powering on. Please also refer to the XML file for additional configuration notes.

#### Nutanix AHV

For detailed installation and deployment guidance, please refer to our [Nutanix blog](#).

#### XEN

The following steps should be followed on the XEN host:

1. Extract the archive.
2. Import the **xva** file into XEN.

## Cloud Appliance Installation

For details of our cloud based products, please refer to the relevant guide in the [documentation library](#).

## Configuring Initial Network Settings

After power up, the following startup message is displayed on the appliance console:

```

Welcome to the Loadbalancer.org appliance.

To perform initial network configuration, log in to the console as
Username: setup
Password: setup

To access the web interface and wizard, point your browser at
http://192.168.2.21:9080/
or
https://192.168.2.21:9443/

lbmaster login:
```

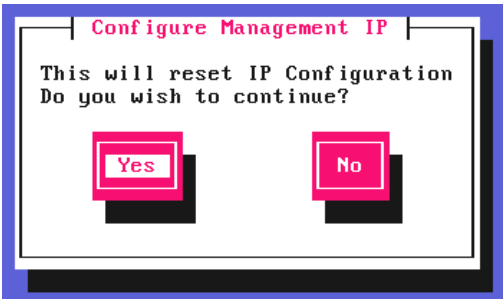
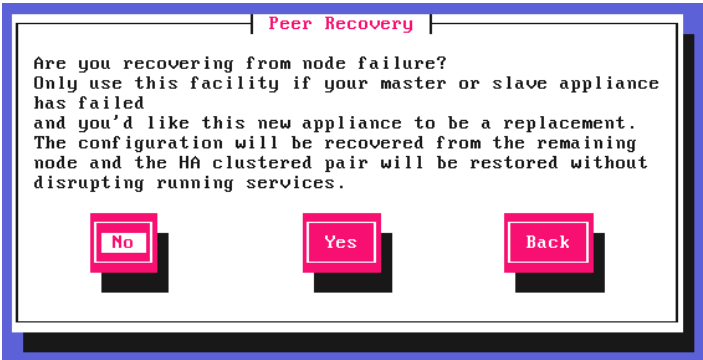
As mentioned in the text, to perform initial network configuration, login as the 'setup' user at the appliance console.

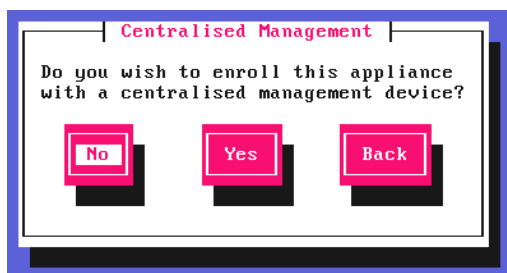
Once logged in, the Network Setup Wizard will start automatically. This will enable you to configure the management IP address and other network settings for the appliance.

login to the console:

Username: setup  
Password: setup

A series of screens will be displayed that allow network settings to be configured:

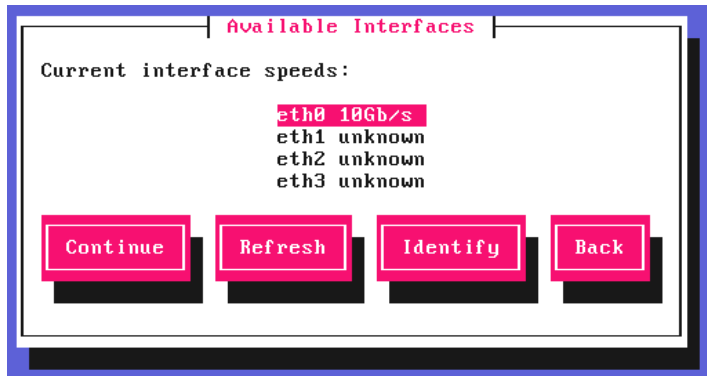
	To continue with the Network Setup Wizard select <b>Yes</b> and hit <ENTER> to continue.
	<div>Next, you'll be asked if you're recovering from node (i.e. Primary or Secondary) failure. If you're simply deploying a new appliance, select <b>No</b> and hit &lt;ENTER&gt; to continue.</div> <div>Note<div>For more details on node recovery using this option please refer to <b>Disaster Recovery After Node (Primary or Secondary) Failure</b>.</div></div>




Next, you'll be asked if you want to enroll the appliance for Centralised Management. If you have been provided with this information select **Yes**, otherwise leave **No** selected, then hit <ENTER> to continue.

Note


For information on how to modify Centralised Management settings via the WebUI, please refer to [Portal Management](#).



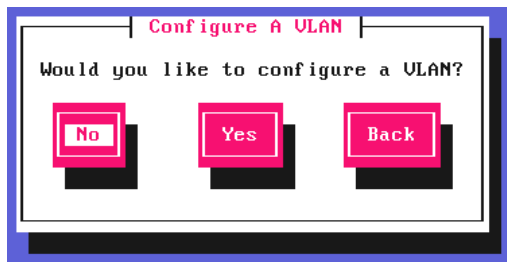
A list of available interfaces will be displayed, hit <ENTER> to continue.



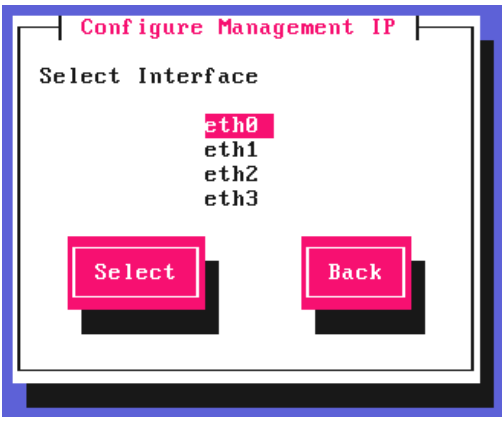
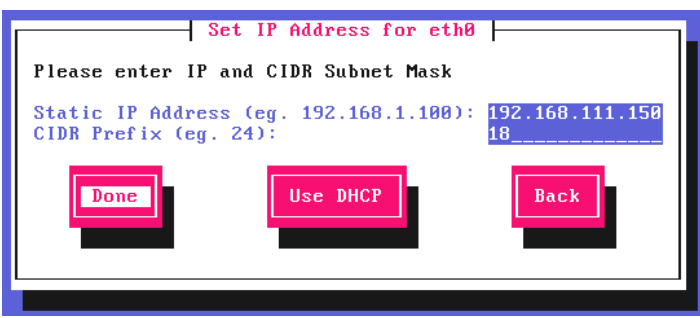

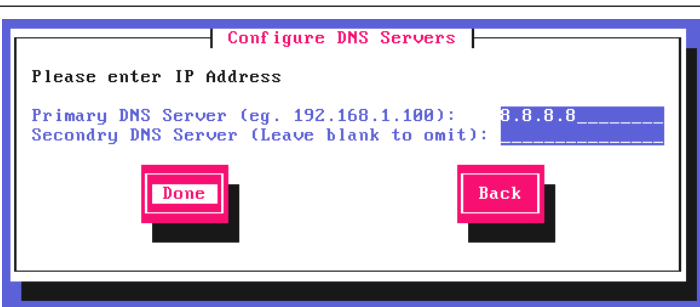
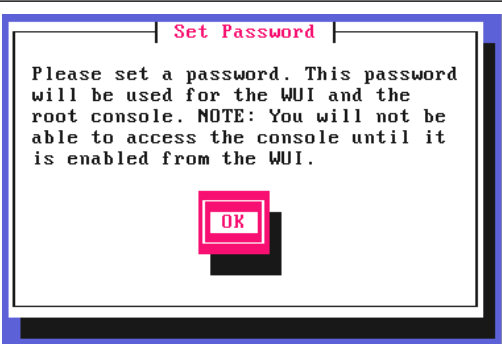
Select **Yes** If you want to configure a bonded interface, if not leave **No** selected, then hit <ENTER> to continue. If you select **Yes**, the screen shown below will be displayed:



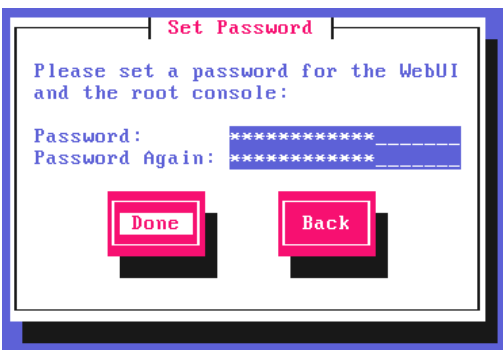
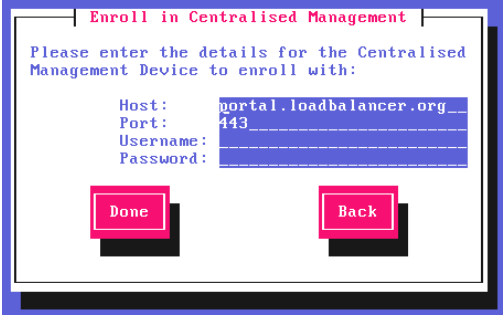
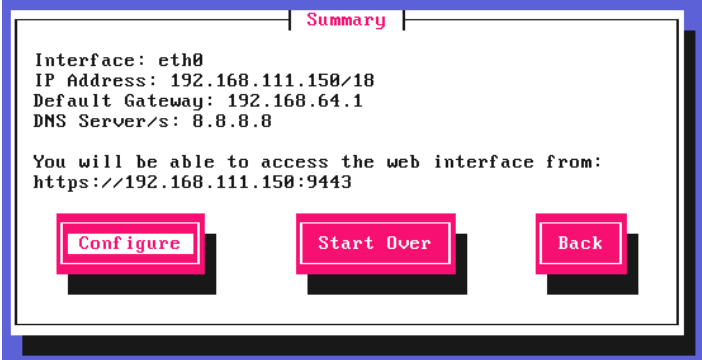
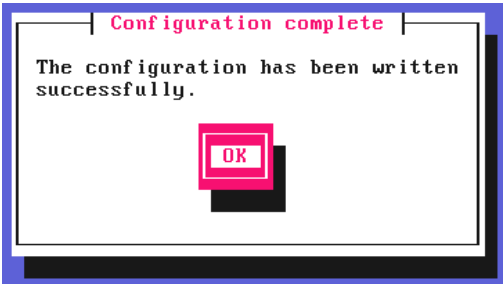
Using the space bar, select the interfaces you'd like to include in the bond, then click **Create**.



Select **Yes** If you want to configure a VLAN, if not leave **No** selected, then hit <ENTER> to continue. If you select **Yes** you'll be prompted to enter a VLAN Tag ID.

	<p>Select the interface that will be used to manage the appliance, then hit &lt;ENTER&gt; to continue.</p>
	<p>Either enter the required management IP address &amp; CIDR prefix and select <b>Done</b> or select <b>Use DHCP</b> to request an address and then hit &lt;ENTER&gt; to continue.</p> <div><div>Note</div><div>A subnet mask such as 255.255.255.0 is not valid, in this case enter 24 instead.</div></div>
	<p>Enter the default gateway address, select <b>Done</b> and hit &lt;ENTER&gt; to continue.</p>
	<p>Configure the required DNS server(s), select <b>Done</b> and hit &lt;ENTER&gt; to continue.</p>
	<p>Hit &lt;ENTER&gt; to continue.</p>



	Enter the password you'd like to use for the 'loadbalancer' WebUI user account and the 'root' Linux user account, select <b>Done</b> and hit <ENTER> to continue.
	If you selected <b>Yes</b> when asked if you want to enroll for Centralised Management, you'll now be prompted for the details. Default values for the Host and Port are set and can be changed if required. Enter a Username and Password, select <b>Done</b> and hit <ENTER> to continue.
	A summary of all settings is displayed, if everything looks good select <b>Configure</b> and hit <ENTER> to continue, all settings will then be applied.
	Once the configuration has been written, the Configuration Complete message will be displayed. Click <b>OK</b> to exit the wizard and return to the command prompt.

## Appliance Access & Configuration Methods

Once the initial network settings have been configured, the appliance can be accessed and configured using the Web User Interface (WebUI). In addition, full root access is provided at the console or via an SSH session.

Note

'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

### Accessing the WebUI

The WebUI is accessed using a web browser. By default, user authentication is based on local Apache .htaccess files. User administration tasks such as adding users and changing passwords can be performed using the WebUI menu option: *Maintenance > Passwords*.

- Note** | A number of compatibility issues have been found with various versions of Internet Explorer and Edge. The WebUI has been tested and verified using both Chrome & Firefox.
- Note** | If required, users can also be authenticated against LDAP, LDAPS, Active Directory or Radius. For more information, please refer to [External Authentication](#).

1. Using a browser, access the WebUI using the following URL:

**https://<IP-address-configured-during-network-setup-wizard>:9443/lbadmin/**

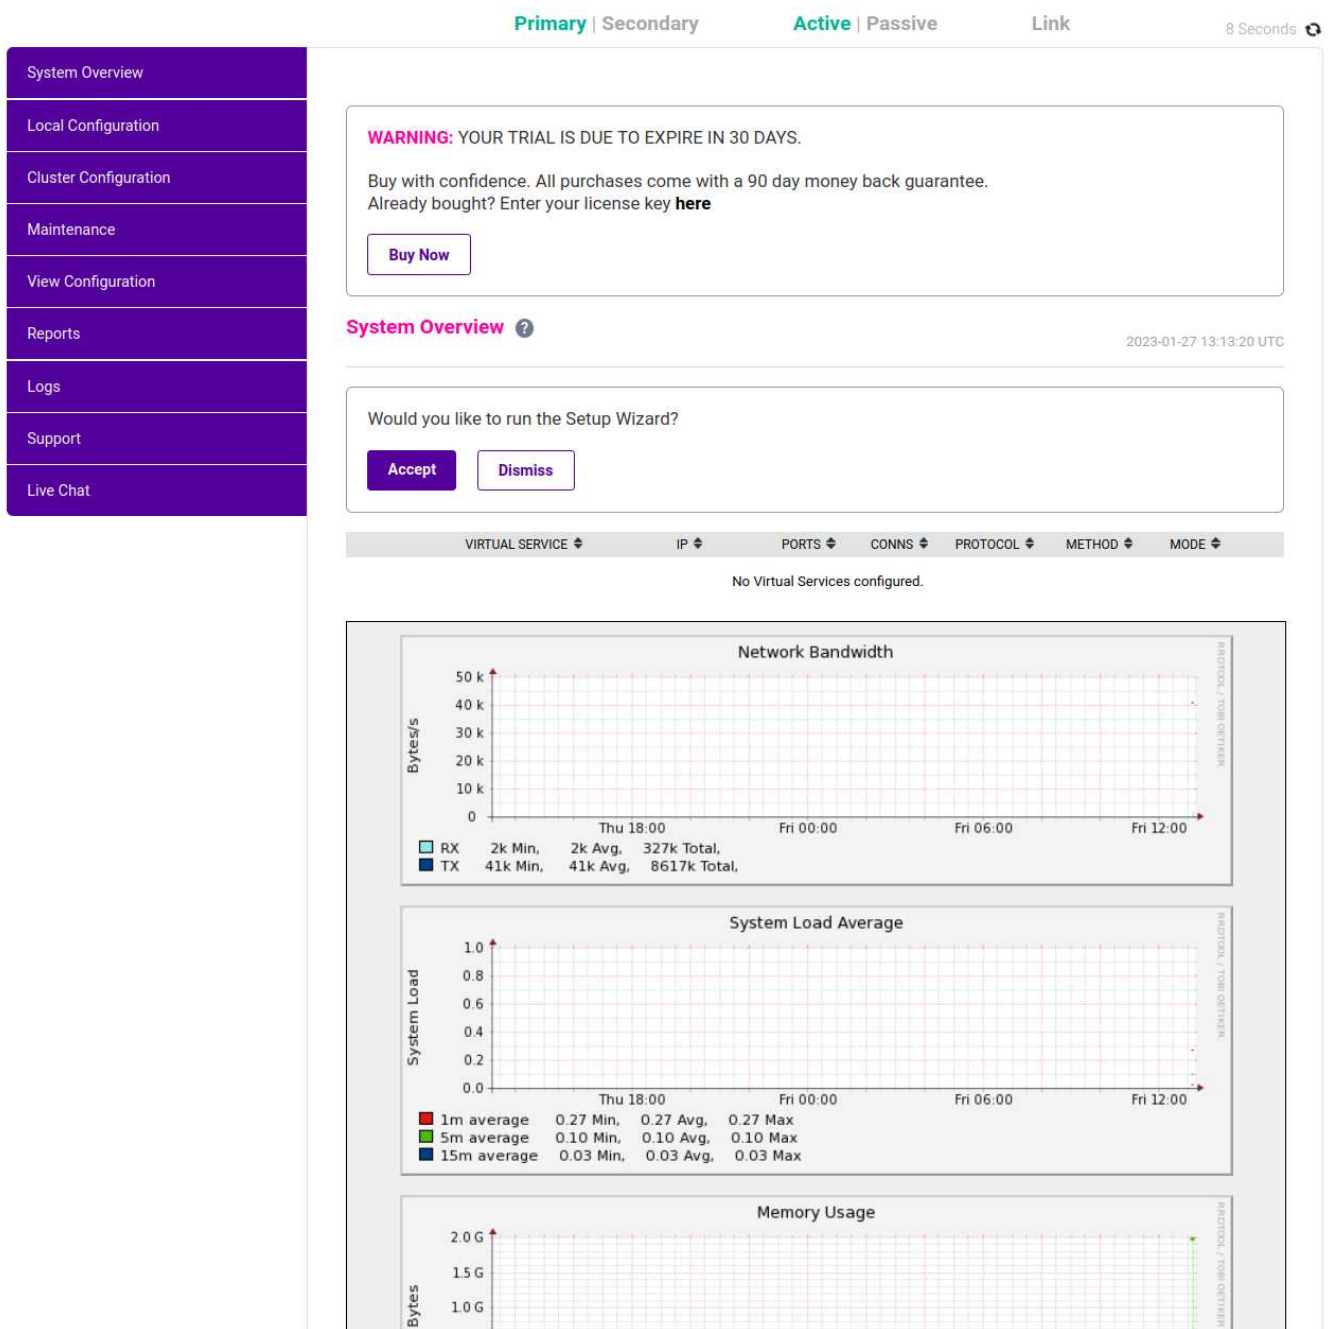
2. Log in to the WebUI:

**Username:** loadbalancer

**Password:** <configured-during-network-setup-wizard>

- Note** | To change the password, use the WebUI menu option: *Maintenance > Passwords*.

Once logged in, the WebUI will be displayed as shown below:

**Note**

The WebUI for the VA is shown, the hardware and cloud appliances are very similar. The licensing related message is platform & model dependent.

- You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

**Note**

The Setup Wizard can only be used to configure Layer 7 services.

**Main Menu Options**

**System Overview** - Displays a graphical summary of all VIPs, RIPs and key appliance statistics

**Local Configuration** - Configure local host settings such as IP address, DNS, system time etc.

**Cluster Configuration** - Configure load balanced services such as VIPs & RIPs

**Maintenance** - Perform maintenance tasks such as service restarts and taking backups

**View Configuration** - Display the saved appliance configuration settings

**Reports** - View various appliance reports & graphs

**Logs** - View various appliance logs

**Support** - Create a support download, contact the support team & access useful links

**Live Chat** - Start a live chat session with one of our Support Engineers

Note

The setup wizard can be used for configuring Layer 7 services only. For details on using the wizard, please refer to [Creating Layer 7 Virtual Services \(Using the Wizard\)](#).

Note

For full details of setting up layer 4 and layer 7 load balanced services, please refer to [Chapter 6 - Configuring Load Balanced Services](#).

## Root Access

Log in to the console or via an SSH session:

**Username:** root

**Password:** <configured-during-network-setup-wizard>

Note

'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

Note

For Windows hosts [PuTTY](#) can be used for SSH access and [WinSCP](#) can be used for secure file transfer.

## Keyboard Layout

By default the appliance is configured with a US keyboard layout. The layout can be changed by editing the file `/etc/sysconfig/keyboard`. For example, to change the layout from US to UK:

1. edit `/etc/sysconfig/keyboard` using a text editor such as **vi** or **vim** for Linux or WinSCP under Windows.
2. replace `KEYTABLE="us"` with `KEYTABLE="uk"`.
3. replace `Layout="us"` with `Layout="uk"`.
4. save the file and re-boot the appliance.

# Chapter 5 - Appliance Management

## Network Configuration

### Physical/Virtual Adapters

The Enterprise 1G, Enterprise 10G and all virtual models have 4 network adapters. The Enterprise 50G also has 4 adapters; 2 of these can be customized to suit your connectivity requirements when purchased. The Enterprise 100G supports up to 6 adapters depending on your choice of interface cards when purchased. Comprehensive information on all models is available [here](#).

If multiple logical interfaces are required, these can be added by specifying multiple IP addresses as shown below. If multiple cables must be connected, an external switch can be used.

Typically, the main reason for using all 4 or 6 interfaces is when bonding (e.g. 802.3ad) is required in a two-arm NAT or SNAT mode highly available configuration.

#### Note

For the VMware appliance, only the first adapter (eth0) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

### Configuring IP Addresses

As mentioned in the previous chapter, initial network settings can easily be configured using the Network Setup Wizard. For more information on using the wizard please refer to [Configuring Initial Network Settings](#).

IP addresses can also be configured using the WebUI menu option: *Local Configuration > Network Interface Configuration*. If a single interface is required, *eth0* is typically used. If 2 interfaces are required, *eth0* is typically used as the internal interface and *eth1* is used as the external interface. However, unlike other appliances on the market you can use any interface for any purpose.

In a simple one-arm configuration, you would just need to configure the IP address and subnet mask for one interface, e.g. *eth0* and if there are remote clients, the relevant default gateway. Both IPv4 and IPv6 addresses can be configured.

CIDR notation is used to specify IP addresses and subnet masks. For example, to specify an IP address of 192.168.2.100 with a subnet mask of 255.255.255.0, then 192.168.2.100/24 would be entered in the relevant interface field as shown in the example below:

eth0

192.168.2.100/24





#### Note

For information on CIDR notation please refer to [Appliance IPv4 Address Format \(CIDR notation\)](#).

To configure IP address(es):

1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.
2. Assign the required IP address/mask, multiple addresses can be assigned as shown below:

#### IP Address Assignment

				
	eth0	eth1	eth2	eth3
eth0	<div>192.168.10.100/24</div>			
		MTU	1500	bytes
eth1	<div>192.168.20.100/24 192.168.40.100/24</div>			
		MTU	1500	bytes
eth2				
		MTU	1500	bytes
eth3				
		MTU	1500	bytes

Configure Interfaces

3. Click **Configure Interfaces**.

#### Note

If you have Virtual Services defined when making changes to the network configuration, you should verify that your Virtual Services are still up and working correctly after making the changes.

## Configuring Bonding

The appliance supports bonding of multiple network adapters.

*To Configure Bonding:*

1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.
2. Any of the available adapters can be bonded. For example, to bond **eth0** and **eth1**, select (check) the *eth0* and *eth1* check-boxes.

### Create New Bond

eth0:	<input checked="" type="checkbox"/>	Interface Speed:
eth1:	<input checked="" type="checkbox"/>	Interface Speed: Link Down
eth2:	<input type="checkbox"/>	Interface Speed: Link Down
eth3:	<input type="checkbox"/>	Interface Speed: Link Down
Bonding Mode	Mode 1 - (Default) Active/Backup ▾	



Create

3. Select the required *Bonding Mode*:

**Mode 0** - Balance round robin. Transmits packets in a numerical order from the first available Secondary through to the last.

**Mode 1** - Active Backup (default). This places one of the adapters in a backup state and will only become active if the link is lost to the active adapter. This mode provides fault tolerance.

**Mode 4** - 802.3ad. Dynamic link aggregation mode. This mode requires a switch that supports IEEE 802.3ad.

#### Note

After changing the bonding mode the appliance must be restarted for the setting to take effect.

4. Click **Create**.

5. The new bond (bond0) is displayed as shown below:

### Create New Bond

eth2:	<input type="checkbox"/>	Interface Speed: Link Down
eth3:	<input type="checkbox"/>	Interface Speed: Link Down
Bonding Mode	Mode 1 - (Default) Active/Backup ▾	



Create

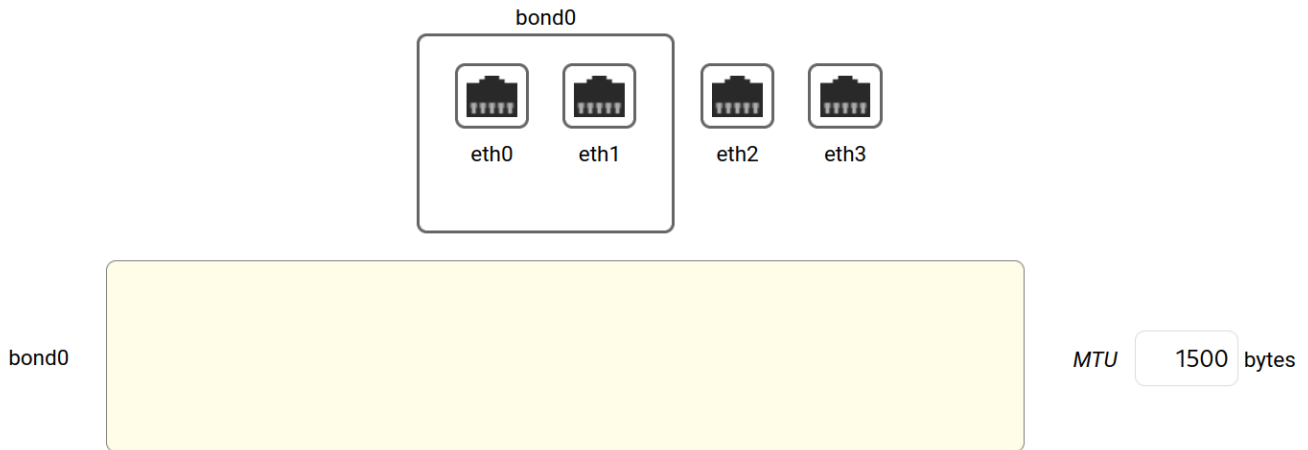
### Active Bonds

bond0	Mode 1	Interfaces: eth0,eth1
-------	--------	-----------------------

Delete

#### Note

At this point the adapters will still have the same IP settings configured previously. Once an IP address is defined for the bond and **Configure Interfaces** is clicked, these addresses will be removed and only the bond address will apply. If bonding is later disabled, these addresses will be re-applied to the adapters.



6. Enter the IP address for bond0 and click **Configure Interfaces**.
7. If the bonding mode has been changed, restart the appliance used the WebUI menu option: *Maintenance > System Control* and clicking **Restart Load Balancer**.

**Note** If you have a Primary and Secondary configured as an HA pair, make sure you configure bonding in the same way on both units. Failure to do this will result in heartbeat (Primary/Secondary communication) related issues.

**Note** If your Real Servers, ESX hosts etc. support network bonding using Broadcom's SLB (Smart Load Balancing), this can cause issues in Layer 4 DR mode if older drivers are used. We have successfully tested SLB (Auto Fallback Disable) with driver version 15.2.0.5. Therefore at least this version is recommended.

## Configuring VLANs

Native 802.1Q VLAN support can be enabled to load balance clusters on multiple VLANs.

In **access mode**, switch ports are dedicated to one VLAN. The switch handles all the tagging and de-tagging of frames - the station connected to the port does not need to be configured for the VLAN at all.

In **trunk mode**, the switch passes on the raw VLAN frames, and the station must be configured to handle them. Trunk mode is usually used to connect two VLAN-carrying switches, or to connect a server or router to a switch.


If the load balancer is connected to an access mode switch port no VLAN configuration is required. If the load balancer is connected to a trunk port, then all the required VLANs will need to be configured on the load balancer.


*To configure a VLAN:*


1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.
2. In the VLAN section select the required adapter (e.g. eth0).
3. Enter the VLAN ID (e.g. 250).
4. Click **Add VLAN**.
5. An extra IP Address Assignment field named eth0.250 will be created as shown below, the required IP address should be entered in this field.





## IP Address Assignment

eth0

eth0.250

eth1

eth2

eth3

eth0

192.168.111.220/18

MTU 1500 bytes

eth0.250

MTU 1500 bytes

Delete eth0.250

6. Click **Configure Interfaces**.

7. To delete the VLAN definition, click the appropriate **Delete** button.

### Note

If you have a clustered pair, don't forget to configure the same VLANs on the Secondary as these will not be replicated/created automatically.

## NIC Offloading

NIC offloading is enabled by default. This will enable (where available) hardware NIC offloading.

*To Configure Offloading:*

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration*.
2. Scroll down to the *Interface Offload* section.
3. configure the required setting.
4. Click **Update**.

## Configuring MTU Settings

*To set the MTU setting for an interface:*

1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*

eth0

192.168.10.100/24

MTU 1500 bytes

2. Enter the required MTU setting.

3. Click **Configure Interfaces**.

### Configuring Default Gateway & Static Routes

To set the Default Gateway for IPv4 and Ipv6:

1. Using the WebUI, navigate to: *Local Configuration > Routing*.
2. In the Default Gateway section define the default gateway as shown in the example below:

#### Routing

Default Gateway				
IP v4	<input type="text" value="192.168.1.254"/>	via interface	<input type="text" value="auto"/>	
IP v6	<input type="text"/>	via interface	<input type="text" value="auto"/>	

3. Click **Configure Routing**.

To configure Static Routes:

1. Using the WebUI, navigate to: *Local Configuration > Routing*.
2. In the Static Routes section configure the subnets & gateway addresses shown in the example below:

Static Routes			
Subnet	<input type="text" value="10.10.0.0/16"/>	via gateway	<input type="text" value="10.10.1.254"/>
Subnet	<input type="text" value="10.20.0.0/16"/>	via gateway	<input type="text" value="10.20.1.254"/>
Subnet	<input type="text"/>	via gateway	<input type="text"/>

3. Click **Configure Routing**.

Note

Unlimited static routes can be defined, additional blank rows will be added to the WebUI screen as they're used.

### Management Gateway

If you want to manage the appliance from a remote subnet that is accessible via a different gateway, the management address and the gateway can be specified here. Once configured, traffic from the management address is routed via the management gateway rather than via the default gateway.

Note

This is not a security feature, only a routing option. The WebUI is still accessible on all appliance IPs as before.

Note

Policy based routing is used to provide this feature. For more information, please refer to the PBR section below.

To configure the Management Gateway:

1. Using the WebUI, navigate to: *Local Configuration > Physical Advanced Configuration*.
2. Scroll down to the **Management Gateway** section.

**Management Gateway**

Management Address	<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">none ▼</div>	?
Via Gateway	<div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>	?

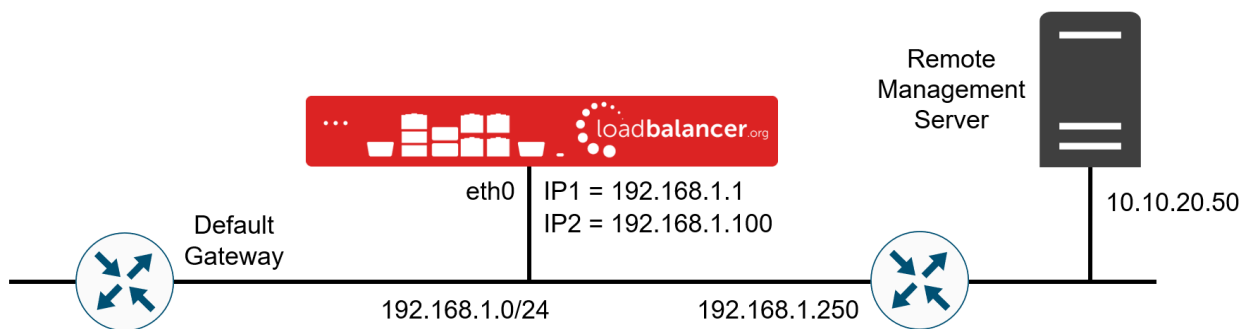
3. Select the required *Management Address* from the drop-down.

**Note** The drop-down is populated with all IP addresses that have been assigned to the various network interfaces.

4. specify the required *Gateway* address.
5. Click **Update**.

### Configuration Example

To enable appliance management on IP **192.168.1.100** from the remote management server shown below:



The following *Management Address* and *Via Gateway* settings would be required:

**Management Gateway**

Management Address	<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">192.168.1.100 ▼</div>	?
Via Gateway	<div style="border: 1px solid #ccc; padding: 2px; display: inline-block;">192.168.1.250</div>	?

**Note** You'll need to ensure that **192.168.1.100** is configured on eth0. This can be done using the WebUI menu option: *Local Configuration > Network Interface Configuration*.

### Policy Based Routing (PBR)

If you require a custom gateway for a particular VIP, this can be achieved using Policy Based Routing.

**Note** If client source addresses are known and predictable, static routes should normally be used to route traffic. In other situations where this is not known or the network is large with many subnets, PBR can be used. Here, return traffic is routed based on the source address of the reply traffic (the VIP/floating IP) rather than on the destination address (the client's IP).

To configure a VIP to return traffic via a custom gateway rather than via the default gateway:

1. Using the WebUI, navigate to: Cluster Configuration > PBR Default Gateways.

Select Floating IP

192.168.111.222 ▾

Gateway Address

192.168.111.254

Submit

Floating IP	Gateway IP	Table Name
-------------	------------	------------

2. Set the Floating IP to the required VIP address, e.g. **192.168.111.222**.
3. Set the Gateway Address to the required value, e.g. **192.168.111.254**.
4. Click **Submit**.
5. Once configured, the new default gateway will be displayed as shown below:

Select Floating IP

Select an IP ▾

Gateway Address

Submit

Floating IP	Gateway IP	Table Name
192.168.111.222	192.168.111.254	lbpbr-1

Delete

To delete the new gateway and configure the VIP to use the appliance's standard default gateway, click **Delete**.

## Configuring Hostname & DNS Configuration

To set the *Hostname*, *Domain* & *DNS servers*:

1. Using the WebUI, navigate to: *Local Configuration* > *Hostname & DNS*.

### Hostname & DNS

Hostname	lbmaster		?
Domain Name	localhost		?
Domain Name Server	Primary	8.8.8.8	?
	Secondary		?
	Tertiary		?

Update

2. Specify the required *Hostname*, by default this is set to **lbmaster**.
3. Specify the Domain name, by default this is set to **localhost**.

- Specify the required DNS server(s).
- Click **Update**.

## Portal Management

To Configure Centralized Portal Management Settings:

- Using the WebUI, navigate to: *Local Configuration > Portal Management*.

### Portal Management

Server Details		
Enabled	<input checked="" type="checkbox"/>	?
Hostname	<input type="text" value="portal.loadbalancer.org"/>	?
Port	<input type="text" value="443"/>	?
<b>Update</b>		

Adopt Node		
Adoption Submitted	no	
Username	<input type="text"/>	?
Password	<input type="password"/>	?
<b>Begin Adoption</b>		

- To enable communication with a Portal Management Server, tick (check) the *Enabled* checkbox.
- Default values for *Hostname* and *Port* are set, change these if required.
- Click **Update**.
- Specify the *Username* & *Password* to provide to the Portal when adopting this appliance.
- Click **Begin Adoption**.

**Note** | These settings can also be configured during the Network Setup Wizard.

## System Date & Time Configuration

### Auto Configuration using NTP Servers

To configure NTP:

- Using the WebUI, navigate to: *Local Configuration > System Date & Time*.

## System Date & Time

Current system time

2023-01-31 15:42:33 UTC

System Timezone

UTC ▼

NTP Servers

Set Timezone & NTP

Date

2023 ▼ – Jan ▼ – 31 ▼

Time

15 : 42

Set Date & Time

2. Select the required *System Timezone*.
3. Define your NTP servers using the *NTP Servers* fields.
4. Click **Set Timezone & NTP**.

## Manual Configuration

*To manually set the date & time:*

1. Set the data & time using the *Date & time* fields.
2. Click **Set Date & Time**.

### Note

When using a clustered pair (i.e. Primary & Secondary) date and time changes on the Primary will not be automatically replicated to the Secondary, therefore the date and time on the Secondary must also be set manually.

## Appliance Internet Access via Proxy

The appliance supports the ability to access the Internet via a proxy server.

*To set the Proxy Server's IP address & Port:*

1. Using the WebUI, navigate to: *Local Configuration > Physical Advanced Configuration*.

### Network Proxy

Proxy Server



Port



Username



Password



2. Enter the proxy's IP address in the *Proxy Server* field.
3. Enter the proxy's port in the *Port* field.
4. Enter a *Username* & *Password* if the proxy requires credentials.
5. Click **Update**.

**Note** | For a clustered pair, this setting must also be manually configured on the Secondary.

## SMTP Relay Configuration

The appliance can be configured with an SMTP smart host to receive all mail messages generated by the load balancer. If this field is not configured the address will be auto-configured based on an MX lookup of the destination email address that's configured under *Cluster Configuration > Layer 4 - Advanced Configuration*.

*To configure a smart host:*

1. Using the WebUI, navigate to: *Local Configuration > Physical Advanced Configuration*.
2. Scroll down to the SMTP Relay section.

### SMTP Relay

Smart Host



3. Enter an appropriate IP address or hostname in the *Smart Host* field.
4. Click **Update**.

**Note** | For a clustered pair, this setting must also be manually configured on the Secondary.

## Syslog Server Configuration

The appliance supports the ability to write all logs either locally, to an external Syslog Server or both. The Syslog server may be specified by IP address or hostname.

*To configure a Syslog server:*

1. Using the WebUI, navigate to: *Local Configuration > Physical Advanced Configuration*.
2. Scroll down to the *Logging* section.

Logging

Rate limit interval

?

Rate limit Burst limit

?

Log Destination

☐ Local Files

☐ Remote syslog Server

☒ Both

?

Remote syslog Server IP

?

Remote syslog Server Port

?

Remote syslog Server Protocol

UDP ▾

?

Remote syslog Server Template

?

Update

- Enter the required *Rate Limit Internal*, the default is 5 seconds.
- Enter the required *Rate Limit Burst Limit*, the default is 200 messages.
- Define whether logs should be written to *Local Files*, a *Remote Syslog Server* or *Both*.
- If *Remote Syslog Server* or *Both* is selected, the following options also apply:

Option	Description
Remote Syslog Server IP	The server may be specified by IP address or hostname. If you use a hostname, make sure DNS is correctly configured on the load balancer.
Remote Syslog Server Port	Specify the Remote Syslog Server port.
Remote Syslog Server Protocol	Select the communications protocol, either TCP or UDP. NOTE: If the load balancer has been configured to keep detailed logs of multiple services, and your syslog server is heavily loaded, we recommend that UDP is used and this is the default.
Remote Syslog Server Template	Specify a Remote Syslog Server template (string format).

- Click **Update**

**Note** For a clustered pair, this setting must also be manually configured on the Secondary appliance.



# SNMP Configuration

The appliance supports SNMP v1, v2 and v3.

To Configure SNMP:

1. Using the WebUI, navigate to: *Local Configuration > SNMP Configuration*.

Protocol Versions		
Enable SNMP v1 and v2	<input type="checkbox"/>	?
Enable SNMP v3	<input type="checkbox"/>	?
Details		
SNMP location	<input type="text" value="Server Room 1"/>	?
SNMP contact	<input type="text" value="IT Dept"/>	?
Authentication		
SNMP v1/v2 community string	<input type="text" value="public"/>	?
USM Username	<input type="text"/>	?
USM Authorization Algorithm	<input type="text" value="SHA"/>	?
USM Authorization Passphrase	<input type="text"/>	?
USM Privacy Algorithm	<input type="text" value="AES"/>	?
USM Privacy Passphrase	<input type="text"/>	?

Update

2. Using the checkboxes in the *Protocol Versions* section, enable the required SNMP version(s).
3. Enter a suitable *SNMP location* and *SNMP contact*.
4. For SNMP v1 & v2:
  - Enter a suitable *SNMP v1/v2 community string*.
5. For SNMP v3:
  - Enter a suitable *USM Username*, *USM Authorization Algorithm*, *USM Authorization Passphrase*, *USM Privacy Algorithm* and *USM Privacy Passphrase*.
6. Click **Update**.
7. Restart SNMPD using the **Restart SNMPD** button at the top of the screen.

**Note** | For a clustered pair, this setting must also be manually configured on the Secondary appliance.

**Note** | More information about the various OIDs and associated MIBs for the appliance please refer to [SNMP Reporting](#).

## Installing the License Key

The appliance can be used completely unrestricted for 30 days without installing a license key. After 30 days, the appliance continues to work but it's no longer possible to make changes to the configuration.

### Note

if you are conducting a PoC (Proof of Concept) using the VA and require more time to complete your evaluation, please contact [sales@loadbalancer.org](mailto:sales@loadbalancer.org) who will be able to provide guidance on how to extend the trial.

When first powered up the following message is displayed for a VA:

**WARNING:** YOUR TRIAL IS DUE TO EXPIRE IN 30 DAYS.

Buy with confidence. All purchases come with a 90 day money back guarantee.  
Already bought? Enter your license key **here**

[Buy Now](#)

The following message is displayed for a hardware appliance:

**WARNING:** This appliance is unregistered. **Please enter your license key** within 30 days to activate your appliance.

If you do not have your license key please **Contact Us**

*To install the license key:*

1. Using the WebUI, navigate to: *Local Configuration > License Key*.

### Install License Key

This unit is in evaluation mode. Please enter your license key to remove this restriction.

If you do not have a license key, please contact [sales@loadbalancer.org](mailto:sales@loadbalancer.org).

No file selected.

[Install License Key](#)

2. Click **Choose File** then browse to and select the license file provided when the appliance was purchased.
3. Click **Install License Key**.

### Note

Once the license is applied, these warning messages will no longer be displayed.

## Running OS Level Commands

The appliance supports the ability to run OS level commands directly from the WebUI.

To run an OS level command:

1. Using the WebUI, navigate to: *Local Configuration > Execute Shell Command*.

### Execute shell command

Execute shell command

2. Enter the relevant command in the field.
3. Click **Execute Shell Command**.
4. The results of the command as well as any errors will be displayed at the top of the screen.

#### Note

Commands that run continuously when executed should be run with a specific count to ensure that they will terminate gracefully and the results can be displayed in the WebUI.

For example with ping use:

```
ping -c 4 192.168.100.254
```

#### Note

The "Execute Shell Command" menu option is disabled by default. This can be enabled using the WebUI option: *Local Configuration > Security*. Set *Appliance Security Mode* to **Custom** then click **Update**.

## Restoring Manufacturer's Settings

The load balancer's settings can be reset to factory default values in two ways. In both cases this will remove all custom configuration from the load balancer. All VIPs, RIPs and other settings will be removed and the IP address configured for eth0 will be set to 192.168.2.21/24.

### Using the WebUI

To restore settings:

1. Using the WebUI, navigate to: *Maintenance > Backup & Restore > Restore Tab*.
2. Click **Restore Manufacturer's Defaults**.

Once restored, restart the appliance to complete the process.

### Using the Console / SSH Session

Run the following command:

```
lbrestore
```

Once restored, restart the appliance to complete the process.

## Note

'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

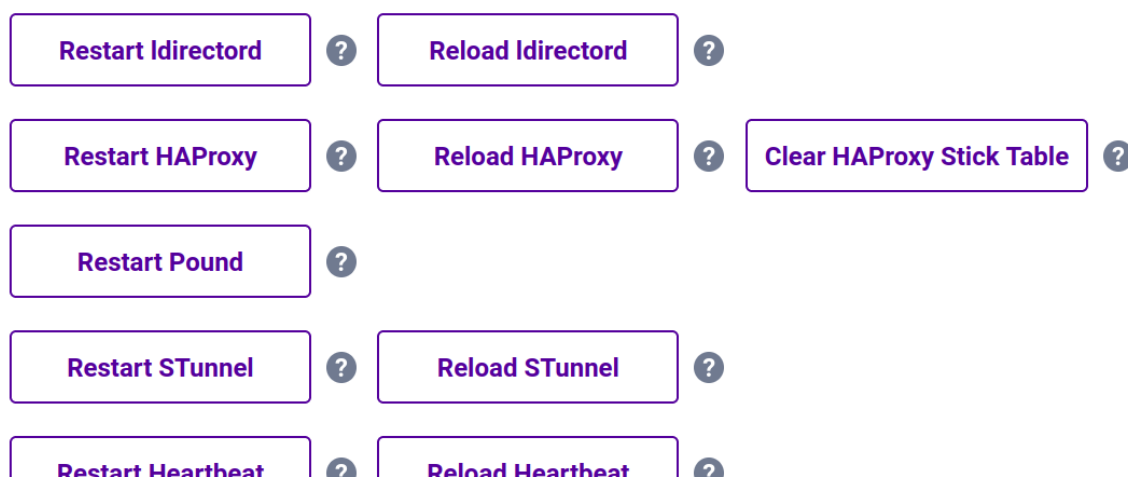
## Restarting & Reloading Services

The various services running on the appliance can be manually reloaded or restarted if required. This is normally only required for HAProxy, Pound, STunnel and Heartbeat when configuration changes are made.

To restart / reload services:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.

### Restart Services



2. Click the relevant restart or reload button.
3. Click **OK** to proceed.

The Following restart & reload options are available:

#### Restart Ldirectord

Restarting Ldirectord will result in a loss of layer 4 services during the restart.

#### Reload Ldirectord

Reloading Ldirectord may result in a loss of layer 4 services during the reload.

#### Restart HAProxy

Restarting HAProxy will result in a loss of layer 7 services during restart. It will cause any persistence tables to be dropped and all connections to be closed.

#### Reload HAProxy

Reloading HAProxy will reload the configuration. If you are using stick tables for persistence the entries will be copied between processes. HAProxy will start a new process (leaving the old one) with the new configuration. New connections will be passed onto this process, the old process will maintain existing connections and eventually terminate when there are no more connections accessing it.

## Note

If you have long lasting TCP connections it can take quite some time for this old process to

terminate, leaving those users running the old configuration. If this is taking too long - See Restart HAProxy.

### **Clear HAProxy Stick Table**

If you are using stick table persistence, this will clear the entries for all tables. Clients may be directed to a different server upon re-connection.

### **Restart Pound**

Restarting Pound will result in a loss of SSL termination services during the restart.

### **Restart STunnel**

Restarting STunnel will result in a loss of SSL termination services during the restart.

### **Reload STunnel**

Restarting STunnel may result in a loss of SSL termination services during the reload.

### **Restart Heartbeat**

Restarting heartbeat will cause a temporary loss of all layer 4, layer 7 and SSL services.

### **Reload Heartbeat**

Reloading heartbeat may cause a temporary loss of all layer 4, layer 7 and SSL services.

### **Restart Firewall**

All firewall rules will be removed, then reloaded from the current configuration. This may result in a temporary loss of service.

### **Restart Syslogd**

Restart Syslogd to load in any changes made to the configuration file.

### **Restart Collectd**

This will not clear the previously collected data. Note that collectd will not start if graphing of all services is disabled.

### **Restart SNMPD**

Restart the SNMP service on the local system.

### **Reload Apache**

Reload Apache performs a graceful restart which causes the parent process to advise the children to exit after their current request (or immediately if they're not serving anything). The parent then reloads its configuration and log files. As each child dies off the parent replaces it with a child with the updated configuration, which begins serving new requests immediately.

### **Restart WAF**

Restarting the WAF will drop all current connections and re-read the config.

### **Reload WAF**

Reload the WAF and re-read the config.

### **Restart GSLB**

Restart GSLB services to make live any changes to configuration. This will impact live services.

## Reload GSLB

Reload GSLB services to make live any changes to configuration. This should not impact live services.

## Restart Gateway

Restart the Portal Management Gateway Service.

# Appliance Restart & Shutdown

The appliance can be restarted or shutdown using the WebUI.

*To restart or shutdown the appliance:*

1. Using the WebUI, navigate to: *Maintenance > System Control*

## System Control

Restart Load Balancer

Halt Load Balancer

2. Select the required option:

Restart Load Balancer - *Shutdown and restart the appliance*

Halt Load Balancer - *Shutdown and halt the appliance*

## Appliance Software Update

We recommend that the appliance is kept up to date to ensure that you benefit from the latest bug fixes, security updates and feature improvements. Both online and offline updates are supported.

### Note

Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

## Determining the Current Software Version

The software version is displayed at the bottom of the WebUI as shown in the example below:

Copyright © Loadbalancer.org Inc. 2002 – 2023

ENTERPRISE VA Max - v8.8.1

English ▼

## Creating a Backup

We recommend that you create a backup before running the update. This can be done using the WebUI option: *Maintenance > Backup & Restore > Backup*. For more information on all backup and restore features, please refer to [Backup & Restore and Disaster Recovery](#).

## Online Update

### Auto-Check for Updates

The appliance periodically contacts the Loadbalancer.org update server ([update.loadbalancer.org](https://update.loadbalancer.org)) and checks for updates. This is the default behavior and can be disabled if preferred. If an update is found, a message will be displayed at the top of the screen as shown in the following example:

**Information:** Update v8.9 is now available for this appliance.

[Online Update](#)

To start the update process click **Online Update**.

### Configuring Online Update

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration*.
2. Scroll down to the *Online Updates* section.

#### Online Updates

Update Server



Auto-check for updates



3. Configure the required settings.
4. Click **Update**.

### Manual Check for Updates

If auto-check for updates is disabled, the update check must be initiated manually.

*To initiate a manual online update check:*

1. Using the WebUI, navigate to: *Maintenance > Software Update*.
2. Select **Online Update**.
3. If the latest version is already installed, a message similar to the following will be displayed:

**Information:** Version v8.8.1 is the current release. No updates are available.

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

**Note** | Do not navigate away whilst the update is ongoing, this may cause the update to fail.

6. Once complete (the update can take several minutes depending on download speed and upgrade version) the following message will be displayed:

**Information:** Update completed successfully.

7. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

## Offline Update

If the load balancer does not have access to the Internet, offline update can be used.

**Note** | Please contact [support@loadbalancer.org](mailto:support@loadbalancer.org) to check if an update is available and obtain the latest offline update files.

*To perform an offline update:*

1. Using the WebUI, navigate to: *Maintenance > Software Update*.
2. Select **Offline Update**.
3. The following screen will be displayed:

### Offline Update

The following steps will lead you through offline update.

1. Contact **Loadbalancer.org support** to obtain the offline update archive and checksum.
2. Save the archive and checksum to your local machine.
3. Select the archive and checksum files in the upload form below.
4. Click *Upload and Install* to begin the update process.

**Archive:**  No file selected.

**Checksum:**  No file selected.

**Upload and Install**

4. Select the Archive and Checksum files.
5. Click **Upload and Install**.

## Updating a Clustered Pair

**Note** | Since services may need to be restarted during the update process, we recommend performing the update during a maintenance window.

*To update a Clustered Pair:*

1. Perform the update on the Secondary (passive) appliance first.
2. Restart/reload services or reboot the appliance as directed.



3. Failover to the Secondary so this is now the active appliance.
4. Now update the Primary appliance in the same way.
5. Restart/reload services or reboot the appliance as directed.

#### Note

For a clustered pair, we recommend fully testing & validating the Primary/Secondary failover process before going live. For more information, please refer to [Testing & Verifying Primary/Secondary Replication & Failover](#).

## Appliance Security Features

The appliance includes a number of security related features that can be used to help ensure the appliance is secure.

### Security Mode

To control how the appliance is accessed and which features are enabled, 3 security modes are provided:

- **Secure - (default)** - in this mode:
  - the WebUI is accessible on HTTPS port **9443**. If you attempt to access the WebUI on HTTP port **9080** you will be redirected to HTTPS port **9443**
  - access to the *Local Configuration > Execute shell Command* menu option is disabled
  - the ability to edit the firewall script & the lockdown wizard is disabled
  - 'root' user console & SSH password access are disabled
- **Custom** - in this mode the following check-box options can be configured to suit your requirements:
  - *Disable Console Access* - prevent 'root' access via the console
  - *Disable SSH Password Access* - prevent 'root' access via SSH
  - *Web User Interface via HTTPS only* - control whether the WebUI is only accessible on HTTPS port **9443** only or via both HTTPS port **9443** and HTTP port **9080**

#### Note

When **Custom** is selected access to the *Local Configuration > Execute shell Command* menu option is automatically enabled.

- **Secure - Permanent** - this mode is the same as **Secure** but the change is *irreversible*

#### Important

Only set the security mode to **Secure - Permanent** if you are 100% sure this is what you want!

*To configure the Appliance Security Mode & related options:*

1. Using the WebUI, navigate to: *Local Configuration > Security*.

Appliance Security Mode	Secure - (Default) ▼	?
HTTPS Port for Web User Interface	9443	?
Web Interface SSL Certificate	Default Self Signed Certificate ▼	?
Ciphers to use	ECDHE-ECDSA-AES256-GCM	?

**Update**

2. Select the required *Appliance Security Mode*.
3. If **Custom** is selected, configure the additional options to suit your requirements.
4. Specify the HTTPS port for the WebUI, the default is **9443**.
5. Select the required SSL certificate for the WebUI. Certificates can be created/uploaded using the WebUI menu option: *Cluster Configuration > SSL Certificate*. If no certificates are available, the appliance's default self-signed certificate will be used.
6. Specify the required cipher, the default is:

ECDHE - ECDSA - AES256 - GCM - SHA384 : ECDHE - ECDSA - AES128 - GCM - SHA256 : DHE - RSA - AES256 - GCM - SHA384 : DHE - RSA - AES128 - GCM - SHA256 : ECDHE - RSA - AES256 - GCM - SHA384 : ECDHE - RSA - AES128 - GCM - SHA256

## Users & Passwords

### Linux 'root' User Account

One of the great advantages of the Loadbalancer.org appliance is that you have full root access. This unlocks the full benefit of the underlying Linux OS. Other vendors tend to lock this down and only provide limited access to certain features and tools.

#### Note

As mentioned in the section above 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, enable the required option(s) and click **Update**.

The root password can be changed at the console, or via an SSH session using the following command:

```
# passwd
```

#### Note

For the AWS and Azure cloud products it's not possible to directly login as root. If root access is required, once you've logged into the console/SSH session using the credentials defined during instance deployment, run the following command:

```
$ sudo su
```

### WebUI User Accounts

By default the appliance includes four predefined WebUI user accounts. The default usernames, passwords, group membership and their primary use are:

Username	Default Password	Default Group	Description (please also refer to the group table below)
configuser	configuser	config	appliance administration account
loadbalancer	(configured during the Network Setup Wizard)	config (*)	appliance administration account
reportuser	reportuser	report	viewing the appliance configuration, reports & logs
maintuser	maintuser	maint	same as reportuser plus can also take servers on/off line & create the support download archive file

(\*) It's not possible to change the default group for the 'loadbalancer' user account.

#### Note

These are Apache .htaccess style accounts and are not related to the local Linux OS level accounts.

The permissions for each group are shown below:

	Menu/Permissions							
Group	System Overview	Local configuration	Cluster Configuration	Maintenance	View Configuration	Reports	Logs	Support
config	Full	Full	Full	Full	View	Full	View	Full
report	View	None	None	None	View	Full	View	View
maint	Full	None	None	None	View	Full	View	Full

It's also possible to define users who will be authenticated by an external LDAP/ADAuth system as described in the *Adding New Users* section below.

## Modifying User Passwords

*To modify a user's password:*

1. Using the WebUI, navigate to: *Maintenance > Passwords*.

## Passwords

configuser	Modify	Delete
loadbalancer	Modify	
maintuser	Modify	Delete
reportuser	Modify	Delete

- Click the **Modify** button next to the relevant user.
- Now change the password for the selected user:

Username	<input type="text" value="loadbalancer"/>
Password *	<input type="password"/>
Re-enter Password *	<input type="password"/>

**Note** | Passwords cannot contain the double quotation mark ( " ).

- Click **Edit User**.

## Adding New Users

*To add new users:*

- Using the WebUI, navigate to: *Maintenance > Passwords*.
- Use the following section:

## Add New User

Username	<input type="text"/>
LDAP/ADAuth User	<input type="checkbox"/>
Password *	<input type="password"/>
Re-enter Password *	<input type="password"/>
Group	report ▼

Add New User

3. Enter the required *Username*.
4. If the user will be authenticated by an external LDAP/ADAuth or RADIUS system, enable (check) the *LDAP/ADAuth User* checkbox.

### Note

For more information on external authentication please refer to the *External Authentication* section below.

5. For locally authenticated users, enter the required *Password*.

### Note

Passwords cannot contain the double quotation mark ( " ).

6. Select the required *Group* for the new user.
7. Click **Add New User**.

## Resetting forgotten Passwords

It's possible to reset passwords via the command line if required. To do this you'll need to login as root to the console/SSH session. The *htpasswd* command can then be used as shown below:

```
htpasswd -b /etc/loadbalancer.org/passwords loadbalancer <new password>
```

### Note

'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

## External Authentication

The appliance supports the following external authentication methods:

- LDAP
- LDAPS
- Active Directory
- Radius

Once a user is configured to use external authentication, they simply enter their credentials for that system to access the appliance. It's important to remember that the username defined in the **Add New User** screen must be the exact same username defined in the external authentication system.

To demonstrate this feature, an Active Directory example is presented below, the steps required to configure against other external authentication systems are similar.

#### Configuring External AD Authentication

To configure external authentication, either at the console or using an SSH session, login as root and run the following command:

```
lbauthconfig
```

**Note** | 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

**Note** | Do not run lbauthconfig from the WebUI.

The following will be displayed:

```
[root@lbmaster ~]# lbauthconfig
#####
#   Loadbalancer.org External Authentication Config Script   #
#####

This script will setup either LDAP or RADIUS authentication for the Load Balancer WUI.
It should be noted that it does not disable the local user access for users such as
"loadbalancer", "maintenance" and "reports" (Or any other additional locally created users)
so it is recommended that once you have tested Auth works successfully that you then assign
long and complex passwords to these users and file them away.

It should also be noted that you will need to add users to the load balancer in order for them
to gain access, authentication is handled by the external authentication method but access is
still controlled by the load balancers default groups so local users identical to those stored
in the external auth provider are required.
```

Step 1. Select authentication type.

1. Local Files (Reset to defaults)
2. LDAP
3. LDAPS
4. Active Directory
5. Radius

[1-5]:

>> To configure Active Directory, enter **4** and hit <ENTER>

Step 2. Enter the hostname/address and port of the authentication server.

Hostname/address: 192.168.112.1

Port [1-65535 (3269)]:

>> Enter the hostname or IP address & port (default is **3269**) of your domain controller, e.g. **192.168.112.1** and hit <ENTER>

Step 3. Enter the LDAP Base Search string.

In the case of AD at a minimum this is your domain name so for "DOMAIN.LOCAL" you would use DC=DOMAIN,DC=LOCAL.

Search Base: DC=lbtestdom,DC=com

>> Enter the LDAP search base string for your domain, e.g. **DC=lbtestdom,DC=com** and hit <ENTER>

Step 4. Enter the LDAP attribute to authenticate against.

In the case of AD this will be "userPrincipalName" or "samAccountName" while OpenLDAP will typically use "uid".

Attrib [userPrincipalName]:

>> Enter the LDAP attribute to authenticate against (default is **UserPrincipalName**) and hit <ENTER>

Step 5. Enter the credentials for a user who can browse the directory.

Username: lbuser

Password:

>> Enter the credentials for a user who can browse AD, e.g. **lbuser** and hit <ENTER>

Step 6. Please add your first user.

This user will be added as a "config" user with full access to the WUI.

Please add additional users via the WUI after this setup process.

Username: tom@lbtestdom.com

Password:

>> Enter the username & password of the first AD authenticated user, e.g. **tom@lbtestdom.com** and hit <ENTER>

The new user will be added as shown below:

Test authentication succeeded.

Creating backup of /var/www/html/lbadmin/.htaccess at /etc/loadbalancer.org/bkup/.htaccess-2019-12-12T15:25:58.621399

Creating backup of /etc/loadbalancer.org/groups at /etc/loadbalancer.org/bkup/groups-2019-12-12T15:25:58.621399

Creating backup of /etc/loadbalancer.org/passwords at /etc/loadbalancer.org/bkup/passwords-2019-12-12T15:25:58.621399

Writing new /var/www/html/lbadmin/.htaccess File.

Writing new /etc/loadbalancer.org/groups File.

Adding tom@lbtestdom.com to /etc/loadbalancer.org/groups File.

Adding password for user tom@lbtestdom.com

Finished.

[root@lbmaster ~]#

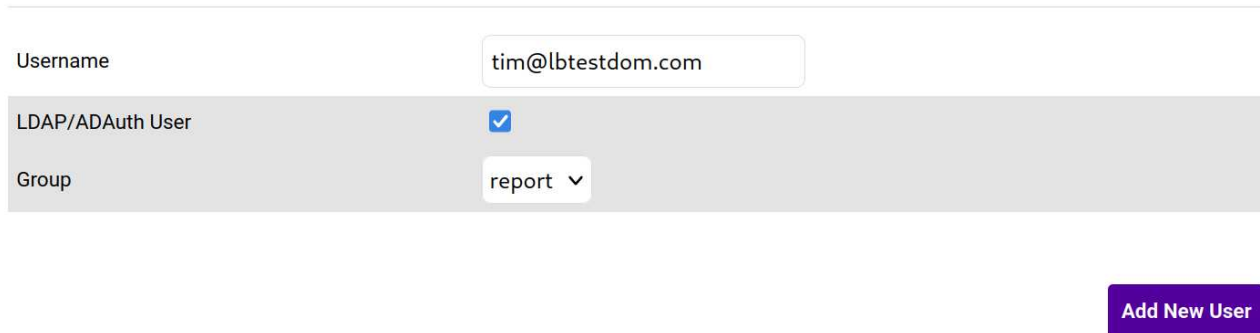
The new user will be added to the appliance and can be viewed using the WebUI menu option: *Maintenance > Passwords*. When the user logs in, they must use their AD credentials.

Adding Additional Users

**Note** | Configuring external authentication using groups is currently not supported.

Once the first user has been added, additional users can be added using the **Add New User** screen which is accessible via the WebUI option: *Maintenance > Passwords* as shown below:

### Add New User



Username	tim@lbtestdom.com
LDAP/ADAuth User	<input checked="" type="checkbox"/>
Group	report ▼

Add New User

1. Specify the same username as the AD / RADIUS user to be added, e.g. **tim@lbtestdom.com**.
2. Enable (check) the *LDAP/ADAuth User* checkbox.
3. Select the required security *Group*.
4. Click **Add New User**.

User **tim@lbtestdom.com** will now be able to login to the appliance with report user access rights using his AD credentials.

## Firewall Configuration

### Note

Whilst the load balancer is capable of supporting complex firewall rules, we do not recommend using the load balancer as your main bastion host. We recommend that the load balancer is deployed behind your external firewall.

If you want to configure firewall rules, here are some points to consider:

- All Virtual Service connections are dealt with on the INPUT chain not the FORWARD chain
- The WebUI runs on HTTP port 9080 (disabled by default) and HTTPS port 9443
- SSH on the load balancer listens on the standard port (22)
- SNAT & DNAT is handled automatically for all layer 4 NAT mode (LVS) and layer 7 (HAProxy) based Virtual/Real load balanced services
- You can use the standard Linux filters against spoofing attacks and syn floods
- LVS has built in DOS attack filters that can be implemented

### Note

Plenty of extra information is available on the Internet relating to Linux Netfilter and LVS, if you need any assistance please contact [support@loadbalancer.org](mailto:support@loadbalancer.org).

## Manual Firewall Configuration

The firewall can be configured manually using the WebUI based script editor. This enables iptables rules and any other required commands to be easily defined. The form allows you to directly edit `/etc/rc.d/rc.firewall`.

Custom rules can be configured, or for belt & braces security your external firewall settings can be replicated on to the load balancer for multi-layer security.



If you're planning to use NAT mode you may want to use the load balancer as your main firewall but we recommend that it is better and simpler to keep your firewall separate from the load balancer, especially if you want to set up VPNs etc. You can also use the firewall script to group ports together using Firewall Marks. For more information please refer to [Firewall Marks](#).

To configure custom firewall rules:

1. Using the WebUI, navigate to: *Maintenance > Firewall Script*.
2. The following screen will be displayed:

### Firewall Script

```
1  #!/bin/sh
2  # $Id$
3
4  #
5
6  # User firewall script for Loadbalancer.org appliance.
7  #
8
9
10
11
12  # Please note:
13  #      Most configurations will not require any changes to be made to
14  #      this script.
15  #
16  #      Administrators will only need to modify this script if their
17  #      needs are not met by the lock-down wizard, auto-NAT, and
18  #      automatic firewall mark functions of the web interface.
19  #
20
21
22
23  ##### One-arm NAT Mode #####
24  # For one-arm NAT, ICMP re-directs will need to be disabled.
25  # (1 = on, 0 = off)
26  #echo "0" >/proc/sys/net/ipv4/conf/all/send_redirects
27  #echo "0" >/proc/sys/net/ipv4/conf/default/send_redirects
28
29
30
31
32  ##### Manual Firewall Marks #####
```

Update

3. Define additional rules anywhere in the script above the last two lines:

```
echo "Firewall Activated"
exit 0;
```

4. Click **Update**.

#### Note

For a clustered pair, firewall script changes must also be manually configured on the Secondary.

#### Note

Be careful !! Make a backup before changing this script so that you know you can roll everything back if you cause a problem. A backup can be created using the WebUI menu option: *Maintenance > Backup & Restore > Make Local Firewall Script Backup*.

### Firewall Lock-down Wizard

The firewall lock down wizard can be used to automatically configure the load balancer to allow access to the various admin ports from one specific IP address or subnet. The wizard automatically detects the IP address of the

client running the WebUI and inserts this into the Admin IP field. The default mask is set to 255.255.255.0 which can be changed as required.

The firewall lockdown wizard uses two files:

- **rc.lockdownwizard** - this file contains the script that can be changed.
- **rc.lockdownwizard.conf** - this file contains a set of variable definitions and is written automatically when **Update firewall lock down** is clicked. The file depends on the rc.lockdownwizard script and the load balancer's configuration. This file should not be changed manually.

When run, rc.lockdownwizard loads the settings from the definitions file rc.lockdownwizard.conf and uses them to generate the firewall rules. You can modify rc.lockdownwizard via ssh or from the WebUI using the **Modify the firewall lock down wizard script** button.

**Note** 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

The default script does not depend on the configured Virtual Services or Real Servers, so the wizard does not need to be re-run when services are changed. However, it does depend on the IP addresses of the Primary and Secondary, and the admin related ports used by the WebUI, heartbeat, and HAProxy. If those settings are changed, the firewall lockdown wizard will need to be re-run in order to reflect the changes. Re-running the firewall lockdown wizard will adapt the rc.lockdownwizard.conf definitions file automatically - any changes made to the script rc.lockdownwizard will remain when you re-run the firewall lockdown wizard.

*To run the lock-down wizard:*

**Note** If you plan to configure an HA pair, either run the lockdown script on each appliance after the pair has been configured, or if it has already been run, temporarily disable the script (on both appliances) whilst performing the pairing process.

1. Using the WebUI, navigate to: *Maintenance > Firewall Lock Down Wizard*.
2. The following screen will be displayed:

### Firewall Lock Down Wizard

**WARNING:** Once the lock-down wizard is enabled, administration access to the load balancer will only be allowed from the Administration Subnet specified below.

Enable lock down script

☐

Administration subnet

192.168.64.7/24



Update firewall lock down

Modify the firewall lock down wizard script

3. Define your administration subnet/host in the *Administration subnet* field.

#### Note

Make sure that the subnet mask is correct - by default a /24 mask is used. To lock down access to a single host use <IP address>/32, e.g. 192.168.2.1/32.

4. Click **Update firewall lock down**.

*To disable the lock-down script:*

1. To disable the lock-down script uncheck the *Enable lock down script checkbox* and click the **Update Firewall lock down** button.

#### Note

If you accidentally block your own access to the appliance you will need to clear the current firewall rules and try again. To clear the firewall tables completely use the following command at the console: `/etc/rc.d/rc.flush-iptables`.

### Conntrack Table Size

By default the connection tracking table size is set to 524288 and is fine in most cases. For high traffic deployment using NAT mode, or when using connection tracking in the firewall script, this value may need to be increased. If the connection tracking table fills up, the following error will be reported in the log:

```
nf_conntrack: table full, dropping packet
```

*To modify this setting:*

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration*.
2. Scroll down to the *Firewall* section.
3. Set *Connection Tracking table size* to the required value.
4. Click **Update**.

#### Note

For a clustered pair, this setting must also be manually configured on the Secondary.

### Appliance Security Lockdown Script

To ensure that the appliance is secure it's recommended that a number of steps should be carried out. These steps have been incorporated into a lockdown script which can be run at the console (recommended) or via an SSH session. When run on the Primary of a correctly configured clustered HA pair, both appliance's will be updated. The script locks down the following:

- the password for the 'loadbalancer' WebUI account
- the password for the Linux 'root' account
- from which subnet/host WebUI and SSH access is permitted

It also regenerates the SSH keys that are used to secure communicating between the Primary and Secondary appliance. To start the script, at the console or via an SSH terminal session run the following command:

```
lbsecure
```

#### Note

'root' user console and SSH password access are disabled by default. These options can be

enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

The following image illustrates how the script works for a single appliance:

```
[root@lbmaster ~]# lbsecure

Loadbalancer.org security lock-down

This script enhances the security of a single or high-availability pair of
load balancers.

You will be asked to provide new passwords for the web interface and the
console root account, plus an IP subnet that should be allowed remote access
to the load balancer's web interface and ssh console.

Please enter a new password for the web interface 'loadbalancer' user. The
password will not be displayed as you type.
    New web interface password:
    Confirm password:

Please enter a new password for the console 'root' user. The password will not
be displayed as you type.
This password will also be used for the console 'setup' user.
    New console password:
    Confirm password:

Please enter an IP subnet that should be allowed remote access to the web
interface and ssh console.
Note that any host outside of this subnet will immediately lose access to the
load balancer. If you are running this script remotely, that includes the
current console.
    Administration subnet: 192.168.64.0/18

Working...
    Generating new SSH keys...
    SSH keys replaced.

    Setting web interface password...
    Setting console root password on local machine...
    Setting console 'setup' password on local machine...
    Passwords set.

    Setting up firewall...
    Firewall enabled.

Security enhancement complete.
```

Once the script has finished, the "Security enhancement complete" message is displayed as shown above.

**Note**

If `lbsecure` is run on the Primary of a correctly configured HA pair, the passwords, firewall rules and SSH keys will also be updated on the Secondary appliance.

You should run `lbsecure` **after** configuring the HA pair to ensure the correct HA related ports are configured in the firewall rules. To reverse the action of `lbsecure`, the command `ibinsecure` can be used.

For a clustered pair, run *lbinsecure* on both Primary and Secondary to completely reverse the configuration applied by running *lbsecure*.

## SSH Keys

This menu option enables SSH keys to be managed.

### Note

Since SSH keys are managed by the appliance and under normal circumstances do not require user intervention, this menu option will not normally be used.

To view/manage SSH keys:

1. Using the WebUI, navigate to: *Local Configuration > SSH Keys*.

SSH Keys

SSH Authentication

Host Keys ?

Create new key pair

Upload key pair

Type	Length (bits)	Date	
DSA	1024	2020-01-15 16:44	<div>Delete</div> <div>Download public key</div>
ECDSA	256	2023-01-20 17:38	<div>Delete</div> <div>Download public key</div>
ED25519	256	2023-01-20 17:38	<div>Delete</div> <div>Download public key</div>
RSA	2048	2020-01-15 16:44	<div>Delete</div> <div>Download public key</div>

User Keys ?

Create new key pair

Upload key pair

Username	Type	Length (bits)	Date	
root	RSA	2048	2020-01-15 16:44	<div>Delete</div> <div>Download public key</div>

Synchronise keys with peer

- The first tab (SSH Keys) enables the following keys to be viewed & managed:
  - **Host Keys** - the host identification key(s) of the local host

- **User Keys** - the public key(s) of the user presented to remote hosts
- The second tab (SSH Authentication) enables the following keys to be viewed & managed:
  - **Host Keys (known\_hosts)** - the known key(s) of hosts that have been previously connected to or have been preconfigured. In an HA pair you will see the peer appliance keys.
  - **User Keys (authorized\_keys)** - the public key(s) of remote hosts that can log in as the specified user. In an HA pair you will see the peer appliance keys.

## Appliance Configuration Files & Locations

The various configuration files used by the appliance are listed in the table below:

Configuration	File & Location
Network	/etc/sysconfig/network-scripts/ifcfg-eth*
Firewall	/etc/rc.d/rc.firewall
Firewall Lockdown Wizard	/etc/rc.d/rc.lockdownwizard.conf
System XML File	/etc/loadbalancer.org/lb_config.xml
Layer 4	/etc/ha.d/conf/loadbalancer.cf
Layer 7	/etc/haproxy/haproxy.cfg
Layer 7 (manual)	/etc/haproxy/haproxy_manual.cfg
Pound SSL	/etc/pound/pound.cfg
STunnel SSL	/etc/stunnel/stunnel.conf
SSL Certificates	/etc/loadbalancer.org/certs
Heartbeat	/etc/ha.d/ha.cf
Heartbeat Resources	/etc/ha.d/haresources
GSLB	/opt/polaris/etc/polaris-lb.yaml
GSLB Topology	/opt/polaris/etc/polaris-topology.yaml
WAF	/etc/httpd/waf.conf.d/90-wafs.conf
SNMP	/etc/snmp/snmpd.conf

# Chapter 6 - Configuring Load Balanced Services

## Introduction

As discussed [here](#), a fundamental choice when setting up load balanced services, is whether to configure the services at layer 4 or at Layer 7.

## Layer 4 Services

### The Basics

Layer 4 services are based on LVS (*Linux Virtual Server*). LVS implements transport layer load balancing inside the Linux kernel. It is used to direct requests for TCP/UDP based services to the Real Servers, and makes services on the Real Servers appear as a Virtual Service on a single IP address.

With the exception of Layer 4 SNAT mode, Layer 4 services are transparent by default, i.e. the source IP address is maintained through the load balancer.

Layer 4 persistence is based on source IP address by default. The time out value is in seconds and each time the client makes a connection the timer is reset, so even a 5 minute persistence setting could last for hours if the client is active and regularly refreshes their connection.

When a VIP is added the load balancer automatically adds a corresponding floating IP address which is activated instantly. Check *View Configuration > Network Configuration* to ensure that the floating IP address has been activated correctly. They will show up as secondary addresses/aliases.

Multiple ports can be defined per VIP, for example 80 & 443. In this case persistence is useful to ensure that clients hit the same backend server for both HTTP & HTTPS traffic and also to prevent the client having to renegotiate the SSL connection.

#### Note

It's not possible to configure a VIP on the same IP address as any of the network interfaces. This ensures services can 'float' (move) between Primary and Secondary appliances when using an HA Pair.

## Creating Layer 4 Virtual Services

Virtual services (VIPs) can be configured either by creating a new VIP and configuring all required settings, or by using the duplicate VIP feature.

Each Virtual Service can have an unlimited number of Real Servers. Typically you'll need one Virtual Service for each distinct cluster (group of load balanced servers). For example, you'd create a VIP for a web cluster, another for an FTP cluster and a third for a SIP cluster. Multiple ports can also be specified for each VIP.

### Defining a New Layer 4 VIP

*To add a new layer 4 VIP:*

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services*.
2. Click **Add a new Virtual Service**.

## Layer 4 - Add a new Virtual Service

Virtual Service		
Label	<input type="text" value="VIP Name"/>	<a href="#">?</a>
IP Address	<input type="text" value="10.0.0.20"/>	<a href="#">?</a>
Ports	<input type="text" value="80"/>	<a href="#">?</a>
Protocol		
Protocol	<input type="text" value="TCP"/>	<a href="#">?</a>
Forwarding		
Forwarding Method	<input type="text" value="NAT"/>	<a href="#">?</a>

[Cancel](#) [Update](#)

3. Enter an appropriate *Label* (name) for the new Virtual Service.
4. Enter the required IP address in the *Virtual Service IP address* field.
5. Enter the required port(s) in the *Virtual Service Ports* field, separate multiple ports with commas, specify a range with a hyphen and specify all ports using an asterisk (\*).

### Note

Several ports are used by the appliance and therefore cannot be used for Virtual Services. For full details please refer to [Ports Used by the Appliance](#).

6. Select the required *Protocol*:
  - **TCP** - Transmission Control Protocol is the default and most common option
  - **UDP** - User Datagram Protocol - used for DNS, SIP, etc.
  - **TCP/UDP** - enable both TCP and UDP on the port(s) specified
  - **One Packet Scheduling** - used for UDP SIP connections
  - **Firewall Marks** - For use when traffic has been tagged in the firewall script using the MARK target
7. Select the required *Forwarding Method*:
  - **Direct Routing (DR)** - This is the default mode for new Layer 4 VIPs. To use this mode, the **ARP Problem** must be solved on each Real Server. For more information on DR mode, please refer to [Layer 4 DR Mode](#).
  - **NAT** - With this mode, the Real Server's default gateway must be changed to be the load balancer. Because the load balancer also handles the return traffic, NAT mode is slower than DR mode. For more information on NAT mode please refer to [Layer 4 NAT Mode](#).
  - **Tunneling** - This is for WAN links (Tunneling). Tunneling has somewhat limited use as it requires an IP tunnel between the load balancer and the Real Server as the VIP is the target address many routers will drop the packet assuming that it has been spoofed. However, it is useful for private networks with Real Servers on multiple subnets.
  - **SNAT** - The mode requires no Real Server changes but is not as fast as DR mode. Also it's non transparent and therefore loses the client source IP information. You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict. For more information on SNAT mode please refer to [Layer 4 SNAT Mode](#).



8. Click **Update**

9. Now proceed to define the RIPs (Real Servers) as described [here](#).

### Duplicating an Existing Layer 4 VIP

If you have existing Virtual Services, these can be duplicated using the *Duplicate Service* feature.

#### Note

This option will copy all Virtual Service settings along with all associated Real Servers. After duplicating, you'll need to change either the IP address or port. If this is not done, the new VIP will clash with the original VIP and will not load. All other settings can remain the same if required.

*To duplicate an existing layer 4 VIP:*

1. Click **Modify** next to the VIP you'd like to duplicate.
2. Click the **Duplicate Service** button.
3. Click **OK** at the prompt to confirm you want to duplicate the VIP.
4. The VIP will be duplicated with a new label , all other settings will be identical.
5. Change the *IP Address*, *Port* and any other setting to suit your requirements.
6. Click **Update**.

### Modifying a Layer 4 VIP

When first adding a Virtual Service, only certain settings can be configured, others are set at their default value to simplify initial configuration. These values can be changed after the Virtual Service has been created by clicking **Modify** next to the relevant Virtual Service. Additional settings that can be changed are:

Section	Setting	Description
Connection Distribution Method	Balance Mode	<p>Select the required method to distribute new connections. The options are:</p> <ul style="list-style-type: none"><li>• <b>Weighted Least-Connection</b> - assign more jobs to servers with fewer jobs, relative to the Real Server's weight (the default).</li><li>• <b>Weighted Round Robin</b> - assign jobs to Real Servers proportionally to the Real Server's weight. Servers with higher weights receive new jobs first and get more jobs than servers with lower weights. Servers with equal weights get an equal distribution of new jobs.</li><li>• <b>Destination Hash</b> - assign jobs to servers through looking up a statically assigned hash table by their destination IP addresses. This algorithm is designed for use with web proxies and is supported with Layer 4 DR mode Virtual Services only.</li></ul> <div><div>Note</div><p>When using this mode, the web proxy servers must be configured in transparent mode as the destination remains set as the page a user requested. If the web proxy servers are configured in explicit/routed mode the destination will become the VIP. If the VIP is configured in either NAT or SNAT mode, the destination will be altered when the traffic is DNAT'ed flowing through the load balancer.</p></div>

Section	Setting	Description
Persistence	Enable	<p>Enable (the default) or disable persistence.</p> <p>Sticky or persistent connections are required for some protocols such as FTP and SIP. It is also kind to clients when using SSL and is sometimes required with HTTP if your web application cannot keep state between real servers.</p> <div> <div>Note</div> <div>If <i>Protocol</i> for the Virtual Service is set to 'One Packet Scheduling', persistence will be based on SIP Call-ID.</div> </div> <div> <div>Note</div> <div>If your Real Servers cannot keep session state persistence themselves, then you will obtain performance but not reliability benefits from a load balancer.</div> </div>
	Timeout	How long do you want connections to be sticky? The persistence time is in seconds and is reset on every connection. By default this is set to 300s (5 mins). Persistence will last for ever if the client clicks on a link within that period.
	Granularity	Specify the granularity with which clients are grouped for persistent virtual services. The source address of the request is masked with this netmask to direct all clients from a network to the same real server. The default is 255.255.255.255, that is, the persistence granularity is per client host. Less specific netmasks may be used to resolve problems with non-persistent cache clusters on the client side.
Health Checks	Check Type	<p>Specify the type of health check to be performed on the Real Servers. As the Check Type drop-down is changed, the related field list changes. The options are:</p> <ul style="list-style-type: none"> <li>• <b>Negotiate</b> - Scan the page specified in <i>Request to Send</i>, and check the returned data for the <i>Response Expected</i> string.</li> <li>• <b>Connect to port</b> - Attempt to make a connection to the specified port.</li> <li>• <b>Ping Server</b> - Use a simple ICMP ping to perform health checks.</li> <li>• <b>External script</b> - Use a custom file for the health check. For more information please refer to <a href="#">External Health Check Scripts</a></li> <li>• <b>No checks, always off</b> - all Real Servers are marked offline.</li> <li>• <b>No checks, always on</b> - all Real Servers are marked online.</li> <li>• <b>5 Connects, 1 Negotiate</b> - Repeating pattern of 5 Connect checks followed by 1 Negotiate check.</li> <li>• <b>10 Connects, 1 Negotiate</b> - Repeating pattern of 10 Connect checks followed by 1 Negotiate check.</li> </ul> <div> <div>Note</div> <div>For full details of all layer 4 health check options, please refer to <a href="#">Health Checks for Layer 4 Services</a>.</div> </div>

Section	Setting	Description
	Check Port	If you want the check port to be different to the port specified for the VIP, set it here. For a multi-port VIP, by default the first port in the list will be used as the check port.
Feedback	Feedback Method	<p>The method the load balancer uses to measure to performance of the Real Servers. The options are:</p> <ul style="list-style-type: none"> <li>• <b>Agent</b> - A simple telnet to port 3333 on the Real Server.</li> <li>• <b>HTTP</b> - A simple HTTP GET to port 3333 on the Real Server.</li> <li>• <b>None</b> - No feedback (default setting).</li> </ul> <p>The load balancer expects a 0-99 integer response from the agent, usually relating to the CPU idle; i.e. a response of 92 would imply that the Real Server's CPU is 92% idle. The load balancer will then use the formula <math>((92/10) * \text{requested\_weight})</math> to find the new weight. Using this method an idle Real Server will get 10 times as many new connections as an overloaded server.</p>
Fallback Server	IP Address	<p>The server to route to if all of the Real Servers in the group fail the health check. The local Nginx fallback server is configured for the ports 80 and 9081 (configured to always show the index.html page).</p> <p>When using HAProxy Layer 7 the Nginx server port 80 is automatically disabled. You can also configure the fallback server to be a 'Hot Spare' if required.</p> <p>For example you have one server in the cluster and one fallback they will act as a Primary/Secondary pair.</p>
	Port	Set the fallback server port, for DR mode leave this blank as it must be the same as the VIP.
	MASQ Fallback	Masquerade fallback. When enabled, this enables the fallback server to be set as a Layer 7 Virtual Service. This is especially useful in WAN/DR site environments.
	Email Alert Destination Address	<p>Destination email address for server health check notifications.</p> <div> <div>Note</div> <div>For more information on configuring email alerts, please refer to <a href="#">Configuring Email Alerts for Virtual Services</a>.</div> </div>

**Note** | If you require a custom gateway for a particular VIP, this can be achieved using Policy Based Routing. For more information on using and configuring PBR please refer to [Policy Based Routing \(PBR\)](#).

## Creating Layer 4 Real Servers (RIPs)

You can add an unlimited number of Real Servers to each Virtual Service. In DR mode, since port redirection is not possible the Real Server port field is not available and the port is automatically set to be the same as the Virtual Service, whilst for a NAT mode Real Server, it's possible to configure the port to be the same or different to the Virtual Service's port.

To add a new layer 4 RIP:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Real Servers*.
2. Click **Add a new Real Server** next to the relevant Virtual Service.

Label	<input type="text" value="RIP Name"/>	?
Real Server IP Address	<input type="text" value="IPAddress"/>	?
Real Server Port	<input type="text"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

3. Enter an appropriate *Label* (name) for the new Real Server.
4. Enter the required IP address in the *Real Server IP Address* field.
5. Enter the required port in the *Real Server Port* field. This only applies to NAT mode, in DR mode port redirection is not possible so by default the port is the same as defined in the VIP.
6. Specify the required *Weight*, this is an integer specifying the capacity of a server relative to the others in the pool, valid values are 0 to 65535, the default is 100. The higher the value, the more connections the server will receive. If the weight is set to 0, the server will effectively be placed in drain mode.
7. Specify the *Minimum Connections*, this is an integer specifying the lower connection threshold of a server. The valid values are 0 through to 65535. The default is 0, which means the lower connection threshold is not set.
8. If Minimum Connections is set with other values, the server will receive new connections when the number of its connections drops below its lower connection threshold. If Minimum Connections is not set but Maximum Connections is set, the server will receive new connections when the number of its connections drops below three fourths of its upper connection threshold.
9. Specify the *Maximum Connections*, this is an integer specifying the upper connection threshold of a server. The valid values of Maximum Connections are 0 through to 65535. The default is 0, which means the upper connection threshold is not set.

## DR Mode Considerations

### The ARP Problem

DR mode works by changing the MAC address of the inbound packets to match the Real Server selected by the load balancing algorithm. To enable DR mode to operate:

1. Each Real Server must be configured to accept packets destined for both the VIP address **and** the Real Server's IP address (RIP). This is because in DR mode the destination address of load balanced packets is the VIP address, whilst for other traffic such as health checks, administration traffic etc. it's the Real Server's own IP address (the RIP). The service/process (e.g. IIS) must also respond to both addresses.
2. Each Real Server must be configured so that it does not respond to ARP requests for the VIP address - only the load balancer should do this.

Configuring the Real Servers in this way is referred to as '*Solving the ARP Problem*'. The steps required depend on the OS used as detailed in the following sections.

### Detecting the ARP Problem

Attempt to connect to the VIP and then use *Reports > Layer 4 Current Connections* to check whether the connection state is SYN\_RECV as shown below.

## Layer 4 Current Connections

Check Status

### IPVS connection entries

pro	expire	state	source	virtual	destination
TCP	00:26	SYN_RECV	192.168.64.7:20415	192.168.111.232:80	192.168.110.240:80
TCP	00:26	SYN_RECV	192.168.64.7:20414	192.168.111.232:80	192.168.110.240:80
TCP	04:18	NONE	192.168.64.7:0	192.168.111.232:80	192.168.110.240:80

If it is, this is normally a good indication that the *ARP Problem* has not been correctly solved.

### Solving the ARP Problem for Linux

There are two different approaches on how to configure a Linux server for correct operation when DR mode load balancing is in use:

- Modifying the server's ARP behaviour and adding the relevant VIP addresses to the loopback interface
- Using NAT to convince the server to accept and reply to packets addressed to the relevant VIP addresses

Four independent methods are described below along with instructions. Each method follows one of the two approaches above. The specific method chosen will depend on technical requirements, the Linux distribution in use, and personal preferences.

The first method involves setting kernel parameters to alter the server's ARP behaviour and adding IP addresses to the loopback interface. This method should be universally applicable to any Linux server **making this the preferred method**.

If setting kernel parameters and adding IP addresses is not possible for some reason, the remaining three methods describe setting up a server for DR mode operation by using NAT via the **redirect** target/statement. The specific instructions depend on the packet filtering framework and tooling in use, which varies between Linux distributions. Methods are presented for iptables, nftables, and the `firewall-cmd` tool.

#### Method 1: ARP Behaviour and Loopback Interface Changes

This is the preferred method as it should be applicable to any Linux server and doesn't require any additional packet filtering or NAT considerations.

Each real server needs the loopback interface to be configured with the virtual IP addresses (VIPs) of the relevant load balanced services. This is often just a single VIP address, but the logic described below can be extended to cover multiple VIPs on a server. Having the VIPs on the loopback interface allows the server to accept inbound load balanced packets that are addressed to a VIP.

The server **must not** respond to ARP requests for the VIP addresses. The server also **must not** use ARP to announce the fact that it owns the VIP addresses. This is necessary to prevent IP address conflicts, as *all* of the real servers *and* the load balancer will own the VIP addresses. Only the load balancer should announce ownership of the VIPs.

To configure the behaviour described above, follow all of the steps below on each real server.

#### Step 1 of 4: Re-configuring ARP behaviour

**This step is only applicable if IPv4-based virtual services are in use.**

Add the following lines to the file `/etc/sysctl.conf` (create this file if it does not already exist):

```
net.ipv4.conf.all.arp_ignore=1
net.ipv4.conf.eth0.arp_ignore=1
net.ipv4.conf.eth1.arp_ignore=1
net.ipv4.conf.all.arp_announce=2
net.ipv4.conf.eth0.arp_announce=2
net.ipv4.conf.eth1.arp_announce=2
```

Adjust the commands shown above to suit the server's network configuration, e.g. a different number of network interfaces or a different interface naming convention.

#### Note

For reference, the effect of these kernel parameter changes on the server is as follows:

- `arp_ignore=1`: This configures the server to only reply to an ARP request if the request's target IP address is local to the incoming interface. This can never be true for VIP addresses on the loopback interface, as the loopback interface can never be an incoming interface for ARP requests from other devices. Hence, ARP requests for VIP addresses are always ignored.
- `arp_announce=2`: This prevents the server from sending an ARP request out of an interface *A* where the ARP request's sender/source address is stated to be an IP address that is local to some other interface *B*. For example, this prevents the server from sending an ARP request *from* a VIP address (which is local to the loopback interface) out of `eth0`, which would announce that the server owns the VIP address.

#### Step 2 of 4: Re-configuring duplicate address detection (DAD) behaviour

**This step is only applicable if IPv6-based virtual services are in use.**

Add the following lines to the file `/etc/sysctl.conf` (create this file if it does not already exist):

```
net.ipv6.conf.lo.dad_transmits=0
net.ipv6.conf.lo.accept_dad=0
```

#### Note

For reference, the effect of these kernel parameter changes on the server is as follows:

- `dad_transmits=0`: This prevents a given interface from sending out duplicate address detection probes in order to test the uniqueness of unicast IPv6 addresses. Any IPv6 VIP addresses will *not* be unique, so this mechanism is disabled.
- `accept_dad=0`: This prevents a given interface from accepting duplicate address detection messages. This prevents any IPv6 VIP addresses from being marked as duplicate addresses.

### Step 3 of 4: Applying the new settings

To apply the new settings, either reboot the real server or execute the following command to immediately apply the changes:

```
/sbin/sysctl -p
```

Steps 1, 2, and 3 can be replaced by instead modifying the necessary kernel variables by writing directly to their corresponding files under `/proc/sys/`. Note that changes made in this way *will not persist across reboots*.

Execute the following commands (as root) to implement these temporary changes (adapting the number of interfaces and interface names as needed):

#### Note

```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_ignore
echo 1 > /proc/sys/net/ipv4/conf/eth0/arp_ignore
echo 1 > /proc/sys/net/ipv4/conf/eth1/arp_ignore
echo 2 > /proc/sys/net/ipv4/conf/all/arp_announce
echo 2 > /proc/sys/net/ipv4/conf/eth0/arp_announce
echo 2 > /proc/sys/net/ipv4/conf/eth1/arp_announce
echo 0 > /proc/sys/net/ipv6/conf/lo/dad_transmits
echo 0 > /proc/sys/net/ipv6/conf/lo/accept_dad
```

### Step 4 of 4: Adding the virtual IP addresses (VIPs) to the loopback interface

Each of the VIP addresses must be permanently added to the loopback interface. VIPs must be added with a network prefix of `/32` for IPv4 addresses or `/128` for IPv6 addresses. The IP addresses can be added using the usual configuration files and tools for modifying network interfaces, which vary between different Linux distributions.

**As an alternative**, the `ip` command can be used as a universal way to add IP addresses to any Linux server. Note that addresses added in this way *will not persist across reboots*. To make these addresses permanent, add the `ip` commands to an appropriate startup script such as `/etc/rc.local`.

Execute the following `ip` command for each IPv4 VIP:

```
ip addr add dev lo <IPv4-VIP>/32
```

Execute the following `ip` command for each IPv6 VIP:

```
ip addr add dev lo <IPv6-VIP>/128
```

To check that the VIPs have been successfully added, execute the command:

```
ip addr ls
```

To remove an IPv4 VIP from the loopback adapter, execute the command:

```
ip addr del dev lo <IPv4-VIP>/32
```

To remove an IPv6 VIP from the loopback adapter, execute the command:

```
ip addr del dev lo <IPv6-VIP>/128
```

#### Method 2: NAT "redirect" via iptables

iptables can be used on each real server to identify incoming packets that are addressed to a virtual IP address (VIP) and redirect those packets to the server itself. This is achieved using the **REDIRECT** target in iptables, which performs the necessary NAT to make this possible. This allows a real server to accept packets addressed to a VIP without the server owning the VIP.

Execute the following command to put the necessary iptables rule in place to redirect traffic for a single IPv4 VIP address. Note that iptables rules added in this way *will not persist across reboots*. To make such a rule permanent, either add the rule to an iptables firewall script, if one is provided with the Linux distribution in question, or add the command to an appropriate startup script such as `/etc/rc.local` on each real server.

```
iptables -t nat -A PREROUTING -d <IPv4-VIP> -j REDIRECT
```

The VIP address should be changed to match the virtual service in question, for example:

```
iptables -t nat -A PREROUTING -d 10.0.0.21 -j REDIRECT
```

The example above will redirect any incoming packets destined for 10.0.0.21 (the virtual service) locally, i.e. to the primary address of the incoming interface on the real server.

If a real server is responsible for serving *multiple* VIPs then additional iptables rules should be added to cover each VIP.

For an IPv6 VIP address, a command like the following should be used:

```
ip6tables -t nat -A PREROUTING -d <IPv6-VIP> -j REDIRECT
```

The VIP address should be changed to match the virtual service in question, for example:

```
ip6tables -t nat -A PREROUTING -d 2001:db8::10 -j REDIRECT
```

#### Note

Method 2 may not be appropriate when using IP-based virtual hosting on a web server. This is because an iptables **REDIRECT** rule will redirect incoming packets to the *primary address* of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 1 instead.

#### Method 3: NAT "redirect" via nftables

nftables is the modern Linux kernel packet filtering framework. It is supported on all major Linux distributions and has replaced iptables as the default framework on most major distributions.

nftables can be used on each real server to identify incoming packets that are addressed to a virtual IP address (VIP) and redirect those packets to the server itself. This is achieved using the **redirect** statement in nftables, which performs the necessary NAT to make this possible. This allows a real server to accept packets addressed to a VIP without the server owning the VIP.



Use a script like the following to put the necessary nftables structures in place to redirect traffic for both IPv4 and IPv6 VIP addresses. To make such a configuration permanent, either add the `inet nat` table to an nftables firewall script, if one is provided with the Linux distribution in question, or configure a script like the following to execute as a startup script on each real server.

```
#!/usr/sbin/nft -f

table inet nat {
    chain prerouting {
        comment "Allow server to accept packets destined for VIP addresses";
        type nat hook prerouting priority -100; policy accept;
        ip daddr <IPv4-VIP> redirect comment "Description"
        ip6 daddr <IPv6-VIP> redirect comment "Description"
    }
}
```

The VIP addresses and comments should be changed to match the virtual services in question, for example:

```
#!/usr/sbin/nft -f

table inet nat {
    chain prerouting {
        comment "Allow server to accept packets destined for VIP addresses";
        type nat hook prerouting priority -100; policy accept;
        ip daddr 10.0.0.21 redirect comment "VIP 1: HTTP"
        ip6 daddr 2001:db8::10 redirect comment "VIP 2: HTTPS"
    }
}
```

The example above will redirect any incoming packets destined for 10.0.0.21 or 2001:db8::10 (the virtual services) locally, i.e. to the primary address of the incoming interface (for each IP version) on the real server.

Note that **Linux kernels prior to 5.2** may not support performing NAT (which is required for the **redirect** statement) in an `inet` family table. In this scenario, use either an `ip` or an `ip6` family table instead, or both if a mixture of IPv4 and IPv6 VIPs are in use on the same server. Also note that older kernels may not support the use of comments in chains.

Note that **Linux kernels prior to 4.18** require explicitly registering both prerouting and postrouting chains in order for the implicit NAT of the **redirect** statement to be correctly performed in both the inbound and outbound directions.

A legacy-friendly setup may look like the following:

```
#!/usr/sbin/nft -f

table ip nat {
    chain prerouting {
        type nat hook prerouting priority -100; policy accept;
        ip daddr 10.0.0.21 counter redirect comment "VIP 1: HTTP"
    }

    chain postrouting {
        type nat hook postrouting priority 100; policy accept;
    }
}

table ip6 nat {
    chain prerouting {
        type nat hook prerouting priority -100; policy accept;
        ip6 daddr 2001:db8::10 counter redirect comment "VIP 2: HTTPS"
    }
}
```

```
chain postrouting {
    type nat hook postrouting priority 100; policy accept;
}
```

#### Note

Method 3 may not be appropriate when using IP-based virtual hosting on a web server. This is because an nftables **redirect** statement will redirect incoming packets to the *primary address* of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 1 instead.

#### Method 4: NAT "redirect" via firewall-cmd

Some recent versions of Linux distributions make use of firewalld as a high-level firewall configuration framework. In this case, while it may actually be iptables performing the work at a lower level, it may be preferred to implement the iptables NAT solution described in [method 2](#) in firewalld, as opposed to directly manipulating iptables. This is achieved by using the `firewall-cmd` tool provided by firewalld and executing a command like the following on each real server:

```
firewall-cmd --permanent --direct --add-rule ipv4 nat PREROUTING 0 -d <IPv4-VIP> -j REDIRECT
```

The VIP address should be changed to match the virtual service in question, for example:

```
firewall-cmd --permanent --direct --add-rule ipv4 nat PREROUTING 0 -d 10.0.0.50 -j REDIRECT
```

To apply the new configuration, reload the firewall rules like so:

```
firewall-cmd --reload
```

Configuration applied in this way will be permanent and will persist across reboots.

#### Note

Method 4 may not be appropriate when using IP-based virtual hosting on a web server. This is because an iptables **REDIRECT** rule will redirect incoming packets to the *primary address* of the incoming interface on the web server rather than any of the virtual hosts that are configured. Where this is an issue, use method 1 instead.

### Solving the ARP Problem for Solaris

With Solaris the loopback interface does not respond to ARP requests so you just add your VIPs to it:

```
ifconfig lo0:1 plumb
ifconfig lo0:1 <VIP> netmask 255.255.255.255 up
```

1. You'll need to add this to the startup scripts on all of your Real Servers.

For Solaris v11 and later, a new command is used:

```
ipadm create-addr -a <VIP>/32 lo0
```

The configuration survives a reboot so there is no need to add this command to a startup script, just run it on each Real Server.

## Solving the ARP Problem for Mac OS X/BSD

OS X is BSDish, so you need to use BSDish syntax:

```
ifconfig lo0 alias <VIP> netmask 255.255.255.255 -arp up
```

You'll need to add this to the startup scripts on all of your Real Servers.

### Note

Don't forget that the service on the Real Servers needs to listen on both the RIP address and VIP address as mentioned previously.

### Note

Failure to correctly configure the Real Servers to handle the **ARP Problem** is the most common mistake in DR mode configurations.

## Solving the ARP Problem for Windows Servers

Windows Server 2012 & Later

Windows Server 2012 and later support Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter that must be installed and configured on each load balanced (Real) Server. The IP address configured on the Loopback Adapter must be the same as the Virtual Service (VIP) address.

### Note

If a Real Server is included in multiple DR mode VIPs, IP addresses for each VIP must be added to the Loopback Adapter.

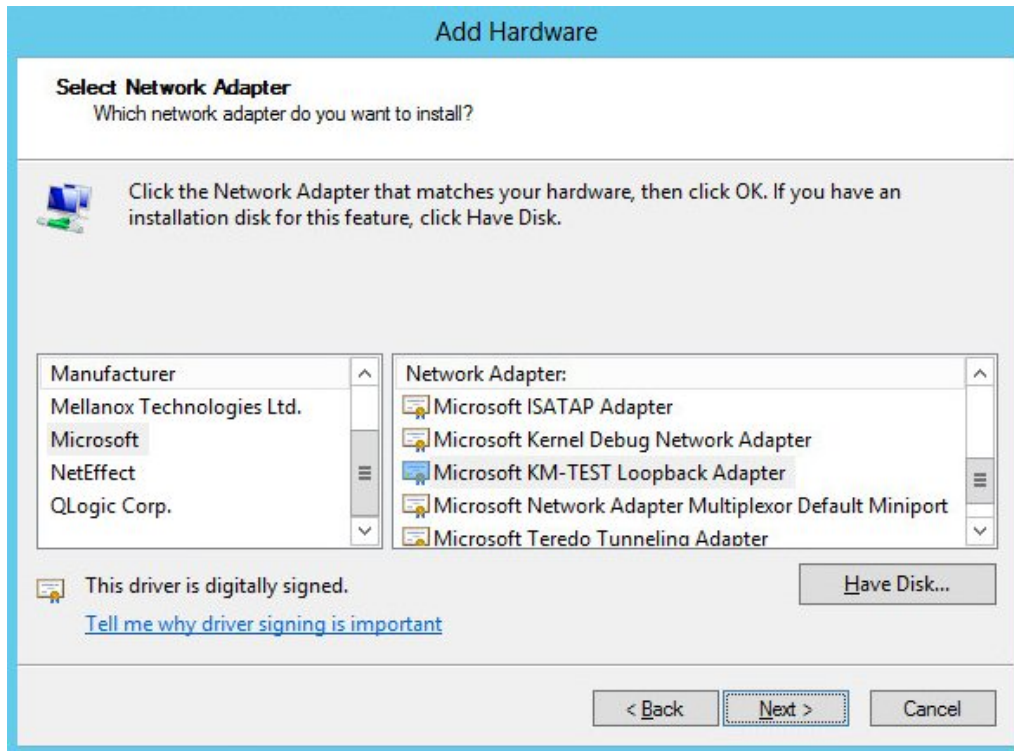
In addition, steps must be taken to set the strong/weak host behavior on each Real Server. This is used to either block or allow interfaces to receive packets destined for a different interface on the same server.

### Important

The following 3 steps must be completed on **each** Real Server.

### Step 1 of 3: Install the Microsoft Loopback Adapter

1. Click **Start**, then run **hdwwiz** to start the Hardware Installation Wizard.
2. Once the Wizard has started, click **Next**.
3. Select **Install the hardware that I manually select from a list (Advanced)**, click **Next**.
4. Select **Network adapters**, click **Next**.



5. Select **Microsoft & Microsoft KM-Test Loopback Adapter**, click **Next**.
6. Click **Next** to start the installation, when complete click **Finish**.

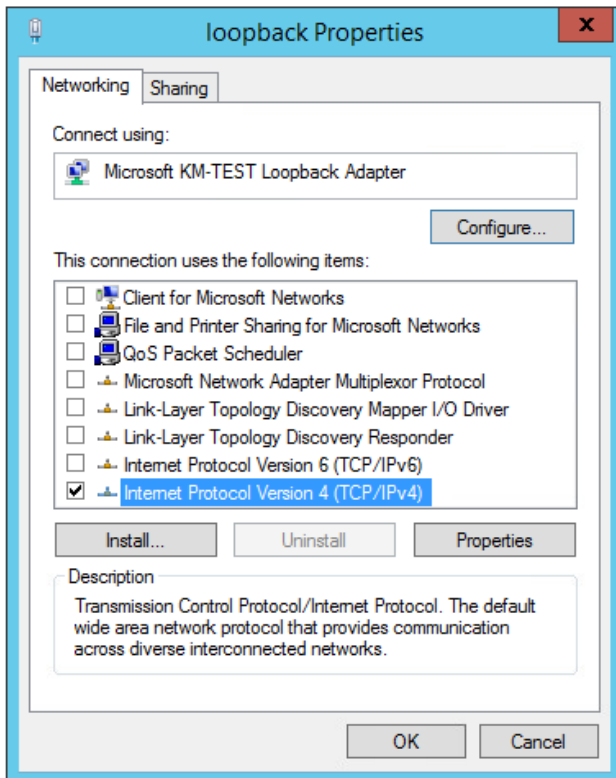
#### Step 2 of 3: Configure the Loopback Adapter

1. Open Control Panel and click **Network and Sharing Center**.
2. Click **Change adapter settings**.
3. Right-click the new Loopback Adapter and select **Properties**.

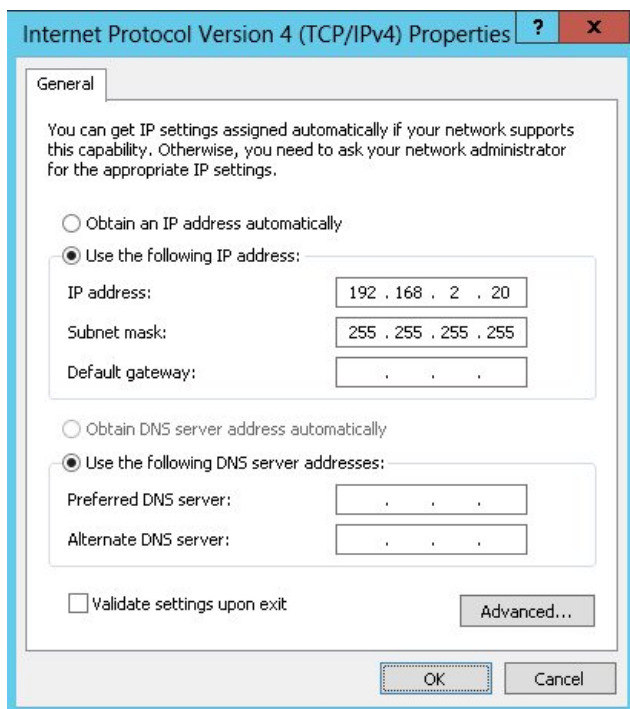
**Note** | You can configure IPv4 or IPv6 addresses or both depending on your requirements.

#### IPv4 Addresses

1. Uncheck all items except **Internet Protocol Version 4 (TCP/IPv4)** as shown below:



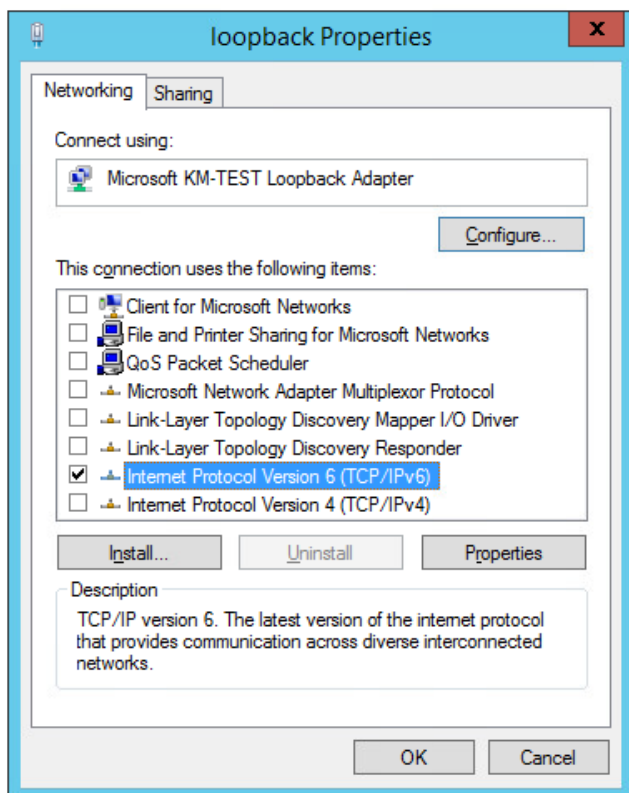
2. Ensure that **Internet Protocol Version (TCP/IPv4)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service address (VIP) with a subnet mask of **255.255.255.255**, e.g. **192.168.2.20/255.255.255.255** as shown below:



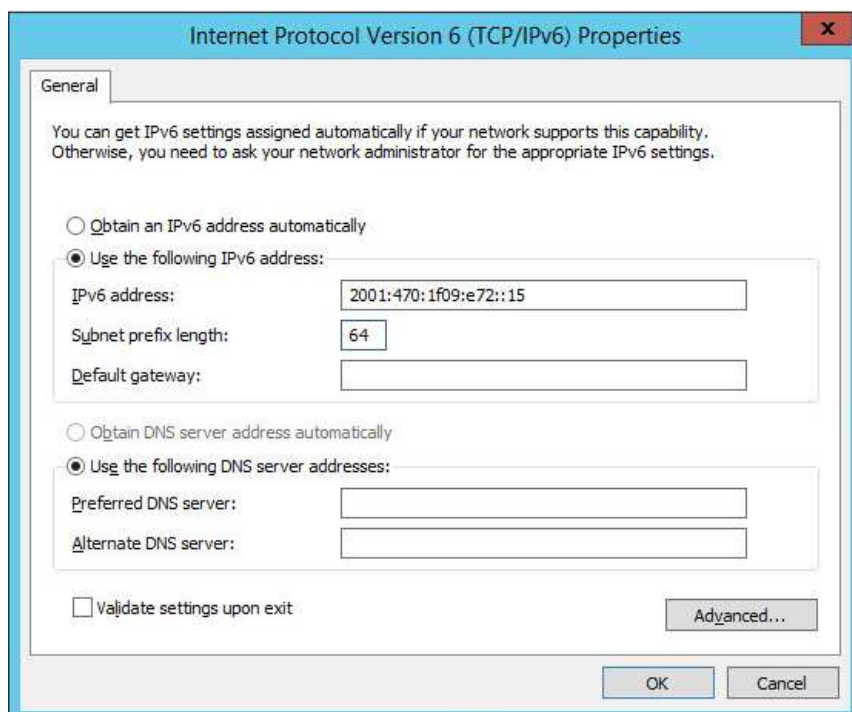
3. Click **OK** then click **Close** to save and apply the new settings.

## IPv6 Addresses

1. Uncheck all items except **Internet Protocol Version 6 (TCP/IPv6)** as shown below:



2. Ensure that **Internet Protocol Version (TCP/IPv6)** is selected, click **Properties** and configure the IP address to be the same as the Virtual Service (VIP) and set the *Subnet Prefix Length* to be the same as your network setting, e.g. **2001:470:1f09:e72::15/64** as shown below:



3. Click **OK** then click **Close** to save and apply the new settings.

### Step 3 of 3: Configure the strong/weak host behavior

Either Network Shell (**netsh**) commands or PowerShell cmdlets can be used to set the required strong/weak host behavior.

The commands in this section assume that the LAN Adapter is named "net" and the Loopback Adapter is named "loopback" as shown in the example below:



#### Important

Either adjust the commands to use the names allocated to your LAN and loopback adapters, or rename the adapters before running the commands. Names are case sensitive so make sure that the interface names used in the commands match the adapter names exactly.

#### Option 1 - Using Network Shell (netsh) Commands

To configure the correct strong/weak host behavior run the following commands:

For IPv4 addresses:

```
netsh interface ipv4 set interface "net" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostreceive=enabled
netsh interface ipv4 set interface "loopback" weakhostsend=enabled
```

For IPv6 addresses:

```
netsh interface ipv6 set interface "net" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostreceive=enabled
netsh interface ipv6 set interface "loopback" weakhostsend=enabled
netsh interface ipv6 set interface "loopback" dadtransmits=0
```

#### Option 2 - Using PowerShell Cmdlets

For IPv4 addresses:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled
-DadTransmits 0 -AddressFamily IPv4
```

```
Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled -AddressFamily IPv4
```

For IPv6 Addresses:

```
Set-NetIpInterface -InterfaceAlias loopback -WeakHostReceive enabled -WeakHostSend enabled
-DadTransmits 0 -AddressFamily IPv6
```

```
Set-NetIpInterface -InterfaceAlias net -WeakHostReceive enabled -AddressFamily IPv6
```

Verifying Strong/Weak Host Settings

To verify that settings have been configured correctly, run the following command on each Real Server to clearly list the settings that have been applied to the interface:

```
netsh interface ipv4 show interface <interface name>
```

For the 'loopback' adapter run:

```
netsh interface ipv4 show interface loopback
```

For the 'net' adapter run:

```
netsh interface ipv4 show interface net
```

The following image shows the output for the loopback adapter:

```
C:\Users\Administrator>netsh interface ipv4 show interface loopback
```

```
Interface loopback Parameters
-----
IfLuid           : ethernet_9
IfIndex          : 15
State            : connected
Metric           : 30
Link MTU         : 1500 bytes
Reachable Time   : 28500 ms
Base Reachable Time : 30000 ms
Retransmission Interval : 1000 ms
DAD Transmits    : 3
Site Prefix Length : 64
Site Id          : 1
Forwarding       : disabled
Advertising      : disabled
Neighbor Discovery : enabled
Neighbor Unreachability Detection : enabled
Router Discovery : dhcp
Managed Address Configuration : enabled
Other Stateful Configuration : enabled
Weak Host Sends  : enabled
Weak Host Receives : enabled
Use Automatic Metric : enabled
Ignore Default Routes : disabled
Advertised Router Lifetime : 1800 seconds
Advertise Default Route : disabled
Current Hop Limit : 0
Force ARPND Wake up patterns : disabled
Directed MAC Wake up patterns : disabled
```

```
C:\Users\Administrator>
```

**Note** | For IPv6, simply replace 'ipv4' with 'ipv6' in the above commands.

For Windows 2012 & later you can also use the following PowerShell Cmdlets to verify the settings:

To view both IPv4 and IPv6:

```
Get-NetIpInterface -InterfaceAlias loopback | FL
```

for IPv4 only:



```
Get-NetIpInterface -InterfaceAlias loopback -AddressFamily IPv4 | FL
```

for IPv6 only:

```
Get-NetIpInterface -InterfaceAlias loopback -AddressFamily IPv6 | FL
```

#### Note

Failure to correctly configure the Real Servers to handle the **ARP Problem** is the most common problem in DR configurations.

### Solving the ARP Problem - Possible Side Effect for Windows 2012 & Later

With DR Mode, the source IP address of return traffic from a Real Server will be the IP address assigned to the loopback adapter, which is the same as the VIP address that the client connected to. For traffic initiated by a Real Server, the source IP address should under normal circumstances be the Real Server's own IP address, i.e. the address assigned to the standard network adapter.

However, due to the way the network adapters are configured to solve the **ARP Problem**, and the way that Windows selects the source IP address, it's possible under certain circumstances for the source IP address of traffic initiated by a Real Server to be the IP address configured on the loopback adapter rather than the Real Server's own IP address. Please refer to [this Microsoft article](#) for more information on how Windows selects the source IP address.

To prevent the IP address(es) assigned to the loopback adapter being used in this way, the following two PowerShell commands should be run on each Windows 2008 R2 & later Real Server to set the **SkipAsSource** flag for all IPs assigned to the loopback adapter:

```
[array]$IPs = Get-NetIPAddress -InterfaceAlias loopback
```

```
Set-NetIPAddress -IPAddress $IPs.IPAddress -InterfaceAlias loopback -SkipAsSource $true
```

- the first command gathers all IP addresses assigned to the loopback adapter
- the second command then sets the **SkipAsSource** flag for all IPs found
- if your loopback adapter is not named 'loopback' modify the commands accordingly

To verify that the flag has been set for all IPs, the following PowerShell command can be used:

```
Get-NetIPAddress -InterfaceAlias loopback
```

#### Note

For more information about these commands, please refer to [this Microsoft article](#).

### Other Windows Settings that May Cause Issues

#### Receive Segment Coalescing (RSC)

RSC is a stateless offload technology that helps reduce CPU utilisation for network processing on the receive side by offloading tasks from the CPU to an RSC-capable network adapter. In rare cases it has been discovered that RSC can adversely effect performance when using DR mode. In these cases the performance issue was addressed by disabling RSC. This can be done in 2 ways:

1) using the NIC's advanced properties tab, disable the following settings:

**Recv Segment Coalescing (IPv4)**

**Recv Segment Coalescing (IPv6)**

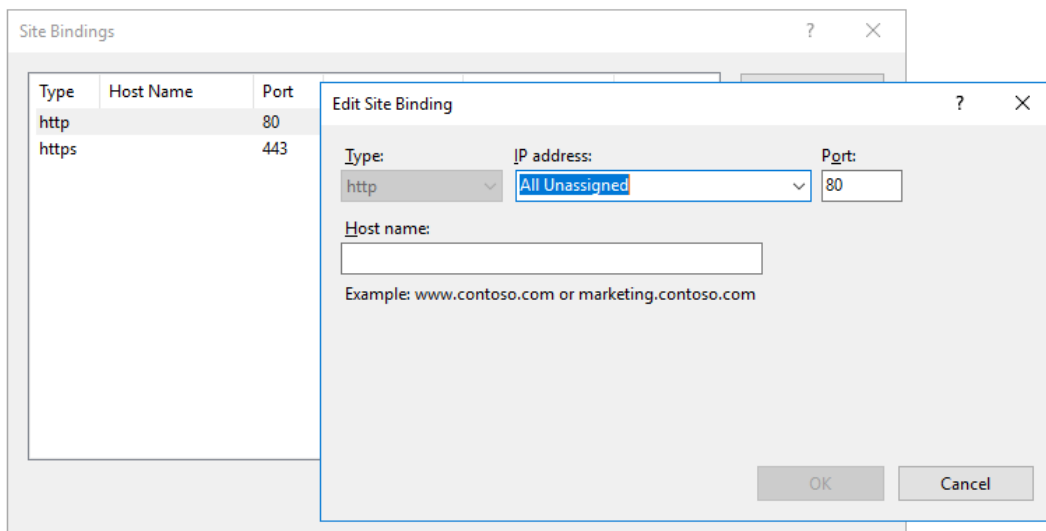
2) using PowerShell:

```
Disable-NetAdapterRsc -Name "MyAdapter" -IPv4  
Disable-NetAdapterRsc -Name "MyAdapter" -IPv6
```

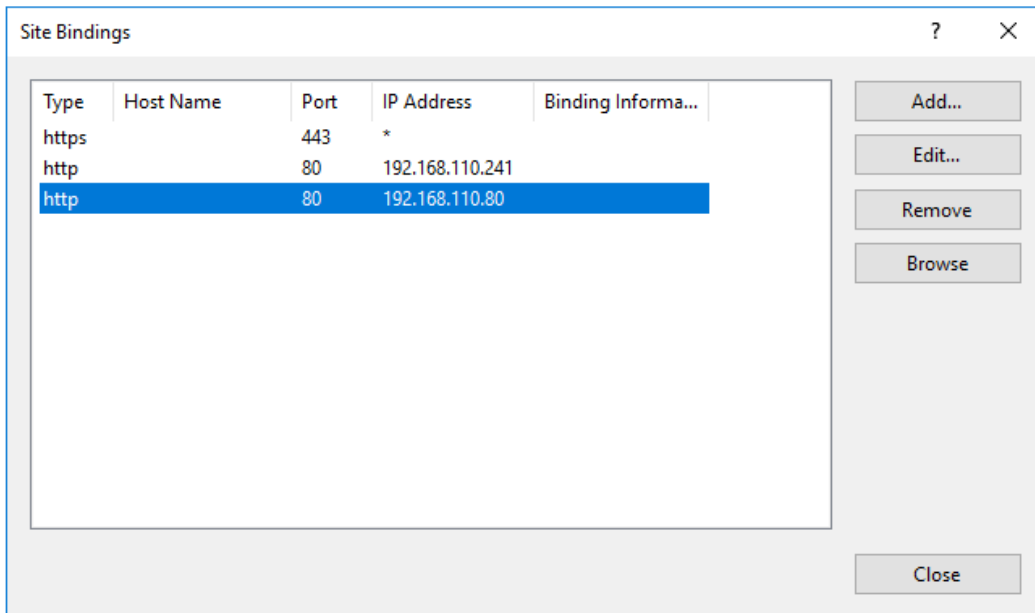
### Configuring Your Application to Respond to Both the RIP and VIP

For DR & TUN modes, it's important to make sure that your application (IIS in this example) responds to both the VIP and RIP.

By default, IIS listens on all configured IP addresses, this is shown in the example below (shows Windows 2016 example). As can be seen the IP address field is set to "All Unassigned".



If the default configuration is left, no further IIS configuration is required. If you do change the IP address in the bindings from "All Unassigned" to a specific IP address, then you need to make sure that you also add a binding for the Virtual Service IP address (VIP) as shown in the example below:



### Important

This example illustrates how IIS must be configured to ensure that its listening on both the RIP and VIP address. Remember that this applies to **ALL** applications when using DR mode.

## Windows Firewall Settings

When IIS is installed, default rules are created that allow inbound traffic on port 80 and 443. By default, these rules apply to the Domain, Public and Private profiles and all interface types. If the rules are changed, make sure that inbound traffic can still reach both the LAN interface and the loopback adapter.

## NAT Mode Considerations

Layer 4 NAT mode requires Real Server return traffic to pass back via the load balancer. This is achieved by setting the Real Server's default gateway to be the load balancer. For an HA Pair, an additional floating IP address should be used to allow failover. Whilst NAT mode is fairly straight forward, a few points need to be considered.

### NAT Mode Potential Issues

1. By default your Real Servers won't be able to access the Internet through the new default gateway (except when replying to requests made through the external VIP).
2. Non-load balanced services on the Real Servers (e.g. RDP for management access to Windows servers) will not be accessible since these have not been exposed via the load balancer.

### Enabling Real Server Internet Access Using Auto-NAT

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Advanced Configuration*.
2. Change Auto-NAT from **off** to the external interface being used - typically **eth1**.
3. Click **Update**.

This activates the **rc.nat** script that forces external network traffic to be MASQUERADED to and from the external network. The iptables masquerade rule that's used for this is shown below:

```
iptables -t nat -A POSTROUTING -o eth1 -j MASQUERADE
```

### Enabling Access to non Load-Balanced Services

If you want specific services to be exposed on your Real Servers you have two choices:

- Setup a Virtual Service with a single Real Server for each service **Or**
- Setup a floating IP address and individual SNAT/DNAT rules for each service as shown in the example below. These lines can be added to the firewall script using the WebUI menu option *Maintenance > Firewall Script*.

```
INT_ADDR="192.168.110.240"  
EXT_ADDR="10.200.110.240"  
iptables -t nat -A POSTROUTING -p tcp -s $INT_ADDR -j SNAT --to-source $EXT_ADDR  
iptables -t nat -A PREROUTING -p tcp -d $EXT_ADDR -j DNAT --to-destination $INT_ADDR
```

Once the above SNAT/DNAT rules have been configured, the following firewall entries will be listed under *View Configuration > Firewall Rules*:

```
Chain PREROUTING (policy ACCEPT 2 packets, 120 bytes)  
pkts bytes target prot opt in out source destination  
0 0 DNAT tcp -- * * 0.0.0.0/0 10.200.110.240 to:192.168.110.240  
Chain POSTROUTING (policy ACCEPT 1 packets, 60 bytes)  
pkts bytes target prot opt in out source destination  
0 0 SNAT tcp -- * * 192.168.110.240 0.0.0.0/0 to:10.200.110.240
```

**Note** | The default gateway on the Real Server must be an IP on the load balancer.

**Note** | If Autonat is already enabled, only the DNAT rule (i.e. not the SNAT rule) will be required.

**Note** | Please don't hesitate to contact [support@loadbalancer.org](mailto:support@loadbalancer.org) to discuss any specific requirements you may have.

### One-Arm (Single Subnet) NAT Mode

Normally the VIP is located on a different subnet to the Real Servers. However, it is possible to perform NAT mode load balancing on a single subnet where the VIP is brought up in the same subnet as the Real Servers. For clients located on this subnet, return traffic would normally be sent directly to the client bypassing the load balancer which would break NAT mode. To address this, the routing table on the Real Servers must be modified to force return traffic to pass via the load balancer. The sections below explain how routing must be modified for Windows hosts and Linux hosts.

#### Route Configuration for Windows Servers

To rectify this issue for Windows servers, a route must be added to each Real Server that takes priority over the default Windows routing rules. This is a simple case of deleting the default On-link route and adding a permanent route via the load balancer using the following commands on each real server:

```
netsh interface ipv4 delete route 192.168.2.0/24 "LAN"  
netsh interface ipv4 add route 192.168.2.0/24 "LAN" 192.168.2.21
```

**Note** | Ensure you specify your local subnet address. Replace "192.168.2.21" with the IP address of your load balancer. Replace "LAN" with the name of your Interface.

**Note** | After running the above commands, reboot the server and check if the updated routing rules have remained. Depending on the specific version of Windows, it may be necessary to add the commands to a startup script. This is because under certain circumstances routing rules for on-link, directly accessible addresses can get reset to defaults after a reboot.

Verify routing rules using the following command:

```
netsh interface ipv4 show route
```

#### Route Configuration for Linux Servers

To rectify this issue for Linux servers, we need to modify the local network route by changing to a higher metric:

```
route del -net 192.168.2.0 netmask 255.255.255.0 dev eth0
route add -net 192.168.2.0 netmask 255.255.255.0 metric 2000 dev eth0
```

**Note** | Ensure you specify your local subnet address.

Then we need to make sure that local network access uses the load balancer as its default route:

```
route add -net 192.168.2.0 netmask 255.255.255.0 gateway 192.168.2.21 metric 0 dev eth0
```

**Note** | Replace 192.168.2.0 & 255.255.255.0 with your local subnet address. Replace 192.168.2.21 with the IP address of your load balancer.






Any local traffic (same subnet) is then handled by the manual route and any external traffic is handled by the default route (which also points at the load balancer).

## Firewall Marks

Using firewall marks enables multiple ports and/or multiple IP addresses to be combined into a single Virtual Service. A common use of this feature is to aggregate port 80 (HTTP) and port 443 (HTTPS) so that when a client fills their shopping cart via HTTP, then moves to HTTPS to give their credit card information, they will remain on the same Real Server.

### Firewall Marks - Auto Configuration

When defining a layer 4 VIP with multiple ports, firewall marks are used automatically in the background to enable this functionality. For example, to configure an HTTP & HTTPS NAT mode Virtual Service, port 80 & 443 must be specified separated by a comma in the 'Virtual Service Ports' field as shown below:

Virtual Service		
Label	<input type="text" value="HTTP-Cluster"/>	
IP Address	<input type="text" value="192.168.115.100"/>	
Ports	<input type="text" value="80,443"/>	
Protocol		
Protocol	<input type="text" value="TCP"/>	
Forwarding		
Forwarding Method	<input type="text" value="NAT"/>	
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

This will automatically configure the load balancer for firewall marks.

**Note** | Persistence will be enabled automatically.

For NAT mode VIPs, leave the *Real Server Port* field blank as shown below:

Label	<input type="text" value="IIS1"/>	?
Real Server IP Address	<input type="text" value="192.168.30.22"/>	?
Real Server Port	<input type="text"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

CancelUpdate

Packets will then be forwarded to the Real Servers on the same port as received by the VIP.

**Note** | For Layer 4 DR mode VIPs, there is no Real Server Port field since port translation is not possible in this mode. Packets will be forwarded to the same port as specified for the VIP.

**Note** | To create an auto firewall mark VIP that listens on **all ports**, simply specify \* in the ports field rather than a specific port number.

**Note** | The Health check port is automatically set to be the first port in the list, e.g. if ports 80 & 443 are defined for the VIP, the check port is automatically set to port 80. This can be changed if required using the *Check Port* field.






## Firewall Marks - Manual Configuration

Firewall Marks can also be configured manually. The basic concept is to create a firewall rule that matches incoming packets to a particular IP address/port and mark them with an arbitrary integer. A Virtual Service is also configured specifying this firewall mark integer instead of the IP address.

### EXAMPLE 1 - Setup a new DR Mode Firewall Mark when no Initial VIP has been Created

#### Step 1: Create the New VIP

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services*.
2. Click **Add a new Virtual Service**.

Virtual Service		
Label	<input type="text" value="Cluster-1"/>	
Firewall Mark Identifier	<input type="text" value="1"/>	
Ports	<input type="text" value="80"/>	
Protocol		
Protocol	<input type="text" value="Firewall Marks"/>	
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Define the required *Label* (name) for the VIP, e.g. **Cluster-1**.
4. Instead of entering an IP address, enter a numeric value, e.g. **1** - this is the numeric reference for the Firewall Mark, this reference is used in step 5 below when defining the firewall rules.
5. The *Virtual Service Ports* field does not need to be set as it is not relevant in this case - the actual port(s) used are defined in the firewall script in step 5 below.
6. Set *Protocol* to **Firewall Marks** - at this point the *Virtual Service Ports* field will be grayed out and the Virtual Service *IP Address* field will be renamed as *Firewall Mark Identifier* as shown above.
7. Click **Update**.

**Note** | Persistence will be enabled automatically.

## Step 2: Define a health check Port

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services*.
2. Click **Modify** next to the new Virtual Service.
3. Enter the appropriate value in the *Check Port* field.
4. Click **Update**.

## Step 3: Add the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Real Servers*.
2. Click **Add a new Real Server**.
3. Enter the required details as shown below.

Label	Server1	?
Real Server IP Address	192.168.111.241	?
Weight	100	?
Minimum Connections	0	?
Maximum Connections	0	?

Cancel
Update

4. Click **Update**.

#### Step 4: Add the Associated Floating IP Address for the VIP

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IPs*.
2. Add a floating IP that corresponds to the required VIP, in this example **192.168.111.240**.

New Floating IP

192.168.111.240

Add Floating IP

3. Click **Add Floating IP**.

#### Step 5: Modify the Firewall Script

1. Using the WebUI, navigate to: *Maintenance > Firewall Script*.
2. Scroll down to the Manual Firewall Marks section and add the following lines as shown below:

```
VIP1="192.168.111.240"
iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 8025 -j MARK --set-mark 1
iptables -t mangle -A PREROUTING -p udp -d $VIP1 --dport 8025 -j MARK --set-mark 1
```



## Firewall Script

```
32 ##### Manual Firewall Marks #####
33
34 # Example: Associate HTTP and HTTPS with Firewall Mark 1:
35 #VIP1="10.0.0.66"
36 #iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 80 -j MARK --set-mark 1
37 #iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 443 -j MARK --set-mark 1
38
39 # A Virtual Service may then be created in the web interface, using 1 as the
40 # service address.
41
42 #It is also possible to bind TCP and UDP protocols together with a firewall mark.
43 #VIP1="192.168.64.27"
44 #iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 80 -j MARK --set-mark 1
45 #iptables -t mangle -A PREROUTING -p udp -d $VIP1 --dport 300 -j MARK --set-mark 1
46
47 VIP1="192.168.111.240"
48 iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 8025 -j MARK --set-mark 1
49 iptables -t mangle -A PREROUTING -p udp -d $VIP1 --dport 8025 -j MARK --set-mark 1
50
51 ##### Packet Filtering #####
52
53 # You should always use a network perimeter firewall to lock down all
54 # external access to the load balancer except the required Virtual Services
55 # and the required services from your admin machine / network (SSH & HTTPS)
56
57 # Allow unlimited traffic on the loopback interface:
58 #iptables -A INPUT -i lo -j ACCEPT
59 #iptables -A OUTPUT -o lo -j ACCEPT
60
61
62
63
```

Update

3. Click **Update**.
4. For a clustered pair, make the same changes to the firewall script on the Secondary unit.





*The VIP is now configured and will be accessible on 192.168.111.240, TCP & UDP port 8025*

### EXAMPLE 2 - Setup a Firewall Mark by Modifying an Existing VIP

In this case, the floating IP address associated with the VIP will already exist so does not need to be created manually.

#### Step 1: Modify the Existing Virtual Service

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services*.
2. Click **Modify** next to the relevant VIP.

Virtual Service		
Label	<input type="text" value="Cluster-2"/>	
Firewall Mark Identifier	<input type="text" value="2"/>	
Ports	<input type="text" value="80"/>	
Protocol		
Protocol	<input type="text" value="Firewall Marks"/>	

3. Change the IP address to the chosen 'mark' value as shown above, e.g. **2**.

4. change the *Protocol* field to **Firewall Marks**.

## Step 2: Define a health check Port

1. Enter the appropriate value in the *Check Port* field, e.g. **80**.
2. Click **Update**.

## Step 3: Modify the Firewall Script

1. Using the WebUI, navigate to: *Maintenance > Firewall Script*.
2. Enter the rules to configure the Firewall Mark as shown in the example below:

```
VIP1="192.168.111.240"
iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 80 -j MARK --set-mark 2
iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 443 -j MARK --set-mark 2
```

```
46 #VIP1="192.168.64.2/"
47 #iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 80 -j MARK --set-mark 1
48 #iptables -t mangle -A PREROUTING -p udp -d $VIP1 --dport 300 -j MARK --set-mark 1
49
50 VIP1="192.168.111.240"
51 iptables -t mangle -A PREROUTING -p tcp -d $VIP1 --dport 80 -j MARK --set-mark 1
52 iptables -t mangle -A PREROUTING -p udp -d $VIP1 --dport 443 -j MARK --set-mark 1
53
54 ##### Packet Filtering #####
55
56 # You should always use a network perimeter firewall to lock down all
57 # external access to the load balancer except the required Virtual Services
58 # and the required services from your admin machine / network (SSH & HTTPS)
59
60 # Allow unlimited traffic on the loopback interface:
61 #iptables -A INPUT -i lo -j ACCEPT
62 #iptables -A OUTPUT -o lo -j ACCEPT
63
64 #Do not delete the following 2 lines
```

**Update**

3. Click **Update**.
4. For a clustered pair, make the same changes to the firewall script on the Secondary unit.

*The VIP is now configured and will be accessible on 192.168.111.240, TCP ports 80 & 443*

## Firewall Mark Notes

1. When using firewall marks the load balancer forwards traffic to the selected Real Server without changing the destination port. So, incoming traffic to port 80 on the Virtual IP will be forwarded to port 80 on one of the Real Servers. Likewise, incoming traffic to port 443 will be forwarded to port 443 on the same Real Server.
2. You can only have one health check port assigned, so if you are grouping port 80 and 443 traffic together you can only check one of these ports, typically this would be port 80.
3. You can specify a range of ports rather than a single port as shown below:

```
iptables -t mangle -A PREROUTING -p tcp -d 10.141.12.34 --dport 1024:5000 -j MARK --set-mark 1
```

(this specifies destination ports from 1024 to 5000)

4. You can leave the upper limit blank to use the default upper limit as shown below:

```
iptables -t mangle -A PREROUTING -p tcp -d 10.141.12.34 --dport 1024: -j MARK --set-mark 1
```

(this specifies destination ports from 1024 to 65535)

5. You can specify a range of IP addresses as shown below:

```
iptables -t mangle -A PREROUTING -p tcp -m iprange --dst-range 10.141.12.34-10.141.12.40 --dport 80 -j MARK --set-mark 1
```

(this specifies the destination IP address as a range from 10.141.12.34 to 10.141.12.40)

## Layer 4 - Advanced Configuration

This section allows you to configure the various layer 4 global settings.

Lock Ldirectord Configuration	<input type="checkbox"/>	?
Check Interval	<input type="text" value="5"/>	?
Check Timeout	<input type="text" value="3"/>	?
Negotiate Timeout	<input type="text" value="5"/>	?
TCP FIN Timeout	<input type="text" value="120"/>	?
UDP Timeout	<input type="text" value="300"/>	?
Failure Count	<input type="text" value="2"/>	?
Quiescent	<input type="text" value="no"/> ▼	?
Email Alert Source Address	<input type="text"/>	?
Email Alert Destination Address	<input type="text"/>	?
Auto-NAT	<input type="text" value="off"/> ▼	?
Multi-threaded	<input type="text" value="yes"/> ▼	?

Update

**Lock Ldirectord Configuration** - Prevent the WebUI from writing the Ldirectord configuration file, so that manual changes are retained. Manual changes to the Ldirectord configuration file may be overwritten if settings are edited in the WebUI. Locking the configuration file will prevent the WebUI from modifying the file so that custom edits are preserved. A warning message will be displayed on all Layer 4 configuration pages, and changes will be denied.

**WARNING:** The Layer 4 configuration is set to read-only – changes made on this page will not be saved. Read-only mode may be disabled on the **Advanced Configuration** page.

**Check Interval** - Layer 4 (Ldirectord) health check interval in seconds. If this setting is too low, you may experience unexpected Real Server downtime.

**Check Timeout** - Layer 4 (Ldirectord) health check timeout in seconds. If this setting is too low, you may induce unexpected Real Server downtime.

**TCP FIN Timeout** - The time to remember an TCP session for after seeing a FIN packet.

**UDP Timeout** - The time to remember a session for after seeing a UDP packet. The timeout is reset on every UDP packet received.

**Negotiate Timeout** - Layer 4 (Ldirectord) negotiate health check timeout in seconds. The negotiate checks may take longer to process as they involve more server side processing than a simple TCP socket connect check. If this setting is too low, you may induce unexpected Real Server downtime.

**Failure Count** - Layer 4 (Ldirectord) number of times a check has to fail before taking server offline. The time to detect a failure and take down a server will be (check interval + check timeout) x failure count.

**Quiescent** - When a Real Server fails a health check, do we kill all connections?

When Quiescent is set to **yes**, on a health check failure the Real Server is not removed from the load balancing table, but the weight is set to 0. Persistent connections will continue to be routed to the failed server, but no new connections will be accepted. When Quiescent is set to **no**, the server is completely removed from the load balancing table on a health check failure. Persistent connections will be broken and sent to a different Real Server.

**Note**

Quiescent only applies to health checks - it has no effect on taking Real Servers offline in System Overview. To manually force a Real Server to be removed from the table, set Quiescent to no and arrange for the server to fail its health check. This may be done, for example, by shutting down the daemon or service, changing the negotiate check value, or shutting down the server.

**Email Alert Source Address** - Specify the global source address of the email alerts. When an email alert is sent, the system will use this address as the 'From' field.

**Email Alert Destination Address** - Specify the global destination email alert address. This address is used to send notifications of Real Server health check failures. This can also be configured on a Virtual Service level.

**Auto NAT** - Automatically NAT outbound network connections from internal servers. By default servers behind the load balancer in a NAT configuration will not have access to the outside network. However clients on the outside will be able to access load balanced services. By enabling Auto NAT the internal servers will have their requests automatically mapped to the load balancer's external IP address. The default configuration is to map all requests originating from internal network eth0 to the external IP on eth1. If you are using a different interface for external traffic you can select it here. Manual SNAT and DNAT configurations for individual servers can also be configured in the firewall script.

**Multi-threaded** - Perform health checks with multiple threads. Using multiple-threads for health checks will increase performance when you have a large number of Virtual Services.

# Layer 7 Services

## The Basics

Layer 7 services are based on HAProxy which is a fast and reliable proxying and load balancing solution for TCP and HTTP-based applications.

Since HAProxy is a full proxy, Layer 7 services are not transparent by default, i.e. the client source IP address is lost as requests pass through the load balancer and instead are replaced by an IP address owned by the load balancer. This is the interface IP by default, but can also be set to any other IP address that the load balancer owns, for example the VIP address.

Layer 7 supports a number of persistence methods: Source IP Address, HTTP Cookie (both application based and inserted), Microsoft Connection Broker, RDP Cookie, HTTP Cookie with Fallback to Source IP, X-Forwarded-For with Fallback to Source IP, Stick On Fallback and Last Successful.

When a VIP is added the load balancer automatically adds a corresponding floating IP address which is activated instantly. Check *View Configuration > Network Configuration* to ensure that the Floating IP address has been activated correctly. They will show up as secondary IP addresses under the relevant interface.

Multiple ports can be defined per VIP, for example 80 & 443. In this persistence (aka affinity/stickiness) will probably be required (default setting) to ensure that clients hit the same backend server for both HTTP & HTTPS traffic and also prevent the client having to renegotiate the SSL connection.

With Layer 7, port re-direction is possible, i.e. VIP:80 → RIP:8080 is supported.

Manual configuration of layer 7 services is supported using the WebUI menu option: *Cluster Configuration > Layer 7 - Manual Configuration*. This enables custom layer 7 VIPs to be created that are able to use HAProxy features not directly supported via the WebUI. For more information, please refer to [Layer 7 - Custom Configurations](#).

<b>Note</b>	It's not possible to configure a VIP on the same IP address as any of the network interfaces. This ensures services can 'float' (move) between Primary and Secondary appliances when using an HA Pair.
-------------	--

## Creating Layer 7 Virtual Services (Using the Wizard)

The wizard can be used to configure one or more Layer 7 Virtual Services and associated Real Servers.

*To run the wizard:*

1. Open the WebUI and start the wizard by clicking the **Accept** button shown above, or by using the WebUI menu option: *Cluster Configuration > Setup Wizard* and clicking **General Layer 7 Virtual Service**.
2. Define the required Virtual Service settings as shown in the example below:

## Setup Wizard - General Layer 7 Virtual Service

Load balancer configuration		
	Primary	Secondary
Hostname	lbmaster	Not configured
Static IP Addresses    eth0	192.168.111.235/18	
Floating IP Addresses		

### Create a new Layer 7 Virtual Service

Label	<input type="text" value="VIP Name"/>		
Virtual Service	IP Address	<input type="text" value="10.0.0.20"/>	
	Ports	<input type="text" value="80"/>	
Layer 7 Protocol	TCP Mode <input type="button" value="v"/>		
<input type="button" value="Create Virtual Service"/>			

3. Click **Create Virtual Service**.

4. Now continue and add the associated load balanced servers (Real Servers) as shown below:

### Attach Real Servers

Label	IP Address	Port	Weight	
<input type="text" value="Web1"/>	<input type="text" value="192.168.1.30"/>	<input type="text" value="80"/>	<input type="text" value="100"/>	
<input type="text" value="Web2"/>	<input type="text" value="192.168.1.40"/>	<input type="text" value="80"/>	<input type="text" value="100"/>	<input type="button" value="✖"/>
<input type="button" value="Add Real Server"/>				
<input type="button" value="Attach Real Servers"/>				

- Use the **Add Real Server** button to define additional Real Servers and use the red cross to delete Real Servers.
- Once you're happy, click **Attach Real Servers** to create the new Virtual Service & Real Servers.
- A confirmation message will be displayed as shown in the example below:

**Information:** Real Server Web1 added.

**Information:** Real Server Web2 added.

**Information:** Virtual Service configured successfully

5. Click **Continue**.
6. Finally, reload HAProxy using the **Reload HAProxy** button in the "commit changes" message box at the top of the screen or by using the WebUI menu option: *Maintenance > Restart Services* and clicking **Reload HAProxy**.

- Note

Running the wizard again will permit additional Layer 7 VIPs and associated RIPs to be defined.
- Note

By default, Real Server health checks are set to a TCP port connect. If you need to configure a more robust check, please refer to [Chapter 8 - Real Server Health Monitoring & Control](#).

Creating Layer 7 Virtual Services

Virtual services (VIPs) can be configured either by creating a new VIP and configuring all required settings, or by using the duplicate VIP feature.

Each Virtual Service can have an unlimited number of Real Servers. Typically you'll need one Virtual Service for each distinct cluster (group of load balanced servers). For example, you'd create a VIP for a web cluster, another for an FTP cluster and a third for a SIP cluster. Multiple ports can also be specified for each VIP.

Defining a New Layer 7 VIP

to add a new layer 7 VIP:

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Services*.
2. Click **Add a new Virtual Service**.

Virtual Service

[Advanced +]

Label

VIP Name

?

IP Address

10.0.0.20

?

Ports

80

?

Protocol

Layer 7 Protocol

HTTP Mode

?

Cancel

Update

3. Enter an appropriate *Label* (name) for the new Virtual Service.
4. Enter the required IP address in the *Virtual Service IP Address* field.
5. Enter the required ports(s) in the *Virtual Service Ports* field, separate multiple ports with commas, specify a range with a hyphen.

- Note

Several ports are used by the appliance and therefore cannot be used for Virtual Services. For full details please refer to [Ports Used by the Appliance](#).

6. Select the Layer 7 protocol to be handled by this Virtual Service, either HTTP or TCP.
- **HTTP Mode** - Selected if the Virtual Service will handle only HTTP traffic. Allows more flexibility in the processing of connections. The HTTP Cookie and HTTP application cookie modes, and the X-Forwarded-For header all require HTTP to be selected. In addition, HAProxy logs will show more information on the

client requests and Real Server responses.

- **TCP Mode** - Required for non HTTP traffic such as HTTPS, RPC, RDP, FTP etc.

7. Click **Update**.

8. Now proceed to define the RIPv (Real Servers) as detailed [here](#).

### Duplicating an Existing Layer 7 VIP

If you have existing Virtual Services, these can be duplicated using the "Duplicate Service" feature.

#### Note

This option will copy all Virtual Service settings along with all associated Real Servers. After duplicating, you'll need to change either the IP address or port. If this is not done, the new VIP will clash with the original VIP and will not load. All other settings can remain the same if required.

*To duplicate an existing layer 7 VIP:*

1. Click **Modify** next to the VIP you'd like to duplicate.
2. Click the **Duplicate Service** button.
3. Click **OK** at the prompt to confirm you want to duplicate the VIP.
4. The VIP will be duplicated with a new label, all other settings will be identical.
5. Change the *IP Address*, *Port* and any other setting to suit your requirements.
6. Click **Update**.

### Modifying a Layer 7 VIP

When first adding a Virtual Service, only certain values can be configured, others are set at their default value to simplify initial configuration. These values can be changed after the Virtual Service has been created by clicking **Modify** next to the relevant Virtual Service. Additional settings that can be changed are:

Section	Setting	Description
Virtual Service	<b>Advanced</b> > Manual Configuration	Enabling this option will prevent the HAProxy configuration file being written for this Virtual Service, leaving the user to configure it via the - Layer 7 Manual Configuration page. If the virtual service labels match the one in the manual configuration you will be able to see the status of the virtual service as well as control it via the system overview as you would any other service.  <div><b>Note</b> For more information on creating a manually defined layer 7 VIP, please refer to <a href="#">Layer 7 - Custom Configurations</a>.</div>
	<b>Advanced</b> > Create Backend Only	Enabling this option will stop the automatic creation of a Frontend in the HAProxy configuration file for this Virtual Service's Backend. The pool of Real Servers will not be directly accessible to clients via the network but can instead be made accessible from another Virtual Service by naming this Virtual Service in a Backend ACL rule.  <div><b>Note</b> For more information on creating backend VIPs please refer to <a href="#">HAProxy Backends</a>.</div>



Section	Setting	Description
Protocol	<b>Advanced</b> > HTTP Pipeline Mode (HTTP mode only)	Select how HAProxy should handle HTTP pipelining to client and server. The options are: <ul style="list-style-type: none"> <li>• <b>Keep-alive Both</b> - Enable pipelining from both the client to HAProxy and from HAProxy to the server.</li> <li>• <b>Close both client and server</b> - Disable pipelining, always closing connections to both client and server using HTTP.</li> <li>• <b>Keep-alive client, close server</b> - Allow client to negotiate pipelining, whilst closing the server connection using HTTP.</li> <li>• <b>Close client, force close server</b> - Close the server connection at the TCP layer, as well as sending the Connection: close header. Also close the client connection using HTTP.</li> </ul>
	<b>Advanced</b> > Work around broken Connection: close (HTTP mode only)	Work around Real Servers that do not correctly implement the HTTP Connection:close option.
	<b>Advanced</b> > Accept Invalid HTTP Requests (HTTP mode only)	This allows invalid characters in header names to be passed through to the backend. If a fix is not immediately available, enable this option. However it can hide further application bugs as well as open security breaches and should only be enabled as a last resort. Ultimately fix your application.
	<b>Advanced</b> > HTTP request timeout (DoS Protection) (HTTP mode only)	Enabling this option helps protect against Slowloris type attacks. With this option enabled the client must send the full HTTP header request within 5 Seconds.
	<b>Advanced</b> > Reuse Idle HTTP Connections (HTTP mode only)	It is possible to reuse idle connections to serve requests from the same session which can be beneficial in terms of performance. It is important to note that the first request of a session is always sent over its own connection, and only subsequent requests may be dispatched over other existing connections.
	<b>Advanced</b> > TCP Keep-alive (TCP Mode only)	Enables the transmission of TCP keep-alive on both the client and the server sides of the connection. Its important to note that this has nothing to do with HTTP keep-alive. This Option is enabled by default when using persistence modes - MS Session Broker and RDP Client Cookie.
	<b>Advanced</b> > Redispatch	If a real server becomes un-responsive ignore persistence and send client connection to another available real server. If Unsure leave enabled.

Section	Setting	Description
Connection Distribution Method	Balance Mode	<p>The scheduler used to specify server rotation. Specify the scheduler to utilize when deciding the backend server to use for the next new connection. The options are:</p> <ul style="list-style-type: none"> <li>• <b>Weighted Round Robin</b> - Incoming requests are distributed to Real Servers in a sequential manner relative to each Real Server's weight. Servers with a higher weight receive more requests. A server with a weight of 200 will receive 4 times the number of requests than a server with a weight of 50. Weightings are relative, so it makes no difference if Real Server #1 and #2 have weightings of 50 and 100 respectively or 5 and 10. The default weight for new Real Servers is 100.</li> <li>• <b>Weighted Least Connection (this is the default for new VIPs)</b> - Incoming requests are distributed to Real Servers with the fewest connections relative to each Real Server's weight. Servers with a higher weight receive more requests. A server with a weight of 200 will receive 4 times the number of requests than a server with a weight of 50. Again, weightings are relative, so it makes no difference if Real Server #1 and #2 have weightings of 50 and 100 respectively or 5 and 10. The default weight for new Real Servers is 100. <i>This is the default method for new VIPs.</i></li> <li>• <b>First</b> - The first server with available connection slots receives the connection. The servers are chosen from the lowest numeric identifier to the highest which defaults to the server's position in the farm. Once a server reaches its maxconn value, the next server is used.</li> </ul>

Section	Setting	Description
Persistence	Persistence Mode	<p>Select how the load balancer should track clients so as to direct each request to the same server. The options are:</p> <ul style="list-style-type: none"> <li>• <b>HTTP Cookie (HTTP mode only)</b> - The load balancer will set an HTTP Cookie to track each client.</li> <li>• <b>Application Cookie (HTTP mode only)</b> - Where an existing HTTP Cookie is set by the web application on the Real Servers, use this to track each client.</li> <li>• <b>SSL Session ID (TCP mode only)</b> - Read the Session ID from the SSL connection and use this to track each client.</li> <li>• <b>MS Session Broker (TCP mode only)</b> - Use the server-set msts RDP Cookie to track clients connecting to a Microsoft Terminal Server. The Session Broker service must be enabled on the real servers.</li> <li>• <b>RDP Client Cookie (TCP mode only)</b> - Use the client-set mstshash RDP Cookie to track clients connecting to a Microsoft Terminal Server. If the cookie is missing, source IP persistence will be used instead.</li> <li>• <b>Source IP</b> - The same source IP always hits the same Real Server.</li> <li>• <b>HTTP Cookie and Source IP (HTTP mode only)</b> - As HTTP Cookie, falling back to Source IP if the cookie is missing from the HTTP request.</li> <li>• <b>X-Forwarded-For and Source IP (HTTP mode only)</b> - Use X-Forwarded-For, falling back to Source IP if the X-Forwarded-For header is missing from the request.</li> </ul> <div> <div>Note</div> <div>You cannot use the set X-Forwarded-For header option with this method of persistence. It will be disabled.</div> </div> <ul style="list-style-type: none"> <li>• <b>Stick On Fallback</b> (An external configured fallback server is required) – This option disables automatically failing back to the Real Server from the fallback server when the Real Server comes back online. This method is appropriate where you have one Real Server and one fallback server. If the Real Server fails, traffic will be handled by the fallback server. When the Real Server comes back online, the fallback server will continue to handle all traffic and no automatic fallback to the Real Server will occur. In order to force fallback you will need to clear the stick table.</li> </ul>

Section	Setting	Description
		<ul style="list-style-type: none"> <li>• <b>Last Successful</b> - Clients will use which ever server last successfully served a client.</li> </ul> <p><b>Note</b>   This includes the fallback server if one is configured.</p> <p>The last successful server is recorded in the stick table and clearing this will cause the server selection for the next connection to be determined by the Balance Mode.</p> <ul style="list-style-type: none"> <li>• <b>None</b> - No persistence. The allocation of clients to Real Servers will be determined solely by the Balance Mode.</li> </ul>
Persistence Options	HTTP Cookie Name	Set the name of the HTTP cookie.
	Application Cookie Name	Set the name of the application cookie.
	<b>Advanced</b> > HTTP Cookie Max Idle Duration	Set the max idle time of the cookie.
	<b>Advanced</b> > HTTP Cookie Max Life Duration	Set the max lifetime of the cookie.
	<b>Advanced</b> > Persistence timeout	The time-out period before an idle connection is removed from the connection table. The source IP address will be removed from memory when it has been idle for longer than the persistence timeout. The default units are minutes.
	<b>Advanced</b> > Persistence table size	The size of the table of connections in KB. The size of the table of connections (approximately 50 bytes per entry) where connection information is stored to allow a session to return to the same server within the timeout period. The default units are in KB.
	<b>Advanced</b> > Clear Stick on Drain	Clearing the stick table when draining a real server is particularly useful and recommended if you have long lived connections with large connection timeouts such as RDP or SSH. This will force users onto another node when they attempt to reconnect and the while server they were attached to is in drain mode. Alternatively disabling this option would allow the user to reconnect and they would only be moved when their persistence entry expired.

Section	Setting	Description
	<b>Advanced</b> > XFF IP Position	<p>With XFF headers its possible to have either more than one header or more than one IP in that header. This option gives the user the ability to select a specific IP position inside the header to use for persistence. For example: X-Forwarded-For: 192.168.1.1, 192.168.1.2, 10.10.10.1.</p> <p>In the above example the -1 (default) position is 10.10.10.1 this will always be the last appended value, -2 is 192.168.1.2 and -3 is 192.168.1.1 and so on for as many IPs as you have in your header.</p> <p>It is possible to do the same thing with Multiple XFF headers:</p> <p>X-Forwarded-For: 192.168.1.1</p> <p>X-Forwarded-For: 192.168.1.2</p> <p>X-Forwarded-For: 10.10.10.1</p> <p>It works the same as the previous example -1 is 10.10.10.1 or the most recently added header -2 is 192.168.1.2 and -3 is 192.168.1.1 and so on. The IP address at the position you select will be stored in the stick table and used for persistence on the next request from the user.</p>

Section	Setting	Description
Health Checks	Check Type	<div> <div>Note</div> <div>For full details of all layer 7 health check options, please refer to <a href="#">Health Checks for Layer 7 Services</a>.</div> </div> <p>Specify the type of health check to be performed on the Real Servers. The options are:</p> <ul style="list-style-type: none"> <li>• <b>Negotiate HTTP/HTTPS (GET)</b> - Scan the page specified in Request to Send, and check the returned data for the Response Expected string</li> <li>• <b>Negotiate HTTP/HTTPS (HEAD)</b> - Request the page headers of the page specified in Request to Send</li> <li>• <b>Negotiate HTTP/HTTPS (OPTIONS)</b> - Request the options of the page specified in Request to Send</li> <li>• <b>Connect to port</b> - Attempt to make a connection to the specified port.</li> <li>• <b>External script</b> - Use a custom file for the health check. For more information please refer to <a href="#">External Health Check Scripts</a>.</li> <li>• <b>MySQL</b> - The check consists of sending two MySQL packets, one Client Authentication packet, and one QUIT packet, to correctly close the MySQL session. It then parses the MySQL Handshake Initialisation packet and/or Error packet. It is a basic but useful test and does not produce error nor aborted connect on the server. However, it requires adding an authorisation in the MySQL table, like this: <div> <pre>USE mysql; INSERT INTO user (Host,User) values ('',''); FLUSH PRIVILEGES;</pre> </div> </li> <li>• <b>No checks, Always on</b> - No health checks, all real servers are marked online.</li> </ul>
Health Check Options	Request to send	Specify a specific location/file for the health check. Open the specified location and check for the Response Expected. Useful for checking a server sided script to check the health of the backend application.
	Response expected	<p>The content expected for a valid health check on the specified file. The Response Expected can be any valid regular expression statement.</p> <div> <div>Note</div> <div>When the <i>Check Type</i> is set to <b>Negotiate HTTP (GET)</b> the result can be inverted by selecting 'not equal' option in the additional drop-down.</div> </div>
	<b>Advanced</b> > Check Port	If specified, this setting overrides the default check port, useful when you are balancing multiple ports.
	<b>Advanced</b> > Username	If authentication is required specify the username here.
	<b>Advanced</b> > Host Header	If the real server's web server is configured to require a Host header, the value to be used in health checks may be set here.
	<b>Advanced</b> > Password	If authentication is required specify the password here.

Section	Setting	Description
ACL Rules	Configure Content Redirects	Enables ACLs to be configured. For more information on configuring ACLs please refer to <a href="#">ACLs (aka Content Switching) and URL Rewriting</a> .
Header Rules	Configure Headers	Enables HTTP headers to be added, set or deleted. For more information on configuring HTTP headers please refer to <a href="#">Modifying HTTP Header Fields</a> .
Feedback Method	Feedback Method	<p>Select whether HAProxy should query each Real Server for its load level. The options are:</p> <ul style="list-style-type: none"> <li>• <b>None</b> - HAProxy will not modify the Real Server's weight.</li> <li>• <b>Agent</b> - The Real Server is queried every health check interval for the real server's percent CPU idle. This is used to set each Real Server's weight to a value proportional to its initial weight. For example, if the initial weight is 100 and the percentage CPU idle is 34, the weight will be set to 34. Remember lower numbers mean lower priority for traffic, when compared with other real servers in the pool.</li> </ul>
Fallback Server	IP Address	IP address of server where to direct requests if all RIPs are down.
	Port	Port of server where to direct requests if all RIPs are down.
	<b>Advanced</b> > Fallback Persistence	Configure the Fallback server to be persistent. During a health check failure users can be forwarded to a fallback server. Setting this to on will make this server persistent so that when the Real Servers are put back in the pool, they will remain on the fallback server until their persistence times out. Setting this to off will move users to a Real Server as soon as one is available.
	<b>Advanced</b> > Encrypt Connection	Enable SSL encryption to the fallback server.
SSL	Enable Backend Encryption	Enabling this option will enable by default the use of HTTPS for all new Backend Servers. This options can then be disabled per backend server under the Real Server settings.
Other	<b>Advanced</b> > Maximum Connections	Specifies the maximal number of concurrent connections that will be sent to this server. If the number of incoming concurrent requests goes higher than this value, they will be queued, waiting for a connection to be released.
	<b>Advanced</b> > Timeout	<p>Use this option to override the default client &amp; server timeouts in the Layer 7 advanced section.</p> <p><b>Client Timeout</b> - The inactivity timeout applies when the client is expected to acknowledge or send data.</p> <p><b>Real Server Timeout</b> - The inactivity timeout applies when the server is expected to acknowledge or send data.</p>

Section	Setting	Description
	<b>Advanced &gt;</b> Set X-Forwarded-For Header	Instruct HAProxy to add an X-Forwarded-For (XFF) header to all requests, showing the client's IP Address. If HTTP is selected under Layer 7 Protocol, HAProxy is able to process the header of incoming requests. With this option enabled, it will append a new X-Forwarded-For header containing the client's IP Address. This information may be extracted by the Real Server for use in web applications or logging.
	<b>Advanced &gt;</b> Force to HTTPS	<p>If set to 'Yes' any HTTP connections that are made on this VIP will be forced to reconnect using HTTPS. This will keep any entered URL. If you are terminating the SSL on the Load balancer you should use the same VIP address for both the SSL Termination and Layer 7 configurations.</p> <p><b>HTTPS Redirect Code</b> - this indicates which type of HTTP redirection is desired. Codes 301, 302, 303, 307 and 308 are supported, with 302 used by default if no code is specified. The options are:</p> <ul style="list-style-type: none"> <li>• 301 means "Moved permanently", and a browser may cache the Location.</li> <li>• 302 means "Moved permanently" and means that the browser should not cache the redirection.</li> <li>• 303 is equivalent to 302 except that the browser will fetch the location with a GET method.</li> <li>• 307 is just like 302 but makes it clear that the same method must be reused.</li> <li>• 308 replaces 301 if the same method must be used.</li> </ul> <p><b>HTTPS Redirect Port</b> - Setting this option to a port other than 443 will cause a port to be specified in the Force-to-HTTPS redirection emitted when clients connect via HTTP.</p>
	<b>Advanced &gt;</b> Use RIP name as Host Header	When set to 'Yes' the Host Header is set to the name allocated to the RIP.
	<b>Advanced &gt;</b> Accept Proxy Protocol	<p>If you wish to use this VIP with STunnel for SSL off-load or another supported proxy such as Amazons ELB whilst passing the client's IP address to the real servers this option needs to be enabled (checked). If using with STunnel please ensure that the 'Enable Proxy Protocol' is enabled in your STunnel VIP.</p> <div> <p><b>Note</b></p> <p>When used with STunnel, the preferred method is to use the 'Enable Proxy Protocol' option in the STunnel VIP's configuration in conjunction with the 'Bind Proxy Protocol to L7 VIP' option. This will configure both the STunnel VIP and the HAProxy VIP in a single step and allows a single HAProxy VIP to support both HTTP and HTTPS. For more information please refer to <a href="#">Transparency at Layer 7</a>.</p> </div>



Section	Setting	Description
	<b>Advanced</b> > Send Proxy Protocol	<p>Enable Proxy Protocol to the backend servers. This option allows the back end servers to see the client's IP address. It should only be enabled if your server supports Proxy Protocol and is configured to use it. The options are:</p> <ul style="list-style-type: none"> <li>• <b>Send V1</b> - This uses the first version of the Proxy Protocol and send the headers in a human readable format.</li> <li>• <b>Send V2</b> - This is the newer version of the Protocol and sends the headers in binary.</li> <li>• <b>Send V2 SSL</b> - This is used to show the client was connected over SSL/TLS.</li> <li>• <b>Send V2 SSL CN</b> - This is the same as V2 SSL but also provides the Common Name from the client certificate if set.</li> </ul>
	<b>Advanced</b> > Enable Compression	Enable gzip HTTP compression. The following MIME types will be compressed when this is enabled: text/html , text/plain , text/css , text/xml , text/javascript , application/javascript , application/xml
	<b>Advanced</b> > Set Source Address	Allows the setting of the source IP address that your backend server will see the traffic coming from. This is useful when you wish to only allow a known IP Address to access your Real Servers or need to allow access through a public gateway.
	<b>Advanced</b> > Enable HSTS	HSTS specifies a period of time during which the users browser (agent) should only access the server in a secure fashion. The recommended duration should be 3 months or more.
	<b>Advanced</b> > Tunnel Timeout	Timeout for the websocket protocol tunnel when no data is passed between client and server. Can be specified as s/m/h for seconds/minutes/hours.
	<b>Advanced</b> > Transparent Proxy	<p>This will set the source address of the data stream leaving the load balancer destined for the real server to be that of the client. As a result you will need to set the default gateway for the real servers to be that of the load balancer. Other wise, depending on the source address the return traffic will route round the load balancer and no responses will be received. Note: Enabling this feature will enable the TProxy system. Once enabled if no longer required it will need to be disabled from: <i>Layer7 - Advanced Configuration &gt; Transparent Proxy</i>.</p>

#### Note

If you require a custom gateway for a particular VIP, this can be achieved using Policy Based Routing. For more information please refer to [Policy Based Routing \(PBR\)](#).

## ACLs (aka Content Switching) and URL Rewriting

The WebUI supports the ability to create ACLs (access-control lists) which can be used to control and direct traffic based on a set of defined rules. This option can be accessed under the *ACL Rules* section by clicking the **Add Rule** button when modifying a VIP:

ACL Rules

Type	Bool	URL/Text	Action	Redirect
<div>Add Rule ?</div>				

This brings up the ACLs pop-up menu:

HAProxy

ACL Rule:

Cancel

Ok

Type

Select

▼

Select the desired ACL type from the *Type* drop-down list, fill in the details as appropriate, and create the ACL by clicking the **Ok** button.

HAProxy

ACL Rule:

Cancel

Ok

Type

path\_beg

▼

Bool

Equals

▼

URL/Text

/example

Action

Drop

▼

An ACL can be modified by clicking on it in the ACL list. ACLs can be reordered by clicking and dragging them in the ACL list:

ACL Rules


Type	Bool	URL/Text	Action	Redirect
header(Content-Type)	Equals	application/json	Drop	<div>Remove</div>
hdr_host	Equals	example.com	Drop	<div>Remove</div>
hdr_beg(Host)	Equals	www	Drop	<div>Remove</div>

Different types of ACLs can be created. Some ACLs are dependent on information from the application layer and so are only available for *HTTP mode* virtual services. The ACL types and their supported modes are listed below.

ACL Type	TCP Mode Support	HTTP Mode Support
path	✗	✓
path_beg	✗	✓
path_end	✗	✓
hdr	✗	✓

ACL Type	TCP Mode Support	HTTP Mode Support
hdr_host	✗	✓
hdr_beg(Host)	✗	✓
Query	✗	✓
SNI	✓	✗
Flags	✓	✓
IP Address	✓	✓
Port	✓	✓
Free Type	✓	✓

#### URL-Based ACLs



#### path:

- Match against the request's full URL path, which starts at the first slash and ends before the (optional) question mark (signifying the start of the query string).
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the full URL path to, e.g. `static/page.html`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

#### path\_beg:

- Match against the *beginning* of the request's URL path.
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the beginning of the URL path to, e.g. `/static/` or `/level_1/level_2/`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

#### path\_end:

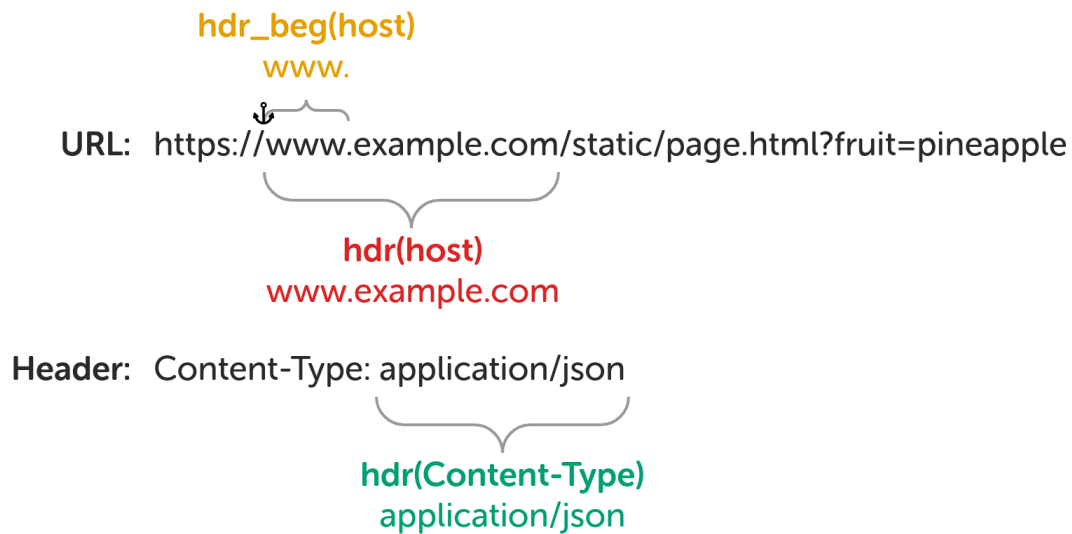
- Match against the *end* of the request's URL path.

- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the end of the URL path to, e.g. `/page.html` or `.png`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

**Query** (aka `url_param`):

- Match against a specified query string parameter.
- **Header/Param:** Name of the query string parameter to match against, e.g. `fruit`.
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the specified query string parameter to, e.g. `pineapple`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

Header-Based ACLs



**hdr\_host** (aka `hdr(host)`):

- Match against the request's full Host header.
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the full Host header to, e.g. `www.example.com` or `en.wikipedia.org`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

**hdr\_beg(host):**

- Match against the *beginning* of the request's Host header.

- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the beginning of the Host header to, e.g. *www.* or *en.*
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

#### hdr:

- Match against a specified request header.
- **Header/Param:** Name of the request header to match against, e.g. *Content - Type*.
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the specified request header to, e.g. *application/json*.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

#### Network/Transport Layer-Based ACLs

##### IP Address (aka src):

- Match against the source IP address of the request.
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** IP address to compare the request's source IP address to. CIDR notation can be used and multiple comma-separated values can be given, e.g. *10.0.0.0/8 123.45.67.8*.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

##### Port (aka dst\_port):

- Match against the destination TCP port of the request.
- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** Integer to compare the request's destination TCP port to.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

#### Miscellaneous ACLs

##### SNI:

- Match against the Server Name Indication (SNI) TLS extension field of the request.

- **Bool:**
  - **Equals:** Perform action if *URL/Text* value matches.
  - **Not equal:** Perform action if *URL/Text* value does not match.
- **URL/Text:** String to compare the SNI field to, e.g. `www.example.com`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

#### Flags:

- Match based on the status of flags set by other ACL rules.
- **Bool:**
  - **Equals:** Perform action if condition described by *URL/Text* field evaluates to *true*.
  - **Not equal:** Perform action if condition described by *URL/Text* field does not evaluate to *true*.
- **URL/Text:** Condition to test, e.g. `flag_a || flag_b` or `protected_path internal_src_addr`.
- **Action:** Action to perform if condition is met. See [ACL Actions](#).

**Note** | See [ACL Examples](#) for an example of how to use the flags feature.

#### Free Type:

- Free-form custom ACL rule(s) to write into the HAProxy configuration verbatim.
- **Freetype:** ACL configuration line(s) to write into the HAProxy configuration file.

**Note** | Multiple ACLs can be defined in a single block. Enter 1 ACL per line.

**Tip** | Full and detailed documentation on how to write ACLs can be found in the HAProxy Configuration Manual, [here](#).

#### ACL Actions

When an ACL rule matches an action is taken (or, alternatively, an action is taken when the rule *doesn't* match, depending on the "bool" setting of the rule). Some actions are dependent on information from the application layer and so are only available for *HTTP mode* virtual services. The different actions and their supported modes are listed below.

ACL Type	TCP Mode Support	HTTP Mode Support
Drop	✓	✓
Deny	✗	✓
Set Flag	✓	✓
URL Location	✗	✓
URL Prefix	✗	✓
Use Backend	✓	✓
Use Server	✓	✓

- **Drop** (aka reject): Stop and immediately close the connection without sending a response.
- **Deny**: Stop and immediately deny the request, emitting the chosen HTTP status code as a response.
  - **Status code**: Status code to use as a response.
    - 200 OK
    - 400 Bad Request
    - 403 Forbidden
    - 405 Method Not Allowed
    - 408 Request Timeout
    - 425 Too Early
    - 429 Too Many Requests
    - 500 Internal Server Error
    - 502 Bad Gateway
    - 503 Service Unavailable
    - 504 Gateway Timeout
- **Set Flag**: Set a flag for use in subsequent ACL rules of type "Flag".
  - **Location/Value**: Name of the flag to set, e.g. `flag_a` or `protected_path`.
- **URL Location** (aka redirect location): Redirect the request to the exact location specified, using a 301 Moved Permanently status code.
  - **Location/Value**: Exact location to redirect to, e.g.  
`https://en.wikipedia.org/wiki/User_Datagram_Protocol`.
- **URL Prefix** (aka redirect prefix): Redirect the request to the *URL path* originally requested prefixed with a specified string, using a 301 Moved Permanently status code.
  - **Location/Value**: String to prefix to the requested URL path to create the redirect location, e.g.  
`https://www.example.com`.
- **Use Backend**: Use the specified backend, or another virtual service, to handle the request.
  - **Location/Value**: Name of a valid backend, or another virtual service, to use, e.g.  
`apache_srv_cluster_b`.
- **Use Server**: Use the specified server to handle the request.
  - **Location/Value**: Name of a valid real server, in the same virtual service or backend, to use, e.g.  
`apache_srv_7`.

## ACL Examples

### Example 1

A virtual service occasionally sees requests for a retired domain, `www.foo.com`. These requests need to be redirected to the new domain: `www.bar.com`. For example, a request for `https://www.foo.com/static/diagram.svg` must be redirected to `https://www.bar.com/static/diagram.svg`.

ACL type to use: `hdr_host`, matching against `www.foo.com`.

ACL action to use: *URL Prefix*, with the prefix `https://www.bar.com`.

## Example 2

A web service has been moved: previously, all of its resources were located under `/web-service/`, but now everything is located under `/legacy/web-service/`. Any requests for old locations, whose paths start with just `web-service`, must be redirected to the correct new locations.

ACL type to use: *path\_beg*, matching against `/web-service/`.

ACL action to use: *URL Prefix*, with the prefix `/legacy`.

## Example 3

A web service hosts an administration panel which is located under `/admin/`. The only legitimate use of this panel should be from users on the local network, which is `10.0.0.0/8`. Any non-local users attempting to access the administration panel should be redirected to a branded page, located at `https://example.com/restricted.html`, which explains that they have attempted to access a restricted part of the service.

First ACL type to use: *path\_beg*, matching against `/admin/`.

First ACL action to use: *Set Flag*, setting the flag `is_admin_panel`.

Second ACL type to use: *IP Address*, matching against the network `10.0.0.0/8`.

Second ACL action to use: *Set Flag*, setting the flag `is_local_user`.

Third ACL type to use: *Flags*, matching against `is_admin_panel !is_local_user`.

Third ACL action to use: *URL Location*, with the location `https://example.com/restricted.html`.

The two flag names together, `is_admin_panel !is_local_user`, create a logical AND (a logical OR could be achieved, instead, by explicitly placing `||` between the flag names). The exclamation mark negates the match on `is_local_user`. The resulting expression of the third ACL will match, and cause the redirection action to be carried out, when `is_admin_panel` is *true* **and** `is_local_user` is *false*.

**Important** | The ACL that evaluates the two flags **must** be placed *after* the two ACLs that set the flags.

## Example 4

It is necessary to force the use of a particular real server in some scenarios. This must be achieved by setting a particular query string parameter to the name of the server. For example, `http://192.168.0.10/?server_override=apache_srv_dev` should trigger the override ACL condition and send the request to the special server.

ACL type to use: *Query*, looking for the `server_override` parameter and matching against `apache_srv_dev`.

ACL action to use: *Use Server*, with the server name `apache_srv_dev`.



When an ACL is created and the *Action* is set to **Use Backend** it's possible to set the *Location/Value* field to either a backend only VIP, a standard VIP or a manually defined VIP.

### 1 - Backend only VIP (Recommended)

A Backend VIP does not have a frontend in the HAProxy configuration, only a backend that defines the associated Real Servers. As a result, there is no floating IP address associated with the VIP. This kind of VIP is used exclusively for ACLs where access is only needed from another VIP and not directly from clients over the network.

*To create a Backend VIP:*

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Services*.
2. Click **Add a New Virtual Service**.
3. Click **[Advanced]** and check (tick) the *Create Backend Only* checkbox.
4. Enter a suitable *Label* (name), e.g. **backend-1**.
5. Select the required *Layer 7 Protocol*.
6. Click **Update**.
7. Now navigate to: *Cluster Configuration > Layer 7 - Real Servers*.
8. Click the **Add a New Real Server** button next to the newly created VIP.
9. Define all the Real Servers that make up the backend.

This new backend VIP can now be referenced as **backend-1** in ACLs.

#### Note

You can modify the Backend Virtual Service in the normal way if additional settings must be configured.

### 2 - As a Manually Defined Backend

Manually defined backends allow additional custom settings that are not directly supported via the WebUI to be configured.

*To create a manually defined backend:*

Using the WebUI menu option: *Cluster Configuration > Layer 7 - Manual Configuration*, the backend 'Blog' can be defined as shown below:

```
backend Blog
mode http
balance roundrobin
option forwardfor
server rip3 192.168.110.242:80 weight 1 check
server rip4 192.168.110.243:80 weight 1 check
```

### 3 - As a Standard VIP with the Required Real Servers

Standard VIPs can also be used. Here, 'Blog' has been defined as an additional standard VIP with 2 Real Servers:

	<b>Blog</b>	<b>192.168.112.116</b>	<b>80</b>	<b>0</b>	<b>HTTP</b>	<b>Layer 7</b>	<b>Proxy</b>	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	<i>BlogRIP1</i>	<i>192.168.110.240</i>	<i>80</i>	<i>100</i>	<i>0</i>	<i>Drain</i>	<i>Halt</i>	
	<i>BlogRIP2</i>	<i>192.168.110.243</i>	<i>80</i>	<i>100</i>	<i>0</i>	<i>Drain</i>	<i>Halt</i>	

#### Note

When defining ACLs that have their *Action* set to **Use Backend** or **Use Server**, the relevant Backend VIP or Real Server must exist before HAProxy can be successfully restarted. Note also that names used are case sensitive.

#### Note

For more details on configuring ACLs please also refer to the HAProxy online documentation available [here](#).

### Using Regular Expressions to Rewrite Requests

Regular expressions can be used to rewrite HTTP requests. This is often used to maintain compatibility between old and new URLs or to turn user-friendly URLs into CMS-friendly URLs, etc. This is achieved using the 'reqrep' and 'reqirep' keywords within a manual layer 7 VIP. The following examples illustrate how these commands can be used:

#### Example 1

Replace "/static/" with "/" at the beginning of any request path.

```
reqirep ^([^\ :]*)\ /static/(.*) \1\ /\2
```

#### Example 2

Replace any host name in the HTTP header with "www.mywebsite.com".

```
reqirep ^Host:\ Host:\ www.mysite.com
```

#### Example 3

Replace /jpg/ with /images/ while maintaining the components before and after the folder.

```
reqrep ^([^\ ]*)\ /jpg/(.*) \1\ /images/\2
```

#### Note

HAProxy uses PCRE compatible regular expressions. For more information about PCRE syntax, see [Regex Quick Start](#) and [Regex Cheat Sheet](#).

#### Note

The "reqrep" keyword is strictly case-sensitive, while "reqirep" is case insensitive. For more details on configuring manually defined layer 7 VIPs please refer to [Configuring Manual Virtual Services](#).

### Modifying HTTP Header Fields

**For HTTP mode virtual services**, the WebUI supports the ability to add, set, delete, and replace HTTP header fields. This option can be accessed under the *Header Rules* section by clicking the **Add Rule** button when modifying a

VIP:

Header Rules				
Phase	Action	Header	Value	Flags
<div>Add Rule ?</div>				

This brings up the headers pop-up menu:

HAProxy

Header Rule:

Cancel

Ok

Type

Request

Option

Select

Select the desired header type from the *Type* drop-down list, fill in the details as appropriate, and create the header rule by clicking the **Ok** button.

HAProxy

Header Rule:

Cancel

Ok

Type

Request

Option

Set

Header

Via-Loadbalancer

Value

Yes

Flags

A header rule can be modified by clicking on it in the header rules list. Header rules can be reordered by clicking and dragging them in the header rules list:

Header Rules				
Phase	Action	Header	Value	Flags
Request	Set	Via-Loadbalancer	Yes	<div>Remove</div>
<div>Request</div>	<div>Delete</div>	<div>Forwarded</div>		<div>Remove</div>
Request	Delete	TE		<div>Remove</div>
<div>Add Rule ?</div>				

Two different types of header rules can be created:

- **Request:** Affect specified HTTP header fields in HTTP *request* messages.
- **Response:** Affect specified HTTP header fields in HTTP *response* messages.

Four different header field manipulation options are supported:

Option	Description
Add	Append an HTTP header field. If a header field of the same name already exists then an additional header field will still be appended.
Set	Append an HTTP header field. If a header field of the same name already exists then it is first removed before a new one is appended. This is useful for handling security information which external users <b>must not</b> be able to set themselves.
Delete	Remove all HTTP header fields of a specified name.
Replace	Perform a regular expression powered "find and replace" operation on all HTTP header field values of a specified name.

The specifics of the header field manipulation options are as follows:

#### Add:

- **Header:** Field name of the HTTP header to append.
- **Value:** Field value of the HTTP header to append.
- **Flags (optional):** Conditionally execute the header manipulation rule based on the status of flags set by ACL rules, e.g. `is_external_request`.

#### Set:

- **Header:** Field name of the HTTP header to append.
- **Value:** Field value of the HTTP header to append.
- **Flags (optional):** Conditionally execute the header manipulation rule based on the status of flags set by ACL rules, e.g. `is_external_request`.

#### Delete:

- **Header:** Field name of the HTTP header to delete.
- **Flags (optional):** Conditionally execute the header manipulation rule based on the status of flags set by ACL rules, e.g. `is_external_request`.

#### Replace:

- **Header:** Field name of the HTTP header to replace.
- **Value:** "Find and replace" operation to execute, of the form `<matching-regular-expression><replacement>`.
- **Flags (optional):** Conditionally execute the header manipulation rule based on the status of flags set by ACL rules, e.g. `is_external_request`.

**Note** | See [HTTP Header Field Modification Examples](#) for an example of how to use the replace option.

**Tip**

It is possible to include dynamic information in header fields, for example the IP address of a request. This is achieved by using log format variables, e.g. %ci for the client IP address, and sample expressions, e.g. %[fe\_name] for the name of the virtual service handling the request. This functionality is explained in detail in the HAProxy online documentation: the relevant section can be found [here](#).

## HTTP Header Field Modification Examples

### Example 1

A secure deployment requires that the X-Forwarded-For header field in client requests never be trusted. The header is under the control of the client and could potentially contain misinformation set by a malicious client. All header fields of this name should be deleted from incoming requests.

Header rule type to use: *Request*.

Header rule option to use: *Delete*.

Header rule "Header" value to use: *X-Forwarded-For*.

### Example 2

A poorly designed web service behind the load balancer leaks sensitive information through an HTTP response header field. The vulnerable response field looks like so: Database-Engine: MongoDB\_3.4.14. All header fields of this name should be deleted from outgoing responses.

Header rule type to use: *Response*.

Header rule option to use: *Delete*.

Header rule "Header" value to use: *Database-Engine*.

### Example 3

A web service behind the load balancer expects to receive information about client requests via HTTP header fields. The service expects to receive the source IP address, destination IP address, and destination port of the client's initial connection to the load balancer, which it expects to find in header fields named X-Client-Source, X-Client-Dest, and X-Client-Dest-Port, respectively. The load balancer should add these header fields to incoming requests and populate their values appropriately. The load balancer should also delete any pre-existing header fields that use the field names that the web service is logging, to prevent clients from tampering with or injecting arbitrary data into the web service's logs.

First header rule type to use: *Request*.

First header rule option to use: *Set*.

First header rule "Header" value to use: *X-Client-Source*.

First header rule "Value" value to use: %ci.

Second header rule type to use: *Request*.

Second header rule option to use: *Set*.

Second header rule "Header" value to use: *X-Client-Dest*.

Second header rule "Value" value to use: *%fi*.

Third header rule type to use: *Request*.

Third header rule option to use: *Set*.

Third header rule "Header" value to use: *X-Client-Dest-Port*.

Third header rule "Value" value to use: *%fp*.

#### Example 4

Any requests containing the HTTP header field "Fruit" with a corresponding value of "pineapple" should have this value changed to "kiwi".

Header rule type to use: *Request*.

Header rule option to use: *Replace*.

Header rule "Header" value to use: *Fruit*.

Header rule "Value" value to use: *pineapple kiwi*.

#### Creating Layer 7 Real Servers (RIPs)

You can add an unlimited number of Real Servers to each Virtual Service. For layer 7 VIPs port redirection is possible so the Real Server port field can be set to a different value to the VIP port. Real Servers in a Layer 7 configuration can be on any subnet in any network as long as they are accessible from the load balancer.

*To add a new layer 7 RIP:*

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Real Servers*.
2. Click **Add a new Real Server** next to the relevant Virtual Service.

Label	<input type="text" value="RIP Name"/>	?
Real Server IP Address	<input type="text"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

CancelUpdate

3. Enter an appropriate *Label* (name) for the new Real Server.
4. Enter the required settings in the *Real Server IP Address* field and *Real Server Port* field.
5. Enable the *Re-encrypt to backend* option if required.

Note

For more details about this option please refer to [SSL Termination on the Load Balancer with Re-encryption \(SSL Bridging\)](#).

6. Enable the *Enable Redirect* option if required. When enabled, this option allows a particular Real Server to respond to GET or HEAD requests with a redirect to a specified location. Requests will be redirected to a URL made up of the prefix specified and the path of their GET or HEAD request. Other request types will continue to be handled by the Real Server.

Note

This option is only available when the associated VIP's *Layer 7 Protocol* is set to **HTTP Mode**.

7. Specify the required *Weight*, this is an integer specifying the capacity of a server relative to the others in the pool, valid values are 0 to 256, the default is 100. The higher the value, the more connections the server will receive. If the weight is set to 0, the server will effectively be placed in drain mode.

Note

The configuration options *Minimum Connections* and *Maximum Connections* are available when the Real Server is modified using **Modify** after the RIP has been created.

## Layer 7 - Custom Configurations

Custom, manually configured Layer 7 services are useful when your configuration requires advanced HAProxy settings that are not directly supported by the WebUI when creating and modifying VIPs & RIPs.

### Configuring Manually Defined Virtual Services

Often, the best approach is to first create a VIP and it's associated RIPs in the normal way, configuring as many settings as you can using the available WebUI configuration options, then convert this to a manual VIP and add the custom HAProxy configuration settings that you require. The steps are as follows:

- **Step 1** - Define the VIP in the normal way using the WebUI menu option: *Cluster Configuration > Layer 7 - Virtual Services*.
- **Step 2** - Define the required Real Servers in the normal way using the WebUI menu option: *Cluster Configuration > Layer 7 - Real Servers*.
- **Step 3** - Now, using the WebUI menu option: *View Configuration > Layer 7*, scroll down to the newly created VIP and copy the whole configuration section for the VIP - from the initial **Listen <VIP name>** line, right to the end of the Real Server definitions.
- **Step 4** - Next, modify the VIP created in Step 1 and click **[Advanced]** in the *Virtual Service* section. Check (tick) the *Manual Configuration* checkbox, then click **Update**.
- **Step 5** - Finally, navigate to: *Cluster Configuration > Layer 7 - Manual Configuration* and paste the configuration copied in Step 3 into the editor window, add the additional custom configuration settings for the VIP and click **Update** when complete.

Note

When you click update in step 5, a syntax check will be done to ensure that the lines you have added are valid.

## Manual Config Example 1 - Simple HTTP Redirect

This example illustrates how a VIP can be created, converted to a manual VIP and then modified to include a custom configuration that forces requests that start with `/staff/` or `/staff` to be redirected to `https://login.domain.com`.

### Note

This example is for demonstration purposes only. Complex ACLs can be configured in the WebUI without the need for a manual configuration. For more information on configuring ACLs please refer to [ACLs \(aka Content Switching\) and URL Rewriting](#).

The lines that need to be manually inserted to achieve this are:

```
acl ACL-1 path_beg /staff/ (see note 1)
acl ACL-2 path_beg /staff (see note 1)
redirect location https://login.domain.com if ACL-1 or ACL-2 (see note 2)
```

## Notes

1. These lines configure 2 ACLs named **ACL-1** & **ACL-2** where the criteria for a match is that the URL starts with either `/staff/` or `/staff`.
2. This line causes a redirect to `https://login.domain.com` to occur when either ACL is matched.

## Configuration Steps:

1. Using the WebUI menu option: *Cluster Configuration > Layer 7 - Virtual Services* create a Layer 7 VIP with the required Label (name), IP Address and Port. At this point leave the *Manual Configuration* checkbox unchecked, e.g.

### Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="VIP1"/>	?
IP Address	<input type="text" value="192.168.2.110"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

2. Using the WebUI menu option: *Cluster Configuration > Layer 7 - Real Servers* define the associated RIPS in the normal way, e.g.



## Layer 7 Add a new Real Server - VIP1

Label	<input type="text" value="RIP1"/>	?
Real Server IP Address	<input type="text" value="192.168.110.111"/>	?
Real Server Port	<input type="text" value="80"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

- Using the WebUI menu option: *View Configuration > Layer 7* scroll down to the newly created VIP. Now copy the entire configuration for the VIP.

```
listen VIP1
bind 192.168.2.110:80 transparent
mode http
balance leastconn
cookie SERVERID maxidle 30m maxlife 12h insert nocache indirect
server backup 127.0.0.1:9081 backup non-stick
option http-keep-alive
timeout http-request 5s
option forwardfor
timeout tunnel 1h
option redispatch
option abortonclose
maxconn 40000
option httplog
server RIP1 192.168.2.111:80 weight 100 cookie Rip1 check inter 4000 rise 2 fall 2
slowstart 8000 minconn 0 maxconn 0 on-marked-down shutdown-sessions
server RIP2 192.168.2.112:80 weight 100 cookie Rip1 check inter 4000 rise 2 fall 2
slowstart 8000 minconn 0 maxconn 0 on-marked-down shutdown-sessions
```

- Using the WebUI menu option: *Cluster Configuration > Layer 7 - Virtual Services*, modify the VIP and check the *Manual Configuration* checkbox and click **Update**.
- Select the WebUI menu option: *Cluster Configuration > Layer 7 - Manual Configuration* and paste the VIP's configuration into the editor window, then add the extra manual config lines:

```
listen VIP1
bind 192.168.2.110:80 transparent
mode http
balance leastconn
acl ACL-1 path_beg /staff/
acl ACL-2 path_beg /staff
redirect location https://login.domain.com if ACL-1 or ACL-2
cookie SERVERID maxidle 30m maxlife 12h insert nocache indirect
server backup 127.0.0.1:9081 backup non-stick
option http-keep-alive
timeout http-request 5s
option forwardfor
timeout tunnel 1h
option redispatch
option abortonclose
maxconn 40000
option httplog
server RIP1 192.168.2.111:80 weight 100 cookie Rip1 check inter 4000 rise 2 fall 2
slowstart 8000 minconn 0 maxconn 0 on-marked-down shutdown-sessions
server RIP2 192.168.2.112:80 weight 100 cookie Rip1 check inter 4000 rise 2 fall 2
slowstart 8000 minconn 0 maxconn 0 on-marked-down shutdown-sessions
```

6. Click **Update**.
7. Now reload HAProxy using the **Reload HAProxy** button in the "commit changes" message box at the top of the screen.

## HAProxy Error Codes

For reference, HAProxy's own error codes are as follows:

Code	When/Reason
200	access to stats, and when replying to monitoring requests.
301	when performing a redirection, depending on the configured code.
302	when performing a redirection, depending on the configured code.
303	when performing a redirection, depending on the configured code.
400	for an invalid or too large request.
401	when an authentication is required to perform the action (when accessing the stats page).
403	when a request is forbidden by a "block" ACL or "reqdeny" filter.
408	when the request timeout strikes before the request is complete.
500	when HAProxy encounters an unrecoverable internal error, such as a memory allocation failure, which should never happen.
502	when the server returns an empty, invalid or incomplete response, or when an "rspdeny" filter blocks the response.
503	when no server was available to handle the request, or in response to monitoring requests which match the "monitor fail" condition.
504	when the response timeout strikes before the server responds.

For a complete HAProxy reference please refer to the [text based manual](#) or the [HTML based manual](#).

## Transparency at Layer 7

HAProxy, Pound and STunnel are all proxies which means that a new connection is established from the proxy out to the backend server in response to an inbound client connection to the proxy. This means that by default the source IP address of the packet reaching the Real Servers will not be the client's IP address, but an IP address owned by the load balancer. The source IP address applied depends on which proxy is in operation:

**HAProxy** - By default the IP address of the network interface is used, but this can also be configured to be any IP address that the load balancer owns using the *Set Source Address* field of the Layer 7 VIP.

**STunnel** - By default the IP address of the STunnel Virtual Service is used, but this can also be configured to be any IP address that the load balancer owns using the *Set Source Address* field of the STunnel VIP.

**Pound** - The IP address of the network interface is used.

## Enabling Transparency

The load balancer can provide the actual client IP address to the Real Servers in 2 ways:

1. By inserting a header that contains the client IP source address. For HTTP traffic the **X-Forwarded-For (XFF)**

header is used, for TCP traffic the **Proxy Protocol Header** is used.

**Note**

For more information about XFF headers please click [here](#), for more information about Proxy Protocol Headers please click [here](#).

2. By modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. The load balancer uses **TProxy** for this purpose.

**Note**

In many cases, option 1 (using Headers) can be used to achieve your objectives. Option 1 is easier to implement because there are no network topology requirements.

These methods can be used independently or in combination to achieve a range of objectives as illustrated in the [Configuration Examples](#) section.

## Inserting Headers

### X-Forwarded-For (XFF) Headers

X-Forward-For headers are inserted by HAProxy when the layer 7 VIP option *Set X-Forwarded-For header* is enabled (the default for new layer 7 VIPs). A new X-Forwarded-For header is appended by the load balancer containing the client's IP address. This information can then be extracted by the Real Servers for use in web applications or logging.

### Proxy Protocol Headers

STunnel & HAProxy can be configured for Proxy Protocol Headers as described below:

**STunnel** - To configure STunnel to **send** Proxy Protocol Headers, the STunnel Virtual Service option *Enable Proxy Protocol* must be enabled.

**Note**

If *Associated Virtual Service* is set when configuring the STunnel Virtual Service *Enable Proxy Protocol* will be enabled by default and cannot be disabled.

**HAProxy** - To configure HAProxy to **send** Proxy Protocol Headers, the layer 7 Virtual Service drop-down *Send Proxy Protocol* must be set to the required header version/type.

To configure HAProxy to **receive** Proxy Protocol Headers, 2 methods can be used:


1. By specifying the Layer 7 Virtual Service where the STunnel VIP will forward its connections when creating / modifying the STunnel Virtual Service. This will also automatically configure the layer 7 VIP to expect Proxy Protocol Headers only for connections from the STunnel VIP where the option was enabled. In this way, the layer 7 VIP will accept traffic with Proxy Protocol Headers from the STunnel VIP as well as standard traffic from other sources that do not present Proxy Protocol Headers. To configure this method:
  - Click **Modify** next to the STunnel VIP
  - Set the *Associated Virtual Service* dropdown to the relevant Layer 7 VIP
  - Click **Update**
2. By enabling the layer 7 Virtual Service option *Accept Proxy Protocol* - this will configure the layer 7 VIP to expect Proxy Protocol Headers for all connections. With this method, the layer 7 VIP will only accept connections from sources that present Proxy Protocol Headers. To configure this method:
  - Click **Modify** next to the HAProxy VIP

- Scroll down to the **Other** section and click **[Advanced]**
- Enable (check) *Accept Proxy Protocol*
- Click **Update**

### Using TProxy to modify the Source IP Address

Loadbalancer.org appliances utilize TProxy to modify the source IP address of each packet. TProxy can be used in conjunction with HAProxy and Pound. When TProxy is enabled, it's important to be aware of the topology requirements for TProxy to operate correctly. Both one-arm and two-arm topologies are supported:

#### TProxy Topology Requirements - One-arm Deployments




- Here, the VIP is brought up in the same subnet as the Real Servers.
- To support remote clients, the default gateway on the Real Servers must be an IP address on the load balancer and routing on the load balancer must be configured so that return traffic is routed back via the router.

#### Note

For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can 'float' (move) between Primary and Secondary appliances.

- To support local clients, return traffic would normally be sent directly to the client bypassing the load balancer which would break TProxy. To address this, the routing table on the Real Servers must be modified to force return traffic to go via the load balancer in the same way as one-arm NAT mode. For more information please refer to [One-Arm \(Single Subnet\) NAT Mode](#).

#### TProxy Topology Requirements - Two-arm Deployments



- Here, 2 subnets are used. The VIP is located in one subnet and the load balanced Real Servers are located in the other. The load balancer requires 2 interfaces, one in each subnet.

#### Note

This can be achieved by using two network adapters, or by creating VLANs on a single adapter.

- The default gateway on the Real Servers must be an IP address on the load balancer.

#### Note

For an HA clustered pair, a floating IP should be added to the load balancer and used as the Real Server's default gateway. This ensures that the IP address can 'float' (move) between Primary and Secondary appliances.

- Clients can be located in the same subnet as the VIP or any remote subnet provided they can route to the VIP.

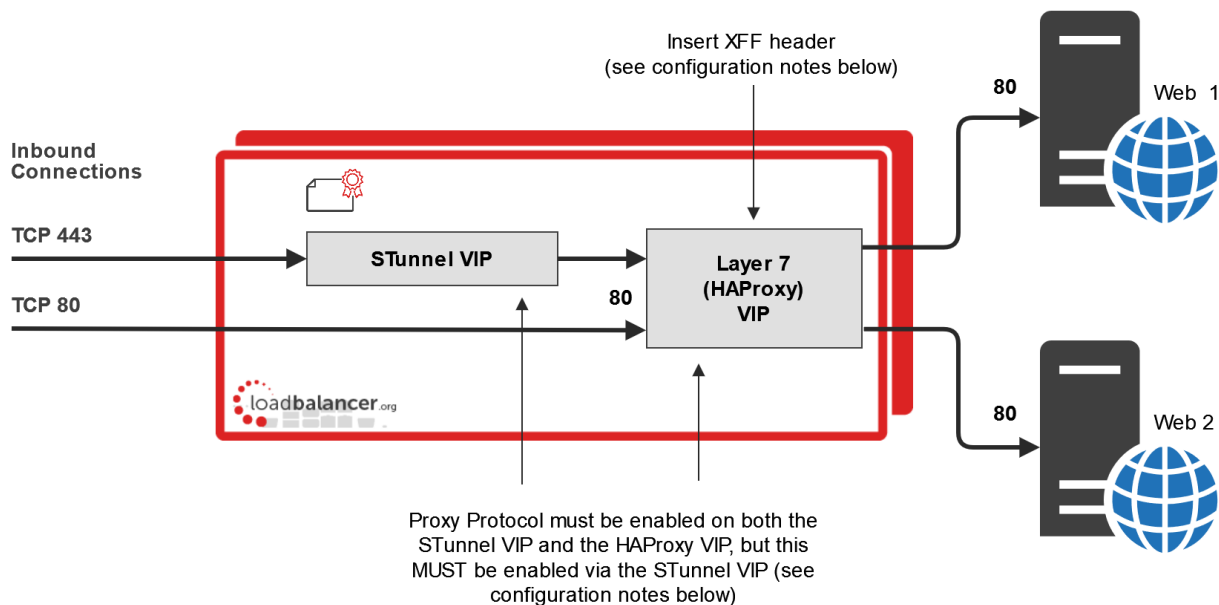
To enable TProxy for a particular layer 7 VIP:

- Click **Modify** next to the HAProxy VIP.
- Scroll down to the **Other** section and click **[Advanced]**.
- Enable (check) *Transparent Proxy*.
- Click **Update**.

## Configuration Examples

### Example 1 - Using Proxy Protocol & X-Forwarded-For Headers

In this example, Proxy Protocol Headers are used with STunnel and HAProxy to present the original client source IP address to the load balanced servers in an XFF header inserted by HAProxy.



## Configuration Notes

1. Configure the STunnel VIP and the HAProxy VIP on the same IP address. Clients then connect to a single IP address for HTTP and HTTPS.
2. Proxy Protocol must be enabled via the STunnel VIP, not via the Layer 7 (HAProxy) VIP. In this way, the HAProxy VIP where STunnel forwards its traffic is automatically configured to accept traffic *with* Proxy Protocol Headers from the STunnel VIP, and also standard traffic *without* Proxy Protocol Headers from other sources, i.e. the direct HTTP connections.

Configuring the STunnel VIP:

Label	<input type="text" value="SSL_VIP1"/>	<a href="#">?</a>
Associated Virtual Service	<input type="text" value="VIP1"/>	<a href="#">?</a>
Virtual Service Port	<input type="text" value="443"/>	<a href="#">?</a>
SSL Operation Mode	<input type="text" value="High Security"/>	
SSL Certificate	<input type="text" value="Default Self Signed Certificate"/>	<a href="#">?</a>
Source IP Address	<input type="text"/>	<a href="#">?</a>
Enable Proxy Protocol	<input checked="" type="checkbox"/>	<a href="#">?</a>
Bind Proxy Protocol to L7 VIP	<input type="text" value="VIP1"/>	<a href="#">?</a>

The STunnel VIP option *Associated Virtual Service* must be set to the backend HAProxy VIP where STunnel will forward its traffic. Then, both the STunnel VIP and the associated HAProxy VIP will be configured automatically.

These STunnel settings will:

1. Configure the STunnel VIP to send Proxy Protocol Headers.
2. Configure the HAProxy VIP to expect Proxy Protocol Headers only from traffic that comes from the STunnel VIP.

*Configuring the Layer 7 (HAProxy) VIP:*

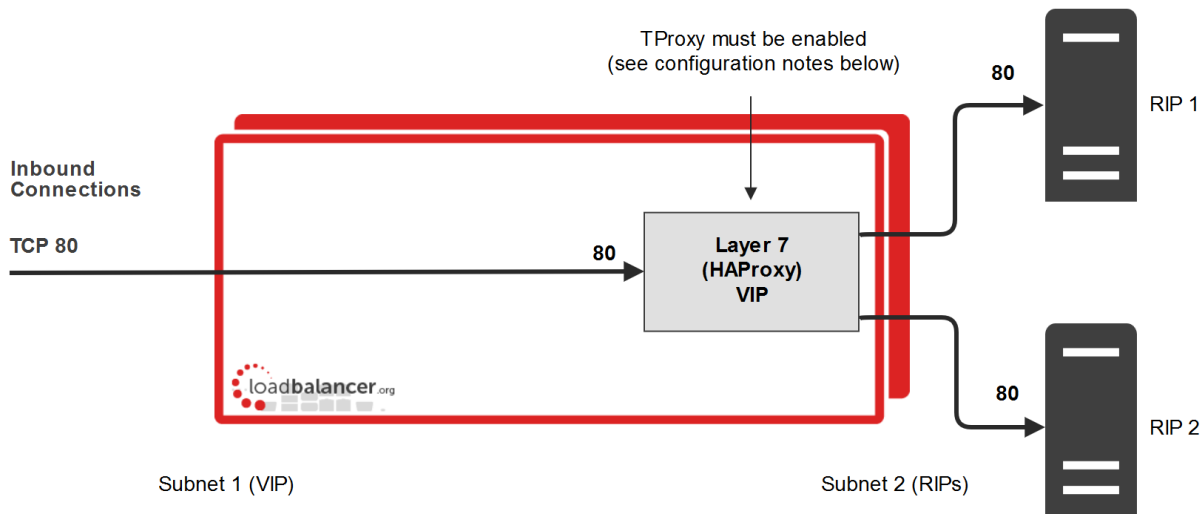
Other	[Advanced -]
Maximum Connections	<input type="text" value="40000"/> <a href="#">?</a>
Timeout	<input type="checkbox"/> <a href="#">?</a>
Set X-Forward-For header	<input checked="" type="checkbox"/> <a href="#">?</a>

X-Forwarded-For Headers must be enabled for HAProxy (this is the default setting).

Once all settings are configured, the **X-Forwarded-For** header received by the load balanced servers Web 1 & Web 2 will contain the source IP address of the client.

#### Example 2 - Using HAProxy & TProxy

In this example, TProxy is enabled for HAProxy so that the source IP address in IP packets is modified by the load balancer to be the clients IP address.

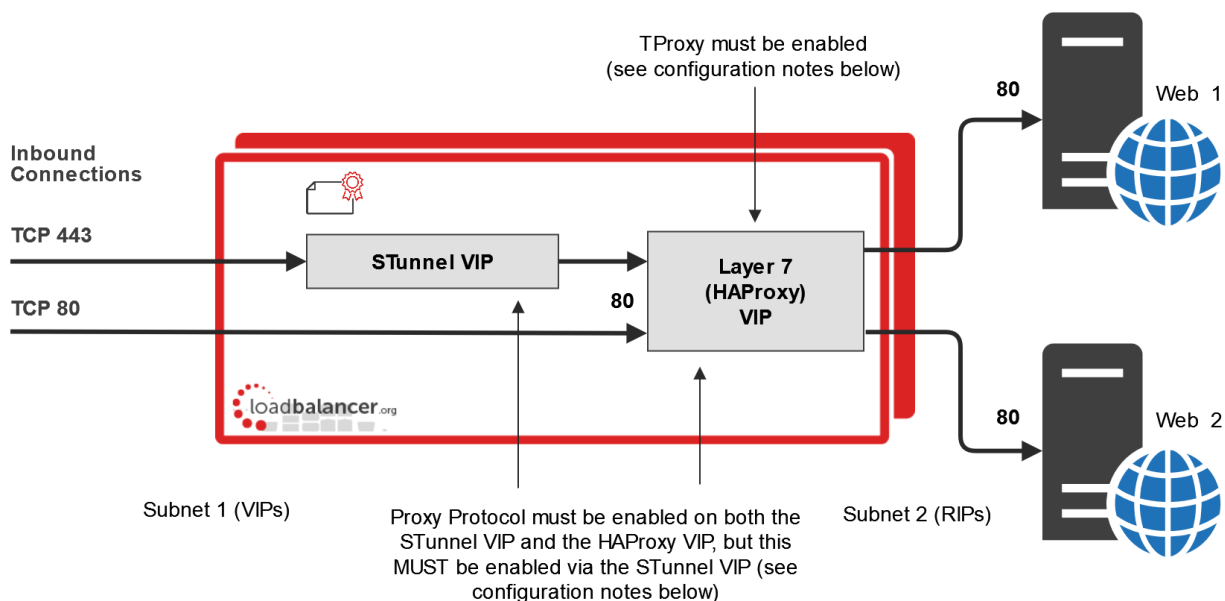


### Configuration Notes & Topology Requirements

1. Certain topology requirements must be met when using TProxy. For details please refer to [Using TProxy to modify the Source IP Address](#).
2. TProxy for HAProxy must be enabled. This is done at the VIP level rather than globally as in previous versions. To enable TProxy at the VIP level, click **Modify** next to the VIP in question, scroll down to the *Other* section and click **[Advanced]**, then enable (check) *Transparent Proxy*.
3. On the Real Servers, the default gateway must be configured to be an IP address on the load balancer. When using a clustered pair, this should be a floating IP to allow failover to the Secondary.

#### Example 3 - Using STunnel, HAProxy & TProxy

In this example, Proxy Protocol Headers are used to pass the client IP address from STunnel to the Layer 7 HAProxy VIP. TProxy is enabled for HAProxy so that the source IP address in packets sent to the Real Servers is modified by the load balancer to be the clients IP address.



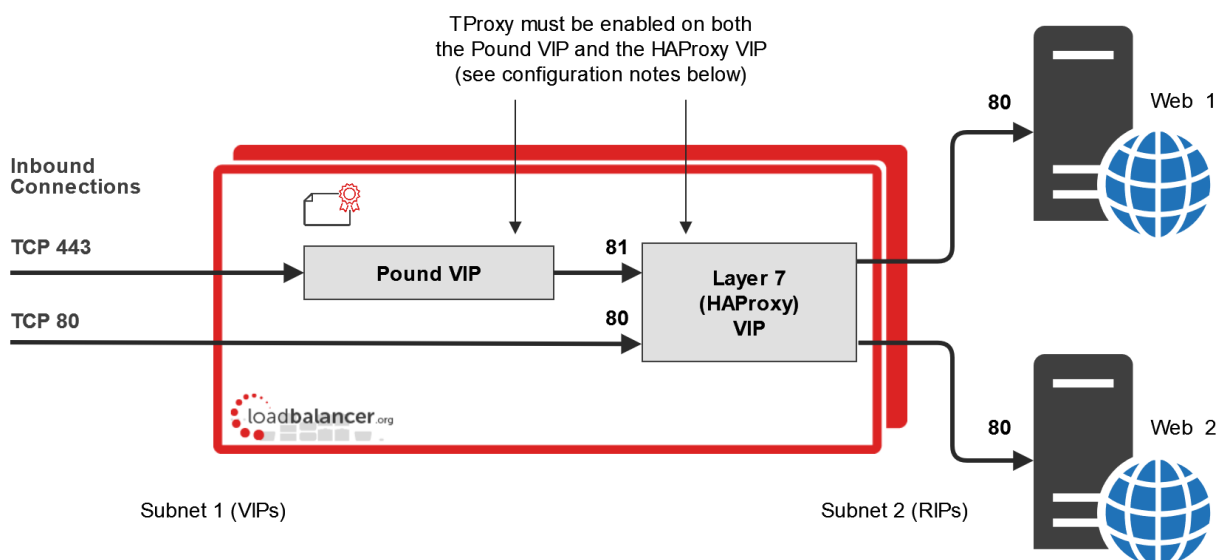
### Configuration Notes & Topology Requirements

1. Certain topology requirements must be met when using TProxy. For details please refer to [Using TProxy to modify the Source IP Address](#).

2. Configure the STunnel VIP and the HAProxy VIP on the same IP address. Clients then connect to a single IP address for HTTP and HTTPS.
3. TProxy for HAProxy must be enabled. This is done at the VIP level rather than globally as in previous versions. To enable TProxy at the VIP level, click **Modify** next to the VIP in question, scroll down to the *Other* section and click **[Advanced]**, then enable (check) *Transparent Proxy*.
4. On the Real Servers, the default gateway must be configured to be an IP address on the load balancer. When using a clustered pair, this should be a floating IP to allow failover to the Secondary.
5. Proxy Protocol must be enabled via the STunnel VIP, not via the Layer 7 (HAProxy) VIP. This is done by checking the *Enable Proxy Protocol* option when either creating or modifying the STunnel VIP (please refer to configuration example 1). This automatically configures the HAProxy VIP to accept traffic with Proxy Protocol Headers from the STunnel VIP and also standard traffic from other sources (i.e. the direct HTTP connections) that do not present Proxy Protocol Headers.
6. If you want to enable HTTP to HTTPS redirection, enable *Force to HTTPS* on VIP1.

#### Example 4 - Using Pound, HAProxy & TProxy

In this example, TProxy is enabled for HAProxy and Pound so that the source IP address is modified by the load balancer to be the clients IP address.



#### Configuration Notes & Topology Requirements

1. Certain topology requirements must be met when using TProxy. For details please refer to [Using TProxy to modify the Source IP Address](#).
2. Configure the Pound VIP and the HAProxy VIP on the same IP address. Clients then connect to a single IP address for HTTP and HTTPS.
3. Configure the Layer 7 HAProxy VIP to listen on 2 ports - e.g. 80 & 81, then use port 80 for client connections on HTTP and port 81 for the Pound backend.
4. When defining Real Servers for HAProxy VIP, ensure that the *Real Server Port* field is set and not left blank.
5. TProxy for HAProxy must be enabled. This is done at the VIP level rather than globally as in previous versions. To enable TProxy at the VIP level, click **Modify** next to the VIP in question, scroll down to the *Other* section and click **[Advanced]**, then enable (check) *Transparent Proxy*.



6. TProxy for Pound must be enabled using the WebUI menu option: *Cluster Configuration > SSL - Advanced Configuration* and *Transparent Proxy* to On.
7. On the load balanced backend Servers, the default gateway must be configured to be an IP address on the load balancer. When using a clustered pair, this should be a floating IP to allow failover to the Secondary appliance.
8. If you want to enable HTTP to HTTPS redirection, you'll need to split the Layer 7 HAProxy VIP into 2 separate VIPs, one on port 80 with *Force to HTTPS* enabled and the other configured to accept traffic from Pound.

## Layer 7 - Advanced Configuration

This section allows you to configure the various layer 7 global settings.

**Lock HAProxy Configuration (Deprecated)** - Prevent the WebUI writing to the HAProxy configuration file. Manual changes to the HAProxy configuration file may be overwritten if settings are edited in the WebUI. Locking the configuration file will prevent the WebUI from modifying the file, so that custom edits are preserved. A warning message will be displayed on all Layer 7 configuration pages, and changes will be denied.

### Note

It's now possible to configure each virtual service as read-only. The manual configuration can then be created using the WebUI option: *Cluster Configuration > Layer 7 - Manual Configuration*. For More information on configuring manual layer VIPs please refer to [Configuring Manual Virtual Services](#).

**Logging** - Set the required logging level for layer 7 services. Logs are written to `/var/log/haproxy.log`.

**Redispatch** - Allows HAProxy to break persistence and redistribute to working servers should failure occur. Normally this setting should not require changing.

**Connection Timeout** - HAProxy connection timeout in milliseconds. This setting should normally not require changing.

**Client Timeout** - HAProxy client timeout in milliseconds. This setting should normally not require changing.

**Real Server Timeout** - HAProxy Real Server timeout in milliseconds. This setting should not require changing.

**Maximum Connections** - HAProxy maximum concurrent connections. This setting should not require changing, unless you are running a high volume site. See also Maximum Connections for a Virtual Service (HAProxy).

**Abort on Close** - Abort connections when users close their connection. Recommended as the probability for a closed input channel to represent a user hitting the 'STOP' button is close to 100%.

**Transparent Proxy** - Enable TProxy support for Layer 7 HAProxy. TProxy support is required in order for the Real Servers behind a layer 7 HAProxy configuration to see the client source IP address. The load balancer must be in a NAT configuration (internal and external subnets) with the Real Servers using an IP address on the load balancer (preferably a floating IP) as their default gateway. Can be used on its own or in combination with Pound TProxy.

### Note

TProxy must be enabled at the VIP level. This is a change from previous versions where enabling it here would enable it for ALL layer 7 VIPs. To enable TProxy at the VIP level, click Modify next to the VIP in question, scroll down to the Other section and click **[Advanced]**, then enable (check) Transparent Proxy. Setting this at the VIP level will also automatically set the Transparent Proxy option here. For more information on using TProxy please refer to [Transparency at Layer 7](#).

Note

Since the load balancer must be in a NAT configuration (i.e. VIPs & RIPv6s in different subnets and default gateway on the Real Servers set as an IP on the load balancer) to utilize TProxy, it's not always an appropriate solution. In situations such as this, it's also possible to use the X-forwarded-for header with layer 7 Virtual Services. Most web servers can then be configured to record the X-Forwarded-For IP address in the log files.

Note

For details on how to enable X-Forwarded-For headers, please click [here](#). For details on how to enable X-Forwarded-For headers in Apache please refer to our [Apache blog](#), For details on how to enable X-Forwarded-For headers in IIS please refer to our [IIS blog](#).

**Disable On Start** - HAProxy brings up all real servers in the UP state after the restart. Enabling this option will bring the real servers up in MAINT mode stopping any connections to them. The init script will then return the real servers back to their previous state pre reload/restart. The init script can do this without this option enabled but while waiting for the init script to get to each service to set the state the real server will be accepting traffic. So it's recommended that you use this with large deployments, or if you just want to stop connections before the previous state has been returned.

**Interval** - Interval between health checks. This is the time interval between Real Server health checks in milliseconds.

**Rise** - Number of health checks to Rise. The number of positive health checks required before re-activating a Real Server.

**Fall** - Number of health checks to Fall. The number of negative health checks required before deactivating a Real Server.

**Slow Start Time** - To minimize the thundering heard effect of a real server recovering from a health check failure getting overwhelmed with all its old users attempting to reconnect at once. This timer will gradually increase the connections for a period set by this value until the end of the timer is reached at which point the server will be running at normal capacity.

Note

If the feed back agent is enabled, the slowstart time MUST be greater than the Interval value.

**Feedback Agent Interval** - The time in milliseconds between each feedback agent check from HAProxy to the feedback agent.

**Advanced Stats** - Enable/disable additional actions available on the HAProxy stats page.

**Request Buffer Length** - Set the health check buffer length in bytes.

Note

Changing this value will effect the performance of HAProxy. Do not make changes unless you know exactly what you are doing.

Lower values allow more sessions to coexist in the same amount of RAM, and higher values allow some applications with very large cookies to work. The default value is 16384 bytes. It is strongly recommended not to change this from the default value, as very low values will break some services such as statistics, and values larger than the default size will increase memory usage, possibly causing the system to run out of memory. Administrators should consider reducing the Maximum Connections parameter if the request buffer is increased.

**Header Buffer Length** - Set the header buffer length, in bytes The header buffer is a section of the request buffer, reserved for the addition and rewriting of request headers. The default value is 1024 bytes. Most applications will

only require a small header buffer, as few headers are added or rewritten.

**Persistence Table Replication** - When enabled, HAProxy's persistence tables are replicated to the Secondary device.

**Replication Port** - Set the TCP port to use for persistence table replication. The default port is TCP 7778.

**eMail Alert From** - Set the 'from address' for email alerts.

**eMail Alert To** - Set the 'to address' for email alerts.

**eMail Server Address** - Set the email server address as either an IP address or FQDN.

**eMail Server Port** - Set the email server TCP port.

Note

For more information on configuring email alerts, please refer to [Configuring Email Alerts for Virtual Services](#).

**Enable Multi-threading** - This can improve performance if limits are being reached.

**Default Number of Threads** - Let the appliance choose a sensible number of worker threads. By default this will be the same as the number of cores available when HAProxy starts or reloads.

**Number of Threads** - Be aware that starting too many threads will have a detrimental affect on performance. Leaving this field blank will have the same effect as selecting default number of threads.

Note

Multi-threading is enabled by default and the number of threads is auto set based on the number of detected CPUs / vCPUs.

**Enable Prometheus Exporter** - Enable the Prometheus exporter for HAProxy. The end point will be mapped to `/lbadmin/stats/l7prometheus/` on the WebUI.

## Floating IPs

In order for the load balancer to function, the unit must physically own the Virtual IP address that the clients are accessing before they get re-directed to a Real Server in the cluster. When new layer 4 or layer 7 Virtual Services (VIPs) are created, Floating IPs are added automatically and can be viewed using the WebUI menu option: *Cluster Configuration > Floating IPs*.

It's also possible to manually define floating IPs if required, this is normally only required when manually configuring firewall marks or when using layer 4 NAT mode or TProxy where in both cases the load balancer must be the default gateway for the Real Servers.

The Floating IPs are controlled by heartbeat to ensure that only one of the load balancer appliance's (normally the Primary) owns the Floating IP(s) at any time.

*To manually add a Floating IP:*

1. Using the WebUI, navigate to: *Cluster Configuration > Floating IPs*.

## Floating IPs

192.168.111.42	<button>Disable</button> <button>Delete</button>
192.168.111.40	<button>Disable</button> <button>Delete</button>

New Floating IP

Add Floating IP

2. Specify the new floating IP.

3. Click **Add Floating IP**.

### Note

When using a clustered pair, ensure that the Secondary also has a static IP address assigned that's in the same subnet as the floating IP being added. Failure to do so will result in heartbeat issues during a failover.

### Note

To disable a floating IP address and bring the IP down, use the relevant **Disable** button. The **Disable** button will be replaced with an **Enable** button to bring it back up when required.

### Note

Floating IPs are not deleted automatically when Virtual Services are removed or the IP address is changed, this must be done manually.

## SSL Termination

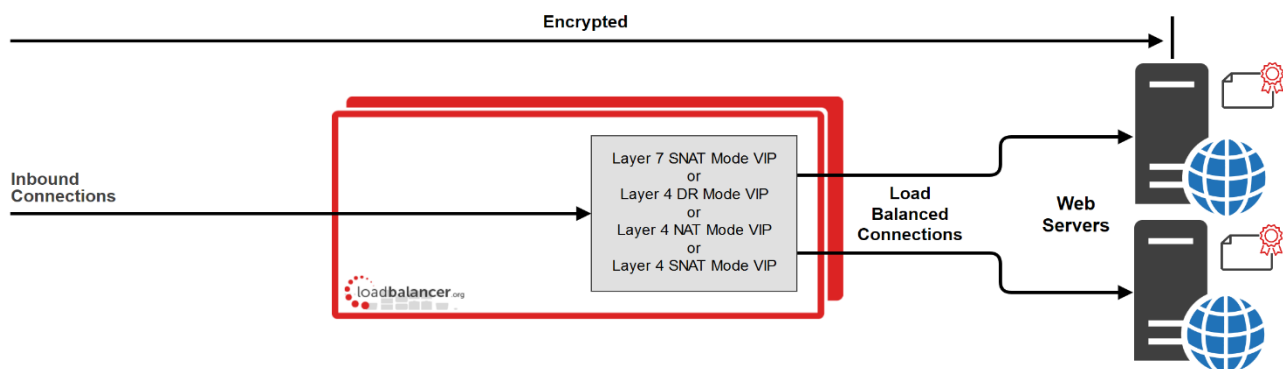
### Concepts

SSL termination can be handled in the following ways:

1. On the Real Servers (recommended) - aka **SSL Pass-through**.
2. On the load balancer - aka **SSL Offloading**.
3. On the load balancer with re-encryption to the backend servers - aka **SSL Bridging**.

The following sections describe each method.

### SSL Termination on the Real Servers (SSL Pass-through)



In this case SSL certificates are installed on each Real Server in the normal way. Data is encrypted from client to server. This provides full end-to-end data encryption as shown in the diagram above.

## Notes

1. This is our recommended solution. SSL termination on the load balancer (SSL Offload) can be very CPU intensive and in most cases, for a scalable solution, terminating SSL on the Real Servers is the best option.
2. It's not possible to use HTTP cookie persistence as well as other layer 7 techniques that control how traffic is sent to the Real Servers because all data is encrypted as it passes through the load balancer.

The load balancer is configured with a VIP that listens on HTTPS port 443 and distributes inbound requests to the Real Servers on port 443 as shown below:

SSL	192.168.110.50	Port 443/tcp	Direct Routing	Add a new Real Server	
SSL1	192.168.110.51	443	Weight 100	Modify	Delete
SSL2	192.168.110.52	443	Weight 100	Modify	Delete

A fairly common configuration is to include port 80 in the VIPs definition and also enable persistence. This ensures that both HTTP and HTTPS requests from a particular client are always sent to the same Real Server as shown below:

SSL	192.168.110.50	Ports 80,443/tcp	Direct Routing	Add a new Real Server	
SSL1	192.168.110.51	80,443	Weight 100	Modify	Delete
SSL2	192.168.110.52	80,443	Weight 100	Modify	Delete

## SSL Termination on the Load Balancer (SSL Offloading)

### Note

SSL termination on the load balancer can be very CPU intensive. In most cases, for a scalable solution, terminating SSL on the Real Servers is the best option.

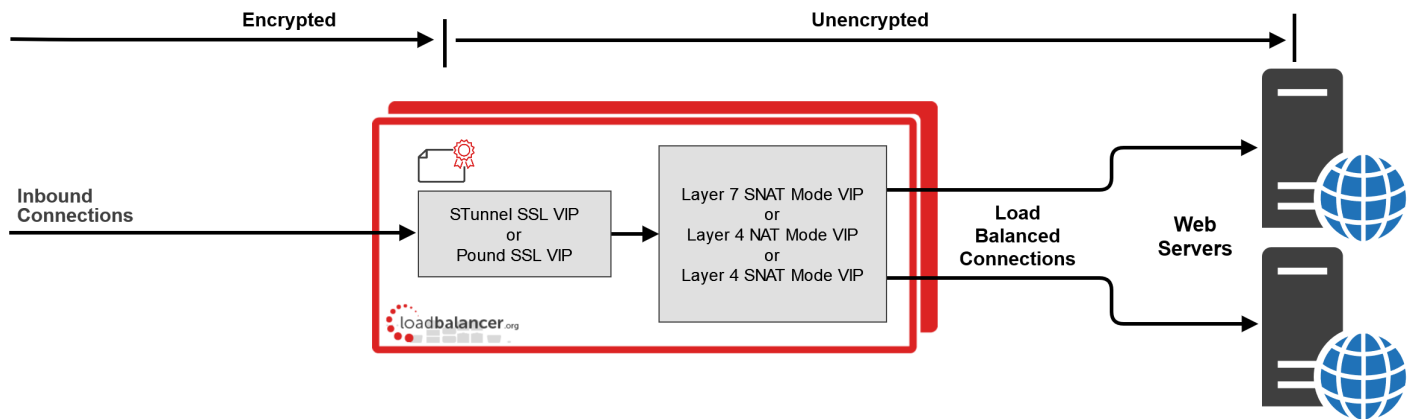
### Note

If you require SSL bridging where the data is re-encrypted from the load balancer to the backend servers, please refer to [SSL Termination on the Load Balancer with Re-encryption \(SSL Bridging\)](#).

SSL Termination can be configured using STunnel, Pound and HAProxy.

## Using STunnel or Pound to Terminate SSL

Here, either an STunnel or Pound SSL Virtual Service is used to decrypt the traffic. Once decrypted, traffic is passed to a second VIP which handles the load balancing between the backend servers as shown below.

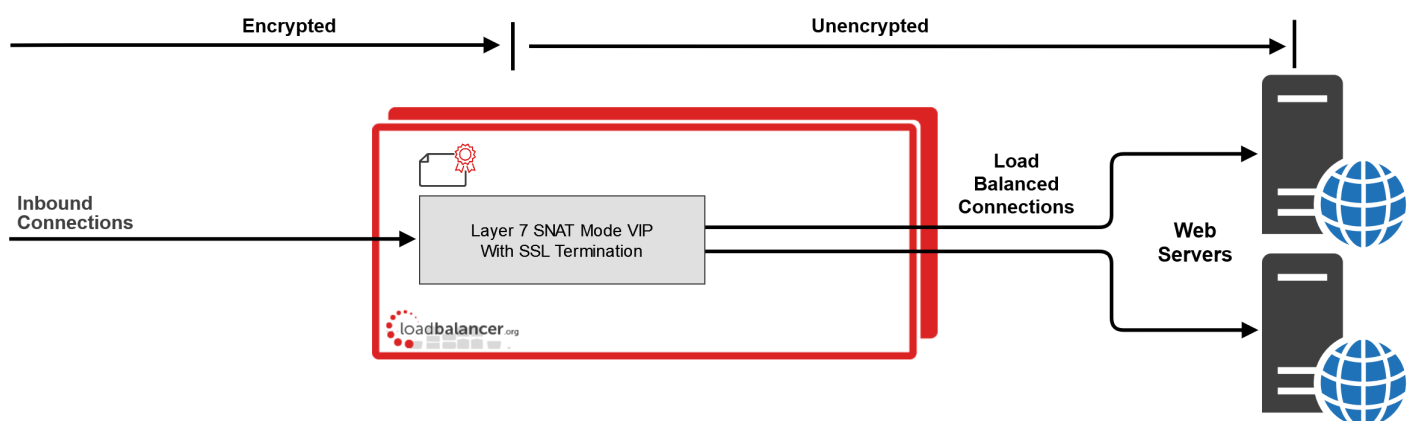


## Notes

1. By default, a self-signed certificate is used for the new VIP. Certificates can be created or uploaded as explained in [Certificates](#).
2. The backend for the STunnel / Pound VIP can be a Layer 7 SNAT mode VIP, a Layer 4 NAT mode VIP or a layer 4 SNAT mode VIP. Layer 4 DR mode cannot be used since Pound & STunnel act as a proxy, and the real servers see requests with a source IP address of the VIP. However, since the Real Servers believe that they own the VIP (due to the loopback adapter configured to handle the **ARP Problem**) they are unable to reply to Pound.
3. If a layer 7 SNAT mode VIP is used as the backend for the STunnel or Pound VIP, cookie based persistence as well as all other layer 7 techniques can be used to control traffic flow to the Real Servers.

## Using HAProxy to Terminate SSL

Here, a single layer 7 SNAT mode VIP handles both the decryption and the load balancing between the backend servers as shown below.



## Notes

1. By default, a self-signed certificate is used for the new VIP. Certificates can be created or uploaded as explained in [Certificates](#).
2. Cookie based persistence as well as all other layer 7 techniques can be used to control traffic flow to the Real Servers.

## Certificates

If you already have an SSL certificate in either PFX or PEM file format, this can be uploaded to the Load balancer

using the certificate upload option as explained in [Uploading Certificates](#). Alternatively, you can create a Certificate Signing Request (CSR) and send this to your CA to create a new certificate, or you can create a locally signed custom certificate.

Generating a CSR on the Load Balancer

CSRs can be generated on the load balancer to apply for a certificate from your chosen CA.

To generate a CSR:

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Certificates*.
2. Click **Add a new SSL Certificate** & select *Create a New SSL Certificate (CSR)*.

I would like to:

- ☐ Upload prepared PEM/PFX file
- ☒ Create a new SSL Certificate Signing Request (CSR) ?
- ☐ Create a new Self-Signed SSL Certificate.

Label:  ?

Domain (CN):  ?

Subject Alternate Name:  ?

Organisation (O):  ?

Organisation unit (OU):  ?

City (L):  ?

State or Province (ST):  ?

Country code (C):  ?

Email address:  ?

CSR Key Length:  ?

**Create** =

3. Enter a suitable *Label* (name) for the certificate.
4. Populate the remaining fields according to your requirements.

**Note** | To specify multiple SANs, separate each name with a comma.

5. Once all fields are complete click **Create**.
6. To view the CSR click **Modify** next to the new certificate, then expand the Certificate Signing Request (CSR) section.
7. Copy the CSR and send this to your chosen CA.
8. Once received, copy/paste your signed certificate into the *Your Certificate* section.

- Intermediate and root certificates can be copied/pasted into the *Intermediate Certificate* and *Root Certificate* sections as required.
- Click **Update** to complete the process.

#### Generating a Self Signed Custom Certificate on the Load Balancer

##### *To generate a Self Signed Certificate:*

- Using the WebUI, navigate to: *Cluster Configuration > SSL Certificates*.
- Click **Add a new SSL Certificate** & select *Create a new Self-Signed SSL Certificate*.
- Enter a suitable *Label* (name) for the certificate.
- Populate the remaining fields according to your requirements.

**Note** | To specify multiple SANs, separate each name with a comma.

- Once all fields are complete click **Create**.

#### Uploading Certificates

Certificates in either PEM or PFX formats can be uploaded to the load balancer.

##### *To upload a Certificate:*

- Using the WebUI, navigate to: *Cluster Configuration > SSL Certificates*.
- Click **Add a new SSL Certificate** & select *Upload prepared PEM/PFX file*.
- Enter a suitable *Label* (name) for the certificate.
- Browse to and select the certificate file to upload (PEM or PFX format).
- Enter the password , if applicable.
- Click **Upload Certificate**, if successful, a message similar to the following will be displayed:

**Information:** cert1 SSL Certificate uploaded successfully.

#### **Note**

If your Primary & Secondary are correctly configured as a clustered pair, when you upload the certificate file to the Primary, the file will be automatically copied over to the Secondary unit.

#### **Note**

It's important to backup all your certificates. This can be done via the WebUI from *Maintenance > Backup & Restore > Download SSL Certificates*.

#### Exporting PFX Certificates from Windows Servers

When exporting certificates from Windows servers, make sure that *Yes, export the private key* is selected, this will enable the output format to be PFX. Also make sure that *Include all certificates in the certification path if possible* is selected.

#### Creating a PEM file



Using a text editor such as vi or vim under Linux or Notepad under Windows, create an empty file (e.g. pem.txt) then copy/paste the entire contents of each of the following items into this file in the order listed below:

- Private Key
- SSL Certificate
- Intermediate Certificate
- Root CA Certificate

Make sure you include the beginning and end tags. The resulting file should look similar to the following:

```
-----BEGIN PRIVATE KEY-----  
(the contents of your Private Key goes here)  
-----END PRIVATE KEY-----  
-----BEGIN CERTIFICATE-----  
(the contents of your SSL Certificate goes here)  
-----END CERTIFICATE-----  
-----BEGIN CERTIFICATE-----  
(the contents of your Intermediate Certificate goes here)  
-----END CERTIFICATE-----  
-----BEGIN CERTIFICATE-----  
(the contents of your Root Certificate goes here)  
-----END CERTIFICATE-----
```

Once created, the file can be uploaded to the load balancer.

#### Converting between certificate formats

In some circumstances it may be required to manually convert certificates between formats. In these cases, OpenSSL can be used. OpenSSL is included on the appliance and is typically included by default in Linux distributions. For Windows, it can be downloaded [here](#). Once installed, you'll have an OpenSSL directory located on your filesystem (by default C:\OpenSSL). To use the program, open a command window, navigate to the location where it was installed (by default C:\OpenSSL\bin) then run the required command as shown below.

#### Converting PFX certificates to PEM format

```
openssl pkcs12 -in file.pfx -nodes -out file.pem
```

#### Converting .cer certificates to PEM format

```
openssl x509 -in file.cer -inform DER -out file.pem -outform PEM
```

#### Converting an Encrypted Private Key to an Unencrypted Key

If a password has been included in the private key, this should be removed before it is used with your PEM file. This can be done using the following command:

```
openssl rsa -in encrypted-server.key -out unencrypted-server.key
```

#### Let's Encrypt

Let's Encrypt is a zero cost Certificate Authority for HTTPS encryption, now trusted by all major root programs, including Google, Microsoft, Apple, Mozilla and Oracle. Used in conjunction with freely available tools it provides

automatic enrollment/renewal, simple cert creation, negating validation emails and manual configuration.

For much more information, please refer to our Let's encrypt [introductory blog](#) and also the [follow up blog](#) that details the `lb-letsencrypt.sh` script and how to use it.

## Creating a SSL Termination

STunnel is used by default for all new SSL terminations. If you want to use Pound or HAProxy, you'll need to set the *SSL Operating Mode* to **Custom** as described in [SSL Operation Mode - Custom](#).

To add a SSL Termination:

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Termination*.
2. Click **Add a new Virtual Service**.

### SSL Termination - Add a new Virtual Service

Label	<input type="text" value="SSL"/>	?
Associated Virtual Service	<input type="text" value="None"/>	?
Virtual Service IP Address	<input type="text"/>	
Virtual Service Port	<input type="text" value="443"/>	?
Backend IP Address	<input type="text"/>	
Backend Virtual Service Port	<input type="text" value="80"/>	?
SSL Operation Mode	<input type="text" value="High Security"/>	
SSL Certificate	<input type="text" value="Default Self Signed Certificate"/>	?
Source IP Address	<input type="text"/>	?
Enable Proxy Protocol	<input checked="" type="checkbox"/>	?
Bind Proxy Protocol to L7 VIP	<input type="text" value="VIP1"/>	?

3. Set the *Associated Virtual Service* to the required value:
  - a. If you leave it set to **None**:
    - Enter the required *Label* (name) for the Virtual Service.
    - Enter the required *Virtual Service IP Address*.
    - Enter the required *Virtual Service Port* - typically **443**.
    - Enter the required *Backend IP Address* - This is normally the same IP address as the Virtual Service IP address but can be any valid IP. The IP address specified must correspond to a Layer 7 SNAT mode VIP or a Layer 4 NAT / SNAT mode VIP. Unencrypted traffic will be sent here for load balancing.
    - Enter the required *backend Virtual Service Port* - typically **80**.

- Select the required *SSL Operation Mode*:

**Note** | For more information see [SSL Operation Mode](#).

- **High Security** - Configure the STunnel VIP for high security
- **FIPS Compliant** - Configure the STunnel VIP for FIPS compliance
- **High Compatibility** - Configure the STunnel VIP for high compatibility
- **Custom** - All settings can be configured manually

**Note** | For more information see [SSL Operation Mode - Custom](#).

- Select the required *SSL Certificate*.
- Set the required *Source IP Address* - by default the *Virtual Service IP Address* is used but this can be changed to any other address owned by the load balancer if required.
- Configure *Enable Proxy Protocol* - If you wish to use HAProxy and the Proxy Protocol this option needs to be enabled (checked) to allow SSL termination on the load balancer whilst passing the client's IP address to the Real Servers. This option only enables a Proxy ACL Rule on a Single STunnel VIP.
- Configure *Bind Proxy Protocol to L7 VIP* - If *Enable Proxy Protocol* is enabled, selecting a layer 7 Virtual service here configures the layer 7 service to expect the proxy protocol from this STunnel service. This enables the layer 7 service to pass the clients IP in a X-Forwarded-For header or with TProxy while still accepting HTTP traffic on the same port (for more information please refer to [Transparency at Layer 7](#)). Note that manually defined layer 7 configurations are not included in the drop-down.

b. If you select a particular VIP where you want to forward the unencrypted STunnel traffic:

- Enter the required *Label* (name) for the Virtual Service.

**Note** | The label will be auto configured based on the *Associated Virtual Service* selected. This can be edited if required.

- Enter the required *Virtual Service Port* - typically **443**.
- Select the required *SSL Operation Mode*:

**Note** | For more information see [SSL Operation Mode](#).

- **High Security** - Configure the STunnel VIP for high security
- **FIPS Compliant** - Configure the STunnel VIP for FIPS compliance
- **High Compatibility** - Configure the STunnel VIP for high compatibility
- **Custom** - All settings can be configured manually

**Note** | For more information see [SSL Operation Mode - Custom](#).











- Select the required *SSL Certificate*.

4. Click **Update** to create the new STunnel Virtual Service.

Once the SSL Termination is created, the associated VIP will be displayed with a padlock symbol in the system overview as shown below:

## System Overview

2023-02-01 16:29:09 UTC

VIRTUAL SERVICE 		IP 	PORTS 	CONNS 	PROTOCOL 	METHOD 	MODE 	
	 <b>VIP1</b>	192.168.111.232	80	0	HTTP	Layer 7	Proxy	

### SSL Operation Mode

The following STunnel settings are auto-configured for each SSL Operation Mode:

STunnel Setting	High Security	FIPS Compliant	High Compatibility
Disable SSLv3 Ciphers	✓	✓	✓
Disable TLSv1.0 Ciphers	✓	✓	✗
Disable TLSv1.1 Ciphers	✓	✓	✗
Disable TLSv1.2 Ciphers	✗	✗	✗
Disable TLSv1.3 Ciphers	✗	✗	✗
Do Not Insert Empty Fragments	✓	✓	✓
Delay DNS Lookups	✓	✓	✓
Honor Cipher Order	✓	✓	✓
Disable SSL Renegotiation	✓	✓	✓

The following SSL Ciphers are auto-configured for each SSL Operation Mode:

### High Security

ECDSA - ECDSA - AES256 - GCM - SHA384 : ECDSA - ECDSA - AES128 - GCM - SHA256 : DHE - RSA - AES256 - GCM - SHA384 : DHE - RSA - AES128 - GCM - SHA256 : ECDSA - RSA - AES256 - GCM - SHA384 : ECDSA - RSA - AES128 - GCM - SHA256

### FIPS Compliant

ECDSA - RSA - AES256 - GCM - SHA384 : ECDSA - ECDSA - AES256 - GCM - SHA384 : ECDSA - RSA - AES256 - SHA384 : ECDSA - ECDSA - AES256 - SHA384 : DHE - DSS - AES256 - GCM - SHA384 : DHE - RSA - AES256 - GCM - SHA384 : DHE - RSA - AES256 - SHA256 : DHE - DSS - AES256 - SHA256 : AES256 - GCM - SHA384 : AES256 - SHA256 : ECDSA - RSA - AES128 - GCM - SHA256 : ECDSA - ECDSA - AES128 - GCM - SHA256 : ECDSA - RSA - AES128 - SHA256 : ECDSA - ECDSA - AES128 - SHA256 : DHE - DSS - AES128 - GCM - SHA256 : DHE - RSA - AES128 - GCM - SHA256 : DHE - RSA - AES128 - SHA256 : DHE - DSS - AES128 - SHA256 : AES128 - GCM - SHA256 : AES128 - SHA256 : AES256 - SHA : AES128 - SHA

### High Compatibility

ECDSA - RSA - AES256 - GCM - SHA384 : ECDSA - RSA - AES128 - GCM - SHA256 : AES128 - GCM - SHA256 : AES256 - SHA256 : AES128 - SHA256 : AES256 - SHA : AES128 - SHA : DHE - RSA - AES256 - SHA256

### Custom (Initial Setting)

```
ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-GCM-SHA256:AES128-GCM-SHA256:AES256-SHA256:AES128-SHA256:AES256-SHA:AES128-SHA:DHE-RSA-AES256-SHA256
```

SSL Operation Mode - Custom

If you set the *SSL Operating Mode* to **Custom** the following SSL/TLS settings can be configured manually:

1. *Ciphers to use* - the default is:

**ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-GCM-SHA256:AES128-GCM-SHA256:AES256-SHA256:AES128-SHA256:AES256-SHA:AES128-SHA:DHE-RSA-AES256-SHA256**

This can be modified as required, or the field can be cleared (blank) to allow all available ciphers (not recommended)

2. *Disable SSLv3 Ciphers* - When checked this option disables all SSLv3 Ciphers.
3. *Disable TLSv1.0 Ciphers* - When checked this option disables all TLSv1.0 Ciphers.
4. *Disable TLSv1.1 Ciphers* - When checked this option disables all TLSv1.1 Ciphers.
5. *Disable TLSv1.2 Ciphers* - When checked this option disables all TLSv1.2 Ciphers.
6. *Disable TLSv1.3 Ciphers* - When checked this option disables all TLSv1.3 Ciphers.
7. *SSL Terminator* - the options are **STunnel**, **HAProxy** or **Pound**.

a. If **STunnel** is selected:

- *Do not Insert Empty Fragments* - This option should be enabled (checked) to ensure mitigation of both the BEAST and CRIME MITM attacks. It is also required for PCI Testing.
- *Delay DNS Lookups* - This option is useful for dynamic DNS, or when DNS is not available during STunnel startup (road warrior VPN, dial-up configurations).
- *Honor Cipher Order* - When choosing a cipher during an SSLv3 or TLSv1 handshake, normally the client's preference is used. If this directive is enabled, the server's preference will be used instead.
- *Disable SSL Renegotiation* - Applications of the SSL renegotiation include some authentication scenarios, or re-keying long lasting connections. On the other hand this feature can facilitate a trivial CPU-exhaustion DoS attack. This option should be enabled (checked) to mitigate the BEAST Attack.
- *Time to Close* - Configure the global client response timeout in seconds. This setting should not require changing.

b. If **HAProxy** is selected:

Note

To use HAProxy for SSL termination, the *Associated Virtual Service* field must be set to a layer 7 VIP.

- *CA Certificate* - If you want to use client certificate authentication, set the relevant CA here.

Note

To configure a *CA Certificate*, create a PEM file with the complete CA certificate chain and private key, then use the WebUI menu option: *Cluster Configuration > CA Certificate Families* to define a new CA Certificate Family and upload the PEM file. Multiple certificates can be associated with each family, the first certificate is activated by default.

c. If Pound is selected:

- *Enable WebDAV Verbs* - Selecting this option permits the use of the following commands:
  - Extended HTTP Requests: PUT, DELETE
  - Standard WebDAV verbs: LOCK, UNLOCK, PROPFIND, PROPPATCH, SEARCH, MKCOL, MOVE, COPY, OPTIONS, TRACE, MKACTIVITY, CHECKOUT, MERGE, REPORT
  - Microsoft WebDAV extensions: SUBSCRIBE, BPROPPATCH, POLL, BMOVE, BCOPY, BDELETE, CONNECT
- *Header Field Name & Header Field Value* - Use to define a custom Pound header.
- *Rewrite HTTP Redirects* - If they point to the backend itself or to the listener (but with the wrong protocol) the response will be changed to show the virtual host in the request.
- *Honor Cipher Order* - When choosing a cipher during an SSLv3 or TLSv1 handshake, normally the client's preference is used. If this directive is enabled, the server's preference will be used instead. This option should be enabled to mitigate the BEAST attack.
- *Client Cipher Renegotiation* - Sets whether the client is allowed to renegotiate the cipher order. This option should be set to "No Client Renegotiation" to mitigate the BEAST attack. The options are:
  - No Client Renegotiation - no client renegotiation will be honored
  - Secure Renegotiation - secure renegotiation will be honored
  - Insecure Renegotiation - insecure renegotiation will be honored

## Server Name Indication (SNI)

Server Name Indication (SNI) rules can be configured for STunnel VIPs. SNI is an extension to the TLS protocol which allows a client to indicate which hostname it is attempting to connect to at the start of the handshaking process. This allows the load balancer to present multiple secure websites on the same IP address and port, but with different certificates. The following section provides more details on configuring SNI.

### Configuring Server Name Indication (SNI) Rules

SNI matching allows traffic to be sent to different backend Virtual Services based on the FQDN requested. SNI rules can be configured after the STunnel VIP is created.

*To configure SNI rules:*

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Termination*.
2. Click **Modify** next to the relevant STunnel VIP.
3. Scroll to the bottom of the screen and click **New SNI Rule**.

Friendly Name	<input type="text" value="rule1"/>
SNI to match	<input type="text" value="www.loadbalancer.org"/>
SSL Certificate	<input type="text" value="Default Self Signed Certificate"/> ▼
Associated Virtual Service	<input type="text" value="VIP1"/> ▼

Add Rule

4. Enter a suitable *Friendly Name* for the new Rule, e.g. **rule1**.
5. Enter the required *SNI to Match* e.g. **www.loadbalancer.org**.
6. Select the required *SSL Certificate*.
7. Use the *Associated Virtual Service* drop-down to select the required Virtual Service, or select **Custom** and enter the required *Backend IP Address* and *Backend Virtual Service Port* and configure the *Enable Proxy Protocol* checkbox according to your requirements.
8. Click **Add Rule**.
9. Repeat the above steps to add additional rules.
10. Once the rules are added, they're displayed in a list under the *Current SNI Rules* section as shown in the example below:

#### Current SNI Rules

Search SNI Rules

Delete

SNI Name	SNI to Match	Certificate	Service	Proxy Protocol	Modify	Delete
rule1	www.loadbalancer.org	server	VIP1	<input checked="" type="checkbox"/>	Modify	<input type="checkbox"/>
rule2	www.lbtestdomain.com	server	VIP1	<input checked="" type="checkbox"/>	Modify	<input type="checkbox"/>

Delete

11. Modify or Delete SNI rules using the buttons provided.
12. To apply the new settings, restart STunnel using the **Reload STunnel** button at the top of the screen.
13. Once SNI rules have been configured for a particular STunnel VIP, this is indicated next to the STunnel VIP name as shown below:

Service Name	IP & Port	Backend & Port	Options
 SSL	192.168.110.235:443	192.168.110.235:80	<input type="button" value="Modify"/> <input type="button" value="SSL Info"/> <input type="button" value="Delete"/>

## SSL Termination on the Load Balancer with Re-encryption (SSL Bridging)

Note

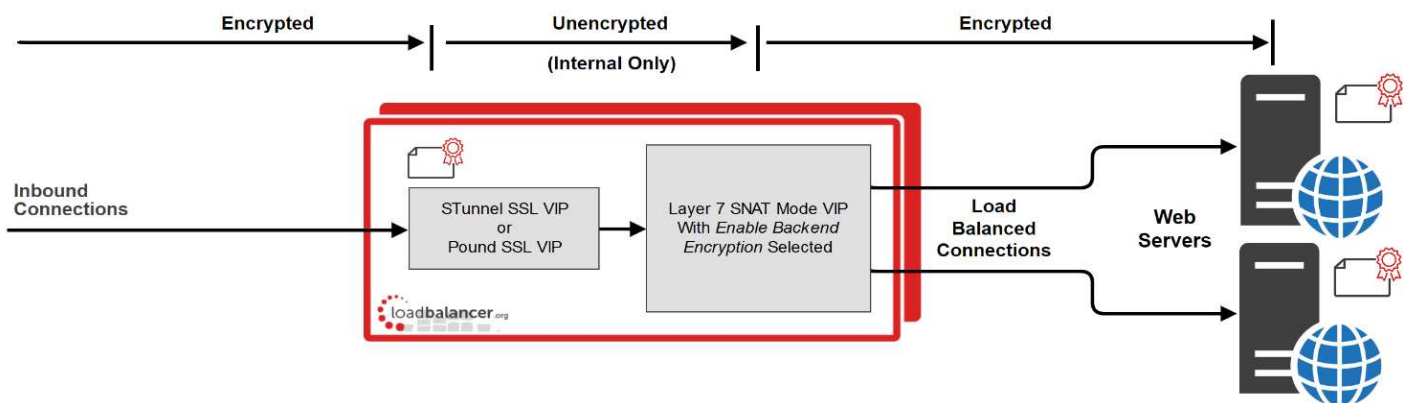
SSL termination on the load balancer can be very CPU intensive. In most cases, for a scalable solution, terminating SSL on the Real Servers is the best option.

Note

For information on configuring SSL termination, please refer to [SSL Termination on the Load Balancer \(SSL Offloading\)](#).

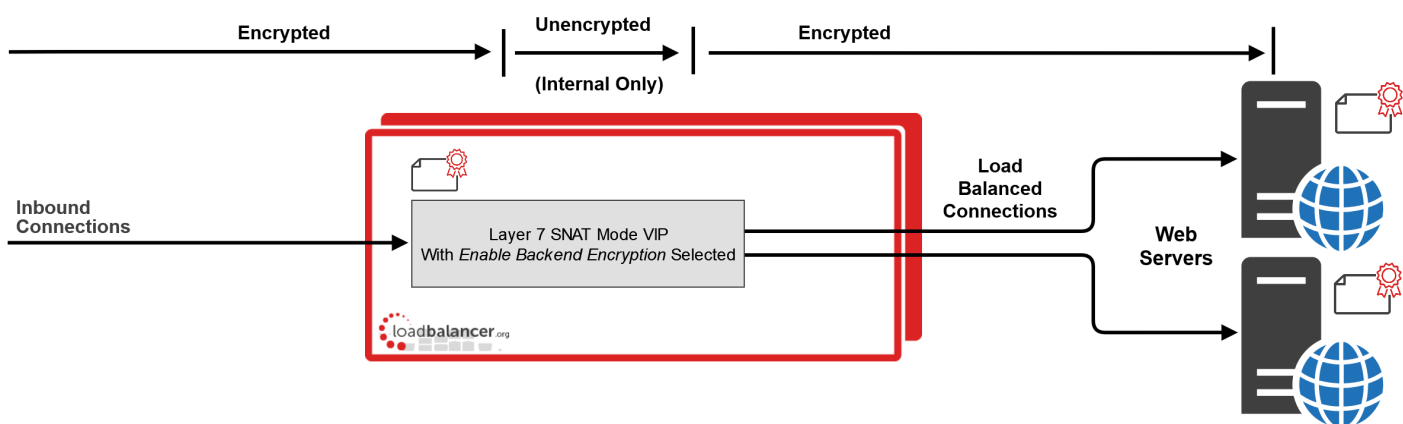
## Using STunnel or Pound to Terminate SSL

Here, SSL is terminated using STunnel or Pound and is then passed to a layer 7 VIP. Traffic is encrypted from the client to the load balancer and is encrypted from the load balancer to the backend servers as shown below.



## Using HAProxy to Terminate SSL

Here, SSL is terminated using HAProxy. In this case a single layer 7 SNAT mode VIP is all that is required to provide encryption from the client to the load balancer and from the load balancer to the backend servers as shown below.



## Notes

- Backend encryption can be enabled for the entire VIP and all associated Real Servers using the VIP option *Enable Backend encryption* or per Real Server using the *Re-Encrypt to Backend* option as detailed below.

*To enable re-encryption at the Virtual Service level:*

- Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Servers*.
- Click **Modify** next to relevant Virtual Service.

### SSL

Enable Backend Encryption



- Scroll down to the *SSL* section and check (enable) the *Enable Backend Encryption* checkbox.
- Click **Update** - you'll be asked if you want to apply the setting to all existing backend servers - click **OK** or **Cancel** as required.

**Note** | The new setting will automatically apply to all new Real Servers added.

*To enable re-encryption at the Real Server level:*



1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Real Servers*.
2. Click **Modify** next to relevant Real Server.

### Layer 7 Add a new Real Server

Label	<input type="text" value="IIS1"/>	?
Real Server IP Address	<input type="text" value="192.168.210.240"/>	?
Real Server Port	<input type="text" value="443"/>	?
Re-Encrypt to Backend	<input checked="" type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

Cancel
Update

3. Check (enable) the option *Re-Encrypt to Backend*.
4. Click **Update**.

## SSL - Advanced Configuration

### Pound Global Settings

**Lock Pound Configuration** - When enabled it will stop the user interface overwriting the configuration files so manual changes can be made.

**Logging** - Activate detailed logging of the Pound SSL termination service. When activated the Pound log is written to `/var/log/poundssl.log`.

**Client Timeout** - Configure the global client response timeout in seconds. This setting should not require changing. The default is 30 seconds.

**Global Server Timeout** - Configure the global Real Server response timeout in seconds. This setting should not require changing.

**Ulimit** - This setting will change the maximum number of file descriptors available to the pound process. The default is 81000.

**Process Threads** - Start the Pound process with X number of threads. Note that these threads are allocated at start so if you're not using them they will take up memory needlessly. The default is 250.

**Transparent Proxy** - Enable TProxy support in Pound SSL. The combination of Pound, TProxy, and HAProxy allows SSL termination on the load balancer whilst passing the client's IP address to the Real Servers. This option also automatically enables TProxy for HAProxy.

#### Note

One consequence of using Transparent Proxy with both Pound and HAProxy is that you can no longer access the HAProxy Virtual Service directly. With transparency turned on, HAProxy will only accept traffic from Pound. One way to get around this is to configure the HAProxy VIP to listen on 2 ports. One will listen on port 80, and be your standard HTTP service. The other will listen on a different port - 81 for example, and will be the destination for traffic from Pound. For

more information please refer to [Transparency at Layer 7](#).

## STunnel Global Settings

STunnel Global Settings		
Debug Level	Emergency (0) ▾	?
Disable Nagle Algorithm	<input type="checkbox"/>	?
Enable FIPS 140-2 mode	<input type="checkbox"/>	?
		<a href="#">Update</a>

**Debug Level** - Option to set the debugging level for all STunnel Services. The Debug Level is a one of the syslog level names or numbers emergency (0), Alert (1), Critical (2), err (3), Warning (4), Notice (5), Information (6), or Debug (7). The higher the number the more detail will be contained in the STunnel Logs.

**Disable Nagle Algorithm** - With this option ticked (enabled) the Nagle Algorithm will be disabled. More details can be found in RFC 896.

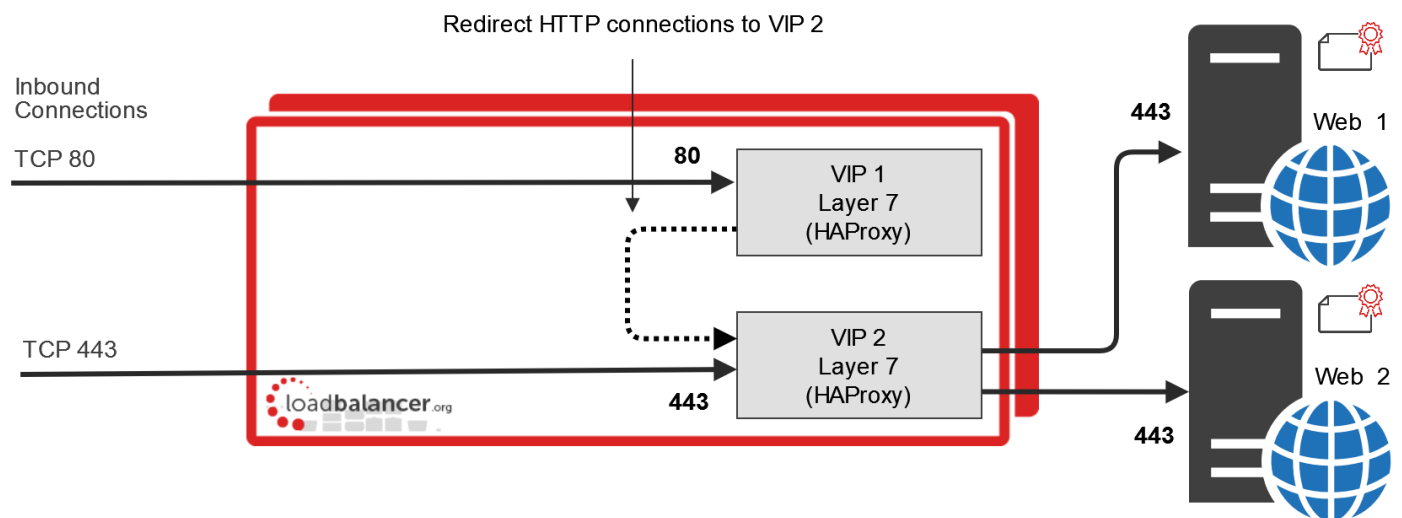
**Enable FIPS 140-2 Mode** - FIPS (Federal Information Processing Standards) are a set of standards that describe document processing, encryption algorithms and other information technology standards for use within non-military government agencies and by government contractors and vendors who work with the agencies. Check to enable FIPS 140-2 mode for STunnel.

## HTTP to HTTPS Redirection

The appliance supports the ability to force HTTP to HTTPS redirection. This can be achieved both when terminating SSL on the Real Servers and when offloading SSL on the load balancer as described in the following sections.

### When Terminating SSL on the Real Servers

This method requires 2 VIPs.



VIP 1 & VIP 2 are configured on the same IP address for HTTP/HTTPS client connections

- **VIP 1** - This is a layer 7 HTTP mode VIP that listens on port 80 and redirects all requests to VIP2. It has the option *Force to HTTPS* enabled.

**Note** | VIP1 does not require any Real Servers to be configured.

**Note** | VIP1 will show purple/green in the System Overview. This when *Force to HTTPS* is enabled and when no Real Servers are configured.

- **VIP 2** - This is a layer 7 TCP mode VIP that listens on port 443 and load balances connections between Web 1 & Web 2.

## VIP 1 Redirect Configuration

Click **Modify** next to the VIP, enable the *Other (Advanced) > Force to HTTPS* option, and set the redirect code as required as shown in the example below:

Force to HTTPS	<input checked="" type="radio"/> Yes <input type="radio"/> No	?
HTTPS Redirect Code	301 (Moved Permanently) ▼	?
HTTPS Redirect Port	<input type="text"/>	?

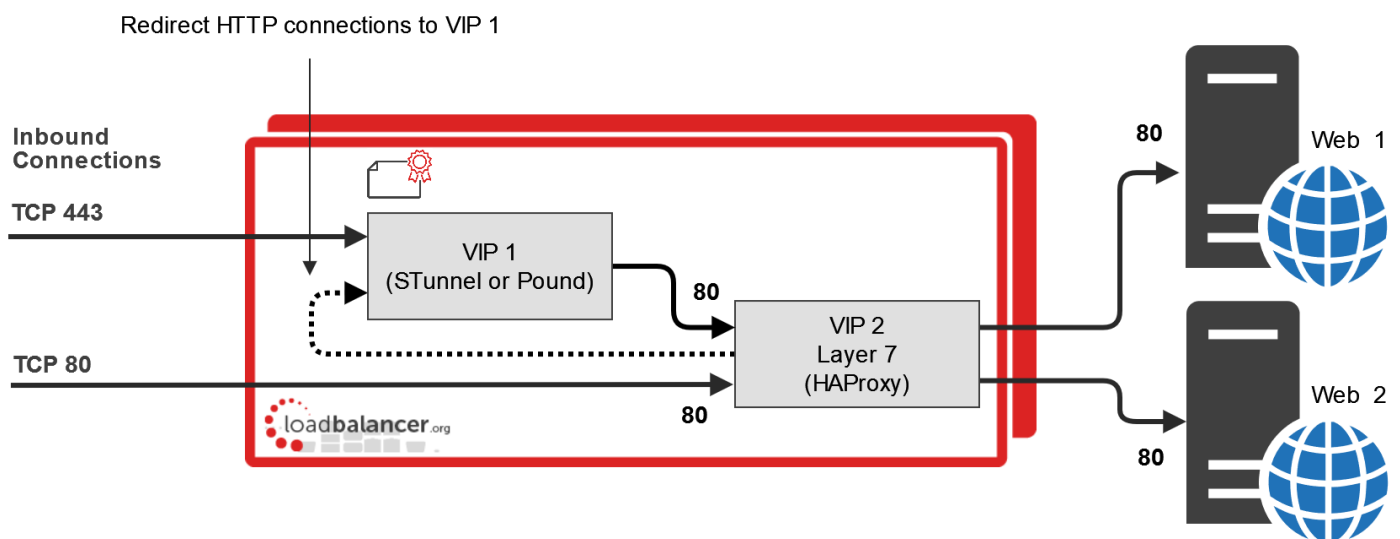
**Note** | The *Force to HTTPS* option is only available when the VIP is in HTTP mode.

**Note** | It's not possible to enable TProxy when using this configuration.

## When Terminating SSL on the Load Balancer

### Using STunnel or Pound

This method requires 2 VIPs.



VIP 1 & VIP 2 are configured on the same IP address for HTTP/HTTPS client connections

- **VIP 1** - This is a Pound or STunnel VIP that listens on port 443, terminates the SSL connection and then forwards the decrypted HTTP connections to VIP2 on port 80.

- **VIP 2** - This is a layer 7 HTTP mode VIP that listens on port 80 and load balances connections between Web 1 & Web 2. It has the option *Force to HTTPS* enabled which redirects the HTTP client connections (see below).

## VIP 2 Redirect Configuration

Click **Modify** next to the VIP, enable the *Other (Advanced)* > *Force to HTTPS* option, and set the redirect code as required as shown in the example below:

Force to HTTPS	<input checked="" type="radio"/> Yes <input type="radio"/> No	?
HTTPS Redirect Code	301 (Moved Permanently) ▼	?
HTTPS Redirect Port	<input type="text"/>	?

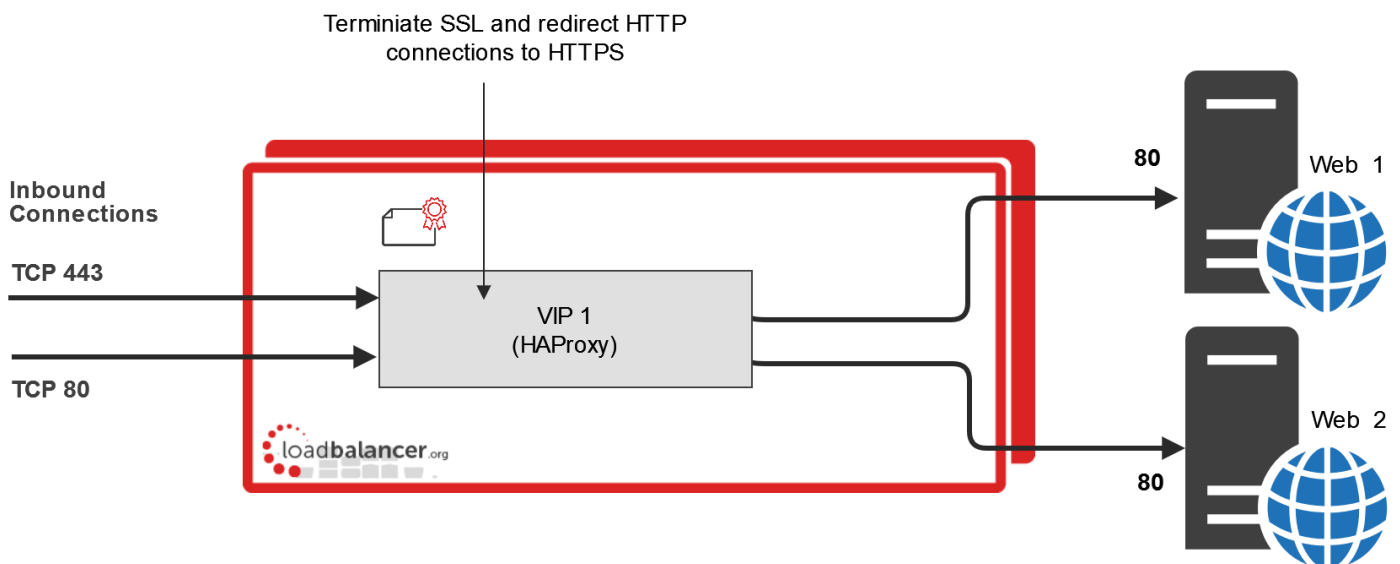
**Note** | The *Force to HTTPS* option is only available when the VIP is in HTTP mode.

**Note** | It's not possible to enable TProxy when using this configuration.

**Note** | If you want to re-encrypt the data from the load balancer to the Real Server, enable the *Re-encrypt to Backend* option for the each Real Server. For more information on using this option please refer to [SSL Termination on the Load Balancer with Re-encryption \(SSL Bridging\)](#).

## Using HAProxy




This method requires 1 VIP.



- **VIP 1** - This is a Layer 7 SNAT mode VIP configured in HTTP mode. It listens on port 80 and load balances the decrypted HTTP requests to Web 1 & Web 2. It is also configured to handle the SSL termination.

## VIP 1 Redirect Configuration

Click **Modify** next to the VIP, enable the *Other (Advanced)* > *Force to HTTPS* option, and set the redirect code as required as shown in the example below:

Force to HTTPS	<input checked="" type="radio"/> Yes <input type="radio"/> No	
HTTPS Redirect Code	301 (Moved Permanently) ▼	
HTTPS Redirect Port	<input type="text"/>	

**Note** | The *Force to HTTPS* option is only available when the VIP is in HTTP mode.

**Note** | It is possible to enable TProxy when using this configuration.

**Note** | If you want to re-encrypt the data from the load balancer to the Real Server, enable the *Re-encrypt to Backend* option for the each Real Server. For more information on using this option please refer to [SSL Termination on the Load Balancer with Re-encryption \(SSL Bridging\)](#).

## Server Feedback Agent

The load balancer can modify the weight (amount of traffic) of each server by gathering data from either a custom agent or an HTTP server. For layer 4 VIPs the feedback method can be set to either agent or HTTP, for Layer 7 VIPs, only the agent method is supported.

A telnet to port 3333 on a Real Server with the agent installed will return the current idle stats as an integer value in the range 0 - 100. The figure returned can be related to CPU utilization, RAM usage or a combination of both. This can be configured using the XML configuration file located in the agents installation folder (by default C:\ProgramData\LoadBalancer.org\LoadBalancer).

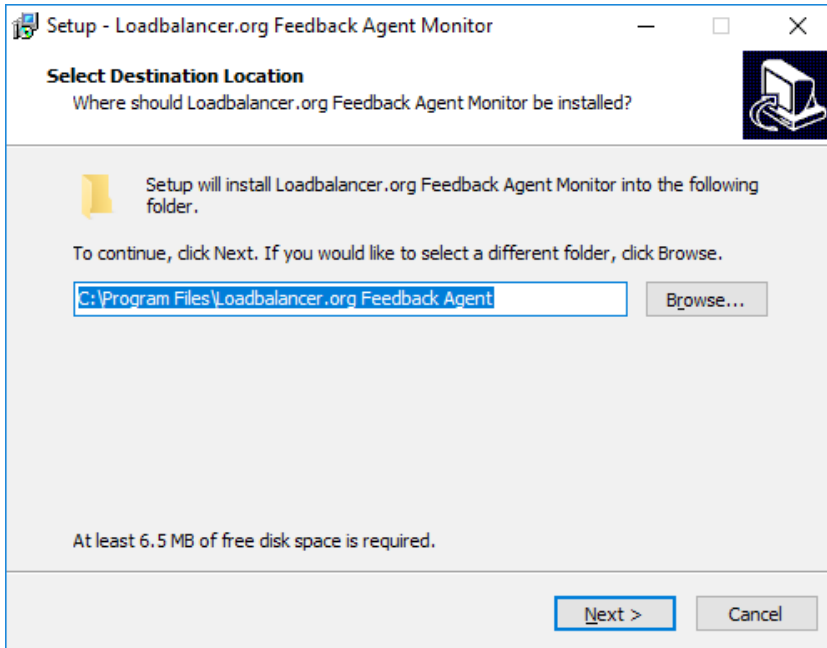
The load balancer typically expects a 0-99 integer response from the agent which by default relates to the current CPU idle state, e.g. a response of 92 would imply that the Real Servers CPU is 92% idle. The load balancer will then use the formula  $(92/100 * \text{requested\_weight})$  to find the new optimized weight.

**Note** | The 'Requested Weight' is the weight set in the WebUI for each Real Server.

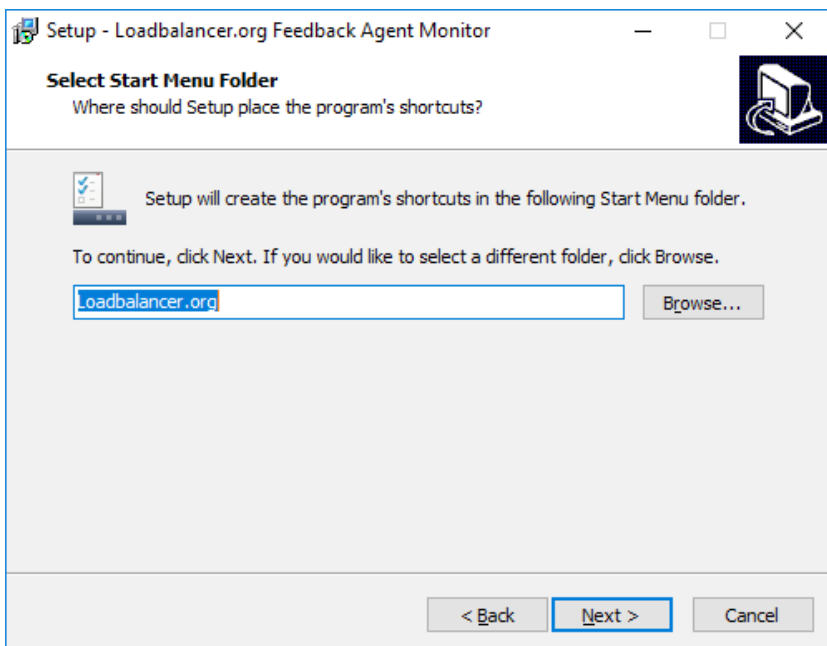
For more information about the feedback agent please refer to [this blog](#).

## Windows Agent


The latest Windows feedback agent can be downloaded from [here](#). To install the agent, run **loadbalanceragent.msi** on each Real Server:




Leave the default location or change according to your requirements, click **Next**.



Leave the default location or change according to your requirements, click **Next**.



Click **Install** to start the installation process.

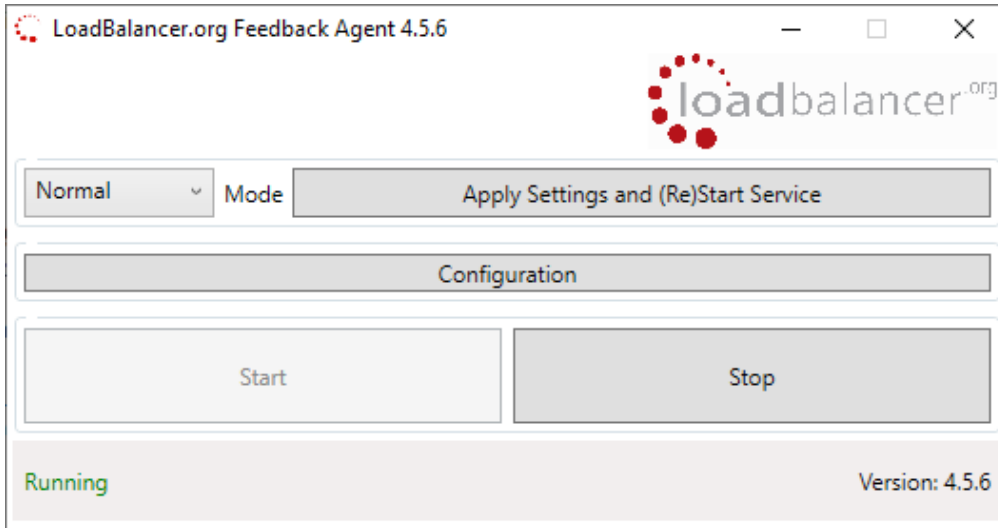


Click **Finish**.

**Note** | The agent should be installed on all Real Servers in the cluster.

## Starting the Agent

Once the installation has completed, you'll need to start the service on the Real Servers. The service is controlled by the Feedback Agent monitor & control program that is also installed along with the Agent. This can be accessed on the Windows server from: *Start> Loadbalancer.org > Loadbalancer.org Feedback Agent*. It's also possible to start the service using the services snap-in - the service is called **LBCPUMon**.



- To start the service, click the **Start** button
- To stop the service, click the **Stop** button

## Linux/Unix Agent

The Linux feedback agent files can be downloaded using the following links:

readme file: <https://downloads.loadbalancer.org/agent/linux/v4.1/readme.txt>

xinetd file: <https://downloads.loadbalancer.org/agent/linux/v4.1/lb-feedback>

feedback script: <https://downloads.loadbalancer.org/agent/linux/v4.1/lb-feedback.sh>

## Installation & Testing

Install xinetd:

```
apt-get install xinetd (if not already installed)
```

insert this line into /etc/services:

```
lb-feedback 3333/tcp # Loadbalancer.org feedback daemon
```

then run the following commands:

```
cp lb-feedback.sh /usr/bin/lb-feedback.sh
chmod +x /usr/bin/lb-feedback.sh
cp lb-feedback /etc/xinetd.d/lb-feedback
chmod 644 /etc/xinetd.d/lb-feedback
/etc/init.d/xinetd restart
```

to test:

```
telnet 127.0.0.1 3333
Trying 127.0.0.1...
Connected to 127.0.0.1.
Escape character is '^]'.
95%
```



Connection closed by foreign host.

**Note** | The agent files must be installed on all Real Servers, not the load balancer.




## Custom HTTP Agent

You can use any HTTP server responding on port 3333 to give feedback information to the load balancer. The format of this information must be an integer number of 0-100 without any header information. Using this method, you can generate a custom response based on your applications requirements.

## Configuring VIPs To Use The Agent

As mentioned, both layer 4 and layer 7 VIPs can be configured to use the feedback agent. To Configure Virtual Services to use Agent/HTTP Feedback follow the steps below:

1. Using the WUI, navigate to:
  - *Cluster Configuration > Layer 4 - Virtual Services* **Or**
  - *Cluster Configuration > Layer 7 - Virtual Services*
2. Click **Modify** next to the relevant Virtual Service

Feedback Method	
Feedback Method	Agent  
Feedback Agent Port	3333 

3. Change the Feedback Method to either **Agent** or **HTTP** for layer 4 VIPs
4. Change the Feedback Method to **Agent** for layer 7 VIPs
5. Click **Update**
6. Reload/restart services as prompted

## Global Server Load Balancing (GSLB)

GSLB enables traffic to be load balanced across multiple geographically dispersed servers. When used in conjunction with the failover and high availability features of the appliance, GSLB enables highly available, multi-site load balanced environments to be created. GSLB functionality is provided using the Open Source **Polaris GSLB**.

### Key Concepts

GSLB enables the load balancer(s) to provide intelligent DNS responses to inbound client queries for one or more sub domains. The responses given depend on the health of each endpoint and if Topology is configured, the location of those endpoints relative to the client (see note below) making the request. Where GSLB is deployed along side application load balancing, the endpoints are usually the VIPs that are configured at each site. Where application load balancing is not used and only GSLB is configured, the endpoints are normally the Real Servers.

DNS delegation is used to delegate responsibility for the sub domain(s) to the GSLB service on the load balancers. Once delegated, it is the GSLB service on the load balancers that is responsible for providing the response to DNS queries for that sub domain.

In a 2 site setup with an HA pair of load balancers in each site, once GSLB and DNS delegation are correctly configured, the 4 load balancers act as intelligent name servers for the sub domains in question.

#### Note

When topology is configured, it is the IP address of the client's local DNS server that is used to determine relative location. Therefore, if multiple client sites use the same DNS server, it's not possible for GSLB to distinguish between locations. If this is required, a local DNS server for each client site must be configured.

## Key features

- Reliable health checking service supporting both TCP, HTTP(S) and custom external checks so that only healthy members/endpoints are returned on lookups
- Failover, round robin and also a topology method that directs clients to servers in the same location
- Can return single or multiple (up to 1024) answers at once
- Option to fallback to any healthy server or refuse the query

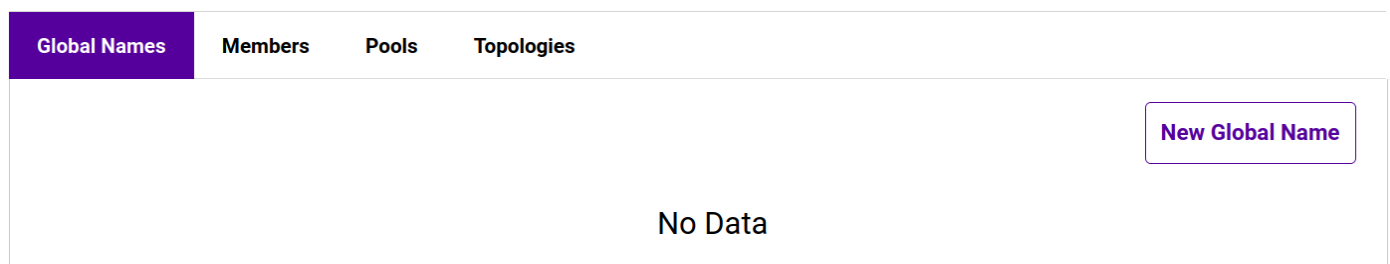
#### Note

For additional background and configuration information about GSLB, please refer to our [blog](#).

## GSLB Configuration

GSLB is configured using the WebUI menu option: *Cluster Configuration > GSLB Configuration*. 4 tabs are used to configure GSLB as illustrated below:

### GSLB Configuration



### Global Names

Global Names are used to define the FQDN's that GSLB responds to.

### Members

Members, also known as 'endpoints' are returned to clients in DNS responses.

### Pools

A Pool links together a Global Name and the relevant members and also defines the health checks, timeouts and other settings that should be used.

### Topologies

Topologies define how network subnets map to sites. In a multi-site deployment, this is used to define which site clients should connect to under normal conditions.

The following table describes the options in each tab.

Tab	Setting	Description
Global Names	Name	Name can be a combination of '0-9', 'a-z', 'A-Z', '-' (dash), '_' (underscore) or a '.' (dot).
	Hostname	A valid RFC 1123 hostname, for example www.example.com
	TTL	TTL is how long to cache the (hostname) DNS response in seconds. For example 3600 would be equal to 1 hour, minimum value is "1" (1s).
Members	Name	Name can be a combination of '0-9', 'a-z', 'A-Z', '-' (dash), '_' (underscore) or a '.' (dot).
	IP	A valid IPv4 address for example 10.0.1.1
	Monitor IP	A valid IPv4 address for example 10.0.1.1
	Weight	Weight of the server, min: 0 (server is disabled), max: 10
Pools	Name	A Name can be a combination of '0-9', 'a-z', 'A-Z', '-' (dash), '_' (underscore) or a '.' (dot).
	Monitor	The type of health check to use. The options are:
	Monitor - HTTP	<p>Perform HTTP(S) GET, succeeds if response HTTP status is 200 (or one of the codes specified in <i>Monitor Expected codes</i>).</p> <ul style="list-style-type: none"> <li>• <i>Monitor use SSL</i> - Whether to use SSL, default is 'No' (false).</li> <li>• <i>Monitor Hostname</i> - Hostname to supply in HTTP Host: header, when using SSL this will also be supplied in SNI, default is "none".</li> <li>• <i>Monitor URL Path</i> - A url path to request, appended after the member's IP address, default is "/".</li> <li>• <i>Monitor Port</i> - An integer between 1 and 65535, if value is not provided, port 80 will be used with use_ssl set to false, port 443 will be used with use_ssl set to true.</li> <li>• <i>Monitor Expected Codes</i> - An array of HTTP codes to match in a response for example "200", "301". Input range is between 100 and 599.</li> </ul>
	Monitor - TCP	<p>Perform a TCP connect. (Optionally: send text, read response, match a reg exp pattern).</p> <ul style="list-style-type: none"> <li>• <i>Monitor Port</i> - see above.</li> <li>• <i>Monitor Send String</i> - A string to send after connecting to a socket, for example 'check'.</li> <li>• <i>Monitor Match Return</i> - A reg exp to match in response, for example 'up' (the reg exp compiles with re.IGNORECASE flag set).</li> </ul>

Tab	Setting	Description
	Monitor - Forced	<p>Forces a member to be either UP or DOWN effectively disabling the health checking.</p> <ul style="list-style-type: none"> <li>• <i>Monitor Status</i> - A string, one of "up" or "down", default is "up".</li> </ul>
	Monitor - External	<p>Will run the script selected in the <i>Monitor Script</i> drop-down. The check will receive the IP address of the member the port and any additional arguments that are passed. If a 'monitor result' is set the check will be deemed a success if the script returns the configured string. If there is no 'monitor result' the exit code will be used.</p> <ul style="list-style-type: none"> <li>• <i>Monitor Port</i> - see above.</li> <li>• <i>Monitor Script</i> - Specify the script to use.</li> <li>• <i>Monitor Parameters</i> - A single or double quoted, comma separated list of additional parameters you may wish to pass in. This value is not required.</li> </ul> <div> <div>Note</div> <div> <p>The parameters that are passed into the shell script are as follows:</p> <ul style="list-style-type: none"> <li>▪ \$1 = pool member IP address</li> <li>▪ \$2 = monitor port</li> <li>▪ \$3 and onwards are taken from <i>Monitor Parameters</i></li> </ul> </div> </div> <ul style="list-style-type: none"> <li>• <i>Monitor Result</i> - A string to match in the response, for example 'success'.</li> </ul> <div> <div>Note</div> <div> <p>For more details on adding additional healthcheck scripts, please refer to <a href="#">External Health Check Scripts (GSLB)</a>.</p> </div> </div>
	Monitor - External Dynamic Weight	<ul style="list-style-type: none"> <li>• This will dynamically adjust the weight based on the output of the health check script. It should output between 0 and 10, 10 being of the highest priority and 0 being offline and removed from the pool. The exit code should be 0 at all times, anything else will report as a health check failure.</li> </ul> <ul style="list-style-type: none"> <li>▪ <i>Monitor Port</i> - see above.</li> <li>▪ <i>Monitor Script</i> - see above.</li> <li>▪ <i>Monitor Parameters</i> - see above.</li> </ul> <div> <div>Note</div> <div> <p>For more details on adding additional healthcheck scripts, please refer to <a href="#">External Health Check Scripts (GSLB)</a>.</p> </div> </div>

Tab	Setting	Description
	LB Method	<p>The load balancing method to use. The options are:</p> <ul style="list-style-type: none"> <li>• <b>wrr</b> - Weighted round robin, round robin with weighting</li> <li>• <b>twrr</b> - Topology weighted round robin, as above but the topology file is also considered to direct clients to end-points in the same region/data center</li> <li>• <b>fogroup</b> - Failover group (active-backup), with this method the first healthy end-point is handed out continuously unless it becomes unhealthy, then, the next healthy end-point is used, etc.</li> </ul>
	Global Names	A Pool can be associated with one or more global names, a pool requires at least one global name. Press "CTRL" and "click" to select multiple "globalnames".
	Members	<p>A pool must have at least one endpoint member. Drag and drop the endpoints.</p> <div> <div>Note</div> <div>If the load balancing method is set to fogroup, the order that the members are added should be the same as the order in which they are used.</div> </div>
Pools (Advanced)	Monitor Interval	In seconds, min: 1, max: 3600
	Monitor Timeout	In milliseconds, min: 100 (0.1ms), max: 10000 (10 seconds)
	Monitor Retries	Retry min: 0, max: 5, default is '0' no retries.
	Fallback	<p>Resolution behavior when all members of the pool are DOWN. The options are:</p> <ul style="list-style-type: none"> <li>• <b>any</b> - default, perform distribution among all the configured members with non-0 weight(ignore health status)</li> <li>• <b>refuse</b> - refuse all queries Note: fallback is set to "any" with all member weights set to 0 will result in a NOERROR response with no answer section data.</li> </ul>
	Max Addresses Returned	Maximum number of A records to return in response, large responses will go over TCP min: "1", max: "1024", default: "1".

## External Health Check Scripts (GSLB)

Custom GSLB health checks can be created and modified using the WebUI.

### Adding Health Check Scripts




New scripts can be created either by using the script templates or by uploading files from an external source.

### Using Script Templates

*To Create a new Script From Template:*

1. Using the WebUI, navigate to *Cluster Configuration > Health Check Scripts* and click **Add New Health Check**.

#### Health Check Details

Name:	<input type="text" value="GSLB-Custom-Check"/>	
Type:	<input type="text" value="GSLB"/>	
Template:	<input type="text" value="Example"/>	

#### Primary Node Health Check Contents

2. Specify an appropriate *Name* for the health check, e.g. **GSLB-Custom-Check**.
3. Set *Type* to **GSLB**.
4. Using the *Template* dropdown select an appropriate template from the **GSLB** section of the list, e.g. **Example**.
5. Modify the script to suit your requirements.
6. Click **Update**.

Once the health check has been added, it will appear in the Health Check Scripts list as shown below:






#### Health Check Scripts

Add New Health Check

Upload Existing Health Check

Health Check Name	Type	In-use	
SMTP	VIP	-	<div>Modify</div> <div>Delete</div>
Ping_IPv4_or_IPv6	VIP	-	<div>Modify</div> <div>Delete</div>
POP3_or_IMAP	VIP	-	<div>Modify</div> <div>Delete</div>
Exchange	VIP	-	<div>Modify</div> <div>Delete</div>
GSLB-Custom-Check	GSLB	-	<div>Modify</div> <div>Delete</div>

The new script will also appear in the *Monitor Script* dropdown when *Monitor* for a Pool is set to **External**:

Monitor	<input type="text" value="External"/>	
Monitor Port	<input type="text" value="80"/>	
Monitor Script	<div><div>GSLB-Custom-Check</div><div>GSLB-Custom-Check</div></div> 	
Monitor Parameters	<input type="text" value="`arg1`,`arg2`"/>	

To Create a new Script by Uploading an External File:

1. Using the WebUI, navigate to *Cluster Configuration > Health Check Scripts* and click **Upload Existing Health Check**.

### Health Checks - Upload Script

Health check Details

Name:

GSLB-London-Paris-Check

?

Type:

☐ Virtual Service
 ☒ GSLB

?

Contents:

Browse...

London-Paris-Check.sh

?

Secondary node contents:

Browse...

No file selected.

?

File is binary:

☐

?

Cancel

Update

2. Specify an appropriate *Name* for the health check, e.g. **GSLB-London-Paris-Check**.
3. Set *Type* to **GSLB**.
4. Click the **Choose File** button next to *Contents*.
5. Browse to and select the required file, e.g. **London-Paris-Check.sh**.

#### Note

If you have an HA Pair and the secondary node requires a different health check, click the **Choose File** button next to *Secondary Node Contents* and browse to and select the required file.

6. If the file is binary, enable the **File is Binary** checkbox - this will prevent the editor window being displayed.
7. Click **Update**.

Once the health check has been added, it will appear in the Health Check Scripts list and in the *Monitor Script* dropdown as explained in [Using Script Templates](#) above.

## GSLB Multi-site Example

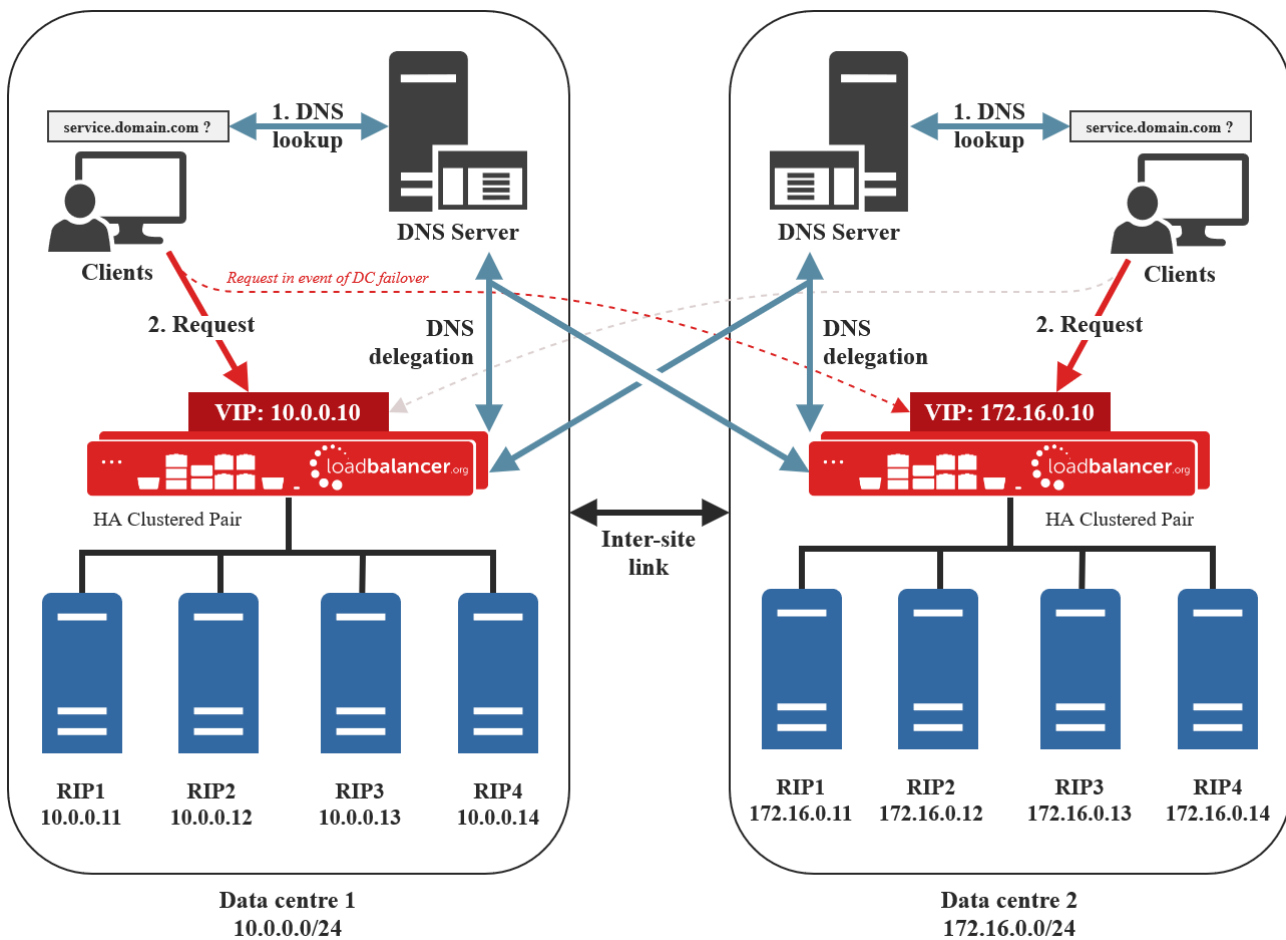
This example demonstrates the steps required to configure GSLB for a 2 site deployment. It's possible for this to be extended to encompass up to 16 sites.

### Conceptual Overview

For multi-site deployments, GSLB functionality can be used to provide high availability and location affinity across multiple sites.

- Clients across multiple sites use the same FQDN to access the load balanced service(s)
- Under normal operation, clients are directed to their site's local load balanced cluster (configured using Topology)

- In the event that a site's load balanced service(s) and/or load balancers are offline, then local clients are automatically directed to a functioning load balanced cluster at another site



#### Explanation:

1. A client tries to access the load balanced service by using the service's FQDN, in this example `service.domain.com`.
2. The client sends a DNS lookup request for `service.domain.com` to its local DNS server.
3. The local site's DNS server has the domain `service.domain.com` delegated to the load balancers.
4. The DNS server sends a delegated DNS lookup request for `service.domain.com` to one of the load balancers.
5. The load balancer that received the delegated DNS lookup request replies to the DNS server by serving up the appropriate, local VIP address. For example, if the request originated from the 10.0.0.0/24 subnet then the VIP in that subnet is served up. Likewise, if the request originated from the 172.16.0.0/24 subnet then the VIP in that subnet is served up. As such, clients are always directed to their local, on-site load balanced service, provided that the on-site instance is online and available and topology has been correctly configured.
6. The DNS server sends the delegated reply to the client.
7. The client connects to the load balanced service at `service.domain.com` by using the local VIP address.

#### Note

In the event that the load balanced cluster and/or load balancers at one site should completely fail then local clients will be directed to the load balanced cluster at the other site and the service will continue to be available. This style of multi-site failover is possible because the load balancer's GSLB functionality continuously health checks the service at each site. When the service at a site is observed to be unavailable then that site's IP address is no longer served



when responding to DNS queries.

## Appliance Configuration

GSLB must be configured on the Primary appliance at each site. The GSLB configuration must be identical across all sites to ensure consistent DNS responses irrespective of which load balancer responds. The following steps assume that an HA pair is already configured in each site. For more information on configuring HA, please refer to [Chapter 9 - Appliance Clustering for HA](#).

### Step 1 – Configure the Global Name

1. Using the WebUI on the Primary appliance for data center 1, navigate to: *Cluster Configuration > GSLB Configuration*.
2. Select the *Global Names* tab.

The screenshot shows the 'Global Names' configuration page. The 'Global Names' tab is selected. A 'New Global Name' button is located in the top right corner. Below this button is a form titled 'New Global Name'. The form contains three rows of input fields: 'Name' (value: service.domain.com), 'Hostname' (value: service.domain.com), and 'TTL' (value: 30, unit: seconds). Each row has a help icon (?) to its right. At the bottom of the form are 'Submit' and 'Cancel' buttons. Below the form, the text 'No Data' is displayed.

3. Define the required *Name* and *Hostname*, in this example both are set to **service.domain.com**.
4. Set the TTL as required, the default of **30** seconds is appropriate in many cases.
5. Click **Submit**.

### Step 2 – Configure the Members

1. Select the *Members* Tab.
2. Click the **New Member** button.

Global Names
Members
Pools
Topologies

New Member

New Member

Name	<input type="text" value="nodes-dc1"/>	?
IP	<input type="text" value="10.0.0.10"/>	?
Monitor IP	<input type="text" value="10.0.0.10"/>	?
Weight	<input type="text" value="1"/> ▾	?

Submit
Cancel

No Data

- In this example, 2 members must be created - one for the VIP in Data Center 1 and the second for the VIP in Data Center 2. Enter a *name* for the first member, e.g. **nodes-dc1** and specify the IP and Monitor IP, in this example **10.0.0.10**.
- Leave the *weight* set to **1**.
- Click **submit**.
- To create the second member, click the **New Member** button again and define the second member (**nodes-dc2**) in the same way.

### Step 3 – Configure the Pool

- Select the *Pools* Tab.

Global Names
Members
**Pools**
Topologies

New Pool

New Pool

Name	service-nodes	?
Monitor	TCP	?
Monitor Port	80	?
Monitor Send String	check	?
Monitor Match Return	up	?
LB Method	twrr	?
Global Names	service.domain.com	?
Members	<div>Available Members</div> <div>Members In Use</div> <div>nodes-dc1</div> <div>nodes-dc2</div>	?

Advanced
Submit
Cancel

No Data

- Enter a suitable name for the Pool, in this example **service-nodes**.
- Set the *monitor* to **TCP** – this will perform a basic TCP port connect to verify each member.
- Set the *monitor port* to the required value, in this example **80**.
- Set the *LB Method* to **twrr** (Topology Weighted Round Robin).
- Select the required *Global Name*.
- Drag the required Members from the Available Members list to the Members in Use list, in this example **nodes-dc1** and **nodes-dc2**.
- Click **Submit**.

## Step 4 – Configure the Topology

1. Select the *Topologies* Tab.

### GSLB Configuration

Global Names Members Pools **Topologies**

New Topology

Name

datacenter2

?

IP/CIDR

172.16.0.0/24

?

Submit

Cancel

name	ips	
datacenter1	10.0.0.0/24	<div>EditDelete</div>

2. 2 Topologies must be created, one for Data Center 1 and one for Data Center 2. Enter a *Name* for the first Topology, e.g. **datacenter1** and specify the *IP/CIDR*, in this example **10.0.0.0/24**.
3. Click **Submit**.
4. To create the second Topology, click the **New Topology** button and define the second topology (**datacenter2**) in the same way.

#### Note

Since this example has 2 data centers with an HA pair in each, you now also need to configure GSLB in the same way for the second pair using the WebUI on the Primary appliance in data center 2 ensuring that the configuration is identical.

## DNS Server Configuration

Once GSLB has been configured at both sites and is identical, the local DNS server at each site must then be configured for GSLB.

The DNS server at each site must be configured to delegate DNS requests for the subdomain in question (in this case **service.domain.com**) to the load balancers. The load balancer's GSLB services will then serve the appropriate A records to the DNS servers and then back to the client making the request.

Using the example presented here, the DNS server at each site would be configured with a delegation for the subdomain **service.domain.com**. The subdomain would be delegated to every load balancer across every site,

which provides multi-site redundancy.

The exact steps for creating a DNS delegation vary between different DNS servers and are outside the scope of this document. For further information, a blog post that walks through creating a DNS delegation on a Microsoft DNS server in the context of setting up GSLB on our appliance can be found [here](#) (see the section titled “Delegating your subdomain to your GSLBs using Microsoft’s DNS Server”).

## GSLB Diagnostics

2 reports are available to view the current state of the running GSLB service. These reports are very useful when first setting up GSLB and also when diagnosing any issues. They are available via the WebUI menu option: *Reports*.

**GSLB Generic State** – This report shows information about the running configuration of GSLB and also the health state of each member/endpoint.

Example:

```
{
  "timestamp": 1658309678.8496737,
  "globalnames": {
    "service.domain.com": {
      "pool_name": "service-nodes",
      "name": "service.domain.com",
      "ttl": 30
    }
  },
  "pools": {
    "service-nodes": {
      "last_status": true,
      "max_addrs_returned": 1,
      "members": [
        {
          "name": "nodes-dc1",
          "status": true,
          "ip": "10.0.0.10",
          "weight": 1,
          "region": "None",
          "status_reason": "monitor passed",
          "retries_left": 2,
          "monitor_ip": "10.0.0.10"
        },
        {
          "name": "nodes-dc2",
          "status": false,
          "ip": "172.16.0.10",
          "weight": 1,
          "region": "None",
          "status_reason": "socket.connect()",
          "retries_left": 0,
          "monitor_ip": "172.16.0.10"
        }
      ],
      "name": "service-nodes",
      "monitor": {
        "send_string": "None",
        "timeout": 5,
        "name": "tcp",
        "match_re": "None",
        "retries": 2,
        "interval": 10,
        "port": 80
      },
      "fallback": "any",
      "lb_method": "twrr"
    }
  }
}
```

This shows that:

- Member 10.0.0.10 is passing its health check
- Member 172.16.0.10 is failing its health check
- The running configuration matches the settings configured in the WebUI

**GSLB PPDNS State** – This report shows information about the running configuration of GSLB and also shows which results will be returned to inbound queries based on the current state of all members/endpoints.

Example:

```
{
  "timestamp": 1658309678.9508624,
  "pools": {
    "service-nodes": {
      "lb_method": "twrr",
      "max_addrs_returned": 1,
      "status": true,
      "fallback": "any",
      "dist_tables": {
        "_default": {
          "index": 0,
          "num_unique_addrs": 1,
          "rotation": [
            "10.0.0.10"
          ]
        }
      }
    }
  },
  "globalnames": {
    "service.domain.com": {
      "pool_name": "service-nodes",
      "ttl": 30
    }
  }
}
```

This shows that:

- Only member 10.0.0.10 is currently in the rotation of addresses being returned because as shown in the **GSLB Generic State** example above, member 172.16.0.10 is failing its health check
- The running configuration matches the settings configured in the WebUI

#### Note

If you want to configure a multi-site load balanced deployment using GSLB and require further assistance, please don't hesitate to contact [support@loadbalancer.org](mailto:support@loadbalancer.org).

## Configuring the Appliance via CLI, API & Direct Service Calls

A command line interface (CLI) is included that enables the appliance to be configured and controlled directly from the command line. A JSON based Application Programming Interface (API) is also provided that enables CLI commands to be called via a Web Service. The API is a wrapper around the CLI, so any CLI command can also be called using the API.

It's also possible to directly control layer 4 and layer 7 services, although the disadvantage here is that changes made will not be reflected in the System Overview. If changes are made via the CLI or API, the System Overview is kept in sync.

## Command Line Interface (CLI)

All CLI actions are invoked using the **lbcli** command. To obtain help for lbcli, run the following command:

```
lbcli --help
```

A summary of all help options is displayed:

```
lbcli --help sections    list help sections.
lbcli --help all         show help for all actions.
lbcli --help [act]      show help for action [act].
lbcli --help about       show information about LBCLI.
```

## LBCLI Command Reference

### Add an ACL rule

```
lbcli add-acl
or
lbcli --action acl --function add
```

Adds an ACL rule to the specified layer 7 virtual service. If allow-duplicates is on then the new ACL will always be appended even if there is an existing, matching ACL defined already.

#### Required Arguments

--vip	expects an argument of type existing-l7-vip
--allow-duplicates	expects an argument of type bool defaults to off
--bool	argument should be one of equal or notEqual defaults to equal
--denystatus	argument should be one of 200, 400, 403, 405, 408, 425, 429, 500, 502, 503 or 504 defaults to 403
--freetype	expects an argument of type string defaults to
--headername	expects an argument of type string defaults to
--location	expects an argument of type string defaults to
--path	expects an argument of type string defaults to /
--pathtype	argument should be one of path_beg, path_end, path_reg, hdr_host, hdr_beg, query, src_blk, header, path, sni, ssl_sni, dst_port, flag or freetype
--redirecttype	argument should be one of url_loc, url_pre, backend, drop, none, use_server or flag

### Create a Certificate Signing Request

```
lbcli add-csr
or
lbcli --action termination --function csr --type certificate
```

Create a Certificate Signing Request. When --yes-i-am-sure on is specified this command will overwrite existing certificates.

#### Required Arguments

--country	expects an argument of type string
--province	expects an argument of type string
--city	expects an argument of type string
--organisation	expects an argument of type string
--unit	expects an argument of type string
--domain	expects an argument of type string
--email	expects an argument of type string
--csrsz	argument should be one of 2048 or 4096 defaults to 4096
--signalgorithm	argument should be one of sha256 or sha256 defaults to

```
--csrname      sha256
                expects an argument of type string
```

#### Optional Arguments

```
--yes-i-am-sure  expects an argument of type bool
```

#### Add a floating IP address

```
lbcli add-floating-ip
```

Adds a floating IP to the cluster.

#### Required Arguments

```
--ip  expects an argument of type ip
```

#### Add a global name

```
lbcli add-gslb-globalname
or
lbcli --action gslb --function add --section globalnames
```

Adds a GSLB global name.

#### Required Arguments

```
--name      expects an argument of type string
--hostname  expects an argument of type hostname
--ttl       expects an argument of type uint
```

#### Add a member

```
lbcli add-gslb-member
or
lbcli --action gslb --function add --section members
```

Adds a GSLB member.

#### Required Arguments

```
--name      expects an argument of type string
--monitor-ip expects an argument of type ip
--ip        expects an argument of type ip
--weight    expects an argument of type gslb-weight
--add-pool  expects an argument of type string
```

#### Add a pool

```
lbcli add-gslb-pool
or
lbcli --action gslb --function add --section pools
```

Adds a GSLB pool.



#### Required Arguments

--add-globalname	expects an argument of type string
--fallback	argument should be one of any or refuse
--monitor	argument should be one of http, tcp, forced, external or externaldynamicweight
--monitor-interval	expects an argument of type uint
--name	expects an argument of type string

#### Optional Arguments

--add-member	expects an argument of type string
--lb-method	argument should be one of wrr, twrr or fogroup
--max-addrs-returned	expects an argument of type uint
--monitor-expected-codes	expects an argument of type http-code
--monitor-expected-codes	expects an argument of type http-code
--monitor-hostname	expects an argument of type hostname
--monitor-parameters	expects an argument of type string
--monitor-port	expects an argument of type port
--monitor-result	expects an argument of type string
--monitor-retries	expects an argument of type uint
--monitor-return-match	expects an argument of type string
--monitor-script	expects an argument of type string
--monitor-send-string	expects an argument of type string
--monitor-status	argument should be one of up or down
--monitor-timeout	expects an argument of type uint
--monitor-url-path	expects an argument of type string
--monitor-use-ssl	expects an argument of type bool

#### Add a topology

```
lbcli add-gslb-topology
or
lbcli --action gslb --function add --section topologies
```

Adds a GSLB topology.

#### Required Arguments

--name	expects an argument of type string
--------	------------------------------------

#### Optional Arguments

--add-ips	expects an argument of type string
-----------	------------------------------------

#### Add a header

```
lbcli add-header
or
lbcli --action headers --function add
```

Adds a header rule to the specified layer 7 virtual service.

#### Required Arguments

--header-type	argument should be one of http-request or http-response
--vip	expects an argument of type existing-l7-vip
--header-option	argument should be one of add, set, del, delete or replace
--header-name	expects an argument of type string
--header-value	expects an argument of type string

## Add a health check

```
lbcli add-healthcheck
```

Add a health check to the cluster. If the `--slave` argument is omitted then the script passed to the `--master` argument will be used on both appliances.

### Required Arguments

```
--label    expects an argument of type string
--kind     argument should be one of check or gslbcheck defaults to check
--master   expects an argument of type string
```

### Optional Arguments

```
--slave    expects an argument of type string
```

## Add a layer 4 real server

```
lbcli add-layer4-server
or
lbcli --action add-rip --layer 4
```

Adds a layer 4 real server to the specified virtual service.

### Required Arguments

```
--vip      expects an argument of type existing-vip
--ip       expects an argument of type ip
--rip      expects an argument of type rip
```

### Optional Arguments

```
--port     expects an argument of type port
--weight   expects an argument of type rip-weight defaults to 100
--minconns expects an argument of type uint
--maxconns expects an argument of type uint
```

## Add a layer 4 virtual service

```
lbcli add-layer4-service
or
lbcli --action add-vip --layer 4
```

Adds a layer 4 virtual service to the cluster.

### Required Arguments

```
--ip       expects an argument of type ip
--ports    expects an argument of type port-list
--vip      expects an argument of type vip
```

### Optional Arguments

```
--forwarding argument should be one of gate, masq, ipip or snat
              defaults to gate
--protocol  argument should be one of tcp, udp, tcpudp, ops or fwm
              defaults to tcp
```

--slave-ip	expects an argument of type ip
--granularity	expects an argument of type cidr
--fallback-ip	expects an argument of type ip
--fallback-port	expects an argument of type port
--persistent	expects an argument of type bool
--persist-time	expects an argument of type uint
--scheduler	argument should be one of wlc, wrr or dh
--feedback	argument should be one of agent or none
--email	expects an argument of type string
--email-from	expects an argument of type string
--check-service	argument should be one of imap, imaps, ldap, nntp, mysql or radius
--check-vhost	expects an argument of type string
--check-database	expects an argument of type string
--check-login	expects an argument of type string
--check-password	expects an argument of type string
--check-type	argument should be one of negotiate, connect, ping, external, off, on, 5 or 10
--check-port	expects an argument of type port
--check-request	expects an argument of type string
--check-response	expects an argument of type string
--check-secret	expects an argument of type string
--check-command	expects an argument of type string
--feedback-port	expects an argument of type port
--fallback-local	expects an argument of type bool
--autoscale-group	expects an argument of type string

## Add a layer 7 real server

```
lbcli add-layer7-server
or
lbcli --action add-rip --layer 7
```

Adds a layer 7 real server to the specified virtual service.

### Required Arguments

--vip	expects an argument of type existing-vip
--ip	expects an argument of type ip
--rip	expects an argument of type rip

### Optional Arguments

--port	expects an argument of type port
--weight	expects an argument of type rip-weight defaults to 100
--encrypted	expects an argument of type bool defaults to off
--redir	expects an argument of type string
--redir-enabled	expects an argument of type bool defaults to off
--minconns	expects an argument of type uint
--maxconns	expects an argument of type uint

## Add a layer 7 virtual service

```
lbcli add-layer7-service
or
lbcli --action add-vip --layer 7
```

Adds a layer 7 virtual service to the cluster.

### Required Arguments

--ip	expects an argument of type ip
--ports	expects an argument of type port-list
--vip	expects an argument of type vip

## Optional Arguments

```
--mode                argument should be one of tcp, ip or http
                        defaults to http
--fallback-ip          expects an argument of type ip
--fallback-port        expects an argument of type port
--service-type         expects an argument of type string
--slave-ip             expects an argument of type ip
--appsession-cookie    expects an argument of type string
--as-port              expects an argument of type port
--autoscale-group      expects an argument of type string
--backend-encryption   expects an argument of type bool
--backend-only         expects an argument of type bool
--check-host          expects an argument of type string
--check-negate-response expects an argument of type bool
--check-password       expects an argument of type string
--check-port           expects an argument of type port
--check-receive        expects an argument of type string
--check-request        expects an argument of type string
--check-type           argument should be one of connect, external,
                        negotiate_http, negotiate_https,
                        negotiate_http_head, negotiate_https_head,
                        negotiate_http_options, negotiate_https_options,
                        mysql or none
--check-username       expects an argument of type string
--clear-stick-drain    expects an argument of type bool
--compression          expects an argument of type bool
--cookie-maxidle       expects an argument of type haproxy-interval
--cookie-maxlife       expects an argument of type haproxy-interval
--cookiename           expects an argument of type string
--enable-hsts          expects an argument of type bool
--encrypt-all-backends expects an argument of type bool
--external-check-script expects an argument of type existing-healthcheck
--fallback-disabled    expects an argument of type bool
--fallback-encrypt     expects an argument of type bool
--fallback-persist     expects an argument of type bool
--feedback-method      argument should be one of agent or none
--feedback-port        expects an argument of type port
--force-to-https       expects an argument of type bool
--force-to-https-port  expects an argument of type port
--forward-for          expects an argument of type bool
--hsts-month           expects an argument of type uint
--http-pipeline        argument should be one of http_keep_alive,
                        http_close, http_server_close or
                        http_force_close
--http-pretend-keepalive expects an argument of type bool
--http-request         expects an argument of type bool
--invalid-http         expects an argument of type bool
--maxconn              expects an argument of type uint
--no-write             expects an argument of type bool
--persist-table-size   expects an argument of type uint
--persist-time         expects an argument of type uint
--persistence          argument should be one of none, ip, tcp,
                        sslsesid, http, appsession, rdp-session,
                        rdp-cookie, http_ip, waf or last
--proxy-bind           expects an argument of type string
--redirect-code        expects an argument of type http-300-code
--redispatch           expects an argument of type bool
--reuse-idle           expects an argument of type bool
--scheduler            argument should be one of first, roundrobin or
                        leastconn
--send-proxy           argument should be one of none, v1, v2, v2_ssl
                        or v2_ssl_cn defaults to none
--send-rip-name-as-host expects an argument of type bool
--slave-ip             expects an argument of type ip
--source-address        expects an argument of type ip
--stunnel-source       expects an argument of type string
--stunnelproxy         expects an argument of type bool
--tcp-keep-alive       expects an argument of type bool
--timeout              expects an argument of type bool
--timeout-client       expects an argument of type haproxy-interval
--timeout-server       expects an argument of type haproxy-interval
--tproxy              expects an argument of type bool
--tunneltimeout        expects an argument of type haproxy-interval
--use-content-inspection expects an argument of type bool
--xff-ip-pos           expects an argument of type int
```

## Add a PBR rule

```
lbcli add-pbr
or
lbcli --action pbr --function set
lbcli set-pbr
```

Adds a policy based routing rule to the appliance.

### Required Arguments

```
--ip          expects an argument of type ip
--gateway     expects an argument of type ip
```

## Add a static IP address

```
lbcli add-static-ip
or
lbcli --action address --function add
```

Adds a static IP to the specified virtual service.

### Required Arguments

```
--interface  expects an argument of type existing-interface
--address    expects an argument of type ip-with-optional-cidr
```

### Optional Arguments

```
--cidr       expects an argument of type cidr
```

## Add a static route

```
lbcli add-static-route
or
lbcli --action route --function static --type add
```

Adds a route via the given gateway to the specified network.

### Required Arguments

```
--gateway  expects an argument of type ip
--network  expects an argument of type ip-with-optional-cidr
```

### Optional Arguments

```
--cidr     expects an argument of type cidr
```

## Add an SSL Termination

```
lbcli add-termination
or
lbcli --action termination --function add --type stunel
```

Adds an SSL Termination to the cluster.

#### Required Arguments

--port	expects an argument of type port
--disablessl2	expects an argument of type bool defaults to true
--disablessl3	expects an argument of type bool defaults to true
--disabletlsv1	expects an argument of type bool defaults to true
--stunnelnsdelay	expects an argument of type bool defaults to true
--stunnelproxy	expects an argument of type bool defaults to false
--servercipherorder	expects an argument of type bool defaults to true
--emptyfragments	expects an argument of type bool defaults to true
--stunnelrenegotiation	expects an argument of type bool defaults to true
--stunneltimetoclose	expects an argument of type uint defaults to 0
--sslcert	expects an argument of type existing-certificate defaults to server
--proxy-bind	expects an argument of type bool defaults to false
--disabletlsv1-1	expects an argument of type bool defaults to false
--disabletlsv1-2	expects an argument of type bool defaults to false
--disabletlsv1-3	expects an argument of type bool defaults to false
--sslmode	argument should be one of high, fips or compatable defaults to high

#### Optional Arguments

--vip	expects an argument of type string
--associated-to	expects an argument of type existing-l7-vip
--proxy-bind	expects an argument of type existing-l7-vip
--source	expects an argument of type ip
--slave-source	expects an argument of type ip
--ciphers	expects an argument of type cipher-list
--ip	expects an argument of type ip
--slave-ip	expects an argument of type ip
--backend-ip	expects an argument of type ip
--slave-backend-ip	expects an argument of type ip
--backend-port	expects an argument of type port
--stunnel-source	expects an argument of type string

#### Add a WAF

```
lbcli add-waf
```

Adds a Web Application Firewall to the specified layer 7 virtual service.

#### Required Arguments

--waf	expects an argument of type string
--vip	expects an argument of type existing-l7-vip

#### Clone a health check

```
lbcli clone-healthcheck
```

Clone an existing health check on the cluster.

#### Required Arguments

--label	expects an argument of type existing-healthcheck
---------	--

Clone a virtual service.

```
lbcli clone-layer4-service
or
lbcli --action clone-vip --layer 4
```

Create a copy of the specified virtual service.

#### Required Arguments

--vip	expects an argument of type existing-vip
--ip	expects an argument of type ip
--ports	expects an argument of type port-list
--clone	expects an argument of type label

#### Optional Arguments

--forwarding	argument should be one of gate, masq, ipip or snat
--protocol	argument should be one of tcp, udp, tcpudp, ops or fwm
--slave-ip	expects an argument of type ip
--granularity	expects an argument of type cidr
--fallback-ip	expects an argument of type ip
--fallback-port	expects an argument of type port
--persistent	expects an argument of type bool
--persist-time	expects an argument of type uint
--scheduler	argument should be one of wlc, wrr or dh
--feedback	argument should be one of agent or none
--email	expects an argument of type string
--email-from	expects an argument of type string
--check-service	argument should be one of imap, imaps, ldap, nntp, mysql or radius
--check-vhost	expects an argument of type string
--check-database	expects an argument of type string
--check-login	expects an argument of type string
--check-password	expects an argument of type string
--check-type	argument should be one of negotiate, connect, ping, external, off, on, 5 or 10
--check-port	expects an argument of type port
--check-request	expects an argument of type string
--check-response	expects an argument of type string
--check-secret	expects an argument of type string
--check-command	expects an argument of type string
--feedback-port	expects an argument of type port
--fallback-local	expects an argument of type bool
--autoscale-group	expects an argument of type string

Clone a virtual service.

```
lbcli clone-layer7-service
or
lbcli --action clone-vip --layer 7
```

Create a copy of the specified virtual service.

#### Required Arguments

--vip	expects an argument of type existing-vip
--ip	expects an argument of type ip
--ports	expects an argument of type port-list
--clone	expects an argument of type label

#### Optional Arguments

--appsession-cookie	expects an argument of type string
--as-port	expects an argument of type port
--autoscale-group	expects an argument of type string
--backend-encryption	expects an argument of type bool
--backend-only	expects an argument of type bool
--check-host	expects an argument of type string
--check-negate-response	expects an argument of type bool
--check-password	expects an argument of type string
--check-port	expects an argument of type port
--check-receive	expects an argument of type string
--check-request	expects an argument of type string

--check-type	argument should be one of connect, external, negotiate_http, negotiate_https, negotiate_http_head, negotiate_https_head, negotiate_http_options, negotiate_https_options, mysql or none
--check-username	expects an argument of type string
--clear-stick-drain	expects an argument of type bool
--compression	expects an argument of type bool
--cookie-maxidle	expects an argument of type haproxy-interval
--cookie-maxlife	expects an argument of type haproxy-interval
--cookie-name	expects an argument of type string
--enable-hsts	expects an argument of type bool
--encrypt-all-backends	expects an argument of type bool
--external-check-script	expects an argument of type existing-healthcheck
--fallback-disabled	expects an argument of type bool
--fallback-encrypt	expects an argument of type bool
--fallback-persist	expects an argument of type bool
--feedback-method	argument should be one of agent or none
--feedback-port	expects an argument of type port
--force-to-https	expects an argument of type bool
--force-to-https-port	expects an argument of type port
--forward-for	expects an argument of type bool
--hsts-month	expects an argument of type uint
--http-pipeline	argument should be one of http_keep_alive, http_close, http_server_close or http_force_close
--http-pretend-keepalive	expects an argument of type bool
--http-request	expects an argument of type bool
--invalid-http	expects an argument of type bool
--maxconn	expects an argument of type uint
--no-write	expects an argument of type bool
--persist-table-size	expects an argument of type uint
--persist-time	expects an argument of type uint
--persistence	argument should be one of none, ip, tcp, sslsesid, http, appsession, rdp-session, rdp-cookie, http_ip, waf or last
--proxy-bind	expects an argument of type string
--redirect-code	expects an argument of type http-300-code
--redispatch	expects an argument of type bool
--reuse-idle	expects an argument of type bool
--scheduler	argument should be one of roundrobin, leastconn or first
--send-proxy	expects an argument of type bool
--send-rip-name-as-host	expects an argument of type bool
--slave-ip	expects an argument of type ip
--source-address	expects an argument of type ip
--stunnel-source	expects an argument of type string
--stunnelproxy	expects an argument of type bool
--tcp-keep-alive	expects an argument of type bool
--timeout	expects an argument of type bool
--timeout-client	expects an argument of type haproxy-interval
--timeout-server	expects an argument of type haproxy-interval
--tproxy	expects an argument of type bool
--tunneltimeout	expects an argument of type haproxy-interval
--use-content-inspection	expects an argument of type bool
--xff-ip-pos	expects an argument of type int

## Delete an ACL rule

```
lbcli delete-acl
or
lbcli --action acl --function del
lbcli --action acl --function delete
lbcli del-acl
```

Delete the first matching ACL rule from the specified layer 7 virtual service.

## Required Arguments

--vip	expects an argument of type existing-l7-vip
--bool	argument should be one of equal or notEqual defaults to equal
--denystatus	argument should be one of 200, 400, 403, 405, 408, 425, 429, 500, 502, 503 or 504 defaults to 403



```
--freetype      expects an argument of type string defaults to
--headername    expects an argument of type string defaults to
--location      expects an argument of type string defaults to
--path          expects an argument of type string defaults to /
--pathtype      argument should be one of path_reg, path_beg, path_end,
                hdr_host, hdr_beg, query, src_blk, header, path, sni,
                ssl_sni, dst_port, flag or freetype
--redirecttype  argument should be one of url_loc, url_pre, backend, drop,
                none or use_server
```

## Delete a floating IP

```
lbcli delete-floating-ip
or
lbcli del-floating-ip
```

Delete a floating IP from the cluster.

### Required Arguments

```
--ip  expects an argument of type ip
```

## Delete a global name

```
lbcli delete-gslb-globalname
or
lbcli --action gslb --function del --section globalnames
lbcli --action gslb --function delete --section globalnames
lbcli del-gslb-globalname
```

Delete the specified GSLB global name.

### Required Arguments

```
--name  expects an argument of type string
```

## Delete a member

```
lbcli delete-gslb-member
or
lbcli --action gslb --function del --section members
lbcli --action gslb --function delete --section members
lbcli del-gslb-member
```

Delete the specified GSLB member.

### Required Arguments

```
--name  expects an argument of type string
```

## Delete a pool

```
lbcli delete-gslb-pool
or
lbcli --action gslb --function del --section pools
lbcli --action gslb --function delete --section pools
lbcli del-gslb-pool
```

Delete the specified GSLB pool.

#### Required Arguments

```
--name expects an argument of type string
```

#### Delete a member

```
lbcli delete-gslb-topology
or
lbcli --action gslb --function del --section topologies
lbcli --action gslb --function delete --section topologies
lbcli del-gslb-topology
```

Delete the specified GSLB member.

#### Required Arguments

```
--name expects an argument of type string
```

#### Delete a header rule

```
lbcli delete-header
or
lbcli --action headers --function delete
lbcli --action headers --function del
lbcli del-header
```

Delete the first matching header rule from the specified layer 7 virtual service.

#### Required Arguments

```
--vip expects an argument of type existing-l7-vip
--header-option argument should be one of add, set, del, delete or
replace
--header-name expects an argument of type string
```

#### Delete a health check

```
lbcli delete-healthcheck
or
lbcli del-healthcheck
```

Delete an existing health check from the cluster.

#### Required Arguments

```
--label expects an argument of type existing-healthcheck
```

#### Delete a header rule

```
lbcli delete-pbr
or
lbcli --action pbr --function del
lbcli --action pbr --function delete
lbcli del-pbr
```

Delete the first matching header rule from the specified layer 7 virtual service.

#### Required Arguments

```
--ip      expects an argument of type ip
--gateway expects an argument of type ip
```

#### Delete a real server

```
lbcli delete-server
or
lbcli del-rip
lbcli delete-rip
```

Delete the real server from the specified virtual service.

#### Required Arguments

```
--vip expects an argument of type existing-vip
--rip expects an argument of type existing-rip
```

#### Delete a virtual service

```
lbcli delete-service
or
lbcli del-service
lbcli del-vip
lbcli delete-vip
```

Delete a virtual service from the cluster.

#### Required Arguments

```
--vip expects an argument of type existing-vip
```

#### Delete a static IP address

```
lbcli delete-static-ip
or
lbcli --action address --function delete
lbcli --action address --function del
lbcli del-static-ip
```

Delete a static IP from the appliance.

#### Required Arguments

```
--interface expects an argument of type existing-interface
--address    expects an argument of type ip-with-optional-cidr
```

#### Optional Arguments

```
--cidr      expects an argument of type cidr
```

#### Delete a static route

```
lbcli delete-static-route
```

```
or
lbcli --action route --function static --type del
lbcli --action route --function static --type delete
lbcli del-static-route
```

Delete a static route from the appliance.

#### Required Arguments

```
--network expects an argument of type ip-with-optional-cidr
```

#### Optional Arguments

```
--cidr expects an argument of type cidr
```

### Delete an SSL termination

```
lbcli delete-termination
or
lbcli --action termination --function del --type stunnel
lbcli --action termination --function delete --type stunnel
lbcli del-termination
```

Delete an SSL termination from the appliance.

#### Required Arguments

```
--vip expects an argument of type existing-termination
```

### Delete a WAF

```
lbcli delete-waf
```

Delete the specified Web Application Firewall from the cluster.

#### Required Arguments

```
--waf expects an argument of type existing-waf
--vip expects an argument of type existing-vip
```

### Disable the API endpoint

```
lbcli disable-api
or
lbcli --action api --function disable
```

Disable the API endpoint on the appliance.

### Disable Floating IP.

```
lbcli disable-floating-ip
```

Make it so that the given floating IP cannot be made active on this appliance.

## Drain a real server

```
lbcli drain
```

Drain connections from the specified real server. No new connections will be allowed to the real server but existing connections will remain until they are closed.

### Required Arguments

```
--vip    expects an argument of type existing-vip
--rip    expects an argument of type existing-rip
```

## Edit a global name

```
lbcli edit-gslb-globalname
or
lbcli --action gslb --function edit --section globalnames
```

Edit a GSLB global name.

### Required Arguments

```
--name      expects an argument of type string
--hostname  expects an argument of type hostname
--ttl       expects an argument of type uint
```

## Edit a member

```
lbcli edit-gslb-member
or
lbcli --action gslb --function edit --section members
```

Edit a GSLB member.

### Required Arguments

```
--name      expects an argument of type string
```

### Optional Arguments

```
--add-pool  expects an argument of type string
--delete-pool expects an argument of type string
--ip        expects an argument of type ip
--weight    expects an argument of type gslb-weight
--monitor-ip expects an argument of type ip
```

## Edit a pool

```
lbcli edit-gslb-pool
or
lbcli --action gslb --function edit --section pools
```

Edit a GSLB pool.

### Required Arguments

--name	expects an argument of type string
--------	------------------------------------

#### Optional Arguments

--add-globalname	expects an argument of type string
--add-member	expects an argument of type string
--delete-globalname	expects an argument of type string
--delete-member	expects an argument of type string
--fallback-method	expects an argument of type string
--lb-method	argument should be one of wrr, twrr or fogroup
--max-addrs-returned	expects an argument of type string
--monitor	argument should be one of http, tcp, forced, external or externaldynamicweight
--monitor-expected-codes	expects an argument of type http-code
--monitor-hostname	expects an argument of type string
--monitor-interval	expects an argument of type string
--monitor-parameters	expects an argument of type string
--monitor-port	expects an argument of type string
--monitor-result	expects an argument of type string
--monitor-retries	expects an argument of type string
--monitor-return-match	expects an argument of type string
--monitor-script	expects an argument of type string
--monitor-send-string	expects an argument of type string
--monitor-status	argument should be one of up or down
--monitor-timeout	expects an argument of type string
--monitor-url-path	expects an argument of type string
--monitor-use-ssl	expects an argument of type bool

#### Edit a topology

```
lbcli edit-gslb-topology
or
lbcli --action gslb --function edit --section topologies
```

Edit a GSLB topology.

#### Required Arguments

--name	expects an argument of type string
--------	------------------------------------

#### Optional Arguments

--add-ips	expects an argument of type string
--delete-ips	expects an argument of type string

#### Modify a health check

```
lbcli edit-healthcheck
```

Modify an existing health check on the cluster. If the --master argument is omitted then the --slave argument cannot be specified.

#### Required Arguments

--label	expects an argument of type existing-healthcheck
---------	--

#### Optional Arguments

--new-label	expects an argument of type string
--kind	argument should be one of check or gslbcheck
--master	expects an argument of type string

--slave expects an argument of type string

## Edit a layer 4 real server

```
lbcli edit-layer4-server
or
lbcli --action edit-rip --layer 4
```

Edit the specified layer 4 real server.

### Required Arguments

--vip expects an argument of type existing-vip  
--rip expects an argument of type existing-rip

### Optional Arguments

--ip expects an argument of type ip  
--port expects an argument of type port  
--weight expects an argument of type rip-weight  
--minconns expects an argument of type uint  
--maxconns expects an argument of type uint

## Edit a layer 4 virtual service

```
lbcli edit-layer4-service
or
lbcli --action edit-vip --layer 4
```

Edit the specified layer 4 virtual service.

### Required Arguments

--vip expects an argument of type existing-vip

### Optional Arguments

--autoscale-group expects an argument of type string  
--check-command expects an argument of type string  
--check-database expects an argument of type string  
--check-login expects an argument of type string  
--check-password expects an argument of type string  
--check-port expects an argument of type port  
--check-request expects an argument of type string  
--check-response expects an argument of type string  
--check-secret expects an argument of type string  
--check-service argument should be one of imap, imaps, ldap, nntp, mysql or radius  
--check-type argument should be one of negotiate, connect, ping, external, off, on, 5 or 10  
--check-vhost expects an argument of type string  
--email expects an argument of type string  
--email-from expects an argument of type string  
--fallback-ip expects an argument of type ip  
--fallback-local expects an argument of type bool  
--fallback-port expects an argument of type port  
--feedback argument should be one of agent or none  
--feedback-port expects an argument of type port  
--force-to-https-port expects an argument of type port-or-empty  
--forwarding argument should be one of snat, masq, gate or ipip  
--granularity expects an argument of type cidr  
--ip expects an argument of type ip  
--negate-response expects an argument of type bool

--persist-time	expects an argument of type uint
--persistent	expects an argument of type bool
--ports	expects an argument of type port-list
--protocol	argument should be one of ops, udp, tcp, tcpudp or fwm
--scheduler	argument should be one of wlc, wrr or dh
--send-rip-name-as-host	expects an argument of type bool
--slave-ip	expects an argument of type ip

## Edit a layer 7 real server

```
lbcli edit-layer7-server
or
lbcli --action edit-rip --layer 7
```

Edit the specified layer 7 real server.

### Required Arguments

--vip	expects an argument of type existing-vip
--rip	expects an argument of type existing-rip

### Optional Arguments

--ip	expects an argument of type ip
--port	expects an argument of type port
--weight	expects an argument of type rip-weight
--minconns	expects an argument of type uint
--maxconns	expects an argument of type uint
--encrypted	expects an argument of type bool
--redir	expects an argument of type string
--redir-enabled	expects an argument of type bool

## Edit a layer 7 virtual service

```
lbcli edit-layer7-service
or
lbcli --action edit-vip --layer 7
```

Edit the specified layer 7 virtual service.

### Required Arguments

--vip	expects an argument of type existing-vip
-------	--

### Optional Arguments

--appsession-cookie	expects an argument of type string
--as-port	expects an argument of type port
--autoscale-group	expects an argument of type string
--backend-encryption	expects an argument of type bool
--check-host	expects an argument of type string
--check-negate-response	expects an argument of type bool
--check-password	expects an argument of type string
--check-port	expects an argument of type port
--check-receive	expects an argument of type string
--check-request	expects an argument of type string
--check-type	argument should be one of connect, external, negotiate_http, negotiate_https, negotiate_http_head, negotiate_https_head, negotiate_http_options, negotiate_https_options, mysql or none
--check-username	expects an argument of type string
--clear-stick-drain	expects an argument of type bool



--compression	expects an argument of type bool
--cookie-maxidle	expects an argument of type haproxy-interval
--cookie-maxlife	expects an argument of type haproxy-interval
--cookiename	expects an argument of type string
--enable-hsts	expects an argument of type bool
--encrypt-all-backends	expects an argument of type bool
--external-check-script	expects an argument of type existing-healthcheck
--fallback-disabled	expects an argument of type bool
--fallback-encrypt	expects an argument of type bool
--fallback-ip	expects an argument of type ip
--fallback-persist	expects an argument of type bool
--fallback-port	expects an argument of type port
--feedback-method	argument should be one of agent or none
--feedback-port	expects an argument of type port
--force-to-https	expects an argument of type bool
--force-to-https-port	expects an argument of type port
--forward-for	expects an argument of type bool
--hsts-month	expects an argument of type uint
--http-pipeline	argument should be one of http_keep_alive, http_close, http_server_close or http_force_close
--http-pretend-keepalive	expects an argument of type bool
--http-request	expects an argument of type bool
--invalid-http	expects an argument of type bool
--ip	expects an argument of type ip
--maxconn	expects an argument of type uint
--mode	argument should be one of http or other_tcp
--no-write	expects an argument of type bool
--persist-table-size	expects an argument of type uint
--persist-time	expects an argument of type uint
--persistence	argument should be one of none, ip, tcp, sslsesid, http, appsession, rdp-session, rdp-cookie, http_ip, waf or last
--ports	expects an argument of type port-list
--proxy-bind	expects an argument of type string
--redirect-code	expects an argument of type http-300-code
--redispatch	expects an argument of type bool
--reuse-idle	expects an argument of type bool
--scheduler	argument should be one of roundrobin, leastconn or first
--send-proxy	expects an argument of type bool
--send-rip-name-as-host	expects an argument of type bool
--slave-ip	expects an argument of type ip
--source-address	expects an argument of type ip
--stunnel-source	expects an argument of type string
--stunnelproxy	expects an argument of type bool
--tcp-keep-alive	expects an argument of type bool
--timeout	expects an argument of type bool
--timeout-client	expects an argument of type haproxy-interval
--timeout-server	expects an argument of type haproxy-interval
--tproxy	expects an argument of type bool
--tunneltimeout	expects an argument of type haproxy-interval
--use-content-inspection	expects an argument of type bool
--xff-ip-pos	expects an argument of type int

## Edit a SSL termination

```
lbcli edit-termination
or
lbcli --action termination --function edit --type stunnel
```

Edit the specified SSL termination.

## Required Arguments

--vip	expects an argument of type existing-termination
--proxy-bind	expects an argument of type bool defaults to false

## Optional Arguments

--port	expects an argument of type port
--associated-to	expects an argument of type existing-l7-vip

--proxy-bind	expects an argument of type existing-l7-vip
--disablesslv2	expects an argument of type bool
--disablesslv3	expects an argument of type bool
--disableletls1	expects an argument of type bool
--stunnelndnsdelay	expects an argument of type bool
--stunnelproxy	expects an argument of type bool
--servercipherorder	expects an argument of type bool
--emptyfragments	expects an argument of type bool
--stunnelrenegotiation	expects an argument of type bool
--stunneltimetoclose	expects an argument of type uint
--sslcert	expects an argument of type existing-certificate
--disableletls1-1	expects an argument of type bool
--disableletls1-2	expects an argument of type bool
--disableletls1-3	expects an argument of type bool
--sslmode	argument should be one of high, fips or compatable
--source	expects an argument of type ip
--slave-source	expects an argument of type ip
--ciphers	expects an argument of type cipher-list
--ip	expects an argument of type ip
--slave-ip	expects an argument of type ip
--backend-ip	expects an argument of type ip
--slave-backend-ip	expects an argument of type ip
--backend-port	expects an argument of type port

## Edit a WAF

```
lbcli edit-waf
```

Edit the specified Web Application Firewall.

### Required Arguments

--waf	expects an argument of type existing-waf
-------	--

### Optional Arguments

--in-anom-score	expects an argument of type uint
--out-anom-score	expects an argument of type uint
--req-data	expects an argument of type bool
--resp-data	expects an argument of type bool
--audit	expects an argument of type bool
--dlogin	expects an argument of type bool
--dlogin-mode	argument should be one of static, openid_google or ldap
--dlogin-google-clientid	expects an argument of type string
--dlogin-google-clientsecret	expects an argument of type string
--dlogin-google-redirect-uri	expects an argument of type string
--dlogin-google-passphrase	expects an argument of type string
--dlogin-google-allowed-domain	expects an argument of type string
--rule-engine	expects an argument of type bool
--proxytimeout	expects an argument of type uint
--dlogin-location	expects an argument of type string
--dlogin-static-username	expects an argument of type string
--dlogin-static-password	expects an argument of type string
--ldap-server	expects an argument of type string
--ldap-server-port	expects an argument of type port
--search-base	expects an argument of type string
--ldap-attribute	expects an argument of type string
--ldap-group	expects an argument of type string
--bind-user	expects an argument of type string
--bind-password	expects an argument of type string
--disable-waf	expects an argument of type bool
--cacheaccel	expects an argument of type bool
--cache-nocache-files	expects an argument of type string
--cache-force-cache	expects an argument of type bool
--cache-object-size	expects an argument of type uint

## Enable the API endpoint

```
lbcli enable-api
or
lbcli --action api --function enable
```

Enable the API endpoint on the appliance. The apikey must be the unencoded value of the base64 string passed in the X-LB-APIKEY header of all requests.

#### Required Arguments

```
--username expects an argument of type string defaults to loadbalancer
--password expects an argument of type string defaults to loadbalancer
```

#### Optional Arguments

```
--apikey expects an argument of type string
```

### Enable Floating IP.

```
lbcli enable-floating-ip
```

Make it so that the given floating IP can be made active on this appliance.

### Remove and add a floating IP address

```
lbcli fix-floating-ip
```

Remove and re-add a floating IP in order to ensure it is up after altering the kernel state.

#### Required Arguments

```
--ip expects an argument of type ip
```

### Remove all ACLs for a virtual service

```
lbcli flush-acls
```

Remove all ACLS for the specified layer 7 virtual service.

#### Required Arguments

```
--vip expects an argument of type existing-l7-vip
```

### Remove all static IP addresses

```
lbcli flush-static-ips
or
lbcli --action address --function flush
```

Remove all static IP addresses from an appliance.

#### Optional Arguments

```
--interface expects an argument of type existing-interface
```

#### Remove all static routes

```
lbcli flush-static-routes
or
lbcli --action route --function static --type flush
```

Remove all static routes from an appliance.

#### Show the properties of the API endpoint

```
lbcli get-api
or
lbcli --action api --function get
```

Show the username, password and unencoded X-LB\_APIKEY header value expected by the API endpoint.

#### Show the default route

```
lbcli get-default-route
or
lbcli --action route --function default --type get
```

Show the default route of the default table for this appliance.

#### Show the DNS servers

```
lbcli get-dns
or
lbcli --action dns --function get
```

Show the Domain Name Servers for this appliance.

#### Show the firewall settings

```
lbcli get-firewall
or
lbcli --action local-configuration --get firewall
```

Show the settings for the firewall for this appliance.

#### Show the graphing settings

```
lbcli get-graphing
or
lbcli --action local-configuration --get graphing
```

Show the graphing settings for this appliance.

#### Show the GSLB reports

```
lbcli get-gslb-reports
or
```

```
lbcli --action gslb --function get --section reports
```

Show the GSLB reports for this appliance.

#### Required Arguments

```
--report argument should be one of get_generic_state or get_ppdns_state
```

#### Show the hostname

```
lbcli get-hostname  
or  
lbcli --action hostname --function get
```

Show the host and domain name for this appliance.

#### Show the interface offload settings

```
lbcli get-interface-offload  
or  
lbcli --action local-configuration --get interface_offload
```

Show the interface offload settings for this appliance.

#### Show the management gateway

```
lbcli get-management-gateway  
or  
lbcli --action local-configuration --get management_gateway
```

Show the details of the routing rule that traffic from this appliance's WebUI will return via.

#### Show the interface MTU settings

```
lbcli get-mtu  
or  
lbcli --action mtu --function get
```

Show the MTU settings for the interfaces on this appliance.

#### Show the NTP servers

```
lbcli get-ntp  
or  
lbcli --action local-configuration --get ntp
```

Show the Network Time Protocol servers this appliance will use.

#### Show the HTTP proxy server settings

```
lbcli get-proxy  
or  
lbcli --action local-configuration --get proxy
```

Show the HTTP proxy server the WebUI on this appliance will use.

Show the security settings

```
lbcli get-security  
or  
lbcli --action local-configuration --get security
```

Show the security settings for this appliance.

Show the email smarthost

```
lbcli get-smarthost  
or  
lbcli --action local-configuration --get smarthost
```

Show the email smarthost this appliance will use to send emails.

Show the email smarthost

```
lbcli get-smtp  
or  
lbcli --action local-configuration --get smtp
```

Show the email smarthost this appliance will use to send emails.

Show the syslog settings

```
lbcli get-syslog  
or  
lbcli --action local-configuration --get syslog
```

Show the syslog settings for this appliance.

Show the current date and time

```
lbcli get-timedate  
or  
lbcli --action local-configuration --get timedate
```

Show the current date and time for this appliance.

Show the URL of the updates server

```
lbcli get-updates  
or  
lbcli --action local-configuration --get updates
```

Show the URL of the updates server this appliance will use to fetch online updates.

Create a cluster

```
lbcli ha_create  
or  
lbcli create-ha
```

Pair this appliance with another to create a high availability cluster. This appliance will become the primary server.

#### Required Arguments

```
--local-ip      expects an argument of type ip
--peer-ip       expects an argument of type string
--peer-password expects an argument of type string
```

#### Halt a real server

```
lbcli halt
```

Close all connections to the specified real server and forbid and new connections.

#### Required Arguments

```
--vip expects an argument of type existing-vip
--rip expects an argument of type existing-rip
```

#### Clear all HAProxy stick tables

```
lbcli haproxy-clear-stick
```

Clear all HAProxy stick tables.

#### Install a licence

```
lbcli install-licence
```

Install a key to licence this appliance.

#### Required Arguments

```
--key expects an argument of type string
```

#### List ACL rules

```
lbcli list-acls
or
lbcli --action acl --function list
```

List all ACL rules for the specified layer 7 virtual service.

#### Required Arguments

```
--vip expects an argument of type existing-l7-vip
```

#### List SSL certificates

```
lbcli list-certificates
or
lbcli --action termination --function list --type certificate
```

List all SSL certificates on this cluster.

#### List floating IP addresses

```
lbcli list-floating-ips
or
lbcli --action list --function floatingip
```

List all floating IPs on this cluster.

#### List global names

```
lbcli list-gslb-globalnames
or
lbcli --action gslb --function list --section globalnames
```

List all GSLB global names.

#### List members

```
lbcli list-gslb-members
or
lbcli --action gslb --function list --section members
```

List all GSLB members.

#### List pools

```
lbcli list-gslb-pools
or
lbcli --action gslb --function list --section pools
```

List all GSLB pools.

#### List topologies

```
lbcli list-gslb-topologies
or
lbcli --action gslb --function list --section topologies
```

List all GSLB topologies.

#### List header rules

```
lbcli list-headers
or
lbcli --action headers --function list
```

List all header rules for the specified layer 7 virtual service.

#### Required Arguments

```
--vip expects an argument of type existing-l7-vip
```



## List health checks

```
lbcli list-healthchecks
```

List existing health checks on the cluster.

## List PBR rules

```
lbcli list-pbr
or
lbcli --action pbr --function get
lbcli get-pbr
```

List Policy Based Routing rules for this appliance.

### Optional Arguments

```
--ip          expects an argument of type ip
--gateway     expects an argument of type ip
```

## List static IP addresses

```
lbcli list-static-ips
or
lbcli --action address --function get
```

List static IPs configured on the specified interface.

### Required Arguments

```
--interface  expects an argument of type existing-interface
```

## List static routes

```
lbcli list-static-routes
or
lbcli --action route --function static --type get
```

List static routes configured on this appliance.

## List WAFs

```
lbcli list-waf
```

List Web Application Firewalls configured on this appliance.

## Lockdown access to the WebUI

```
lbcli lockdown
```

Lockdown the WebUI on this appliance so that it only responds to clients on the specified subnet.

### Required Arguments

```
--enabled expects an argument of type bool
```

#### Optional Arguments

```
--network expects an argument of type ip-and-cidr
```

#### Report the node statuses

```
lbcli nodestatus
```

Report the status of the nodes in this cluster as this appliance understands them.

#### Bring a real server online

```
lbcli online
```

Bring a real server that is halted or drained back online by allowing it to receive connections once more.

#### Required Arguments

```
--vip expects an argument of type existing-vip  
--rip expects an argument of type existing-rip
```

#### Reboot appliance

```
lbcli reboot  
or  
lbcli --action power --function restart
```

Reboot this appliance.

#### Regenerate default certificate

```
lbcli regenerate-default-certificate  
or  
lbcli --action termination --function regenerate_local --type certificate
```

Regenerate the default self-signed certificate used by the webui and SSL-terminations.

#### Reload WebUI

```
lbcli reload-apache  
or  
lbcli reload-webui
```

Reload the WebUI service on this appliance.

#### Reload GSLB

```
lbcli reload-gslb
```

Reload the GSLB service on this appliance.

### Reload HAProxy

```
lbcli reload-haproxy
```

Reload the HAProxy service on this appliance.

### Reload Heartbeat

```
lbcli reload-heartbeat
```

Reload the Heartbeat service on this appliance.

### Reload LDirectord

```
lbcli reload-ldirectord
```

Reload the LDirectord service on this appliance.

### Reload STunnel

```
lbcli reload-stunnel
```

Reload the STunnel service on this appliance.

### Reload Syslog

```
lbcli reload-syslog
```

Reload the Syslog service on this appliance.

### Reload WAF

```
lbcli reload-waf
```

Reload the WAF service on this appliance.

### Break a cluster

```
lbcli reset-cluster-config  
or  
lbcli --action reset --function cluster
```

Break a high-availability pair.

### Restart Autoscaled

```
lbcli restart-autoscaling
```

Reload the Autoscaled service on this appliance.

### Restart AZHA

```
lbcli restart-azha
```

Reload the AZHA service on this appliance.

#### Restart Collectd

```
lbcli restart-collectd
```

Reload the Collectd service on this appliance.

#### Restart Firewall

```
lbcli restart-firewall
```

Reload the Firewall service on this appliance.

#### Restart HAProxy

```
lbcli restart-haproxy
```

Reload the HAProxy service on this appliance.

#### Restart Heartbeat

```
lbcli restart-heartbeat
```

Reload the Heartbeat service on this appliance.

#### Restart LDirectord

```
lbcli restart-ldirectord
```

Reload the LDirectord service on this appliance.

#### Restart Pound

```
lbcli restart-pound
```

Reload the Pound service on this appliance.

#### Restart SNMP

```
lbcli restart-snmp
```

Reload the SNMP service on this appliance.

#### Restart STunnel

```
lbcli restart-stunnel
```

Reload the STunnel service on this appliance.

#### Restart Syslog

```
lbcli restart-syslog
```

Reload the Syslog service on this appliance.

#### Restart WAF

```
lbcli restart-waf
```

Reload the WAF service on this appliance.

#### Restore factory defaults

```
lbcli restore
```

Restore this appliance to factory defaults.

#### Required Arguments

```
--yes-i-am-sure  expects an argument of type bool
```

#### Set default route

```
lbcli set-default-route  
or  
lbcli --action route --function default --type set
```

Set the default gateway for this appliance.

#### Required Arguments

```
--interface  expects an argument of type existing-interface-or-auto  
              defaults to auto  
--gateway    expects an argument of type ip
```

#### Set DNS servers

```
lbcli set-dns  
or  
lbcli --action dns --function set
```

Set the Domain Name Servers for this appliance.

#### Optional Arguments

```
--dns0  expects an argument of type ip-or-empty  
--dns1  expects an argument of type ip-or-empty  
--dns2  expects an argument of type ip-or-empty
```

#### Set firewall properties

```
lbcli set-firewall
or
lbcli --action local-configuration --set firewall
```

Set the properties for the firewall for this appliance.

#### Required Arguments

```
--conntrack-size expects an argument of type uint
```

#### Set graphing properties

```
lbcli set-graphing
or
lbcli --action local-configuration --set graphing
```

Set the graphing properties for this appliance.

#### Optional Arguments

```
--layer4      expects an argument of type bool
--layer7      expects an argument of type bool
--interfaces  expects an argument of type bool
--interfaces  expects an argument of type bool
--load        expects an argument of type bool
--memory      expects an argument of type bool
--disk        expects an argument of type bool
--interval    expects an argument of type uint
--timeout     expects an argument of type uint
```

#### Set hostname

```
lbcli set-hostname
or
lbcli --action hostname --function set
```

Set the host and domain name for this appliance.

#### Required Arguments

```
--hostname expects an argument of type string
--domain   expects an argument of type string defaults to localdomain
```

#### Set interface offload

```
lbcli set-interface-offload
or
lbcli --action local-configuration --set interface_offload
```

Enable or disable TCP offload for this appliance.

#### Required Arguments

```
--hardware-offload expects an argument of type bool
```

#### Set the management gateway

```
lbcli set-management-gateway
or
lbcli --action local-configuration --set management_gateway
```

Set the routing rule which traffic from the WeBUI will follow for this appliance.

#### Required Arguments

```
--iface-ip  expects an argument of type existing-static-ip
--gateway    expects an argument of type ip
```

#### Set the MTU

```
lbcli set-mtu
or
lbcli --action mtu --function set
```

Set the MTU for the specified interface.

#### Required Arguments

```
--interface  expects an argument of type existing-interface
--mtu        expects an argument of type mtu
```

#### Set NTP server

```
lbcli set-ntp
or
lbcli --action local-configuration --set ntp
```

Set the Network Time Protocol servers used by this appliance.

#### Set HTTP proxy

```
lbcli set-proxy
or
lbcli --action local-configuration --set proxy
```

Set the HTTP proxy used by the WebUI of this appliance.

#### Optional Arguments

```
--ip          expects an argument of type hostname-or-ip
--port        expects an argument of type port
--username    expects an argument of type string
--password    expects an argument of type string
```

#### Set security settings

```
lbcli set-security
or
lbcli --action local-configuration --set security
```

Set the security settings for this appliance.

#### Optional Arguments

```
--security-mode    argument should be one of custom, secure or secure-perm
--rootaccess       expects an argument of type bool
--sshpasword       expects an argument of type bool
--httpsonly        expects an argument of type bool
--httpsport        expects an argument of type bool
--certificate       expects an argument of type existing-certificate
--httpsciphers     expects an argument of type cipher-list
```

#### Set the email smarthost

```
lbcli set-smarthost
or
lbcli --action local-configuration --set smarthost
```

Set the email smarthost this appliance will use to send emails.

#### Optional Arguments

```
--smtp-relay  expects an argument of type string
```

#### Set SNMP properties

```
lbcli set-snmp
or
lbcli --action local-configuration --set snmp
```

Set the properties of the SNMP service on this appliance.

#### Optional Arguments

```
--community    expects an argument of type string
--location      expects an argument of type string
--contact       expects an argument of type string
--user          expects an argument of type string
--auth-method   expects an argument of type string
--auth          expects an argument of type string
--priv-method   expects an argument of type string
--priv          expects an argument of type string
--snmpv3-enable expects an argument of type string
--snmpv2-enable expects an argument of type string
```

#### Set syslog properties

```
lbcli set-syslog
or
lbcli --action local-configuration --set syslog
```

Set the properties of the syslog service on this appliance.

#### Optional Arguments

```
--interval    expects an argument of type uint
--burst        expects an argument of type uint
--destination  expects an argument of type string
--ip           expects an argument of type ip
--port         expects an argument of type port
--protocol     argument should be one of tcp or udp
--template     expects an argument of type string
```



## Set time and date

```
lbcli set-timedate
or
lbcli --action local-configuration --set timedate
```

Set the time and date of this appliance.

### Required Arguments

```
--year      expects an argument of type uint
--month     expects an argument of type month
--day       expects an argument of type day
--hour      expects an argument of type hour
--minute    expects an argument of type minute
```

## Enable or disable update check

```
lbcli set-updates
or
lbcli --action local-configuration --set updates
```

Enable or disable the update check on this appliance.

### Optional Arguments

```
--auto-check-for-updates  expects an argument of type bool
--update-server           expects an argument of type string
```

## Show configuration.

```
lbcli show-full-config
or
lbcli --action list --function dumpconfig
```

Show configuration for this appliance.

## Show configuration for LDirectord.

```
lbcli show-layer4-advanced
or
lbcli --action list --function advanced --layer 4
```

Show configuration for LDirectord for the cluster.

## Show configuration for Layer 4 virtual services.

```
lbcli show-layer4-services
or
lbcli --action list --function virtual --layer 4
```

Show the configuration for Layer 4 virtual services on this cluster.

## Show configuration for HAProxy.

```
lbcli show-layer7-advanced
or
```

```
lbcli --action list --function advanced --layer 7
```

Show configuration for HAProxy for the cluster.

Show configuration for Layer 7 virtual services.

```
lbcli show-layer7-services  
or  
lbcli --action list --function virtual --layer 7
```

Show the configuration for Layer 7 virtual services on this cluster.

Shutdown appliance.

```
lbcli shutdown  
or  
lbcli --action power --function shutdown
```

Shutdown this appliance.

Show GSLB status.

```
lbcli status-gslb
```

Show the status of the GSLB service for this appliance.

Create a technical support download.

```
lbcli support-download
```

Create a technical support download for this appliance.

Get uptime.

```
lbcli uptime
```

Get the uptime for this appliance.

## Running lbcli from a remote Linux Host

These commands can also be run from a remote Linux host. This example halts VIP1/RIP1:

```
ssh root@192.168.111.42 "lbcli --action halt --vip VIP1 --rip RIP1"
```

## Running lbcli from a remote Windows Host

These commands can also be run from a remote Windows host. This example halts VIP1/RIP1:

```
plink -pw loadbalancer root@192.168.111.42 "lbcli --action halt --vip VIP1 --rip RIP1"
```

## Notes

1. **PuTTY** must be installed to use the *plink* command.
2. The password for the root user is set during the Network Setup Wizard.
3. 192.168.111.42 is the IP address of the load balancer.

## Application Programming Interface (API)

### Enabling the API

To enable the API, run the following command:

```
lbcli enable-api --username (username) --password (password) --apikey (apikey)
```

e.g. The following command sets the username and password to "loadbalancer" and sets the apikey to "ApIk3y2345118".

```
lbcli --action api --function enable --username loadbalancer password loadbalancer --apikey ApIk3y2345118
```

**Note** | To automatically generate a 32 character random key exclude the --apikey option.

Once enabled, the API enables all CLI commands to be executed using HTTP POST requests.

### API Endpoint

API calls must be posted to the following URL on the load balancer:

```
https://<appliance IP address>:9443/api/v2
```

### Creating the JSON Request

To illustrate how JSON requests are formed, the following examples show the CLI command and the corresponding JSON request in each case.

#### Example 1 - Halt a Real Server

This example shows how RIP1 of VIP1 can be halted.

*lbcli command:*

```
lbcli --action halt --vip VIP1 --rip RIP1
```

*JSON equivalent:*

```
{
  "lbcli": [{
    "action": "halt",
    "vip": "VIP1",
    "rip": "RIP1"
  }]
}
```

## Example 2 - Add a Layer 7 VIP

This example shows how to add a Layer 7 HTTP mode VIP.

*lbcli command:*

```
lbcli --action add-vip --layer 7 --vip VIP1 --ip 192.168.1.1 --ports 80 --mode http
```

*JSON equivalent:*

```
{
  "lbcli": [{
    "action": "add-vip",
    "layer": "7",
    "vip": "VIP1",
    "ip": "192.168.1.1",
    "ports": "80",
    "mode": "http"
  }]
}
```

## Example 3 - Restart HAProxy

This example shows how to restart HAProxy.

*lbcli command:*

```
lbcli --action restart-haproxy
```

*JSON equivalent:*

```
{
  "lbcli": [{
    "action": "restart-haproxy"
  }]
}
```

## JSON Syntax Validation

The website <https://jsonlint.com/> is a useful resource for JSON syntax validation. Paste the JSON into the window provided and click **Validate JSON**.

## Sending API Requests

Using Postman

**Postman** is a great tool for testing and verifying API functionality. For more details on using Postman, please refer to [our blog](#).

Using script/code

### Example 1 - using bash

This example shows how bash can be used to make API calls.

```
#!/bin/bash
# loadbalancer.org api/v2 curl api call
while true; do
  case "$1" in
    -l | --loadbalancer ) loadbalancer="$2"; shift 2 ;;
    -u | --username ) username="$2"; shift 2 ;;
    -p | --password ) password="$2"; shift 2 ;;
    -j | --json ) json="$2"; shift 2 ;;
    -a | --apikey ) apikey=$(echo $2 | base64); shift 2 ;;
    * ) break ;;
  esac
done
if [ $loadbalancer != "" ] || [ $username != "" ] || [ $password != "" ] || [ $json != "" ] || [ $apikey != "" ]; then
  curl -u ${username}:${password} -X POST \
    --header "X-LB-APIKEY: ${apikey}" \
    --header Content-Type:application/json \
    -d @${json} https://${loadbalancer}:9443/api/v2/ -k
else
  echo "./apicall.sh --loadbalancer 192.168.111.220 --username loadbalancer --password loadbalancer --json /api-json/add-vip.json --apikey aj1hf5GzSgYwMvNBUn6T3s4pIXo0lF7R"
fi
```

## Example 2 - Using Microsoft PowerShell

This example shows how PowerShell can be used to make API calls.

```
# Configure Variables
$user = "loadbalancer"           # Appliance API Username
$pass = "loadbalancer"          # Appliance API Password
$apikey = "aj1hf5GzSgYwMvNBUn6T3s4pIXo0lF7R" # Appliance Non base64 encoded APIKEY
$ip = "192.168.111.220"         # Appliance Management IP address
$jsonfile = "c:\api-json\add-vip.json" # Local path to JSON configuration file

# Disable certificate verification checks to allow for the self signed WebUI certificate on the load balancer
if (-not ([System.Management.Automation.PSTypeName]'TrustAllCertsPolicy').Type)
{
  add-type @"
    using System.Net;
    using System.Security.Cryptography.X509Certificates;
    public class TrustAllCertsPolicy : ICertificatePolicy {
      public bool CheckValidationResult(
        ServicePoint srvPoint, X509Certificate certificate,
        WebRequest request, int certificateProblem) {
        return true;
      }
    }
"@
}
[System.Net.ServicePointManager]::CertificatePolicy = New-Object TrustAllCertsPolicy

# Base64 encode the apkey
$bytes_api = [System.Text.Encoding]::ASCII.GetBytes("$apikey")
$base64_api = [System.Convert]::ToBase64String($bytes_api)
$APIAuthValue = "${base64_api}"

# Set the auth headers
$headers = @"Authorization"=$basicAuthValue; "X-LB-APIKEY"=$APIAuthValue

# Fetch the json configuration file
$json = Get-Content $jsonfile -Raw

# Send the API call
Invoke-WebRequest -Uri ("https://$ip:9443/api/v2/") -Method Post -Body $json -ContentType "application/json" -Headers $headers
```

**Note** | Change the variable declarations to suit your environment.

The add-vip.json file referred to in both examples:

```
{
  "lbcli": [{
    "action": "add-vip",
    "layer": "7",
    "vip": "VIP1",
    "ip": "192.168.111.228",
    "ports": "80",
    "mode": "http"
  }]
}
```

#### Note

For more examples of making API call using other languages please refer to our blog [How to automate load balancer deployments, Part 2!](#).

## Using ipvsadm to configure Layer 4 Services

For layer 4 services, the ipvsadm command can be used. Several examples are provided below.

Add a TCP based Virtual Service & use weighted round robin scheduling:

```
ipvsadm -A -t 192.168.65.192:80 -s wrr
```

Add a TCP based Real Server in DR mode:

```
ipvsadm -a -t 192.168.65.192:80 -g -r 192.168.70.196:80
```

Add a TCP based Real Server in NAT mode:

```
ipvsadm -a -t 192.168.65.192:80 -m -r 192.168.70.196:80
```

Add a UDP based Virtual Service & use weighted least connection scheduling:

```
ipvsadm -A -u 192.168.65.192:80 -s wlc
```

Add a UDP based Real Server in DR mode:

```
ipvsadm -a -u 192.168.65.192:80 -g -r 192.168.70.196:80
```

Delete a TCP based Virtual Service:

```
ipvsadm -D -t 192.168.65.180:80
```

Delete a TCP based Real Server:

```
ipvsadm -d -t 192.168.65.122:80 -r 192.168.70.134:80
```

View the current running config:

```
ipvsadm -ln
```

Command output:

```
IP Virtual Service version 1.2.1 (size=4096)
Prot LocalAddress:Port Scheduler Flags
-> RemoteAddress:Port Forward Weight ActiveConn InActConn
TCP 192.168.65.120:80 rr
-> 192.168.70.130:80 Route 1 0 0
-> 192.168.70.131:80 Route 1 0 0
TCP 192.168.65.122:80 rr
-> 192.168.70.132:80 Mass 1 0 0
-> 192.168.70.133:80 Mass 1 0 0
```

#### Note

Please note that since these changes are being made directly to the running configuration, the services that are displayed in the System Overview will no longer match the running configuration when ipvsadm/socat commands are used. Using the **lbcli** command or the API does not have this disadvantage since the System Overview will show the correct VIP and RIP status.

## Using Linux socket commands to configure Layer 7 Services

For layer 7 HAProxy VIPs, the socat socket command can be used as shown in the examples below.

To take a server offline:

```
echo "disable server VIP_Name/RIP_Name" | socat unix-connect:/var/run/haproxy.stat stdio
```

To bring a server online:

```
echo "enable server VIP_Name/RIP_Name" | socat unix-connect:/var/run/haproxy.stat stdio
```

To set the weight of a Real Server:

```
echo "set weight VIP_Name/RIP_Name 0" | socat unix-connect:/var/run/haproxy.stat stdio
```

To view HAProxy's running configuration:

```
echo "show info" | socat unix-connect:/var/run/haproxy.stat stdio
```

To clear HAProxy's statistics:

```
echo "clear counters all" | socat unix-connect:/var/run/haproxy.stat stdio
```

#### Note

Other Linux Socket command examples can be found [here](#).

#### Note

Since these changes are being made directly to the running configuration, the services that are displayed in the System Overview will no longer match the running configuration when ipvsadm/socat commands are used. Using the **lbcli** command or the API does not have this

disadvantage since the System Overview will show the correct VIP and RIP status.



# Chapter 7 - Web Application Firewall (WAF)

## Introduction

A web application firewall (WAF) filters, monitors, and blocks HTTP traffic to and from a web application.

The load balancer includes a built-in WAF. It can be deployed in front of a web application to provide an additional layer of security, where required. It is based on the free and open-source ModSecurity WAF engine and includes the **OWASP ModSecurity Core Rule Set** (CRS) by default. The CRS is a set of generic attack detection rules. It aims to protect web applications from a wide range of attacks, including the OWASP Top 10, while keeping false positives (false alerts) to a minimum.

The OWASP Top 10 represents a broad consensus about the most critical security risks to web applications. These risks are broken down into ten categories, as shown in the table below:

Category	Description
<b>A01</b> - Broken Access Control	Access control failures typically lead to unauthorized information disclosure, modification, or destruction of all data or performing a business function outside the user's limits.
<b>A02</b> - Cryptographic Failures	Previously known as <i>Sensitive Data Exposure</i> , the focus is on failures related to cryptography (or lack thereof). Which often lead to exposure of sensitive data.
<b>A03</b> - Injection	An application is vulnerable to injection attack when, for example, user-supplied data is not validated, filtered, or sanitized by the application.
<b>A04</b> - Insecure Design	A new category which focuses on risks related to design and architectural flaws, with a call for more use of threat modeling, secure design patterns, and reference architectures.
<b>A05</b> - Security Misconfiguration	The application might be vulnerable if the application is, for example, missing appropriate security hardening across any part of the application stack.
<b>A06</b> - Vulnerable and Outdated Components	You are likely vulnerable, for example, if you do not know the versions of all components you use (both client-side and server-side), including nested dependencies.
<b>A07</b> - Identification and Authentication Failures	Previously known as <i>Broken Authentication</i> , confirmation of the user's identity, authentication, and session management is critical to protect against authentication-related attacks.
<b>A08</b> - Software and Data Integrity Failures	A new category which focuses on making assumptions related to software updates, critical data, and CI/CD pipelines without verifying integrity.
<b>A09</b> - Security Logging and Monitoring Failures	Detecting and responding to breaches is critical. This category is to help detect, escalate, and respond to active breaches.
<b>A10</b> - Server-Side Request Forgery (SSRF)	SSRF flaws occur whenever a web application is fetching a remote resource without validating the user-supplied URL.

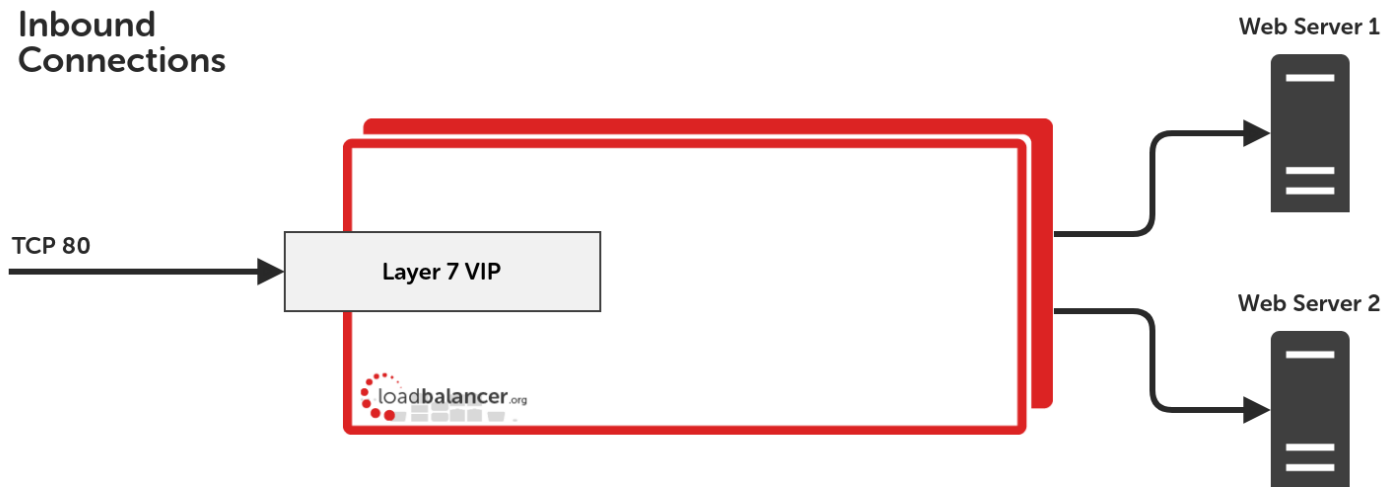
More details can be found at the [OWASP Top 10 website](#).

## Implementation Concept

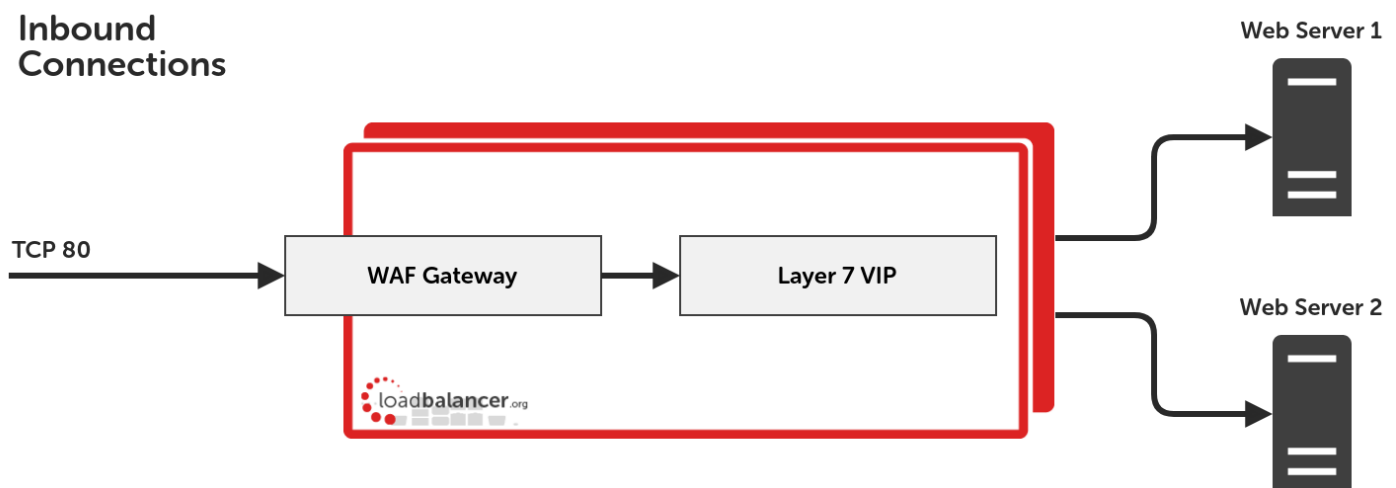
The load balancer supports the ability to define multiple WAF gateways, one for each web application to be protected. Each WAF gateway is associated with a layer 7 VIP when created. On creation, the data path is automatically modified so that the WAF becomes the initial connection point for inbound client connections, as

illustrated below:

Data flow before WAF is deployed:



Modified data flow once WAF is deployed:



## Notes

- When defining a WAF gateway on the load balancer, the associated layer 7 VIP must be selected from a drop-down list. This enables the WAF to be automatically configured to listen on the same TCP socket as the original layer 7 VIP. The WAF gateway is then automatically configured to forward packets to the original layer 7 VIP.
- Each WAF gateway is associated with one layer 7 VIP.
- Once the WAF gateway is defined, the *Label*, *IP Address*, *Port* and *Protocol* of the associated layer 7 VIP cannot be edited to ensure the association remains intact. If changes to these settings are required, take a backup copy of the WAF gateway's manual configuration (if one exists), remove the WAF, make the changes, and then recreate the WAF.
- Each WAF gateway is actually comprised of two component parts: an additional layer 7 VIP, which acts as the frontend to the WAF, and an Apache instance, of which ModSecurity is a module. Both are automatically created when the WAF gateway is configured.

## Creating a New WAF Gateway

For reasons mentioned in the previous section, the layer 7 VIP must be created first (if it doesn't already exist), followed by the WAF gateway.

### Step 1 - Create the Layer 7 VIP

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Services* and click **Add a new Virtual Service**.

#### Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="Web-Cluster"/>	?
IP Address	<input type="text" value="192.168.110.46"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

2. Enter a suitable Label (name) for the VIP, e.g. **Web-Cluster**.
3. Enter a valid IP address, e.g. **192.168.110.46**.
4. Enter a valid port, e.g. **80**.
5. Click **Update**.

### Step 2 - Define the associated Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Real Servers* and click **Add a new Real Server** next to the VIP just created.

#### Layer 7 Add a new Real Server - Web-Cluster

Label	<input type="text" value="Web1"/>	?
Real Server IP Address	<input type="text" value="192.168.110.241"/>	?
Real Server Port	<input type="text" value="80"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

2. Enter a suitable Label (name) for the RIP, e.g. **Web1**.
3. Enter a valid IP address, e.g. **192.168.110.241**.
4. Enter a valid port, e.g. **80**.

5. Click **Update**.

### Step 3 - Define the WAF Gateway

1. Using the WebUI, navigate to: *Cluster Configuration > WAF - Gateway* and click **Add a new WAF gateway**.

#### WAF - Add A New Gateway

Select Layer 7 Virtual Service	Web-Cluster ▾	?
WAF Label	WAF-Web-Cluster	?
Ruleset	Core Rule Set 3.3.2 ▾	?

**Cancel** **Update**

2. Select the VIP created in step 1 in the drop down.
3. The WAF label (name) field will be populated automatically, this can be changed if required.
4. The *Rule Set* will automatically choose the latest available stable version of the Core Rule Set. This should not ordinarily require changing.
5. Click **Update**.

### Step 4 - Reload Services to Apply the new Settings

1. Click *System Overview* in the WebUI.
2. Reload the services (WAF and HAProxy) as prompted in the "commit changes" message box.

### Step 5 - View Configured Services

1. The original layer 7 VIP and the auto created layer 7 WAF frontend VIP are now displayed in the system overview as shown below:

#### System Overview ?

2023-02-01 16:53:53 UTC

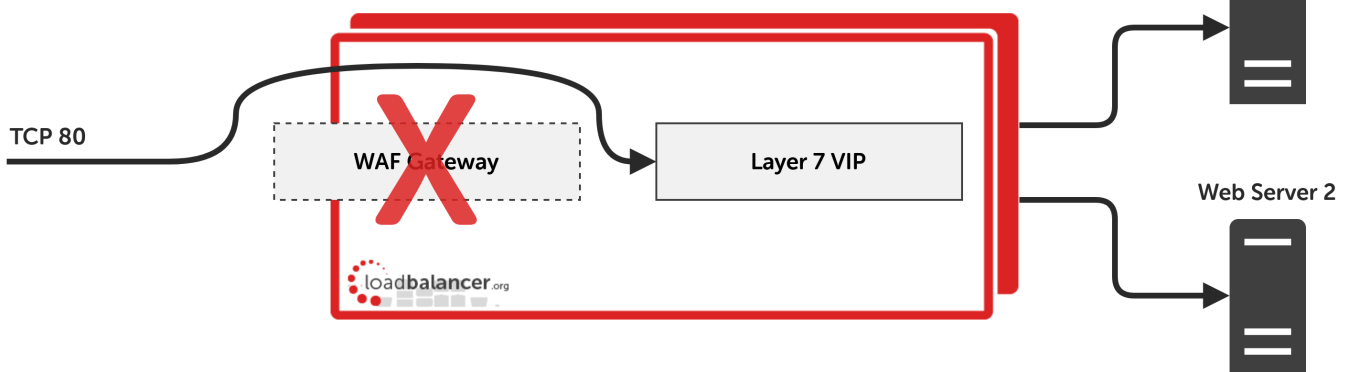
	VIRTUAL SERVICE ▾	IP ▾	PORTS ▾	CONN. ▾	PROTOCOL ▾	METHOD ▾	MODE ▾	
↑	Web-Cluster	192.168.110.46	65435	0	HTTP	Layer 7	Proxy	
↑	WAF-Web-Cluster	192.168.110.46	80	0	HTTP	Layer 7	Proxy	

## Fail Open and Fail Closed Designs

In the event that a WAF gateway should become unavailable for any reason, for example being overwhelmed by a denial of service attack, there are two response strategies that can be employed:


- **Fail open:** Traffic flow is allowed to continue by bypassing the failed WAF gateway, and service availability is maintained (*default*)

## Inbound Connections



- Fail closed: Traffic flow is stopped until the WAF gateway can recover

## Inbound Connections



Loadbalancer.org's top priority is ensuring that services are highly available, first and foremost. As such, the load balancer's default policy is to *fail open* in the event of a WAF gateway failure.

There may be scenarios where security is the *primary* concern, with high availability being a secondary consideration in comparison. For example, if working with highly sensitive web traffic, it may be preferential to experience a service outage rather than allowing the service to be accessed insecurely by bypassing a WAF gateway.

## Configuring a WAF Gateway for Fail Open Operation (Default)

**Note** | This is the default starting configuration for all WAF gateways.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Services*.
2. Find the relevant pair of layer 7 virtual services: the original layer 7 VIP and the automatically created WAF frontend VIP, which is titled "WAF-..." by default.

## Layer 7 - Virtual Services

Search..	<a href="#">Add a new Virtual Service</a>				
Service Name	IP	Port	Config Type		
Web_Service	192.168.85.150	Ports 65435	Auto	<a href="#">Modify</a>	<input type="text"/>
WAF-Web_Service	192.168.85.150	Ports 80	Auto	<a href="#">Modify</a>	<input type="text"/>

3. Click **Modify** next to the "WAF-..." frontend VIP.
4. Scroll down to the *Fallback Server* section.
5. Set the *IP Address* to the address that the original layer 7 VIP is listening on, e.g. **192.168.85.150**.
6. Set the *Port* to the port that the original layer 7 VIP is listening on, e.g. **65435**.
7. Click **Update**.

Fallback Server		[Advanced +]
Disable Fallback	<input type="checkbox"/>	<a href="#">?</a>
IP Address	<input type="text" value="192.168.85.150"/>	<a href="#">?</a>
Port	<input type="text" value="65435"/>	<a href="#">?</a>

8. Click *System Overview* in the WebUI.
9. Reload the services (WAF and HAProxy) as prompted in the "commit changes" message box.

## Configuring a WAF Gateway for Fail Closed Operation

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Services*.
2. Find the relevant pair of layer 7 virtual services: the original layer 7 VIP and the automatically created WAF frontend VIP, which is titled "WAF-..." by default.

## Layer 7 - Virtual Services

Search..	<a href="#">Add a new Virtual Service</a>				
Service Name	IP	Port	Config Type		
Web_Service	192.168.85.150	Ports 65435	Auto	<a href="#">Modify</a>	<input type="text"/>
WAF-Web_Service	192.168.85.150	Ports 80	Auto	<a href="#">Modify</a>	<input type="text"/>

3. Click **Modify** next to the "WAF-..." frontend VIP.
4. Scroll down to the *Fallback Server* section.
5. Set the *IP Address* to the local address of the fallback server, **127.0.0.1**.
6. Set the *Port* to the port that the fallback server is listening on, **9081**.

7. Click **Update**.

Fallback Server		[Advanced +]
Disable Fallback	<input type="checkbox"/>	?
IP Address	<input type="text" value="127.0.0.1"/>	?
Port	<input type="text" value="9081"/>	?

8. Click *System Overview* in the WebUI.

9. Reload the services (WAF and HAProxy) as prompted in the "commit changes" message box.

**Tip**

It's possible to use a dedicated, external fallback server instead of the local one on 127.0.0.1. It's also possible to customise the HTML of the built-in fallback page which is presented to users. For more information, refer to the [Fallback Server](#) section.

## WAF Gateway Settings

Each WAF gateway has a variety of settings which can be configured through the WebUI. To access these settings, navigate to *Cluster Configuration > WAF - Gateway* and click **Modify** next to the WAF gateway in question.

**Important**

After modifying any WAF gateway settings, be sure to save the new configuration to disk by pressing the **Update** button and then apply the new configuration by reloading services as prompted in the "commit changes" message box.

### Disable Web Application Firewall

*Default value: False.*

When checked, this option completely disables all ModSecurity and CRS configuration for a given WAF gateway (this includes disabling any manual WAF configuration, if present, while leaving the manual configuration itself intact). The proxy sandwich remains in place, allowing traffic to continue to flow in the same way, except without any WAF functionality present.

### Ruleset

*Default value: Core Rule Set 3.3.2.*

The WAF rule set to use for a given WAF gateway can be selected from the drop-down list. There are currently two rule sets to choose from:

- **Core Rule Set 3.3.2** (default): the latest stable release of the OWASP ModSecurity Core Rule Set.
- **Core Rule Set 2**: a legacy option provided for backward compatibility with older WAF installations. Provides CRS version 2.2.9.

### Paranoia Level

*Default value: Paranoia Level 1.*

Not available when the Core Rule Set 2 (legacy) is in use.

This can be set to paranoia level 1, 2, 3, or 4. Paranoia level 1 offers a baseline level of security with a minimal, or zero, need to tune away false positives. At the other end of the scale, paranoia level 4 offers the strongest level of

security and features many additional rules, but is extremely likely to cause a large number of false positives, requiring a significant investment of time to tune them away.

See the section on [Paranoia Levels](#) for full details about this key concept.

## Rule Engine Traffic Blocking

*Default value: False.*

Allows the WAF rule engine to take disruptive actions and block malicious looking traffic for this WAF gateway.

By default, the rule engine of a newly created WAF gateway is not able to block traffic. This is sometimes referred to as **detection only mode**. This means that WAF rule logic is processed in order to examine traffic but disruptive actions, e.g. "deny" and "drop", are **never executed**. As such, traffic is never blocked, even if it triggers rules and appears to be malicious.

One approach to configuring and tuning a WAF deployment is to leave the WAF gateway in detection only mode, pass known good traffic through the WAF (e.g. traffic from user testing), and then use the resulting log data to tune the WAF. This "tuning" is accomplished by writing rule exclusions to cover all false positives caused by the known good traffic. Once confident that all false positives have been accounted for, traffic blocking could then be enabled for the WAF gateway's rule engine. This takes the WAF gateway out of detection only mode and allows it to start actively blocking malicious looking traffic.

## Process Request Data

*Default value: True.*

Instructs the WAF engine to buffer and process request bodies. This allows the data in request bodies to be inspected, for example the parameters of a POST request.

## Process Response Data

*Default value: False.*

Instructs the WAF engine to buffer and process response bodies. This allows the data in response bodies to be inspected, for example an HTML response.

By default, a WAF gateway only processes request data, i.e. the data in requests coming in from clients. It's also possible to process response data, i.e. the data passed back *to* clients *from* the back end web application. This can be useful, for example, to catch instances of data leakage, such as an unintended SQL database error being passed back to the client (which may expose information about the type, version, and configuration of database software in use).

## Inbound Anomaly Score

*Default value: 20.*

Sets the inbound anomaly score threshold: the cumulative anomaly score at which an inbound request will be blocked.

See the section on [Anomaly Scoring](#) for full details about this key concept.

## Outbound Anomaly Score



*Default value: 4.*

Sets the outbound anomaly score threshold: the cumulative anomaly score at which an outbound response will be blocked.

See the section on [Anomaly Scoring](#) for full details about this key concept.

## Audit Mode

*Default value: False.*

Enables the audit logging engine for all transactions passing through a WAF gateway.

Audit logs record full transaction data, *including full request bodies*. This information can be invaluable for troubleshooting particularly difficult issues.

### Warning

Audit logs can grow **extremely large** very quickly. As such, it is **strongly recommended not to enable audit logging on a production machine**: doing so is likely to fill the logging disk partition and is **likely to cause disruptive issues on production machines**.

## WAF Proxy Timeout

*Default value: 120.*

The Apache proxy service that hosts a WAF gateway has a 60 second timeout by default. This can be changed if required.

## Enable Cache Acceleration

*Default value: False.*

While not directly related to WAF functionality, the proxy sandwich that hosts the WAF functionality also features a simple object cache. It will only cache objects that are HTML and below 64k in size, independent of any cache or no-cache options that your real servers may provide.

Enabling cache acceleration exposes the following cache-specific options:

- Force 'no-cache' override
- Location to exclude from the cache
- Cache object size

## Double Login Enable

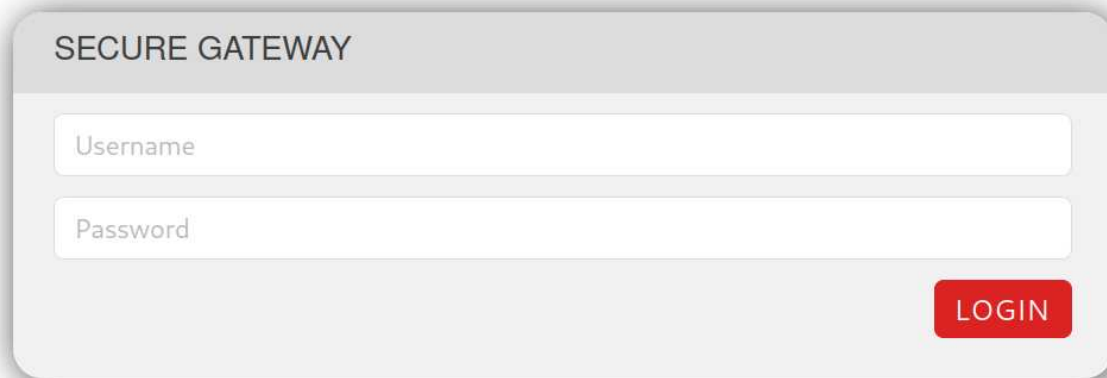
*Default value: False.*

### Web Gateway Authentication

While not directly related to WAF functionality, the proxy sandwich that hosts the WAF functionality also features a simple web gateway / "double login" page. It supports the following authentication methods:

- Locally defined static user
- Google OpenID

Once enabled, users will be prompted for credentials when accessing the WAF:

A login form titled "SECURE GATEWAY" with a light gray background and rounded corners. It contains two input fields: "Username" and "Password", both with light gray borders and placeholder text. A red "LOGIN" button is positioned to the right of the password field.

Secured by Loadbalancer.org

Enabling double login exposes the following double login-specific options:

- Location to protect
- Double login mode (Static user; OpenID Connect - Google)
- Static username
- Static password

## WAF - Advanced Configuration

**Note** | These settings should not typically require changing.

The global WAF service has two advanced settings which can be configured through the WebUI. To access these settings, navigate to *Cluster Configuration > WAF - Advanced Configuration*.

The two advanced settings are *match limits*, which help avoid potential **regular expression denial of service** attacks by preventing PCRE (the underlying pattern matching library that evaluates regular expressions in ModSecurity / the WAF) from consuming huge amounts of system resources. PCRE uses a function called `match()` which it calls repeatedly, sometimes recursively. The match limits are imposed on the number of times this function is called during a match. The default values are both 250000, which is the minimum recommended value. A modern system (i.e. 4+ CPU cores, 8+ GB of RAM) could run without issue in production with match limits of 500000.

**Important** | After modifying any advanced WAF settings, be sure to save the new configuration to disk by pressing the **Set PCRE Match Limits** button and then apply the new configuration by reloading services as prompted in the "commit changes" message box.

### PCRE Match Limit

*Default value: 250000.*

Defines the global value of the `SecPcreMatchLimit` directive, which sets a maximum limit on the number of calls to the underlying match function when evaluating a regular expression in the WAF engine.

### PCRE Match Limit Recursion

*Default value: 250000.*

Defines the global value of the `SecPcreMatchLimitRecursion` directive, which sets a maximum limit on the number of *recursive* calls to the underlying match function when evaluating a regular expression in the WAF engine.

## WAF - Advanced Configuration

PCRE Match Limit

250000



PCRE Match Limit Recursion

250000



Set PCRE Match Limits


## Working With the Core Rule Set

### What is the Core Rule Set?

The OWASP® (Open Web Application Security Project) **CRS (Core Rule Set)** is an open source collection of rules that work with ModSecurity® and compatible web application firewalls (WAFs). These rules are designed to provide easy to use, generic attack detection capabilities, with a minimum of false positives (false alerts), to your web application as part of a well balanced defense-in-depth solution.

### Core Rule Set Map

A map of the Core Rule Set can be [downloaded as a PDF file](#) from our website. It shows each rule along with its ID number, a summary explanation, its phase, its paranoia level, and any relationships it has with other rules.



## Anomaly Scoring

The Core Rule Set 3 is designed as an anomaly scoring rule set. This section explains what anomaly scoring is and how to use it.

### Overview of Anomaly Scoring

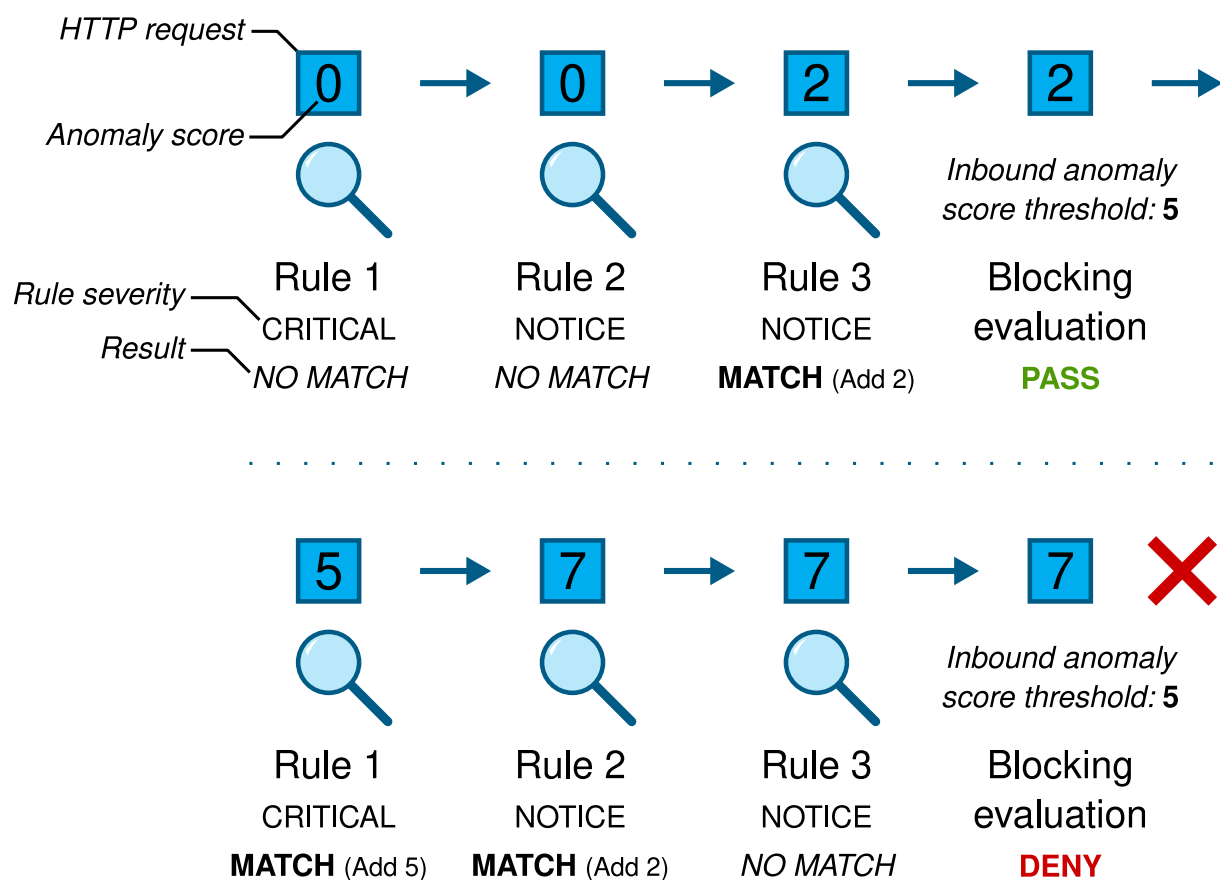
Anomaly scoring, also known as "collaborative detection", is a scoring mechanism used in the Core Rule Set. It assigns a numeric score to HTTP transactions (requests and responses), representing how 'anomalous' they appear to be. Anomaly scores can then be used to make blocking decisions. The default CRS blocking policy, for example, is to block any transaction that meets or exceeds a defined anomaly score threshold.

### How Anomaly Scoring Mode Works

Anomaly scoring mode combines the concepts of *collaborative detection* and *delayed blocking*. The key idea to understand is that **the inspection/detection rule logic is decoupled from the blocking functionality**.

Individual rules designed to detect specific types of attacks and malicious behavior are executed. If a rule matches, no immediate disruptive action is taken (e.g. the transaction is not blocked). Instead, the matched rule contributes to a transactional *anomaly score*, which acts as a running total. The rules just handle detection, adding to the anomaly score if they match. In addition, an individual matched rule will typically log a record of the match for later reference, including the ID of the matched rule, the data that caused the match, and the URI that was being requested.

Once all of the rules that inspect *request* data have been executed, *blocking evaluation* takes place. If the anomaly score is greater than or equal to the inbound anomaly score threshold then the transaction is *denied*. Transactions that are not denied continue on their journey.



Continuing on, once all of the rules that inspect *response* data have been executed, a second round of blocking evaluation takes place. If the *outbound* anomaly score is greater than or equal to the outbound anomaly score threshold then the transaction is *denied*.

#### Note

Having separate inbound and outbound anomaly scores and thresholds allows for request data and response data to be inspected and scored independently.

### Summary of Anomaly Scoring Mode

To summarize, anomaly scoring mode in the CRS works like so:

1. Execute all *request* rules
2. Make a blocking decision using the *inbound* anomaly score threshold
3. Execute all *response* rules
4. Make a blocking decision using the *outbound* anomaly score threshold

### Anomaly Score Thresholds

An anomaly score threshold is the cumulative anomaly score at which an inbound request or an outbound response will be blocked.

Most detected inbound threats carry an anomaly score of 5 (by default), while smaller violations, e.g. protocol and

standards violations, carry lower scores. An anomaly score threshold of 7, for example, would require multiple rule matches in order to trigger a block (e.g. one "critical" rule scoring 5 plus a lesser-scoring rule, in order to reach the threshold of 7). An anomaly score threshold of 10 would require at least two "critical" rules to match, or a combination of many lesser-scoring rules. **Increasing the anomaly score thresholds makes the CRS less sensitive** and hence less likely to block transactions.

Rule coverage should be taken into account when setting anomaly score thresholds. Different CRS rule categories feature different numbers of rules. SQL injection, for example, is covered by more than 50 rules. As a result, a real world SQLi attack can easily gain an anomaly score of 15, 20, or even more. On the other hand, a rare protocol attack might only be covered by a single, specific rule. If such an attack only causes the one specific rule to match then it will only gain an anomaly score of 5. If the inbound anomaly score threshold is set to anything higher than 5 then attacks like the one described will not be stopped. As such, a CRS installation should aim for an inbound anomaly score threshold of 5.

**Warning** | Increasing the anomaly score thresholds may allow some attacks to bypass the CRS rules.

**Note** | An outbound anomaly score threshold of 4 (the default) will block a transaction if any single response rule matches.

CRS uses two anomaly score thresholds, which can be defined for each WAF gateway. This is done using the WebUI, by navigating to: *Cluster Configuration > WAF - Gateway* and clicking **Modify** next to the relevant WAF. The two score thresholds are:

- Inbound Anomaly Score threshold
- Outbound Anomaly Score threshold

### Severity Levels

Each CRS rule has an associated *severity level*. Different severity levels have different anomaly scores associated with them. This means that different rules can increment the anomaly score by different amounts if the rules match.

The four severity levels and their *default* anomaly scores are:

Severity Level	Anomaly Score
CRITICAL	5
ERROR	4
WARNING	3
NOTICE	2

For example, by default, a single matching CRITICAL rule would increase the anomaly score by 5, while a single matching WARNING rule would increase the anomaly score by 3.

## Paranoia Levels


Paranoia levels are an essential concept when working with the Core Rule Set. This section explains the concept behind paranoia levels and how to work with them on a practical level.

### Introduction to Paranoia Levels

The paranoia level (PL) makes it possible to define how aggressive the Core Rule Set is. Paranoia level 1 (PL 1)

provides a set of rules that hardly ever trigger a false alarm (ideally never, but it can happen, depending on the local setup). PL 2 provides additional rules that detect more attacks (these rules operate *in addition* to the PL 1 rules), but there's a chance that the additional rules will also trigger new false alarms over perfectly legitimate HTTP requests.

This continues at PL 3, where more rules are added, namely for certain specialized attacks. This leads to even more false alarms. Then at PL 4, the rules are so aggressive that they detect almost every possible attack, yet they also flag a lot of legitimate traffic as malicious.



A higher paranoia level makes it harder for an attacker to go undetected. Yet this comes at the cost of more false positives: more false alarms. That's the downside to running a rule set that detects almost everything: your business / service / web application is also disrupted.

When false positives occur they need to be tuned away. In ModSecurity parlance: rule exclusions need to be written. A rule exclusion is a rule that disables another rule, either disabled completely or disabled partially only for certain parameters or for certain URIs. This means **the rule set remains intact** yet the CRS installation is no longer affected by the false positives.

#### Note

Depending on the complexity of the service (web application) in question and on the paranoia level, the process of writing rule exclusions can be a *substantial* amount of work.

For more information on this topic, see the section on [False Positives and Tuning](#).

## Description of the Four Paranoia Levels

The CRS project views the four paranoia levels as follows:

Paranoia Level	Description
1	Baseline security with a minimal need to tune away false positives. This is CRS for everybody running an HTTP server on the internet. Please report any false positives encountered with a PL 1 system to <a href="mailto:support@loadbalancer.org">support@loadbalancer.org</a> .
2	Rules that are adequate when real user data is involved. Perhaps an off-the-shelf online shop. Expect to encounter false positives and learn how to tune them away.
3	Online banking level security with lots of false positives. From the CRS project's perspective, false positives are accepted and expected here, so it's important to learn how to write rule exclusions.

Paranoia Level	Description
4	Rules that are so strong (or paranoid) they're adequate to protect the "crown jewels". To be used at one's own risk: be prepared to face a large number of false positives.

## Choosing an Appropriate Paranoia Level

It's important to think about a service's security requirements. The difference between protecting a personal website and the admin gateway controlling access to an enterprise's Active Directory are very different. The paranoia level needs to be chosen accordingly, while also considering the resources (time) required to tune away false positives at higher paranoia levels.

Running at the highest paranoia level, PL 4, may seem appealing from a security standpoint, but *it could take many weeks to tune away the false positives encountered*. It is crucial to have enough time to fully deal with all false positives.

**Warning** Failure to properly tune an installation runs the risk of exposing users to a vast number of false positives. This can lead to a poor user experience, and might ultimately lead to a decision to completely disable the WAF / Core Rule Set. As such, **setting a high PL in blocking mode *without* adequate tuning to deal with false positives is very risky.**

For an enterprise environment, consider developing an internal policy to map the risk levels and security needs of different assets to the minimum acceptable paranoia level to be used for them, for example:

- **Risk Class 0:** No personal data involved → PL 1
- **Risk Class 1:** Personal data involved, e.g. names and addresses → PL 2
- **Risk Class 2:** Sensitive data involved, e.g. financial/banking data; highest risk class → PL 3

## Setting the Paranoia Level

To set the paranoia level for a WAF gateway, navigate to *Cluster Configuration > WAF - Gateway* and click **Modify** next to the WAF gateway in question.

From the **Paranoia Level** drop-down list, select the desired paranoia level.

**Paranoia Level**

Paranoia Level 1 ▼

**Important** After modifying any WAF gateway settings, be sure to save the new configuration to disk by pressing the **Update** button and then apply the new configuration by reloading services as prompted in the "commit changes" message box.

## How Paranoia Levels Relate to Anomaly Scoring

It's important to understand that paranoia levels and CRS anomaly scoring (the CRS anomaly threshold/limit) are **two entirely different things with no direct connection**. The paranoia level controls the number of rules that are enabled while the anomaly threshold defines how many rules can be triggered before a request is blocked.

At the conceptual level, these two ideas *could* be mixed if the goal was to create a particularly granular security concept. For example, saying "we define the anomaly threshold to be 10, but we compensate for this by running at paranoia level 3, which we acknowledge brings more rule alerts and higher anomaly scores."

This is *technically* correct but it overlooks the fact that there are attack categories where CRS scores very low. For example, there is a plan to introduce a new rule to detect POP3 and IMAP injections: this will be a single rule, so, under normal circumstances, an IMAP injection would never score more than 5. Therefore, an installation running at an anomaly threshold of 10 could never block an IMAP injection, even if running at PL 3. In light of this, it's generally advised to **keep things simple and separate**: a CRS installation should aim for an anomaly threshold of 5 and a paranoia level as deemed appropriate.

## False Positives and Tuning

When a *genuine* transaction causes a rule from the Core Rule Set to match in error it is described as a **false positive**. False positives need to be tuned away by writing *rule exclusions*, as this section explains.

### What are False Positives?

The Core Rule Set provides *generic* attack detection capabilities. A fresh CRS deployment has no awareness of the web services that may be running behind it, or the quirks of how those services work. It is possible that *genuine* transactions may cause some CRS rules to match in error, if the transactions happen to match one of the generic attack behaviors or patterns that are being detected. Such a match is referred to as a *false positive*, or false alarm.

False positives are particularly likely to happen when operating at higher **paranoia levels**. While paranoia level 1 is designed to cause few, ideally zero, false positives, higher paranoia levels are increasingly likely to cause false positives. Each successive paranoia level introduces additional rules, with *higher* paranoia levels adding *more aggressive* rules. As such, the higher the paranoia level is the more likely it is that false positives will occur. That is the cost of the higher security provided by higher paranoia levels: the additional time it takes to tune away the increasing number of false positives.

### Example False Positive

Imagine deploying the CRS in front of a WordPress instance. The WordPress engine features the ability to add HTML to blog posts (as well as JavaScript, if you're an administrator). Internally, WordPress has rules controlling which HTML tags are allowed to be used. This list of allowed tags has been studied heavily by the security community and it's considered to be a secure mechanism.

Consider the CRS inspecting a request with a URL like the following:

```
www.example.com/?wp_post=<h1>Welcome+To+My+Blog</h1>
```

At paranoia level 2, the `wp_post` query string parameter would trigger a match against an XSS attack rule due to the presence of HTML tags. CRS is unaware that the problem is properly mitigated on the server side and, as a result, the request causes a false positive and may be blocked. The false positive may generate an error log line like the following:

```
[Wed Jan 01 00:00:00.123456 2022] [:error] [pid 2357:tid 140543564093184] [client 10.0.0.1:0]
[client 10.0.0.1] ModSecurity: Warning. Pattern match
"<(?:a|abbr|acronym|address|applet|area|audioscope|b|base|basefont|bdo|bgsound|big|blackface|
blink|blockquote|body|bq|br|button|caption|center|cite|code|col|colgroup|comment|dd|del|dfn|di
r|div|dl|dt|em|embed|fieldset|fn|font|form|frame|frameset|h1|head|..." at ARGS:wp_post. [file
"/etc/crs/rules/REQUEST-941-APPLICATION-ATTACK-XSS.conf"] [line "783"] [id "941320"] [msg
"Possible XSS Attack Detected - HTML Tag Handler"] [data "Matched Data: <h1> found within
ARGS:wp_post: <h1>welcome to my blog</h1>"] [severity "CRITICAL"] [ver "OWASP_CRS/3.3.2"] [tag
"application-multi"] [tag "language-multi"] [tag "platform-multi"] [tag "attack-xss"] [tag
"OWASP_CRS"] [tag "capec/1000/152/242/63"] [tag "PCI/6.5.1"] [tag "paranoia-level/2"]
[hostname "www.example.com"] [uri "/" ] [unique_id "Yad-7q03dV56xYsnGhYJlQAAAAA"]
```

This example log entry provides lots of information about the rule match. Some of the key pieces of information



are:

- The message from ModSecurity, which explains what happened and where:

```
ModSecurity: Warning. Pattern match "<(?:a|abbr|acronym ...)" at ARGS:wp_post.
```

- The rule ID of the matched rule:

```
[id "941320"]
```

- The additional matching data from the rule, which explains precisely what caused the rule match:

```
[data "Matched Data: <h1> found within ARGS:wp_post: <h1>welcome to my  
blog</h1>"]
```

## Why are False Positives a Problem?

### Alert Fatigue

If a system is prone to reporting false positives then the alerts it raises may be ignored. This may lead to real attacks being overlooked. For this reason, leaving false positives mixed in with real attacks is dangerous: the false positives should be resolved.

### Sensitive Information and Regulatory Compliance

A false positive alert may contain sensitive information, for example usernames, passwords, and payment card data. Imagine a situation where a web application user has set their password to '/bin/bash': without proper tuning, this input would cause a false positive every time the user logged in, writing the user's password to the error log file in plaintext as part of the alert.

It's also important to consider issues surrounding regulatory compliance. Data protection and privacy laws, like GDPR and CCPA, place strict duties and limitations on what information can be gathered and how that information is processed and stored. The unnecessary logging data generated by false positives can cause problems in this regard.

### Poor User Experience

When working in strict blocking mode, false positives can cause legitimate user transactions to be blocked, leading to poor user experience. This can create pressure to disable the CRS or even to remove the WAF solution entirely, which is an unnecessary sacrifice of security for usability. The correct solution to this problem is to tune away the false positives so that they don't reoccur in the future.

## Tuning Away False Positives

### Directly Modifying CRS Rules

**Warning** | Making direct modifications to CRS rule files is a bad idea and is strongly discouraged.

It may seem logical to prevent false positives by modifying the offending CRS rules. If a detection pattern in a CRS rule is causing matches with genuine transactions then the pattern could be modified. **This is a bad idea.**

*Directly modifying CRS rules essentially creates a fork of the rule set. Any modifications made would be undone by a rule set update, meaning that any changes would need to be continually reapplied by hand. This is a tedious, time consuming, and error-prone solution.*

There are alternative ways to deal with false positives, as described below. These methods sometimes require slightly more effort and knowledge but they do not cause problems when performing rule set updates.

## Rule Exclusions

### Overview

The ModSecurity WAF engine has flexible ways to tune away false positives. It provides several *rule exclusion* (RE) mechanisms which allow rules to be modified *without* directly changing the rules themselves. This makes it possible to work with third-party rule sets, like the Core Rule Set, by adapting rules as needed while leaving the rule set files intact and unmodified. This allows for easy rule set updates.

Two fundamentally different types of rule exclusions are supported:

- **Configure-time rule exclusions:** Rule exclusions that are applied once, at *configure-time* (e.g. when (re)starting or reloading ModSecurity, or the server process that holds it). For example: "remove rule X at startup and never execute it."

This type of rule exclusion takes the form of a ModSecurity directive, e.g. `SecRuleRemoveById`.

- **Runtime rule exclusions:** Rule exclusions that are applied at *runtime* on a per-transaction basis (e.g. exclusions that can be conditionally applied to some transactions but not others). For example: "if a transaction is a POST request to the location 'login.php', remove rule X."

This type of rule exclusion takes the form of a `SecRule`.

#### Note

Runtime rule exclusions, while granular and flexible, have a computational overhead, albeit a small one. A runtime rule exclusion is an extra `SecRule` which must be evaluated for every transaction.

In addition to the two *types* of exclusions, rules can be excluded in two different *ways*:

- **Exclude the entire rule/tag:** An entire rule, or entire category of rules (by specifying a tag), is removed and will not be executed by the rule engine.
- **Exclude a specific variable from the rule/tag:** A *specific variable* will be excluded from a specific rule, or excluded from a category of rules (by specifying a tag).

These two methods can also operate on multiple individual rules, or even entire rule categories (identified either **by tag** or by using a **range of rule IDs**).

The combinations of rule exclusion types and methods allow for writing rule exclusions of varying granularity. Very coarse rule exclusions can be written, for example "remove all SQL injection rules" using `SecRuleRemoveByTag`. Extremely granular rule exclusions can also be written, for example "for transactions to the location 'web\_app\_2/function.php', exclude the query string parameter 'user\_id' from rule 920280" using a `SecRule` and the action `ctl:ruleRemoveTargetById`.

The different rule exclusion types and methods are summarized in the table below, which presents the main ModSecurity directives and actions that can be used for each type and method of rule exclusion:

	Exclude entire rule/tag	Exclude specific variable from rule/tag
Configure-time	<code>SecRuleRemoveById*</code> <code>SecRuleRemoveByTag</code>	<code>SecRuleUpdateTargetById</code> <code>SecRuleUpdateTargetByTag</code>

	Exclude entire rule/tag	Exclude specific variable from rule/tag
Runtime	ctl:ruleRemoveById** ctl:ruleRemoveByTag	ctl:ruleRemoveTargetById ctl:ruleRemoveTargetByTag

\*Can also exclude ranges of rules or multiple space separated rules.

\*\*Can also exclude ranges of rules.

#### Important

It is not currently possible to exclude *phase 1* rules using *runtime* rule exclusions on Loadbalancer.org appliances. Configure-time rule exclusions must be used instead. This is a known limitation of the WAF implementation. The [Core Rule Set Map](#) can be used to lookup the phase of a CRS rule. More information on rule phases can be found in the [ModSecurity Reference Manual](#).

#### Note

There's also a third group of rule exclusion directives and actions, the use of which is discouraged. As well as excluding rules "ById" and "ByTag", it's also possible to exclude "ByMsg" (SecRuleRemoveByMsg, SecRuleUpdateTargetByMsg, ctl:ruleRemoveByMsg, and ctl:ruleRemoveTargetByMsg). This excludes rules based on the message they write to the error log. These messages can be dynamic and may contain special characters. As such, trying to exclude rules by message is difficult and error-prone.

#### Rule Tags

CRS rules typically feature multiple tags, grouping them into different categories. For example, a rule might be tagged by attack type ('attack-rce', 'attack-xss', etc.), by language ('language-java', 'language-php', etc.), and by platform ('platform-apache', 'platform-unix', etc.).

Tags can be used to remove or modify entire categories of rules all at once, but some tags are more useful than others in this regard. Tags for *specific* attack types, languages, and platforms may be useful for writing rule exclusions. For example, if lots of the SQL injection rules are causing false positives but SQL isn't in use anywhere in the back end web application then it may be worthwhile to remove all CRS rules tagged with 'attack-sqli' (SecRuleRemoveByTag attack-sqli).

Some rule tags are *not* useful for rule exclusion purposes. For example, there are generic tags like 'language-multi' and 'platform-multi': these contain hundreds of rules across the entire CRS, and they don't represent a meaningful rule property to be useful in rule exclusions. There are also tags that categorize rules based on well known security standards, like CAPEC and PCI DSS (e.g. 'capec/1000/153/267', 'PCI/6.5.4'). These tags may be useful for informational and reporting purposes but are not useful in the context of writing rule exclusions.

Excluding rules using tags may be more useful than excluding using rule ranges in situations where a category of rules is spread across multiple files. For example, the 'language-php' rules are spread across several different rule files (both inbound and outbound rule files).

#### Rule Ranges

As well as rules being tagged using different categories, CRS rules are organized into files by general category. In addition, CRS rule IDs follow a consistent numbering convention. This makes it easy to remove unwanted types of rules by removing ranges of rule IDs. For example, the file REQUEST-913-SCANNER-DETECTION.conf contains rules related to detecting well known scanners and crawlers, which all have rule IDs in the range 913000-913999. All of the rules in this file can be easily removed using a configure-time rule exclusion, like so:

```
SecRuleRemoveById "913000-913999"
```

Excluding rules using rule ranges may be more useful than excluding using tags in situations where tags are less relevant or where tags vary across the rules in question. For example, a rule range may be the most appropriate solution if the goal is to remove all rules contained in a single file, regardless of how the rules are tagged.

#### Support for Regular Expressions

Most of the configure-time rule exclusion directives feature some level of support for using regular expressions. This makes it possible, for example, to exclude a dynamically named variable from a rule. The directives with support for regular expressions are:

- `SecRuleRemoveByTag`

A regular expression is used for the tag match. For example, `SecRuleRemoveByTag "injection"` would match both "attack-injection-generic" and "attack-injection-php".

- `SecRuleRemoveByMsg`

A regular expression is used for the message match. For example, `SecRuleRemoveByMsg "File Access"` would match both "OS File Access Attempt" and "Restricted File Access Attempt".

- `SecRuleUpdateTargetById`, `SecRuleUpdateTargetByTag`, `SecRuleUpdateTargetByMsg`

A regular expression can optionally be used in the target specification by enclosing the regular expression in forward slashes. This is useful for dealing with dynamically named variables, like so:

```
SecRuleUpdateTargetById 942440 "!REQUEST_COOKIES:/^uid_.*/"
```

This example would exclude request cookies named "uid\_0123456", "uid\_6543210", etc. from rule 942440.

#### Note

The 'ctl' action for writing runtime rule exclusions does **not** support any use of regular expressions. This is a known limitation of the ModSecurity rule engine.

#### Example 1 (`SecRuleRemoveById`)

(Configure-time RE. Exclude entire rule.)

**Scenario:** Rule 933151, "PHP Injection Attack: Medium-Risk PHP Function Name Found", is causing false positives. The web application behind the WAF makes no use of PHP. As such, it is deemed safe to tune away this false positive by completely removing rule 933151.

#### Rule Exclusion:

```
# CRS Rule Exclusion: 933151 - PHP Injection Attack: Medium-Risk PHP Function Name Found
SecRuleRemoveById 933151
```

#### Example 2 (`SecRuleRemoveByTag`)

(Configure-time RE. Exclude entire tag.)

**Scenario:** Several different parts of a web application are causing false positives with various SQL injection rules. None of the web services behind the WAF make use of SQL, so it is deemed safe to tune away these false

positives by removing all the SQLi detection rules.

#### Rule Exclusion:

```
# CRS Rule Exclusion: Remove all SQLi detection rules
SecRuleRemoveByTag attack-sqli
```

Example 3 (*SecRuleUpdateTargetById*)

(Configure-time RE. Exclude specific variable from rule.)

**Scenario:** The content of a POST body parameter named 'wp\_post' is causing false positives with rule 941320, "Possible XSS Attack Detected - HTML Tag Handler". Removing this rule entirely is deemed to be unacceptable: the rule is not causing any other issues, and the protection it provides should be retained for everything apart from 'wp\_post'. It is decided to tune away this false positive by excluding 'wp\_post' from rule 941320.

#### Rule Exclusion:

```
# CRS Rule Exclusion: 941320 - Possible XSS Attack Detected - HTML Tag Handler
SecRuleUpdateTargetById 941320 "!ARGS:wp_post"
```

Example 4 (*SecRuleUpdateTargetByTag*)

(Configure-time RE. Exclude specific variable from rule.)

**Scenario:** The values of request cookies with random names of the form 'uid\_<STRING>' are causing false positives with various SQL injection rules. It is decided that it is not a risk to allow SQL-like content in cookie values, however it is deemed unacceptable to disable the SQLi detection rules for anything apart from the request cookies in question. It is decided to tune away these false positives by excluding only the problematic request cookies from the SQLi detection rules. A regular expression is to be used to handle the random string portion of the cookie names.

#### Rule Exclusion:

```
# CRS Rule Exclusion: Exclude the request cookies 'uid_<STRING>' from the SQLi detection rules
SecRuleUpdateTargetByTag attack-sqli "!REQUEST_COOKIES:/^uid_.*/"
```

Example 5 (*ctl:ruleRemoveById*)

(Runtime RE. Exclude entire rule.)

**Scenario:** Rule 920230, "Multiple URL Encoding Detected", is causing false positives at the specific location '/webapp/function.php'. This is being caused by a known quirk in how the web application has been written, and it cannot be fixed in the application. It is deemed safe to tune away this false positive by removing rule 920230 for that specific location only.

#### Rule Exclusion:

```
# CRS Rule Exclusion: 920230 - Multiple URL Encoding Detected
SecRule REQUEST_URI "@beginsWith /webapp/function.php" \
    id:1000, \
    phase:1, \
    pass, \
    nolog, \
```

```
ctl:ruleRemoveById=920230"
```

Example 6 (*ctl:ruleRemoveByTag*)

(Runtime RE. Exclude entire tag.)

**Scenario:** Several different locations under '/web\_app\_1/content' are causing false positives with various SQL injection rules. Nothing under that location makes any use of SQL, so it is deemed safe to remove all the SQLi detection rules for that location. Other locations *may* make use of SQL, however, so the SQLi detection rules **must** remain in place everywhere else. It has been decided to tune away the false positives by removing all the SQLi detection rules for locations under '/web\_app\_1/content' only.

#### Rule Exclusion:

```
# CRS Rule Exclusion: Remove all SQLi detection rules
SecRule REQUEST_URI "@beginsWith /web_app_1/content" \
    id:1010, \
    phase:1, \
    pass, \
    nolog, \
    ctl:ruleRemoveByTag=attack-sqli"
```

Example 7 (*ctl:ruleRemoveTargetById*)

(Runtime RE. Exclude specific variable from rule.)

**Scenario:** The content of a POST body parameter named 'text\_input' is causing false positives with rule 941150, "XSS Filter - Category 5: Disallowed HTML Attributes", at the specific location '/dynamic/new\_post'. Removing this rule entirely is deemed to be unacceptable: the rule is not causing any other issues, and the protection it provides should be retained for everything apart from 'text\_input' at the specific problematic location. It is decided to tune away this false positive by excluding 'text\_input' from rule 941150 for location '/dynamic/new\_post' only.

#### Rule Exclusion:

```
# CRS Rule Exclusion: 941150 - XSS Filter - Category 5: Disallowed HTML Attributes
SecRule REQUEST_URI "@beginsWith /dynamic/new_post" \
    id:1020, \
    phase:1, \
    pass, \
    nolog, \
    ctl:ruleRemoveTargetById=941150;ARGS:text_input"
```

Example 8 (*ctl:ruleRemoveTargetByTag*)

(Runtime RE. Exclude specific variable from rule.)

**Scenario:** The values of request cookie 'uid' are causing false positives with various SQL injection rules when trying to log in to a web service at location '/webapp/login.html'. It is decided that it is not a risk to allow SQL-like content in this specific cookie's values for the login page, however it is deemed unacceptable to disable the SQLi detection rules for anything apart from the specific request cookie in question at the login page only. It is decided to tune away these false positives by excluding only the problematic request cookie from the SQLi detection rules, and only when accessing '/webapp/login.html'.

#### Rule Exclusion:

```
# CRS Rule Exclusion: Exclude the request cookie 'uid' from the SQLi detection rules
```

```
SecRule REQUEST_URI "@beginsWith /webapp/login.html" \
    "id:1030,\
    phase:1,\
    pass,\
    nolog,\
    ctl:ruleRemoveTargetByTag=attack-sqli;REQUEST_COOKIES:uid"
```

It's possible to write a conditional rule exclusion that tests something other than just the request URI. Conditions can be built which test, for example, the source IP address, HTTP request method, HTTP headers, and even the day of the week.

**Tip**

Multiple conditions can also be chained together to create a logical AND by using ModSecurity's **chain** action. This allows for creating powerful rule logic like "for transactions that are from source IP address 10.0.0.1 AND that are for location '/login.html', exclude the query string parameter 'user\_id' from rule 920280". Extremely granular and specific rule exclusions can be written, in this way.

## Adding Custom WAF Configuration

Rule exclusions and custom rules can be added to a WAF gateway. To access this functionality:

1. Using the WebUI, navigate to: *Cluster Configuration > WAF - Manual Configuration*.
2. From the drop-down list, select the name of the WAF gateway in question.

By default, a WAF gateway's manual configuration contains a set of commented-out examples. The examples can be uncommented and adapted for use. Alternatively, all of the example text can safely be deleted to give a clean slate.

To add rule exclusions or custom rules, write or paste the custom content directly into the text box.

## WAF - Manual Configuration

WAF-Web\_Service ▾

```
1 # CRS Rule Exclusion: 932150 - Remote Command Execution: Direct Unix Command
2 #                               Execution
3 #
4 # The argument "text_input" contains arbitrary user input. Natural text like
5 # "ping pong is great" causes a false positive with rule 932150. Exclude
6 # "text_input" from the rule.
7 SecRuleUpdateTargetById 932150 "!ARGS:text_input"
8
9
10
11 # All traffic hitting the WAF comes from the public internet, so deny all access
12 # to locations under "/app/admin_portal".
13 SecRule REQUEST_URI "@beginsWith /app/admin_portal/" \
14     "id:1000,\
15     phase:1,\
16     deny,\
17     log,\
18     msg:'Deny admin portal access through the WAF, i.e. for external traffic'"
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
```

Update

### Important

After modifying any WAF manual configuration, be sure to save the new configuration to disk by pressing the **Update** button and then apply the new configuration by reloading services as prompted in the "commit changes" message box.

## Per-Transaction Resource Limits

WAFs can be susceptible to simple denial-of-service attacks. Bombarding a WAF with certain types of HTTP requests can cause serious problems:

- Requests with large bodies (excluding file uploads) consume a significant amount of CPU time while the WAF rules are executed against the body data.
- Requests with large bodies consume a large amount of memory while they're buffered for inspection.
- Requests with very large headers consume a significant amount of CPU time while the WAF rules are executed against the header data.
- Large file uploads consume a large amount of disk space while they're buffered during a transaction.

A WAF can thus be attacked from multiple angles: attempting to exhaust its CPU, memory, and disk space. As such, it is imperative to control and limit the resources that a **single HTTP transaction** can consume. This hardens the WAF against these kinds of attacks.

### Default Limits

The following default resource limits are enforced for all WAF gateways:

- Maximum request body size (including file uploads): **13 MB** (13107200 bytes)




- Maximum request body size (excluding file uploads): **131 kB** (131072 bytes)
- Maximum amount of request body to store in memory before switching to disk: **131 kB** (131072 bytes)
- Maximum size of a single HTTP request header field: **8 kB** (8190 bytes)

#### Tip

The default limits for the maximum non-file request body size and the maximum amount of memory to use to buffer a request body are equal. This means that non-file request bodies under the size limit will always be buffered in memory, which is fast. Only request bodies of file uploads exceeding 131 kB will be streamed to disk, which is much slower.

#### Example 1

As an example, under the default limits, it is permissible to submit an HTTP request with up to 131 kB of non-file upload data, e.g. data from filling in form fields:




A request that exceeds the limit will:

- receive a 400 Bad Request response status code if the WAF engine is **enabled**
- pass through the WAF gateway without issue if the WAF engine is **not enabled**, i.e. is in **detection only mode**

#### Example 2

As another example, under the default limits, it is permissible to submit an HTTP request with up to 13 MB of file upload data, either one large file or multiple smaller files:



A request that exceeds the limit will:

- receive a `500 Internal Server Error` response status code if the WAF engine is **enabled**
- pass through the WAF gateway without issue if the WAF engine is **not enabled**, i.e. is in **detection only mode**

## Handling Large HTTP Requests

When working with some web applications, it can be necessary to handle HTTP requests that exceed the default size limits in one or more ways.

### Large File Uploads

The maximum request body size, **including** file uploads, should be as large as the largest file upload that the web application should accept. If the web application should accept file uploads that are *larger* than the default 13 MB limit then the limit should be increased.

As an example, to increase the limit to **20 MB** for a given WAF gateway, add the following directive to the WAF gateway's manual configuration:

```
# Buffer request bodies of up to 20 MB in size.
# Allows for requests to contain up to 20 MB of file upload data.
SecRequestBodyLimit 20000000
```

For additional granularity, a `Location` directive can be used to increase the limit for certain locations only. For example:

```
# For the admin portal upload page only, buffer request bodies of up to 500 MB
# in size. Allows for requests to contain up to 500 MB of file upload data.
<Location "/admin-portal/upload.php">
  SecRequestBodyLimit 500000000
</Location>
```

**Note** | Request body handling has a hard upper limit of 1 GB.

## Requests Containing Large Amounts of Non-File Upload Data

The maximum request body size, **excluding** file uploads, should be as small as is practical. An example of non-file upload data would be data from filling in form fields. Unlike file upload data, **non-file** upload data is inspected: the WAF rules are executed against the data. For large chunks of data, rule execution can use a significant amount of CPU time.

If a web application should accept request bodies containing *more* than the default limit of 131 kB of non-file upload data then the limit should be increased. Some web applications transfer state data or small files in this way, which can often exceed the 131 kB limit where this technique is used.

### Warning

Increasing the non-file upload request body size limit makes the WAF more susceptible to denial-of-service attacks. This limit should **only** be increased if necessary and should still be kept as small as possible.

As an example, to increase the limit to **200 kB** for a given WAF gateway, add the following directive to the WAF gateway's manual configuration:

```
# Buffer request bodies, excluding file uploads, of up to 200 kB in size.
# Allows for requests to contain up to 200 kB of non-file upload data.
SecRequestBodyNoFilesLimit 200000
```

Because increasing this limit makes a WAF more susceptible to denial-of-service attacks, it is a very good idea to increase the limit only where it is needed. For additional granularity, a `Location` directive can be used to increase the limit for certain locations only. For example:

```
# For the user profile avatar upload functionality only, buffer request bodies
# containing up to 200 kB of non-file upload data. Allows users to upload
# avatars up to 200 kB in size.
<Location "/user/profile/avatar-upload.php">
    SecRequestBodyNoFilesLimit 200000
</Location>
```

### Note

The limit for the maximum amount of memory to use to buffer a request body can be set using the `SecRequestBodyInMemoryLimit` directive. The default values of the `SecRequestBodyNoFilesLimit` and `SecRequestBodyInMemoryLimit` directives are equal.

## Requests with Massive Header Fields

If a web application requires the use of massive HTTP request header fields, with any single header field exceeding the default size limit of 8 kB, then the limit should be increased.

As an example, to increase the limit to **13 kB** for a given WAF gateway, which should be large enough to let through many known large authentication headers which are seen very occasionally, add the following directive to the WAF gateway's manual configuration:

```
# Allow HTTP request header fields of up to 13 kB in size.
# Should be large enough to let through many known large authentication headers.
LimitRequestFieldSize 13000
```

Note that the `LimitRequestFieldSize` directive cannot be defined inside a `Location` directive, so being granular with this setting is not possible.

## Note

To handle HTTP request header fields above *15 kB* in size it's necessary to increase HAProxy's **Request buffer length**. This can be found under *Cluster Configuration > Layer 7 - Advanced Configuration*. This is required as HAProxy becomes the limiting factor, not Apache.

# WAF Gateway Error Logs

## Logging Mechanism Overview

If a transaction does not cause any ModSecurity (WAF) rules to match then nothing is written to the WAF error log.

If a transaction *does* cause rules to match:

- A record of each rule match is written to the WAF error log.
- If a transaction's inbound or outbound anomaly score is high enough to reach either the inbound or outbound anomaly score thresholds then a record of this is written to the WAF error log, stating the action taken, if any (e.g. "Access denied with code 403").
- A summary of the transaction's anomaly score is written to the WAF error log, stating the makeup of the anomaly score total per-paranoia level and per-rule category.

## Tip

If a rule is *expected* to match and should *not* log those matches for some reason (e.g. it's a helper rule), the `noLog` action can be used to prevent a rule from creating error log output and cluttering up the log.

## Viewing the Error Logs

To view the error logs for a given WAF gateway:

1. Using the WebUI, navigate to: *Logs > WAF Error*.
2. From the drop-down list, select the name of the WAF gateway in question.

### Default View

After selecting a WAF gateway from the drop-down list, the default view displays error log lines in their raw, unedited format. The most recent 500 log lines are displayed in **reverse chronological order**, with the most recent log entry shown at the top.

Error WAF-Web\_Service ▾

Simple

Breakdown

Fixes

Empty Log

Download

```
[Wed Feb 01 16:35:47.658635 2023] [:error] [pid 6905:tid 140248700925696] [client 192.168.85.1:0]
[client 192.168.85.1] ModSecurity: Warning. Operator GE matched 5 at TX:inbound_anomaly_score. [file
"/etc/modsecurity.d/rules_3-3-2/RESPONSE-980-CORRELATION.conf"] [line "91"] [id "980130"] [msg
"Inbound Anomaly Score Exceeded (Total Inbound Score: 5 -
SQLI=0,XSS=0,RFI=0,LFI=0,RCE=5,PHPI=0,HTTP=0,SESS=0): individual paranoia level scores: 5, 0, 0, 0"]
[ver "OWASP_CRS/3.3.2"] [tag "event-correlation"] [hostname "example.com"] [uri "/login"] [unique_id
"Y9qU4zw4XsSqM8XUvX2LCQAAAAQ"]
[Wed Feb 01 16:35:47.657121 2023] [:error] [pid 6905:tid 140248700925696] [client 192.168.85.1:0]
[client 192.168.85.1] ModSecurity: Access denied with code 403 (phase 2). Operator GE matched 5 at
TX:anomaly_score. [file "/etc/modsecurity.d/rules_3-3-2/REQUEST-949-BLOCKING-EVALUATION.conf"] [line
"100"] [id "949110"] [msg "Inbound Anomaly Score Exceeded (Total Score: 5)"] [severity "CRITICAL"]
[ver "OWASP_CRS/3.3.2"] [tag "application-multi"] [tag "language-multi"] [tag "platform-multi"] [tag
"attack-generic"] [hostname "example.com"] [uri "/login"] [unique_id "Y9qU4zw4XsSqM8XUvX2LCQAAAAQ"]
```

### Simple View

When the **Simple** button is pressed, the most recent 1000 error log lines are shown in chronological order

(matching the log file itself), with the oldest log entry shown at the top.

Log entries are broken up across multiple lines to aid readability. Individual log entries are separated by horizontal lines.

Error WAF-Web\_Service ▾

SimpleBreakdownFixesEmpty LogDownload

[Wed Feb 01 16:35:47.654387 2023]

[ :error]

[pid 6905:tid 140248700925696]

[client 192.168.85.1:0]

[client 192.168.85.1]

• ModSecurity: Warning. Matched phrase "bin/bash" at ARGS:usr.

[file "/etc/modsecurity.d/rules\_3-3-2/REQUEST-932-APPLICATION-ATTACK-RCE.conf"]

[line "500"]

• [id "932160"]

[msg "Remote Command Execution: Unix Shell Code Found"]

[data "Matched Data: bin/bash found within ARGS:usr: /bin/bash"]

[severity "CRITICAL"]

[ver "OWASP\_CRS/3.3.2"]

[tag "application-multi"]

[tag "language-shell"]

[tag "platform-unix"]

[tag "attack-rce"]

[tag "paranoia-level/1"]

[tag "OWASP\_CRS"]

[tag "capec/1000/152/248/88"]

[tag "PCI/6.5.2"]

[hostname "example.com"]

[uri "/login"]

[unique\_id "Y9qU4zW4XsSqM8XUvX2LCQAAAAQ"]

## Breakdown View

When the **Breakdown** button is pressed, a summary of which matched rules are present in the error log is presented. The information presented is:

- Occurrences: The number of times a given event is present in the log, e.g. "how many times rule X has matched for host name Y at location Z".
- Rule ID: The ID number of the rule that matched.
- Hostname: The host name that the matches were against. Useful to break up data if multiple services are sitting behind a single WAF gateway.
- Severity (optional): If present, the severity of the rule that matched (as described in the [Severity Levels](#) section).
- URI: The location that the match was against.

Error WAF-Web\_Service ▾

SimpleBreakdownFixesEmpty LogDownload

Occurrences | Rule ID | Hostname | [-Rule severity (optional)-] | URI

1 942100 example.com [- (5) CRITICAL-] /login

1 932160 example.com [- (5) CRITICAL-] /login

## Fixes View

When the **Fixes** button is pressed, a best efforts, automated list of rule exclusions are generated. These rule exclusions are based on the assumption that only known, good traffic has passed through (and hence been logged by) the WAF gateway. It is thus assumed that the error log contains only *false positives*, which the load balancer attempts to generate rule exclusions to resolve.

### Caution

The *Fixes* view is not a substitute for a proper tuning process. Writing meaningful rule exclusions requires an understanding of the web service behind the WAF gateway, which the script that generates the automated rule exclusions can never provide.

Error WAF-Web\_Service ▾

Simple

Breakdown

Fixes

Empty Log

Download

```
<LocationMatch "(?i)^/login$">
SecRuleRemoveById 932160
SecRuleRemoveById 942100
</LocationMatch>
```

## Breakdown of a Log Entry

A single, full error log entry from the *Simple* view is presented below:

[illegible]

The most important parts and their meanings are as follows:

- [Fri Dec 03 12:14:06.323315 2021]

The date and time that the log entry was written.

- [client 192.168.85.1]

The source IP address of the transaction.

- `ModSecurity: Warning. Pattern match "(?:^|=)..." at ARGS:text_input.`

A summary message from ModSecurity describing what has happened and where.

- `[id "932150"]`

The rule ID of the rule that matched.

- `[msg "Remote Command Execution: Direct Unix Command Execution"]`

A summary message from the matched rule describing what happened.

- `[data "Matched Data: ping found within ARGS:text_input: ping pong is great!"]`

Additional data from the rule describing what happened.

- `[severity "CRITICAL"]`

The severity of the rule that matched (as described in the [Severity Levels](#) section).

- `[ver "OWASP_CRS/3.3.2"]`

The Core Rule Set version of the rule that matched.

- `[tag "platform-unix"]`

A tag used to categorise the type of rule or type of attack it is designed to detect.

- `[tag "paranoia-level/1"]`

A tag used to report the paranoia level of the rule that matched.

- `[tag "OWASP_CRS"]`

A tag used to report that the matched rule was part of the Core Rule Set.

- `[hostname "example.com"]`

The host name used in the transaction.

- `[uri "/"]`

The location requested by the transaction.

- `[unique_id "YaoKDmrMWmz0GZzK9ZmMCAAAAEA"]`

The transaction's unique ID, which ties together all log entries relating to a single transaction.

An error log entry from the *Simple* view showing the outcome of inbound blocking evaluation is presented below:

```
[Fri Dec 03 12:14:06.327738 2021]
```

```

[:error]
[pid 2547:tid 140172380346112]
[client 192.168.85.1:0]
[client 192.168.85.1]
• ModSecurity: Access denied with code 403 (phase 2). Operator GE matched 5 at
TX:anomaly_score.
[file "/etc/modsecurity.d/rules_3-3-2/REQUEST-949-BLOCKING-EVALUATION.conf"]
[line "93"]
• [id "949110"]
[msg "Inbound Anomaly Score Exceeded (Total Score: 5)"]
[severity "CRITICAL"]
[ver "OWASP_CRS/3.3.2"]
[tag "application-multi"]
[tag "language-multi"]
[tag "platform-multi"]
[tag "attack-generic"]
[hostname "example.com"]
[uri "/"]
[unique_id "YaoKDrMwMz0GZzK9ZmMCAAAAEA"]

```

The most important unique parts and their meanings are as follows:

- ModSecurity: Access denied with code 403 (phase 2). Operator GE matched 5 at TX:anomaly\_score.

A summary message from ModSecurity describing what has happened (access denied with status code 403 Forbidden) and why (anomaly score reached threshold of 5).

- [msg "Inbound Anomaly Score Exceeded (Total Score: 5)"]

A summary message from the matched (blocking evaluation) rule describing what happened, in a slightly more readable way.

An error log entry from the *Simple* view showing an event correlation summary is presented below:

```

[Fri Dec 03 12:14:06.329017 2021]
[:error]
[pid 2547:tid 140172380346112]
[client 192.168.85.1:0]
[client 192.168.85.1]
• ModSecurity: Warning. Operator GE matched 5 at TX:inbound_anomaly_score.
[file "/etc/modsecurity.d/rules_3-3-2/RESPONSE-980-CORRELATION.conf"]
[line "91"]
• [id "980130"]
[msg "Inbound Anomaly Score Exceeded (Total Inbound Score: 5 -
SQLI=0,XSS=0,RFI=0,LFI=0,RCE=5,PHPI=0,HTTP=0,SESS=0): individual paranoia level scores: 5, 0,
0, 0"]
[ver "OWASP_CRS/3.3.2"]
[tag "event-correlation"]
[hostname "example.com"]
[uri "/"]
[unique_id "YaoKDrMwMz0GZzK9ZmMCAAAAEA"]

```

The most important unique part and its meaning is as follows:

- [msg "Inbound Anomaly Score Exceeded (Total Inbound Score: 5 - SQLI=0,XSS=0,RFI=0,LFI=0,RCE=5,PHPI=0,HTTP=0,SESS=0): individual paranoia level scores: 5, 0, 0, 0"]

A summary message from the matched (event correlation) rule giving a full breakdown of the transaction's anomaly score, both by paranoia level and category. The paranoia level breakdown is given in ascending paranoia level (i.e. PL 1, PL 2, PL 3, PL 4).



## Chapter 8 - Real Server Health Monitoring & Control

### Configuring Health Checks

The appliance supports a comprehensive range of health check options to check and verify the health of Real Servers. These range from simple ping checks to much more complex negotiate options to determine that the underlying daemon/service is running and responding correctly. The specific options available depend on whether services are deployed at Layer 4 or Layer 7, details of both are covered in the following sections.

#### Health Checks for Layer 4 Services

At layer 4, Real Server health checking is performed by Ldirectord. To configure health checks, use the WebUI menu option: *Cluster Configuration > Layer 4 - Virtual Services*, then click **Modify** next to the VIP to be configured. The health check options available depend on the check type selected.

Health Checks	
Check Type	Connect to port ?
Check Port	? ?
Feedback	
Feedback Method	? ?
Fallback Server	
IP Address	? ?

**Note** | For new Layer 4 VIPs the default check type is *Connect to Port*.

The following Check Types are supported:

**Negotiate** - Sends a request and looks for a specific response.

**Note** | If a Negotiate check is selected and *Response Expected* is left blank, the appliance will check the location specified in Request To Send (if blank the root will be checked) and will consider all HTTP 2xx (usually HTTP 200) and HTTP 3xx response statuses as valid and the server will be marked as up. All other responses including no response or a timeout will be considered invalid and the server will be marked as down.

**Connect to port** - Just do a simple connect to the specified port/service & verify that it's able to accept a connection.

**Ping server** - Sends an ICMP echo request packet to the Real Server.

**External script** - Use a custom file (typically a script but can also be a binary) for the health check. Select the script from the *Check Script* drop-down. For more information on using custom external health check scripts please refer to [External Health Check Scripts](#).

**No checks, always Off** - All Real Servers are assumed to be down.

**No checks, always On** - All Real Servers are assumed to be up.

**5 Connects, 1 Negotiate** - Do 5 connect checks and then 1 negotiate check.

**10 Connects, 1 Negotiate** - Do 10 connect checks and then 1 negotiate check.

The following table describes the health check options associated with each check type:

Option	Sub Option	health check Description
Negotiate		<i>Send a request and matches a receive string</i>
	Check Port	The port to monitor. This can normally be left blank in which case the port checked is the same port defined for the Real Servers. Note that for DR mode, the port cannot be specified at the Real Server level, so the port specified for the VIP is used. Sometimes the check port differs from service port.

Option	Sub Option	health check Description
	Protocol	<p>Set the protocol for the health check. The options are:</p> <ul style="list-style-type: none"> <li>• <b>HTTP</b> - use HTTP as the negotiate protocol (also requires filename, path &amp; text expected)</li> <li>• <b>HTTPS</b> - use HTTPS as the negotiate protocol (also requires filename, path &amp; text expected)</li> <li>• <b>HTTP Proxy</b> - Use an HTTP proxy check</li> <li>• <b>FTP</b> - use FTP as the negotiate protocol (also requires login/password, filename in the default folder)</li> <li>• <b>IMAP (IPv4 only)</b> - use IMAP as the negotiate protocol (requires login/password)</li> <li>• <b>IMAPS (IPv4 only)</b> - use IMAPS as the negotiate protocol (requires login/password)</li> <li>• <b>POP</b> - use POP as the negotiate protocol (also requires login/password)</li> <li>• <b>POPS</b> - use POPS as the negotiate protocol (also requires login/password)</li> <li>• <b>LDAP (IPv4 only)</b> - use LDAP as the negotiate protocol (also requires username/password)</li> <li>• <b>SMTP</b> - use SMTP as the negotiate protocol</li> <li>• <b>NNTP (IPv4 only)</b> - use NNTP as the negotiate protocol</li> <li>• <b>DNS</b> - use DNS as the negotiate protocol</li> <li>• <b>MySQL (IPv4 only)</b> - use MySQL as the negotiate protocol (also requires username/password)</li> <li>• <b>SIP</b> - use SIP as the negotiate protocol (also requires username/password)</li> <li>• <b>Simple TCP</b> - Sends a request string to the server and checks the response</li> <li>• <b>RADIUS (IPv4 only)</b> - use RADIUS as the negotiate protocol (also requires username/password)</li> </ul> <p><b>Additional Negotiate Check Options (depending on type selected) :</b></p> <p><i>Login</i> - the username when authentication is required  <i>Password</i> - the password when authentication is required  <i>Database Name</i> - The database to use for the MySQL check  <i>Radius Secret</i> - the RADIUS secret string for the RADIUS negotiate check</p>

Option	Sub Option	health check Description
	Virtual Host	Used when using a negotiate check with HTTP or HTTPS. Sets the host header used in the HTTP request. In the case of HTTPS this generally needs to match the common name of the SSL certificate. If not set then the host header will be derived from the request url for the real server if present. As a last resort the IP address of the real server will be used.
	Request to Send	This is used with negotiate checks and specifies the <i>Request to Send</i> to the server. The use of this parameter varies with the protocol selected in Negotiate Check Service. With protocols such as HTTP and FTP, this should be the object to request from the server. Bare filenames will be requested from the web or FTP root. With DNS, this should be either a name to look up in an A record, or an IP address to look up in a PTR record. With databases, this should be a SQL SELECT query. The <i>Response Expected</i> field is not used by the SQL health check since the data returned is not read, the answer must simply be 1 or more rows). With LDAP, this should be the search base for the query. The load balancer will perform an (ObjectClass=*) search relative to this base. With Simple TCP, this should be a string to send verbatim to the server.
	Response Expected	<p>This is the response that must be received for check to be a success. The check succeeds if the specified text (response) is found anywhere in the response from the web server when the file specified in the <i>Request to Send</i> field is requested.</p> <p>For example, a file called 'check.txt' could be placed in the default folder of the web server, this text file could just have the text OK in the file, then when the negotiate check runs, it would look for a file called 'check.txt' containing OK. If found, the test would succeed, if not found it would fail and no new sessions will be sent to that server.</p> <div> <p>Note</p> <p>If <i>Response Expected</i> is left blank, the appliance will check the location specified in <i>Request To Send</i> (if blank the root will be checked) and will consider all HTTP 2xx (usually HTTP 200) and HTTP 3xx response statuses as valid and the server will be marked as up. All other responses including no response or a timeout will be considered invalid and the server will be marked as down.</p> </div>
Connect to Port		<i>Attempt to make a connection to the specified port</i>
	Check Port	The port to monitor. This can normally be left blank in which case the port checked is the same port defined for the Real Servers. Note that for DR mode, the port cannot be specified at the Real Server level, so the port specified for the VIP is used. Sometimes the check port differs from service port.
Ping Server		<i>Test Real Server availability using an ICMP ping</i>
External Script		<i>Use a custom file (script or binary) for the health check</i>
	Check Port	(See above)

Option	Sub Option	health check Description
	External Script	<p>Select the required external check script from the drop-down. For more information on creating and using custom external health check scripts please refer to <a href="#">External Health Check Scripts</a>.</p> <div> <p><b>Note</b></p> <p>By default the Microsoft SQL external health check is not available in the drop down. This health check requires several Microsoft related per-requisites such as the Microsoft Linux ODBC driver, and these must first be installed and configured. To install the prerequisites and configure required settings, at the console or via an SSH session, login as root and run the following command:</p> <pre>\$ lb_mssql -i</pre> <p>Once completed, the additional option "ms-sql-check" will appear in the External Script drop-down.</p> </div> <div> <p><b>Note</b></p> <p>'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: <i>Local Configuration &gt; Security</i>. You'll need to Set <i>Appliance Security Mode</i> to <b>Custom</b>, configure the required option(s) and click <b>Update</b>.</p> </div>
<b>No Checks Always Off</b>		<i>No checking will take place and no real or fallback servers will be activated</i>
<b>No Checks Always On</b>		<i>No checking will take place and Real Servers will always be assumed to be up</i>
<b>5 Connects, 1 Negotiate</b>		<i>Repeating pattern of 5 Connect checks followed by 1 Negotiate check</i>
	Check Port	(See above)
	Protocol	(See above)
	Virtual Host	(See above)
	Request to Send	(See above)
	Response Expected	(See above)
<b>10 Connects, 1 Negotiate</b>		<i>Repeating pattern of 10 Connect checks followed by 1 Negotiate check</i>
	Check Port	(See above)
	Protocol	(See above)
	Virtual Host	(See above)
	Request to Send	(See above)
	Response Expected	(See above)

## Health Checks for Layer 7 Services

At layer 7, Real Server health checking is performed by HAProxy. To configure health checks, use the WebUI menu option: *Cluster Configuration > Layer 7 - Virtual Services*, then click **Modify** next to the VIP to be configured. The health check options available depend on the check type selected and whether the **[Advanced]** option is clicked. When clicked, the advanced options for each check type are displayed, when clicked again, they are hidden.

The screenshot shows the 'Health Checks' configuration page. At the top, there's a 'Health Checks' header with a '[Advanced +]' toggle. Below it, a dropdown menu is open, showing various check types. The 'Connect to port' option is currently selected and highlighted in blue. Other options include 'Negotiate HTTP (GET)', 'Negotiate HTTP (HEAD)', 'Negotiate HTTPS (GET)', 'Negotiate HTTPS (HEAD)', 'Negotiate HTTP (OPTIONS)', 'Negotiate HTTPS (OPTIONS)', 'External script', 'MySQL', and 'No checks, always On'. To the right of the dropdown, there's a question mark icon. Below the dropdown, there's an 'Add Rule' button and another question mark icon. The background shows parts of the 'ACL Rules' table with columns 'Type', 'Bool', and 'URL/Text', and the 'Header Rules' section.

**Note** | For new Layer 7 VIPs the default check type is *Connect to Port*.

The following Check Types are supported:

**Negotiate HTTP/HTTPS (GET)** - Scan the page specified in *Request to Send*, and check the returned data for the *Response Expected* string.

**Negotiate HTTP/HTTPS (HEAD)** - Request the page headers of the page specified in *Request to Send*.

**Negotiate HTTP/HTTPS (OPTIONS)** - Request the options of the page specified in *Request to Send*.

**Note** | If a Negotiate (Get) check is selected and *Response Expected* is left blank, the appliance will check the location specified in Request To Send (if blank the root will be checked) and will consider all HTTP 2xx (usually HTTP 200) and HTTP 3xx response statuses as valid and the server will be marked as up. All other responses including no response or a timeout will be considered invalid and the server will be marked as down.

**Connect to Port** - Attempt to make a connection to the specified port.

**External script** - Use a custom file (typically a script but can also be a binary) for the health check. Select the script from the *Check Script* drop-down. For more information on using custom external health check scripts please refer to [External Health Check Scripts](#).

**MySQL** - The check consists of sending two MySQL packets, one Client Authentication packet, and one QUIT packet, to correctly close the MySQL session. It then parses the MySQL Handshake Initialisation packet and/or Error packet. It is a basic but useful test and does not produce error nor aborted connect on the server. However, it requires adding an authorisation in the MySQL table, like this:

```
USE mysql; INSERT INTO user (Host,User) values ('',''); FLUSH PRIVILEGES;
```

**No checks, Always On** - No health checks, all real servers are marked online.

**Note** | By default, a TCP connect health check is used for newly created layer 7 Virtual Services.

The following table describes the health check options associated with each check type:

Option	Sub Option	health check Description
Negotiate HTTP/HTTPS (GET)		<i>Scan the page specified in Request to Send, and check the returned data for the Response Expected string</i>
	Request to Send	Specify a specific location/file for the health check. Open the specified location and check for the <i>Response Expected</i> .
	Response Expected	<p>The content expected for a valid health check on the specified file. The <i>Response Expected</i> can be any valid regular expression statement.</p> <p>For example, if the Real Servers have a virtual directory called /customers, with a default page that contains the word 'welcome', <i>Request to Send</i> would be set to "customers" (without quotes) and <i>Response Expected</i> would be set to "Welcome" (without quotes). Provided that the load balancer can access the page and see the text 'Welcome', the health check would pass.</p> <p><b>Note</b>   If <i>Response Expected</i> is left blank, the appliance will check the location specified in Request To Send (if blank the root will be checked) and will consider all HTTP 2xx (usually HTTP 200) and HTTP 3xx response statuses as valid and the server will be marked as up. All other responses including no response or a timeout will be considered invalid and the server will be marked as down.</p> <p><b>Note</b>   It's possible to escape characters in the response expected. For example, if you wanted to look for "success" (including the quotes), specify \"success\" in <i>Response Expected</i>.</p>
	<b>Advanced</b> > Check Port	The port to monitor. This can normally be left blank in which case the port checked is the same port defined for the Real Servers. However, sometimes the check port differs from service port in which case it can be specified here. Also useful for multiport VIPs where the real server port field is left blank. In this case, the default checkport is the first in the list. This can be changed using this field if required.
	<b>Advanced</b> > Username	Specify a username if authentication is required.
	<b>Advanced</b> > Host Header	If the Real Server is configured to require a Host header, the value to be used in health checks may be set here.
	<b>Advanced</b> > Password	Specify a password if authentication is required.

Option	Sub Option	health check Description
Negotiate HTTP/HTTPS (HEAD)		<i>Request the page headers of the page specified in Request to Send</i>
	Request to Send	(see above)
	<b>Advanced</b> > Check Port	(see above)
	<b>Advanced</b> > Username	(see above)
	<b>Advanced</b> > Host Header	(see above)
	<b>Advanced</b> > Password	(see above)
Negotiate HTTP/HTTPS (OPTIONS)		<i>Request the options of the page specified in Request to Send</i>
	Request to Send	(see above)
	<b>Advanced</b> > Check Port	(see above)
	<b>Advanced</b> > Username	(see above)
	<b>Advanced</b> > Host Header	(see above)
	<b>Advanced</b> > Password	(see above)
Connect to port		<i>Attempt to make a connection to the specified port</i>
	<b>Advanced</b> > Check Port	(See above)
External Script		<i>Use a custom file (script or binary) for the health check</i>



Option	Sub Option	health check Description
	Check Script	<p>Select the required external check script from the drop-down. For more information on creating and using custom external health check scripts please refer to <a href="#">External Health Check Scripts</a>.</p> <div> <p><b>Note</b></p> <p>By default the Microsoft SQL external health check is not available in the drop down. This health check requires several Microsoft related per-requisites such as the Microsoft Linux ODBC driver, and these must first be installed and configured. To install the prerequisites and configure required settings, at the console or via an SSH session, login as root and run the following command:</p> <pre>\$ lb_mssql -i</pre> <p>Once completed, the additional option "ms-sql-check" will appear in the External Script drop-down.</p> <p><b>Note</b></p> <p>'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: <i>Local Configuration &gt; Security</i>. You'll need to Set <i>Appliance Security Mode</i> to <b>Custom</b>, configure the required option(s) and click <b>Update</b>.</p> </div>
MySQL		<i>Check MySQL</i>
	<b>Advanced</b> > Username	<p>perform a Client Authentication check, using &lt;username&gt; This requires an update into the MySql servers, as shown below, using MySQL client software:</p> <pre>USE mysql;</pre> <pre>INSERT INTO user (Host,User) values ('&lt;ip_of_haproxy&gt;','&lt;username&gt;');</pre> <pre>FLUSH PRIVILEGES;</pre> <div> <p><b>Note</b></p> <p>Without the user option, a MySql Handshake is performed</p> </div>
No checks, always On		<i>No checking will take place and Real Servers will always be assumed to be up</i>

## External Health Check Scripts

Custom health checks can be created and modified using the WebUI.


### Default Scripts

By default, 4 external scripts are available:

- SMTP

- Ping\_IPv4\_or\_IPv6
- POP3\_or\_IMAP
- Exchange

These are available in the *External Script* drop-down when the *Check Type* for a VIP is set to **External Script** as shown below:

Health Checks		
Check Type	External script ▼	?
Check Port	<input type="text"/>	?
External script	Ping_IPv4_or_IPv6 ▼ 	?
Feedback		
Feedback Method	SMTP Ping_IPv4_or_IPv6 POP3_or_IMAP Exchange	?
Fallback Server		

## Adding Additional Health Check Scripts

New scripts can be created either by using the script templates or by uploading files from an external source.

### Using Script Templates

*To Create a new Script From Template:*

1. Using the WebUI, navigate to *Cluster Configuration > Health Check Scripts* and click **Add New Health Check**.

Health Check Details		
Name:	Untitled	?
Type:	Virtual Service ▼	?
Template:	Multi-port-check.sh ▼	?
Primary Node Health Check Contents		


2. Specify an appropriate *Name* for the health check, e.g. **Multi-Port-Check**.
3. Set *Type* to **Virtual Service**.
4. Using the *Template* dropdown select an appropriate template from the **Virtual Service** section of the list, e.g. **Multi-port-check.sh**.
5. Modify the script if required.
6. Click **Update**.

Once the health check has been added, it will appear in the Health Check Scripts list below the 4 default scripts as shown below:

## Health Check Scripts

			<a href="#">Add New Health Check</a>
			<a href="#">Upload Existing Health Check</a>
Health Check Name	Type	In-use	
SMTP	VIP	-	<a href="#">Modify</a> <a href="#">Delete</a>
Ping_IPv4_or_IPv6	VIP	-	<a href="#">Modify</a> <a href="#">Delete</a>
POP3_or_IMAP	VIP	-	<a href="#">Modify</a> <a href="#">Delete</a>
Exchange	VIP	-	<a href="#">Modify</a> <a href="#">Delete</a>
Multi_Port_Check	VIP	-	<a href="#">Modify</a> <a href="#">Delete</a>

The new script will also appear in the *External Script* dropdown when the *Check Type* for a VIP is set to **External Script**:

Health Checks		
Check Type	External script ▾	?
Check Port	<input type="text"/>	?
External script	SMTP ▾ 	?
<b>Feedback</b>		
Feedback Method		?
<b>Fallback Server</b>		

SMTP

Ping\_IPv4\_or\_IPv6

POP3\_or\_IMAP

Exchange

Multi\_Port\_Check

### Uploading External Files

To Create a new Script by Uploading an External File:

1. Using the WebUI, navigate to *Cluster Configuration > Health Check Scripts* and click **Upload Existing Health Check**.

Health check Details		
Name:	<input type="text" value="Check-Control-Servers"/>	<a href="#">?</a>
Type:	<input checked="" type="radio"/> Virtual Service <input type="radio"/> GSLB	<a href="#">?</a>
Contents:	<input type="button" value="Browse..."/> control-server-check.sh	<a href="#">?</a>
Secondary node contents:	<input type="button" value="Browse..."/> No file selected.	<a href="#">?</a>
File is binary:	<input type="checkbox"/>	<a href="#">?</a>

- Specify an appropriate *Name* for the health check, e.g. **Check-Control-Servers**.
- Set *Type* to **Virtual Service**.
- Click the **Choose File** button next to *Contents*.
- Browse to and select the required file, e.g. **control-server-check.sh**.

#### Note

If you have an HA Pair and the secondary node requires a different health check, click the **Choose File** button next to *Secondary Node Contents* and browse to and select the required file.

- If the file is binary, enable the **File is Binary** checkbox - this will prevent the editor window being displayed.
- Click **Update**.

Once the health check has been added, it will appear in the Health Check Scripts list and in the *External Script* dropdown as explained in [Using Script Templates](#) above.

#### Note

For additional Layer 4 health check options such as *Check interval* and *Failure Count* please refer to [Layer 4 - Advanced Configuration](#). For Layer 7, please refer to [Layer 7 - Advanced Configuration](#).

## Testing External Health Check Scripts at the Command Line

Health check scripts use 4 passed parameters. These 4 values represent *Virtual Service IP Address*, *Virtual Service Port*, *Real Server IP Address* and *Real Server Port*. If a script does not use all 4 values, for example the ping.sh script, then a zero should be entered as a place-holder.

#### Note

The **Skeleton** and **Example** script templates provide additional information about the 4 passed parameters.

```
# ./<check-script-name> <$1> <$2> <$3> <$4>
```

Examples:

```
# ./SMTP-check.sh 192.168.1.1 25 192.168.1.10 25
```

```
# ./ping.sh 192.168.1.1 0 192.168.1.10 0
```

to check the return value, use the command:

```
# echo $?
```

A return value of 0 means the check has passed, any other value means it has failed.

## Simulating Health Check Failures

It may not always be possible to take a server offline to check that health checks are working correctly. In these cases, firewall rules can be used. The following rules can be configured at the console, using SSH or via the WebUI option *Local Configuration > Execute a Shell Command*.

### Note

'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

**To block access to a particular Real Server port:**

```
iptables -A OUTPUT -p tcp --dport <Check Port> -d <REAL-SERVER-IP> -j DROP
```

e.g. `iptables -A OUTPUT -p tcp --dport 80 -d 192.168.65.60 -j DROP`

**To re-enable access to a particular Real Server port:**

```
iptables -D OUTPUT -p tcp --dport <Check Port> -d <REAL-SERVER-IP> -j DROP
```

e.g. `iptables -D OUTPUT -p tcp --dport 80 -d 192.168.65.60 -j DROP`

### Note

Make sure these rules are cleared after testing & verification is complete!

## Disabling Health Checks

In some cases it may be desirable to completely disable health checking and simply assume that the Real Servers are up and working correctly. This can be configured by setting the health check option to **No Checks, Always On** - this applies to both layer 4 and layer 7 services.

## Fallback Server

The fallback server is activated under the following conditions for both Layer 4 & Layer 7 Virtual Services:

- When all associated Real Servers have failed their health check
- When all associated Real Servers have been taken offline via the WebUI

The fallback page can be provided in the following ways:

- Using the load balancer's built in NGINX fallback page
- Using a separate server to host the fallback page
- Using a Layer 7 VIP

## Local Fallback Server

The appliance has a built in fallback server that uses NGINX. The local fallback page can be modified using the WebUI menu option: *Maintenance > Fallback Page*:

### Fallback Page

```

1  <html>
2  <head>
3  <title>The page is temporarily unavailable</title>
4  <style>
5  body { font-family: Tahoma, Verdana, Arial, sans-serif; }
6  </style>
7  </head>
8  <body bgcolor="white" text="black">
9  <table width="100%" height="100%">
10 <tr>
11 <td align="center" valign="middle">
12 The page you are looking for is temporarily unavailable.<br/>
13 Please try again later.<br/>
14 (WUI port reminder 9080)
15 </td>
16 </tr>
17 </table>
18 </body>
19 </html>
20
21
22
23
24
```

### Notes

1. The local fallback server is an NGINX instance that by default listens on port 9081.
2. If a layer 4 VIP is added that listens on port 80, NGINX is automatically configured to listen on ports 9081 & 80.
3. You can use any valid HTML for the default page, simply copy and paste the required HTML into the Fallback Page.
4. If you are using the load balancer for your holding page and your Real Servers are all offline then the local NGINX server is exposed to hacking attempts, if you are concerned about this you can change the fallback server to be one of your internal servers.

## Using a Separate Dedicated Server

For DR mode the fallback server must be listening on the same port as the VIP (port re-mapping is not possible with DR mode). Also, don't forget to solve the **ARP Problem** for the dedicated fallback server. For more information please refer to [The ARP Problem](#).

For NAT mode don't forget to set the default gateway of the fallback server to the internal IP of the load balancer or when you have 2 appliances in a cluster, to a floating IP.

## Using a Layer 7 VIP

It's possible to set the fallback server to be a layer 7 VIP. This is especially useful in WAN/DR site environments. It also enables an external fallback server to be easily configured for Layer 4 VIPs without having to comply with the requirements mentioned in the previous section. To do this, create a layer 7 fallback VIP and configure your fallback server as an associated RIP. Then enable the MASQ option for the Layer 4 VIP and set the fallback VIP as

its fallback server. If all servers are down, requests will then be routed via the Layer 7 VIP to your fallback server. If the layer 4 VIP is multi-port, specify 0 as the port for the fallback server. Requests will then be forwarded to the correct port.

## Configuring A Real Server as the Fallback Server

It's possible to configure one of the Real Servers as the fallback server. This can be useful for example when all servers are very busy and health checks start to fail simply because the response is taking longer than the configuration allows. In this case, traffic will still be sent to one of the Real Servers rather than to a separate fallback page.

## Configuring Primary/Secondary Real Servers

If you want to setup a VIP that sends all traffic to a primary server and only sends traffic to a secondary server if the primary server fails, configure the VIP with the primary server as a RIP, and the secondary server as the fallback server.

## Configuring Email Alerts for Virtual Services

Email alerts can be configured for layer 4 and layer 7 Virtual Services. This enables emails to be sent when one or more of the associated Real Servers fail their health check and also when they subsequently start to pass their health check.

### Layer 4







For layer 4 Virtual Services, settings can be configured globally for all VIPs or individually per VIP.

#### Global Layer 4 Email Settings

Once configured, these settings apply to all layer 4 VIPs by default.

*To configure global email alert settings for layer 4 services:*

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 Advanced Configuration*.

Email Alert Source Address	<input type="text" value="lb1@loadbalancer.org"/>	
Email Alert Destination Address	<input type="text" value="alerts@loadbalancer.org"/>	
Auto-NAT	<input type="text" value="off"/> 	
Multi-threaded	<input type="text" value="yes"/> 	
		<input type="button" value="Update"/>

2. Enter an appropriate email address in the *Email Alert Source Address* field.

e.g. lb1@loadbalancer.org

3. Enter an appropriate email address in the *Email Alert Destination Address* field.

e.g. alerts@loadbalancer.org

4. Click **Update**.

## VIP Level Settings

**Note** | VIP level settings override the global settings.

Once configured, these settings apply to the individual VIP.

*To configure VIP level email alerts:*

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 Virtual Service* and click **Modify** next to the VIP to be configured.
2. Scroll down to the *Fallback Server* section.

Email Alert Destination Address

alerts@loadbalancer.org



Cancel

Update

3. Enter an appropriate email address in the *Email Alert Destination Address* field.

e.g. alerts@loadbalancer.org

4. Click **Update**.

**Note** | You can set the *Email Alert Source Address* field as explained above if required to configure a default source address.

## Configuring a Smart Host (SMTP relay)

For layer 4 services (and Heartbeat), email alerts are sent to the mail server defined in the destination domain's DNS MX record directly from the load balancer by default. However, configuring a smart host (mail relay server) will instead direct all outgoing mail to the specified relay. This will significantly help to improve the deliverability of the email. A smart host is an email server through which approved devices can send emails. Where possible, we recommend that you use a smart host to send email alerts.

*To configure a Smart Host:*

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration*.
2. Scroll down to the *SMTP Relay* section.
3. Set the smart host using an FQDN or IP address.
4. Click **Update**.

**Note** | Layer 7 services do not use the smart host configured here.





## Layer 7

For layer 7 services, email settings are configured globally for all VIPs.

*To configure global email alert settings for layer 7 services:*



1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 Advanced Configuration*.

eMail Alert From	<input type="text" value="lb1@loadbalancer.org"/>	
eMail Alert To	<input type="text" value="alerts@loadbalancer.org"/>	
eMail Server Address	<input type="text" value="mail.loadbalancer.org"/>	
eMail Server Port	<input type="text" value="25"/>	

2. Enter an appropriate email address in the *eMail Alert From* field.

e.g. lb1@loadbalancer.org

3. Enter an appropriate email address in the *eMail Alert To* field.

e.g. alerts@loadbalancer.org

4. Enter an appropriate IP address/FQDN in the *eMail Server Address* field.

e.g. mail.loadbalancer.org

5. Enter an appropriate port in the *eMail Server Port* field.


e.g. 25

6. Click **Update**.





## Real Server Monitoring & Control using the System Overview

### Real Server Monitoring

The System Overview includes a visual display indicating the health status of all Virtual and Real Servers as shown in the example below:

**System Overview** 

2023-01-31 17:57:01 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy	
	RDS-Cluster	192.168.110.152	3389	0	TCP	Layer 7	Proxy	

Clicking on each Virtual Service expands the view so that the associated Real Servers can also be seen:

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	IIS1	192.168.110.240	80	100	0	Drain	Halt	
	IIS2	192.168.110.241	80	100	0	Drain	Halt	
	RDS-Cluster	192.168.110.152	3389	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	RDS1	192.168.110.240	3389	100	0	Drain	Halt	
	RDS2	192.168.110.241	3389	100	0	Drain	Halt	

The various colors used to indicate status are:

- **Green** - All Real Servers in the cluster are healthy
- **Yellow** - One or more Real Servers in the cluster has failed or has been taken offline using *Halt* or *Drain*
- **Red** - All Real Servers in the cluster have failed
- **Blue** - All Real Servers have been taken offline using *Drain* or *Halt* (see below)
- **Purple/Green** - Used to indicate that a particular VIP is used for HTTP to HTTPS redirection

This information is also displayed when clicking the system overview help button:

System Overview

The following colors and icons are used to show the real-time status of your Load balanced Virtual Services

Colour	Image	Details
Green		Virtual Service / Real Server healthy
Yellow		Virtual Service needs attention
Blue		Real Server taken offline
Red		Virtual Service / Real Server down
Purple		Virtual Service FORCE-TO-HTTPS

The Virtual Services may be sorted using drag and drop, or by clicking on the column headings.






## Real Server Control

The System Overview also enables each Real Server to be taken offline. This can be achieved in 2 ways:

1. **Drain** - This option allows all existing connections to complete & close gracefully. It also prevents any new connections being established.

2. **Halt** - This option drops all existing connections immediately without waiting. It also prevents any new connections being established.

The screen shot below shows that RDS2 has been placed in drain mode:

	<b>RDS-Cluster</b>	<b>192.168.110.152</b>	<b>3389</b>	<b>0</b>	<b>TCP</b>	<b>Layer 7</b>	<b>Proxy</b>	
	<i>REAL SERVER</i>	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
	<b>RDS1</b>	<b>192.168.110.240</b>	<b>3389</b>	<b>100</b>	<b>0</b>	<b>Drain</b>	<b>Halt</b>	
<b>None</b>	<b>RDS2</b>	<b>192.168.110.241</b>	<b>3389</b>	<b>0</b>	<b>0</b>	<b>Online (drain)</b>	<b>Halt</b>	

To bring RDS2 back online, click the *Online (drain)* link. If the server had been halted rather than drained, then the link would be displayed as *Online (Halt)*.

#### Note

If a particular Real Server is used in multiple VIPs you'll be asked if you'd like to apply the offline/online action to all relevant VIPs or only a single VIP. This simplifies taking Real Servers offline for maintenance purposes.

#### Note

Halting or draining all Real Servers in a cluster activates the fallback server.

## Ordering of VIPs



The display order of configured VIPs can be changed either by clicking on the column heading, or by drag and drop.

### Sort by Column

If VIPs are ordered by a particular column, this is indicated using arrows next to the column heading as shown below:

## System Overview ?

2023-01-31 18:00:40 UTC

VIRTUAL SERVICE ⚙		IP ⚙	PORTS ⚙	CONNS ⚙	PROTOCOL ⚙	METHOD ⚙	MODE ⚙		
📈	⬆️	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy	
📉	⬆️	RDS-Cluster	192.168.110.152	3389	0	TCP	Layer 7	Proxy	

In this example, the VIPs are ordered alpha-numerically by Virtual Service name. To change the order, click on the required column heading then click save. If you want to reverse the order for a particular column, click that column heading again. For example, clicking on the IP column heading shows the following:

### EDIT MODE

Cancel

Save

VIRTUAL SERVICE ⚡		IP ⚡	PORTS ⚡	CONNS ⚡	PROTOCOL ⚡	METHOD ⚡	MODE ⚡
⚠	RDS-Cluster	192.168.110.152	3389	0	TCP	Layer 7	Proxy
⬆	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy

Clicking on the IP column heading again changes the order to:

## EDIT MODE

[Cancel](#)[Save](#)

	VIRTUAL SERVICE ⌵	IP ⌵	PORTS ⌵	CONNS ⌵	PROTOCOL ⌵	METHOD ⌵	MODE ⌵
↑	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy
⚠	RDS-Cluster	192.168.110.152	3389	0	TCP	Layer 7	Proxy

Once you've set the required order, click **Save**.

## Drag & Drop

To re-order VIPs by drag and drop, simply click the mouse on any part of the VIP:

## EDIT MODE

[Cancel](#)[Save](#)

	VIRTUAL SERVICE ⌵	IP ⌵	PORTS ⌵	CONNS ⌵	PROTOCOL ⌵	METHOD ⌵	MODE ⌵
⚠	RDS-Cluster ✎	192.168.110.152	3389	0	TCP	Layer 7	Proxy
↑	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy

Then drag it to the required position:

## EDIT MODE

[Cancel](#)[Save](#)

	VIRTUAL SERVICE ⌵	IP ⌵	PORTS ⌵	CONNS ⌵	PROTOCOL ⌵	METHOD ⌵	MODE ⌵
↑	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy
⚠	RDS-Cluster ✎	192.168.110.152	3389	0	TCP	Layer 7	Proxy

And release it. Once you've set the required order, click **Save**.

## Real Server Monitoring & Control using the HAProxy Statistics Report

### Real Server Monitoring

The layer 7 Statistics Report shows the status of all Virtual Servers and Real Servers as shown in the example below:

## HAProxy

### Statistics Report for pid 8570

#### > General process information

pid = 8570 (process #1, nbproc = 1, nbthread = 1)  
uptime = 0d 0h00m04s  
system limits: memmax = unlimited; ulimit-n = 80049  
maxsock = 80049; maxconn = 40000; maxpipes = 0  
current conns = 1; current pipes = 0/0; conn rate = 2/sec  
Running tasks: 1/26; idle = 100 %

active UP  
active UP, going down  
active DOWN, going up  
active or backup DOWN  
active or backup DOWN for maintenance (MAINT)  
active or backup SOFT STOPPED for maintenance  
Note: "NOLB"/"DRAIN" = UP with load-balancing disabled.

backup UP  
backup UP, going down  
backup DOWN, going up  
not checked

Display option:

- Scope :
- Hide 'DOWN' servers
- Refresh now
- CSV export

External resources:

- Primary site
- Updates (v1.8)
- Online manual

Web-Cluster		Queue			Session rate			Sessions					Bytes		Denied		Errors		Warnings		Server									
Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle	
Frontend			0	0	-	0	0	0	40 000	0		0	0	0	0	0	0	0	0	0	OPEN									
backup	0	0	-	0	0	0	0	-	0	0	?	0	0	0	0	0	0	0	0	no check			1	-	Y					-
Web1	0	0	-	0	0	0	0	-	0	0	?	0	0	0	0	0	0	0	0	4s UP	L4OK in 0ms	100	Y	-	0	0	0	0s	-	
Web2	0	0	-	0	0	0	0	-	0	0	?	0	0	0	0	0	0	0	0	4s UP	L4OK in 0ms	100	Y	-	0	0	0	0s	-	
Web3	0	0	-	0	0	0	0	-	0	0	?	0	0	0	0	0	0	0	0	4s UP	L4OK in 0ms	100	Y	-	0	0	0	0s	-	
Backend	0	0	0	0	0	0	0	4 000	0	0	?	0	0	0	0	0	0	0	0	4s UP		300	3	1		0	0s			

ADFS-Cluster		Queue		Session rate		Sessions								Bytes		Denied		Errors		Warnings		Server									
						Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme
Frontend				0	0	-	0	0	40 000	0		0	?	0	0	0	0	0	0	0	0	0	OPEN								
backup		0	0	-	0	0		0	-	0	0	?	0	0	0	0	0	0	0	0	0	no check		1	-	Y					-
ADFS1		0	0	-	0	0		0	-	0	0	?	0	0	0	0	0	0	0	0	0	4s UP	L7OK/200 in 0ms	100	Y	-	0	0	0s	-	
ADFS2		0	0	-	0	0		0	-	0	0	?	0	0	0	0	0	0	0	0	0	4s UP	L7OK/200 in 0ms	100	Y	-	0	0	0s	-	
Backend		0	0		0	0		0	4 000	0	0	?	0	0	0	0	0	0	0	0	0	4s UP		200	2	1		0	0s		

Exchange-HTTPS																															
	Queue			Session rate			Sessions						Bytes		Denied		Errors		Warnings		Server										
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle	
Frontend				0	0	-	0	0	0	40 000	0		0	0	0	0	0	0	0	0	0	OPEN									
backup	0	0	-	0	0		0	0	-	-	0	?	0	0	0	0	0	0	0	0	0	no check		1	-	Y					-
Exch1	0	0	-	0	0		0	0	-	0	0	?	0	0	0	0	0	0	0	0	0	4s UP	L4OK in 0ms	100	Y	-	0	0	0s	-	
Exch2	0	0	-	0	0		0	0	-	0	0	?	0	0	0	0	0	0	0	0	0	4s UP	L4OK in 0ms	100	Y	-	0	0	0s	-	
Backend	0	0		0	0		0	0	0	4 000	0	?	0	0	0	0	0	0	0	0	0	4s UP		200	2	1		0	0s		

stats																																
	Queue			Session rate			Sessions				Bytes		Denied		Errors			Warnings		Server												
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntme	Thrtle		
Frontend				2	2	-	1	2	2	2 000	2				488	262	0	0	0			OPEN										
Backend	0	0		0	0		0	0	0	200	0	0s	488	262	0	0	0	0	0	0	0	4s UP		0	0	0	0		0			

To access the page, navigate to: *Reports > Layer 7 Status* - a new tabbed window will be opened.

Note | From v8.8.1, all statistics are automatically refreshed every 30 seconds.

## Real Server Control

It's also possible to control layer 7 Real Servers using the HAProxy statistics Report. By default this is not enabled.

To enable Real Server Control:

- Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Advanced*.
- Enable (check) the **Advanced Stats** option.
- Click **Update**.
- Reload HAProxy using the button at the top of screen. With this setting, the HAProxy stats page has the ability to control the state of Real Servers as shown below:

# HAProxy

## Statistics Report for pid 7354

### > General process information

pid = 7354 (process #1, nbproc = 1, nbthread = 1)  
 uptime = 0d 0h00m29s  
 system limits: memmax = unlimited; ulimit-n = 80049  
 maxsock = 60049; maxconn = 40000; maxpipes = 0  
 current conns = 1; current pipes = 0/0; conn rate = 5/sec  
 Running tasks: 1/26; idle = 100 %

active UP  
 active UP, going down  
 active DOWN, going up  
 active or backup DOWN  
 active or backup DOWN for maintenance (MAINT)  
 active or backup SOFT STOPPED for maintenance  
 backup UP  
 backup UP, going down  
 backup DOWN, going up  
 not checked  
 Note: "NO LB" "DRAIN" = UP with load-balancing disabled.

Display option:

- Scope:
- [Hide DOWN servers](#)
- [Refresh now](#)
- [CSV export](#)

External resources:

- [Primary site](#)
- [Updates \(v1.8\)](#)
- [Online manual](#)

Web-Cluster																																
	Queue			Session rate			Sessions					Bytes		Denied		Errors		Warnings		Status			Server									
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp									Retr	Redis	LastChk	Wght	Act
<input type="checkbox"/>	Frontend	0	0	-	0	0	-	0	0	0	40 000	0	?	0	0	0	0	0	0	0	0	OPEN										
<input type="checkbox"/>	backup	0	0	-	0	0	-	0	0	0	-	0	?	0	0	0	0	0	0	0	no check		1	-	Y						-	
<input type="checkbox"/>	Web1	0	0	-	0	0	-	0	0	0	-	0	?	0	0	0	0	0	0	0	29s UP	L4OK in 0ms	100	Y	-	0	0	0s		-		
<input type="checkbox"/>	Web2	0	0	-	0	0	-	0	0	0	-	0	?	0	0	0	0	0	0	0	29s UP	L4OK in 0ms	100	Y	-	0	0	0s		-		
<input type="checkbox"/>	Web3	0	0	-	0	0	-	0	0	0	-	0	?	0	0	0	0	0	0	0	29s UP	L4OK in 0ms	100	Y	-	0	0	0s		-		
<input type="checkbox"/>	Backend	0	0	-	0	0	-	0	0	0	4 000	0	?	0	0	0	0	0	0	0	29s UP		300	3	1		0	0s				

Choose the action to perform on the checked servers:

Apply

ADFS-Cluster										Set slate to READY										Set slate to DRAIN										Set slate to MAINT										Server																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							
		Queue			Session rate																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																										

Choose the action to perform on the checked servers:

Apply

Exchange-HTTPS		Agent force UP Agent force DOWN Kill Sessions																														
		Queue			Session rate			Total	LbTot	Bytes			Denied			Req	Errors Conn	Resp	Warnings		Status	LastChk	Server									
		Cur	Max	Limit	Cur	Max	Limit			Last	In	Out	Req	Resp	Retr				Redis	Wght			Act	Bck	Chk	Dwn	Downtime	Thrtht				
<input type="checkbox"/>	Frontend	0	0	-	0	0	-	0	0	40 000	0	0	?	0	0	0	0	0	0	0	OPEN		1	-	Y							
<input type="checkbox"/>	backup	0	0	-	0	0	-	0	0	-	0	0	?	0	0	0	0	0	0	0	no check											
<input type="checkbox"/>	Exch1	0	0	-	0	0	-	0	0	0	-	0	?	0	0	0	0	0	0	0	L4OK in 0ms	L4OK in 0ms	100	Y	-	0	0	0s	-			
<input type="checkbox"/>	Exch2	0	0	-	0	0	-	0	0	0	-	0	?	0	0	0	0	0	0	0	L4OK in 0ms	L4OK in 0ms	100	Y	-	0	0	0s	-			
<input type="checkbox"/>	Backend	0	0	-	0	0	-	0	0	4 000	0	0	?	0	0	0	0	0	0	0	29s UP		200	2	1		0	0s				

Choose the action to perform on the checked servers:

Apply

- Use the checkboxes to select the relevant Real Server(s), then select the required action in the drop-down, then click **Apply**.

# Chapter 9 - Appliance Clustering for HA

## Introduction

Loadbalancer.org appliances can be deployed as single unit or as a clustered pair.

**Note** | We always recommend deploying a clustered pair to avoid introducing a single point of failure.

## Clustered Pair Concepts

When configured as a clustered pair, the appliances work in **Active-Passive** mode. In this mode the active unit (normally the Primary) handles all traffic under normal circumstances. If the active unit fails, the passive unit (normally the Secondary) becomes active and handles all traffic.

Whilst HA clustering can be configured at any time, we generally recommend that the Primary appliance should be fully configured first and then the Secondary appliance should be added to create the clustered pair. Once this is done, all load balanced services configured on the Primary are automatically replicated to the Secondary. Pairing must be performed on the unit that is to be the Primary appliance.

**Note** | For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

## Primary/Secondary Operation

### Pair Communication

To allow the pair to function correctly, the Primary and Secondary must be able to:

- Perform an ICMP echo request (ping) to each other
- Communicate with each other on TCP port 22
- Communicate with each other on UDP port 6694 (or the selected custom port if this has been changed)

### Heartbeat

By default, heartbeat uses ucast over UDP port 6694 to communicate between the Primary and Secondary appliances. The link enables the state of each to be monitored by the other and permits a failover to the passive unit if the active unit should fail. For hardware appliances, it's possible to configure both ucast and serial communication if required.

**Note** | For hardware appliances, if the load balancer pair is located in close proximity, enabling serial communication in addition to ucast is recommended. Once the serial cable is connected between the appliances, serial comms can be enabled using the WebUI menu option: *Cluster Configuration > Heartbeat Configuration*. When serial communication is disabled, console access via the serial port is activated.

Ping checks to a common node such as the default gateway can also be configured. If the active node loses access to the ping node, the system will fail-over to the peer. However, if both nodes lose access, no fail-over will occur.

## Primary Secondary Replication

Once the Primary and Secondary are paired, all load balanced services are automatically replicated from Primary to

Secondary. This ensures that should the Primary unit fail, the Secondary is already configured to run the same services. Replication of the configured load balanced services from the Primary to the Secondary appliance occurs over the network using SSH/SCP.

### Settings that are NOT Replicated to the Secondary Appliance

Local appliance settings such as hostname and DNS, IP address configuration and time/date settings are not replicated and must be manually configured on each appliance. The table below lists all settings by WebUI menu option that are not replicated from the Primary appliance to the Secondary appliance when an HA pair is configured.

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	All network settings including IP address(es), bonding configuration and VLANs
Local Configuration	Routing	Routing configuration including default gateways and static routes
Local Configuration	System Date & time	All time and date related settings
Local Configuration	Physical – Advanced Configuration	Various settings including Internet Proxy, Management Gateway, Firewall connection tracking table size, NIC offloading, SMTP relay, logging and Syslog Server
Local Configuration	Security	Appliance security settings
Local Configuration	SNMP Configuration	Appliance SNMP settings
Local Configuration	Graphing	Appliance graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Software Updates	Appliance software update management
Maintenance	Firewall Script	Appliance firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

### Manually Forcing Appliance Synchronisation

If changes are made to the Primary appliance of an HA Pair whilst the Secondary appliance is down, the HA Pair must be manually re-synchronised once the secondary is back up.

*To Force Primary to Secondary Synchronisation:*

1. Using the WebUI, navigate to: *Maintenance > Backup & Restore*.
2. Select the Synchronisation tab.
3. Click **Synchronize Configuration with peer**.

### To Create an HA Pair (Add a Secondary)


#### Note

If you have already run the firewall lockdown wizard on either appliance, you'll need to ensure that it is temporarily disabled on both appliances whilst performing the pairing process.



1. Power up a second appliance that will be the Secondary and configure initial network settings - for more details on initial deployment and network setup, please refer to [Chapter 4 - Appliance Fundamentals](#).
2. Using the WebUI of the Primary appliance, navigate to: *Cluster Configuration > High-Availability Configuration*.

Create a Clustered Pair

 **LOADBALANCER**

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41


Password for *loadbalancer* user on peer

••••••••••


Add new node

3. Specify the IP address and the *loadbalancer* user's password (the default is 'loadbalancer') for the Secondary (peer) appliance as shown above.
4. Click **Add new node**.
5. A warning will be displayed indicating that the pairing process will overwrite the new Secondary appliance's existing configuration, click **OK** to continue.
6. The pairing process now commences as shown below:

Create a Clustered Pair

 **LOADBALANCER** 192.168.110.40 Primary

Attempting to pair..

 **LOADBALANCER** 192.168.110.41 Secondary

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41

Password for *loadbalancer* user on peer

••••••••••

7. Once complete, the following will be displayed under: *Cluster Configuration > High Availability Configuration* on the Primary appliance:

### High Availability Configuration - primary


 **LOADBALANCER** 192.168.1.40 Primary


 **LOADBALANCER** 192.168.1.41 Secondary

**Break Clustered Pair**

The following will be displayed under: *Cluster Configuration > High Availabuiltiy Configuration* on the Secondary appliance:

### High Availability Configuration - secondary

 **LOADBALANCER** 192.168.1.41 Secondary

 **LOADBALANCER** 192.168.1.40 Primary

**Break Clustered Pair**

**Make Active**

**Promote to Primary**

- Now restart heartbeat as prompted in the "commit changes" message box at the top of the screen.

#### Note

Clicking the **Restart Heartbeat** button on the Primary appliance will also automatically restart heartbeat on the Secondary appliance.

### To Break an HA Pair (Remove a Secondary)

- Using the WebUI of the Primary or Secondary appliance, navigate to: *Cluster Configuration > High-Availability Configuration* (shows the Primary appliance).

### High Availability Configuration - primary


 **LOADBALANCER** 192.168.1.40 Primary


 **LOADBALANCER** 192.168.1.41 Secondary

**Break Clustered Pair**

- To break the pair, click the **Break Clustered Pair** button.
- Click **OK** to confirm you want to proceed.

### High Availability Configuration - primary


 **LOADBALANCER** 192.168.1.40 Primary

 **LOADBALANCER** 192.168.1.41 Secondary

Attempting to break..

- Once the process is complete, the pairing configuration screen will be displayed on both appliances:

### Create a Clustered Pair



**Local IP address**

**IP address of new peer**

**Password for loadbalancer user on peer**

**Add new node**

5. To complete the reconfiguration, restart the system services on both appliances as directed in the "commit changes" message box.

### Notes

1. Load balanced services will be momentarily interrupted as system services are restarted.
2. After the pair is broken, the Secondary appliance will be left configured as a Secondary and any configured load balanced services will remain.
3. If you later want to use the Secondary as a Primary, use the *Cluster Configuration > High Availability Configuration* menu option on the Secondary to setup a new pair. The Secondary will then be re-configured as a Primary, and the added peer will be configured as a Secondary.
4. If you want to reset the configuration to factory defaults, use the WebUI menu option: *Maintenance > Backup & Restore > Restore > Restore Manufacturer's Defaults*.


## Promoting a Secondary to Primary


This is useful if the Primary unit fails and you'd like to change the now active Secondary to be a Primary, and then add the repaired/replaced Primary back as a Secondary unit.

To promote a Secondary unit to become a Primary:

1. Using the WebUI of the Secondary appliance, navigate to: *Cluster Configuration > High-Availability Configuration*.

### High Availability Configuration - secondary

 **Secondary**

 **Primary**

**Break Clustered Pair**


**Make Active**

**Promote to Primary**

2. Click **Promote**.

3. Click **OK** to confirm you want to proceed.

High Availability Configuration - secondary



**Note** | If the Primary is still up and operational, it will not be possible to promote the Secondary.

4. Once complete, the unit will be configured as a Primary.

**Note** | Please refer to [Chapter 14 - Backup & Restore and Disaster Recovery](#) for more details on how to recover from various appliance failure scenarios.

Configuring Heartbeat

The default heartbeat settings are appropriate in most cases, but can be changed if required.

To configure Heartbeat:

1. Using the WebUI of the Primary appliance, navigate to: *Cluster Configuration > Heartbeat Configuration*.

**Note** | The screen shot below shows the configuration screen for a hardware appliance. The virtual appliance does not have the serial option checkbox in the communications method section.

## Heartbeat Configuration

<b>Communication method</b>		
Serial	<input type="checkbox"/>	?
UDP Unicast	<input checked="" type="checkbox"/>	?
UDP Broadcast ( <i>Deprecated</i> )	Off ▾	?
UDP Port for broadcast & unicast	6694	?
<b>Peer Failure Detection</b>		
Keep-alive message interval	3 seconds	?
Dead peer timer	10 seconds	?
Warning timer	5 seconds	?
<b>Routing Failure Detection</b>		
Test IP addresses		?
Test time-out	10 seconds	?
<b>Email Alerts</b>		
Email Alert Destination Address		?
Email Alert Source Address		?
<b>Automatic Fail-back</b>	<input type="checkbox"/>	?

Modify Heartbeat configuration

### Communication Method

**Serial** - Enable or disable heartbeat Primary/Secondary communication over the serial port (hardware appliances only). Ucast is the default heartbeat communication method. However, if the load balancer pair is located in close proximity, enabling serial communication in addition to ucast is recommended. This method requires a null modem cable (supplied with the appliance) to be connected to the serial port of each appliance.

**Note** | When serial communication is disabled, console access via the serial port is activated.

**UDP Unicast** - Enable or disable unicast heartbeat Primary/Secondary communication. This is the default method of heartbeat communication and uses unicast UDP between Primary and Secondary, with a destination port specified by the *UDP Port for broadcast & unicast* parameter. When unicast is enabled, the load balancer determines the correct interface and IP addresses to use based upon the configured Secondary IP address.

**UDP Broadcast (Deprecated)** - Enable or disable broadcast heartbeat Primary/Secondary communication, and choose the interface. This option is deprecated - please migrate to Unicast. This method of heartbeat communication uses broadcast UDP between Primary and Secondary, with a destination port specified by the *UDP Port for broadcast & unicast* parameter. Care must be taken when using broadcast on multiple pairs of load balancers in the same network. Each high-availability pair must operate on a different UDP port if they are not to interfere with each other. If heartbeat communication over the network is required, it is recommended that unicast be used in preference to broadcast.

**UDP Port for unicast & broadcast** - The UDP port number used by heartbeat for network communication over unicast or broadcast. By default, heartbeat uses UDP port 6694 for unicast or broadcast communication. If you have multiple load balancer pairs on the same subnet, and wish to use broadcast, you will need to set each pair to a different UDP port.

## Peer Failure Detection

**Keep-alive message interval** - Specify the number of seconds between keepalive pings. The Keepalive setting must be less than the warntime and deadtime.

**Dead peer timer** - The number of seconds communication can fail before a fail over is performed. A very low setting of deadtime could cause unexpected failovers.

**Warning timer** - If communication fails for this length of time write a warning to the logs. This is useful for tuning your deadtime without causing failovers in production.

## Routing Failure Detection

**Test IP address** - Specify one or more mutually accessible IP address to test network availability. A good ping node to specify is the IP address of a router that both the Primary and Secondary node can access. If the active node loses access to the ping node, the system will fail-over to the peer. However, if both nodes lose access, no fail-over will occur. Multiple IP addresses may be given, separated by spaces or commas. In this case, all addresses must be reachable for the routing test to pass.

**Test time-out** - Specify the time-out, in seconds, for the routing test. If a response is not received from the test address within the time-out period, the route to that host will be considered dead.

## Email Alerts

**Email Alert Destination Address** - Specify the email address where to send heartbeat alerts. In the event of failover or failback the email address specified will receive an alert.

**Email Alert Source Address** - Specify the email address from where to send heartbeat alerts. In the event of failover failback the email specified will send an alert.

**Note** | Both Primary and Secondary appliances will send an email in the event of a failover or failback.

## Fail-back Settings

**Automatic Fail-back** - Enable/disable auto-failback. This option controls the cluster behavior when the Primary returns to service after a failure. With Automatic Fail-back enabled, the Primary will automatically return to active status, taking back the floating IP addresses from the Secondary. With Automatic Fail- back disabled, the Secondary will remain active and will retain the floating IP addresses. Fail-over back to the Primary must then be controlled manually.

**Note** | Automatic Fail-back is disabled by default. Manual intervention is required to force the repaired Primary to become active and the Secondary unit to return to passive mode. For more information on controlling failback please refer to [Forcing Primary/Secondary Failover & Failback](#).

## Configuring a Smart Host (SMTP relay)

For Heartbeat (and layer 4 services), email alerts are sent to the mail server defined in the destination domain's DNS MX record directly from the load balancer by default. However, configuring a smart host (mail relay server) will instead direct all outgoing mail to the specified relay. This will significantly help to improve the deliverability of the

email. A smart host is an email server through which approved devices can send emails. Where possible, we recommend that you use a smart host to send email alerts.

*To configure a Smart Host:*

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration*.
2. Scroll down to the *SMTP Relay* section.
3. Set the smart host using an FQDN or IP address.
4. Click **Update**.

## Connection State & Persistence Table Replication

### Layer 4 VIPs

If you want the current connection state and persistence table to be available when the active appliance (typically the Primary) swaps over to the passive appliance (typically the Secondary), then you can start the synchronisation daemons on both appliances to replicate this data in real time as detailed below.

First, login to the Primary appliance using SSH or at the console, then as root run the following commands:

```
ipvsadm --start-daemon master
ipvsadm --start-daemon backup
```

Next, login to the Secondary appliance using SSH or at the console, then as root run the following commands:

```
ipvsadm --start-daemon master
ipvsadm --start-daemon backup
```

#### Note

'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

To ensure that these sync daemons are started on each reboot, put these commands in the file `/etc/rc.d/rc.firewall`. This can be done using the WebUI menu option: *Maintenance > Firewall Script*. Make sure that the full path is specified in the firewall script, i.e.

```
/usr/local/sbin/ipvsadm --start-daemon master
/usr/local/sbin/ipvsadm --start-daemon backup
```

After a few seconds you can confirm that it is working by viewing the connections report on each appliance which is available in the WebUI by navigating to: *Reports > Layer 4 Current Connections* as shown in the following examples:

*The active appliance:*

#### IPVS connection entries

pro	expire	state	source	virtual	destination
TCP	02:13	NONE	192.168.64.7:0	192.168.111.221:23	192.168.110.240:23
TCP	12:12	ESTABLISHED	192.168.64.7:53177	192.168.111.221:23	192.168.110.240:23
TCP	12:14	ESTABLISHED	192.168.64.7:53180	192.168.111.221:23	192.168.110.240:23

*The passive appliance:*

#### IPVS connection entries

pro	expire	state	source	virtual	destination
TCP	12:08	ESTABLISHED	192.168.64.7:53177	192.168.111.221:23	192.168.110.240:23
TCP	02:12	NONE	192.168.64.7:0	192.168.111.221:23	192.168.110.240:23
TCP	12:12	ESTABLISHED	192.168.64.7:53180	192.168.111.221:23	192.168.110.240:23

You can also run the following command at the command line:

```
ipvsadm -Lc
```

As shown, the state of all current connections as well as the persistence entries (i.e. those in state 'NONE') are replicated to the passive device. This enables current connections to continue through the passive appliance should the active appliance fail.

To stop the replication, run the following commands on both appliances:

```
ipvsadm --stop-daemon master  
ipvsadm --stop-daemon backup
```

#### Note

Setting this option can generate a high level traffic between the Primary and Secondary appliances.

#### Note

Once configured, you'll see multicast UDP traffic from the active appliance to multicast IP address 224.0.0.81 on port 8848.

## Layer 7 VIPs

If you want the current persistent connection table to work when the Primary load balancer swaps over to the Secondary then this can be enabled using the WebUI. Enabling this option will replicate persistence tables for all relevant layer 7 VIPs to the peer load balancer.

*To enable persistence state table replication:*

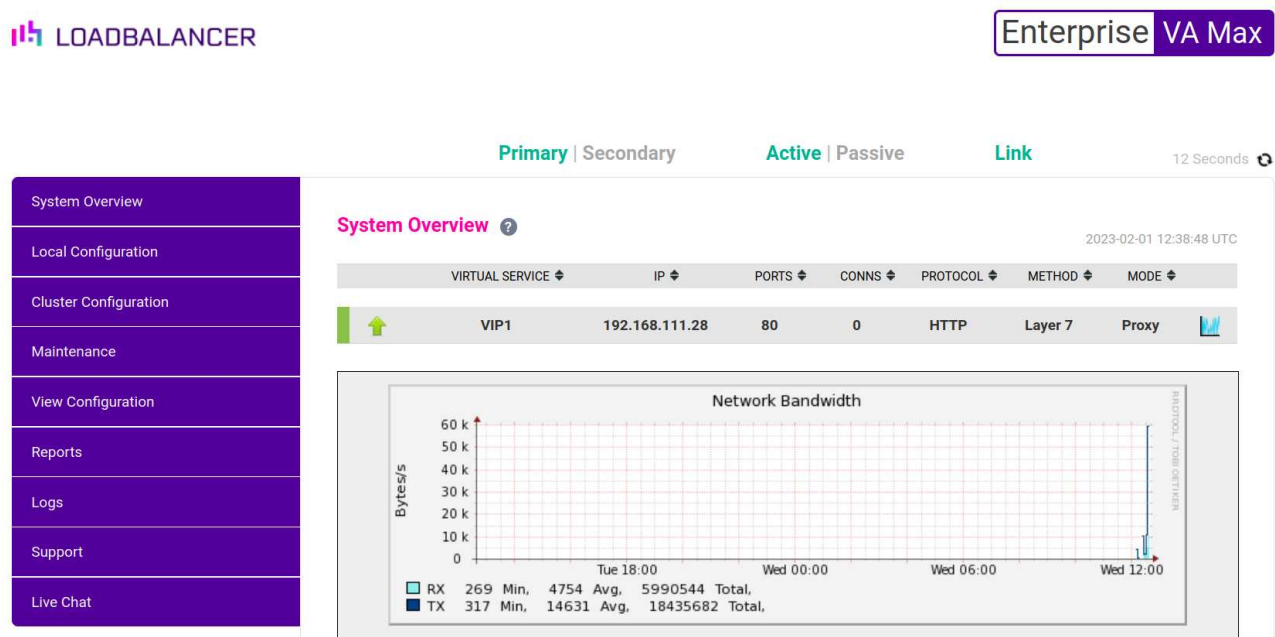
1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Advanced Configuration*.
2. Enable the *Persistence Table Replication*.
3. Click **Update**.
4. Now reload HAProxy using the **Reload** button in the "commit changes" message box at the top of the screen.

## Clustered Pair Diagnostics



Heartbeat State Diagnostics

The status of the appliance is shown at the top of the screen. For a working pair, the normal view is shown below:





Primary   Secondary	Active   Passive	Link	<p>this is a Primary unit, a Secondary unit has been defined, but the link is down (e.g. serial cable unplugged) so the state cannot be determined. In this case the floating IPs may be active on both units.</p> <p><b>Action:</b> check &amp; verify the heartbeat configuration, check the serial cable (if applicable), check heartbeat logs &amp; if required restart heartbeat on both units.</p>
---------------------	------------------	------	--

## Split Brain Scenarios

Split brain can occur if heartbeat on the Primary/Secondary clustered pair can no longer communicate with one another. In this case both units will assume that the other appliance is down and will bring up the Virtual Services. The system status will look similar to the following on both units:


Primary | Secondary
Active | Passive
Link
55 Seconds

 **Error:** The heartbeat link to the secondary node is down

 **Error:** heartbeat may be active on both load balancers

When heartbeat communication is re-established, heartbeat will automatically attempt to resolve the split brain and ensure that only one of the units is active. If heartbeat fails to do this automatically, the system status will show as follows on both units:

Primary | Secondary
Active | Passive
Link
12 Seconds

 **Error:** heartbeat may be active on both load balancers. To force this node to take control, click the button below.

Take over

The **Take over** button can then be used on either Primary or Secondary to attempt to force that unit to become active.

## Forcing Primary/Secondary Failover & Failback

To force the Secondary to become active & the Primary to become passive:

Either use the **Take over** button in the Secondary's system overview:

**Information:** This device is currently passive. Please see the active device for Virtual Service statistics.  
[ Advanced ]

**Take over** Make this node active

**Note** | Click on **[Advanced]** to show this button.

Or run the following command on the Secondary:

```
/usr/local/sbin/hb_takeover.php all
```

To force the Primary to become active & the Secondary to become passive:

Either use the **Take over** button on the Primary as explained above or run the following command on the Primary:

```
/usr/local/sbin/hb_takeover.php all
```

**Note** | These commands can either be run on the console, via an SSH session or via the WebUI menu option: *Local Configuration > Execute shell command*.

**Note** | The "Execute Shell Command" menu option is disabled by default. This can be enabled using the WebUI option: *Local Configuration > Security*. Set *Appliance Security Mode* to **Custom** then click **Update**.

**Note** | 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

## Testing & Verifying Primary/Secondary Replication & Failover

**Important** | Make sure you verify that Primary/Secondary failover occurs correctly before going live. This proves the resilience of the HA cluster and makes you aware of the failover/failback process.

**Note** | When testing appliance fail-over, if heartbeat is configured to use only the serial link, don't just pull the serial cable out. This will not cause a fail-over but will cause a split brain (i.e. both units active) to occur. Testing must be done by pulling both the network and serial cable (if used) as detailed below.

### STEP 1 - Verify Basic Settings for the clustered pair

1. On the Primary unit verify that the system status appears as follows:

**Primary** | **Secondary**      **Active** | **Passive**      **Link**

2. On the Secondary unit verify that the system status appears as follows:

**STEP 2 - Verify Replication**

1. Verify that the load balanced services have been replicated to the Secondary unit, this can be done by using either the *View Configuration* or *Edit Configuration* menus to validate that the same Virtual & Real Servers exist on the Secondary as on the Primary.

**STEP 3 - Verify Failover to the Secondary (using the Take over button)**

1. On the Secondary unit, click the **[Advanced]** option in the information box, then click the **Take Over** button.
2. Verify that the Secondary's status changes to *Active*:

3. Verify that the Primary's status changes to *Passive*:

4. Using the WebUI menu option: *View Configuration > Network Configuration*, verify that the floating IPs associated with the VIPs have been brought up on the Secondary unit and brought down on the Primary.

e.g. the partial screen shot below from the View Network Configuration screen on the Secondary unit shows the status of eth0:

```
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 00:0c:29:92:18:4f brd ff:ff:ff:ff:ff:ff
    inet 192.168.111.223/18 brd 192.168.127.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet 192.168.111.72/18 brd 192.168.127.255 scope global secondary eth0
        valid_lft forever preferred_lft forever
```

This shows the secondary IP address 192.168.111.72 (the VIP address) is up and therefore the Secondary has become active as intended.

**STEP 4 - Verify Fail-back to the Primary (using the Take over button)**

1. On the Primary unit, click the **[Advanced]** option in the information box, then click the **Take Over** button.
2. Verify that the Primary's status has changed to *Active*:

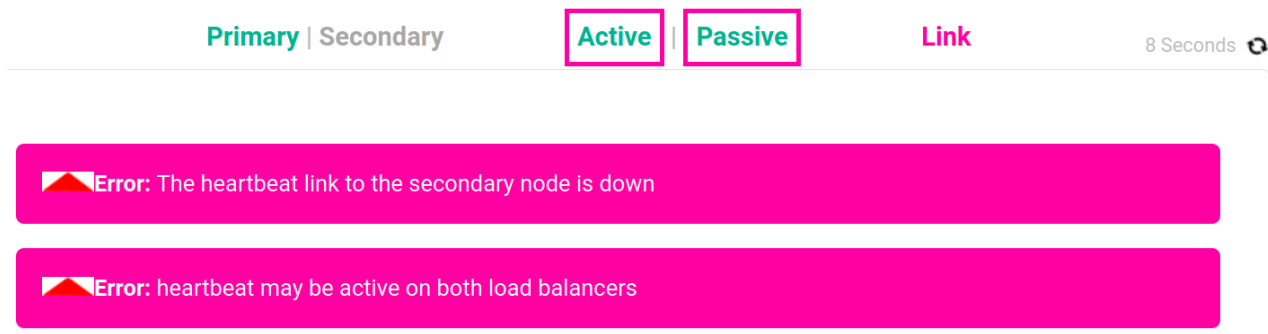
3. Verify that the Secondary's status has changed to *Passive*:

4. Also, using the WebUI menu option: *View Configuration > Network Configuration*, verify that the floating IPs

associated with the VIPs have been brought up on the Primary unit and brought down on the Secondary (see STEP 3 above for more details).

#### STEP 5 - Verify Failover to the Secondary (when removing the network and serial cable from Primary)

1. Remove the network cable and serial cable (if applicable) from the Primary.
2. Verify that the Secondary's status has changed as follows:



This indicates that the Secondary is unable to communicate with the Primary. This means that either the Primary is down, or is still up but is unreachable. In both cases the Secondary will become active.

3. On the Secondary using the WebUI menu option: *View Configuration > Network Configuration*, verify that the floating IPs associated with the VIPs have been brought up (see STEP 3 above for more details).

#### STEP 6 - Verify normal operation resumes (when reconnecting the network & serial cable to Primary)

1. Reconnect the cables to the Primary.
2. Verify that the Primary's status is set to *Active*:

Primary | Secondary      Active | Passive      Link

3. Verify that the Secondary has changed to *Passive*:

Primary | Secondary      Active | Passive      Link

4. Also, using the WebUI menu option: *View Configuration > Network Configuration*, verify that the floating IPs associated with the VIPs have been brought up on the Primary unit and brought down on the Secondary.

#### Note

If the power cable on the Primary had been removed rather than disconnecting the network cable and serial cable (if applicable), once the Primary is brought back up the Secondary would remain active and the Primary would come back up in a passive state. The **Take over** button on the Primary would then need to be used to force the Primary to become active.

#### Heartbeat Log

The log provides a useful insight into the health of an HA pair and can be viewed using the WebUI menu option: *Logs > Heartbeat*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/ha.log`.

## Late Heartbeats

Late Heartbeats are caused when the heartbeat packets either arrive at the other appliance later than the *Warning Timer* configured in the WebUI under *Cluster Configuration > Heartbeat Configuration* or not at all.

The following log snippet shows 3 entries reporting a late heartbeat:

```
Apr 26 21:03:53 lbmaster heartbeat: [2474]: WARN: Late heartbeat: Node lbmaster: interval 5110 ms
Apr 26 21:04:30 lbmaster heartbeat: [2474]: WARN: Late heartbeat: Node lbmaster: interval 6430 ms
Apr 26 21:04:35 lbmaster heartbeat: [2474]: WARN: Late heartbeat: Node lbmaster: interval 5830 ms
```

This suggests that there is either high latency or packet loss and is the first signs of issues that can lead to a split brain.

## Troubleshooting late Heartbeats

- Verify that the system load average is acceptable - An overloaded system can be struggling to send or receive messages.
- Check if the network card is saturated - An overloaded NIC can add latency.

**Note** | For more information, please refer to [Resource Utilisation Graphs](#).

- Confirm that there are no bottlenecks or packet loss further down the wire - Such as links between data centres where there may be traffic congestion/contention.
- Once a possible cause has been identified and addressed, monitor the situation to make sure the issue has been resolved.

# Chapter 10 - Configuration Examples

## Introduction

This section presents a number of examples that illustrate how the appliance can be configured.

### Initial Network Settings

For more information on configuring initial network settings please refer to [Configuring Initial Network Settings](#). For more information on how to access the WebUI please refer to [Accessing the WebUI](#).

### Example 1 - One-Arm DR Mode (Single Appliance)

This DR (Direct Return) mode example has one Virtual Service (VIP) with two Real Servers (RIPs). It's a straight forward deployment mode that can be used in many situations. It also offers the highest performance because return traffic passes directly from the Real Servers to the client rather than passing back via the load balancer.

#### Configuration Overview


- **Configure Network Settings** - a single Interface is needed, eth0 is normally used
- **Define the Virtual Service (VIP)** - all Real (backend) Servers are accessed via this IP address
- **Define the Real Servers (RIPs)** - define the Real Servers that make up the cluster
- **Implement the required changes to the Real Servers** - for DR mode, the **ARP Problem** must be solved


#### Network Settings


Configure the various network settings as outlined below:


1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.
2. Scroll down to the *IP Address Assignment* section.

**IP Address Assignment**

  
eth0

  
eth1

  
eth2

  
eth3

eth0

192.168.2.120/24

MTU

1500

 bytes

3. Specify the IP address & subnet mask for eth0 (normally eth0 is used for single-arm configurations although this is not mandatory), e.g. **192.168.2.120/24**.
4. Click **Configure Interfaces**.
5. Using the WebUI, navigate to: *Local Configuration > Hostname & DNS*.
6. Specify the DNS server(s).

Domain Name Server	Primary	<input type="text" value="192.168.2.254"/>	<a href="#">?</a>
	Secondary	<input type="text"/>	<a href="#">?</a>
	Tertiary	<input type="text"/>	<a href="#">?</a>

- Click **Update**.
- Using the WebUI, navigate to: *Local Configuration > Routing*.

Default Gateway			
IP v4	<input type="text" value="192.168.2.254"/>	via interface	<input type="text" value="auto"/> <a href="#">?</a>
IP v6	<input type="text"/>	via interface	<input type="text" value="auto"/> <a href="#">?</a>

- Specify the Default Gateway.
- Click **Configure Routing**.

## Virtual Service (VIP)

Next, configure the Virtual Service. This is the IP address that is presented to clients. Any packet arriving at the load balancer with that IP address and port number will be forwarded to the the Real Servers associated with the Virtual Service. This example is for Web traffic on TCP port 80.

- Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services* and click **Add a new Virtual Service**.

Virtual Service			
Label	<input type="text" value="ExVIP1"/>		<a href="#">?</a>
IP Address	<input type="text" value="192.168.2.150"/>		<a href="#">?</a>
Ports	<input type="text" value="80"/>		<a href="#">?</a>
Protocol			
Protocol	<input type="text" value="TCP"/>		<a href="#">?</a>
Forwarding			
Forwarding Method	<input type="text" value="Direct Routing"/>		<a href="#">?</a>

- Enter a suitable Label (name) for the VIP, e.g. **ExVIP1**.
- Enter a valid IP address, e.g. **192.168.2.150**.
- Enter a valid port, e.g. **80**.
- Select the required Protocol, .e.g. **TCP**.
- Ensure that the *Forwarding Method* is set to **Direct Routing**.



7. Click **Update**.

## Real Servers (RIPs)

Each Virtual Service requires a cluster of Real Servers (backend servers) to forward the traffic to.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Real Servers* and click **Add a new Real Server** next to the relevant Virtual Service.

Label	<input type="text" value="RIP1"/>	?
Real Server IP Address	<input type="text" value="192.168.2.151"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

2. Enter a suitable Label (name) for the RIP, e.g. **RIP1**.
3. Enter a valid IP address, e.g. **192.168.2.151**.

### Note

There is no option to specify a Real Server port because port redirection is not possible when using DR mode. The port used will be the same as that configured for the VIP.

4. The weight defaults to 100 making the Real Server active immediately.
5. Leave *Minimum Connections* & *Maximum Connections* set to 0 which means unrestricted.
6. Click **Update**.
7. Repeat for the other Real Server.

## Physical Real Server Changes - Solve the ARP Problem

For DR mode, the **ARP Problem** must be solved on each Real Server:

- Each Real Server must be configured to respond to its own IP address **and** the VIP address
- Each Real Server must be configured so that it only responds to ARP requests for its own IP address, it should **not** respond to ARP requests for the VIP address - only the load balancer must respond to these requests

### Note

Failure to correctly configure the Real Servers to handle the **ARP Problem** is the most common problem in DR configurations. For more information on the **ARP Problem** and the solution for various OSs, please refer to [The ARP Problem](#).

## Basic Testing & Verification

Once configured, a few quick checks can be performed to verify the setup:

1. Using *System Overview* check that the VIP & RIPs are shown as active (green).
2. Using a browser, navigate to the VIP address, i.e. **http://192.168.2.150** to verify that you can reach the Real

Servers via the Virtual Service.

3. Check *Reports > Layer 4 Current Connections* to ensure that client connections are reported in state 'ESTABLISHED'. If connections are in state 'SYN\_RECV', this normally indicates that the **ARP Problem** on the Real Servers has not been solved.

## Example 2 - One-Arm Layer 4 SNAT Mode (Single Appliance)

This layer 4 SNAT mode example has one Virtual Service (VIP) with two Real Servers (RIPs). This mode is ideal for example when you want to load balance both TCP and UDP but you're unable to use DR mode or NAT mode due to network topology or Real Server related reasons. Layer 4 SNAT mode is non-transparent by default, i.e. the Real Servers will see the source IP address of the load balancer.

**Note** | In this mode, no changes are required to the Real Servers.

### Configuration Overview


- **Configure Network Settings** - a single Interface is needed, eth0 is normally used
- **Define the Virtual Service (VIP)** - all Real (backend) Servers are accessed via this IP address
- **Define the Real Servers (RIPs)** - define the Real Servers that make up the cluster


### Network Settings


Configure the various network settings as outlined below:


1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.

**IP Address Assignment**

eth0

eth1




eth2

eth3

eth0

MTU  bytes

2. Specify the IP address & subnet mask for eth0 (normally eth0 is used for single-arm configurations although this is not mandatory), e.g. **192.168.2.120/24**.
3. Click **Configure Interfaces**.
4. Using the WebUI, navigate to: *Local Configuration > Hostname & DNS*.
5. Specify the DNS server(s).

Domain Name Server	Primary	<input type="text" value="192.168.2.254"/>	
	Secondary	<input type="text"/>	
	Tertiary	<input type="text"/>	

- Click **Update**.
- Using the WebUI, navigate to: *Local Configuration > Routing*.

Default Gateway			
IP v4	<input type="text" value="192.168.2.254"/>	via interface	auto <span>▼</span> <span>?</span>
IP v6	<input type="text"/>	via interface	auto <span>▼</span> <span>?</span>

- Specify the Default Gateway.
- Click **Configure Routing**.

## Virtual Service (VIP)

Next, configure the Virtual Service. This is the IP address that is presented to clients. Any packet arriving at the load balancer with that IP address and port number will be forwarded to the the Real Servers associated with the Virtual Service. This example is for RDP traffic on TCP/UDP port 3389.

- Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services* and click **Add a new Virtual Service**.

Virtual Service	
Label	<input type="text" value="ExVIP2"/> <span>?</span>
IP Address	<input type="text" value="192.168.2.150"/> <span>?</span>
Ports	<input type="text" value="3389"/> <span>?</span>
Protocol	
Protocol	TCP/UDP <span>▼</span> <span>?</span>
Forwarding	
Forwarding Method	SNAT <span>▼</span> <span>?</span>

Cancel
Update

- Enter a suitable Label (name) for the VIP, e.g. **ExVIP2**.
- Enter a valid IP address, e.g. **192.168.2.150**.
- Enter the required port, e.g. **3389**.
- Select the required *Protocol* , in this example **TCP/UDP** (UDP support was added in RDP v8.0).
- Ensure that *Forwarding Method* is set to **SNAT**.
- Click **Update**.

## Real Servers (RIPs)

Each Virtual Service requires a cluster of Real Servers (backend servers) to forward the traffic to.

- Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Real Servers* and click **Add a new Real Server** next to the relevant Virtual Service.

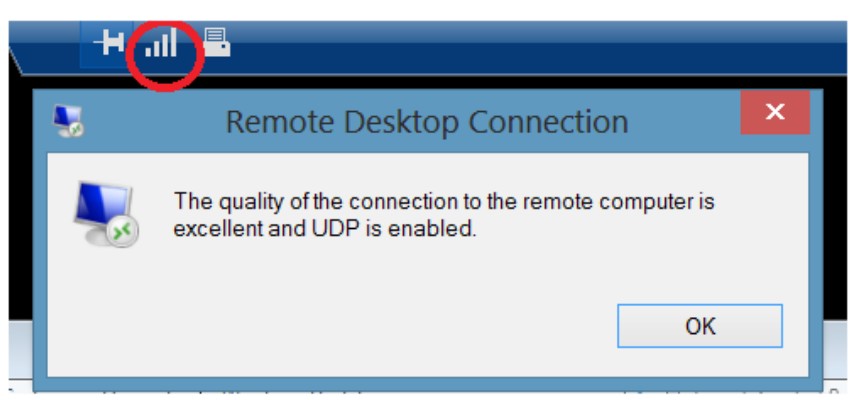
Label	<input type="text" value="RDS1"/>	?
Real Server IP Address	<input type="text" value="192.168.2.151"/>	?
Real Server Port	<input type="text" value="3389"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

2. Enter a suitable Label (name) for the RIP, e.g. **RDS1**.
3. Enter a valid IP address, e.g. **192.168.2.151**.
4. Enter the required port, e.g. **3389**.
5. The weight defaults to 100 making the Real Server active immediately.
6. Leave *Minimum Connections* & *Maximum Connections* set to 0 which means unrestricted.
7. Click **Update**.
8. Repeat for the other Real Server.

## Basic Testing & Verification

Once configured, a few quick checks can be performed to verify the setup:

1. Using *System Overview* check that the VIP & RIPs are shown as active (green).
2. Using Windows RDP client (mstsc.exe) to connect to the VIP address, i.e. **192.168.2.150** to verify that you can start an RDP session.
3. Verify that the RDP session supports TCP & UDP by clicking the connection info button on the RDS Connection Bar:



## Example 3 - Two-Arm NAT Mode (Clustered Pair)

This example shows how to configure two appliances as a clustered pair using layer 4 NAT mode.

### Note

Using two appliances configured as a clustered pair is Loadbalancer.org's recommended

configuration and ensures that no single point of failure is introduced.

Note

When using two-arm NAT mode all Real Servers should be in the same subnet as the internal interface of the load balancer and the default gateway on each Real Server must be set to be an IP on the load balancer, for a clustered pair this should be a floating IP to allow failover.

## Configuration Overview

- **Configure the Primary's Network Settings** - two Interfaces are needed, this can be either two physical interfaces such as eth0 and eth1, or one physical interface and a secondary interface/alias
- **Configure the Secondary's Network Settings** - two Interfaces are needed, this can be either two physical interfaces such as eth0 and eth1, or one physical interface and a secondary interface/alias
- **On the Primary, Define the Virtual Service (VIP)** - all Real Servers are accessed via this IP address
- **On the Primary, Define the Real Servers (RIPs)** - define the Real Servers that make up the cluster
- **Implement the required changes to the Real Servers** - in NAT mode, the Real Servers default gateway must be set to be the load balancer
- **Create the HA Clustered Pair** - pair the Primary & Secondary to synchronize the appliances
- **Verify Heartbeat Settings** - check that the default heartbeat settings are appropriate

## Primary Unit - Network Settings

1. Using the WebUI on the Primary unit, navigate to: *Local Configuration > Network Interface Configuration*.

### IP Address Assignment

eth0

192.168.2.120/24

MTU 1500 bytes

eth1

192.168.20.120/24

MTU 1500 bytes

2. Specify the IP address & mask for eth0 - normally eth0 is configured as the internal interface although this is not mandatory, e.g. **192.168.2.120/24**.
3. Specify the IP address & mask for eth1 - normally eth1 is configured as the external interface although this is not mandatory, e.g. **192.168.20.120/24**.

Note

For a VA make sure that the virtual NIC associated with eth1 is connected to the virtual switch, by default only the first NIC is connected.


4. Click **Configure Interfaces**.


## Secondary Unit - Network Settings


Configure the various network settings as outlined below:


1. Using the WebUI on the Secondary appliance, navigate to: *Local Configuration > Network Interface Configuration*.

**IP Address Assignment**

eth0

eth1

eth2

eth3

eth0

192.168.2.121/24

MTU 1500 bytes

eth1

192.168.20.121/24

MTU 1500 bytes

2. Specify the IP address & mask for eth0 - normally eth0 is configured as the internal interface although this is not mandatory, e.g. **192.168.2.121/24**.
3. Specify the IP address & mask for eth1 - normally eth1 is configured as the external interface although this is not mandatory, e.g. **192.168.20.121/24**.

### Note

For a VA make sure that the virtual NIC associated with eth1 is connected to the virtual switch, by default only the first NIC is connected.

4. Click **Configure Interfaces**.

## Virtual Service (VIP)

Next, configure the Virtual Service. This is the IP address that is presented to clients. Any packet arriving at the load balancer with that IP address and port number will be forwarded to the the Real Servers associated with the Virtual Service. This should be done on the Primary appliance.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Virtual Services* and click **Add a new Virtual Service**.

Virtual Service		
Label	<input type="text" value="ExVIP3"/>	<a href="#">?</a>
IP Address	<input type="text" value="192.168.20.150"/>	<a href="#">?</a>
Ports	<input type="text" value="80"/>	<a href="#">?</a>
Protocol		
Protocol	<input type="text" value="TCP"/>	<a href="#">?</a>
Forwarding		
Forwarding Method	<input type="text" value="NAT"/>	<a href="#">?</a>

2. Enter a suitable label (name) for the VIP, e.g. **ExVIP3**.
3. Enter a valid IP address, e.g. **192.168.20.150**.
4. Enter a valid port, e.g. **80**.
5. Ensure that *Forwarding Method* is set to **NAT**.
6. Click **Update**.

## Real Servers (RIPs)

Each Virtual Service requires a cluster of Real Servers (backend servers) to forward the traffic to. This should be done on the [Primary appliance](#).

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 - Real Servers* and click **Add a new Real Server**.

Label	<input type="text" value="RIP1"/>	<a href="#">?</a>
Real Server IP Address	<input type="text" value="192.168.2.151"/>	<a href="#">?</a>
Real Server Port	<input type="text" value="80"/>	<a href="#">?</a>
Weight	<input type="text" value="100"/>	<a href="#">?</a>
Minimum Connections	<input type="text" value="0"/>	<a href="#">?</a>
Maximum Connections	<input type="text" value="0"/>	<a href="#">?</a>

2. Enter a suitable Label (name) for the RIP, e.g. **RIP1**.
3. Enter a valid IP address, e.g. **192.168.2.151**.
4. Enter a valid port, e.g. **80**.
5. *Weight* defaults to 100 making the Real Server active immediately.
6. Leave *Minimum Connections* & *Maximum Connections* set to 0 which means unrestricted.

7. Click **Update**.
8. Repeat for the other Real Server.

## Physical Real Server Changes - Set the Default Gateway

When using NAT mode, each Real Server's default gateway must be changed to be the load balancer. For a clustered pair, you must define an additional floating IP for this purpose. Then, if failover occurs, the same IP will also be brought up on the Secondary. To add a floating IP to use as the default gateway, use *Cluster Configuration > Floating IPs*.

New Floating IP

192.168.2.254


Add Floating IP

Define the IP address that you'd like to use for the default gateway, then click **Add Floating IP**. Now configure the default gateway on each Real Server to use this address.

## Create the HA Clustered Pair

1. Using the WebUI on the Primary appliance, navigate to: *Cluster Configuration > High Availability Configuration*

### Create a Clustered Pair

 **LOADBALANCER**

**Local IP address**

192.168.2.120

**IP address of new peer**

192.168.2.121

**Password for loadbalancer user on peer**


••••••••••

Add new node

2. Leave the *Local IP address* set as the address assigned to eth0, in this case **192.168.2.120**.
3. Specify the *IP address of new peer* (i.e. the Secondary appliance), in this case **192.168.2.121**.
4. Specify the *loadbalancer* users password (the default is 'loadbalancer') for the Secondary appliance.
5. Click **Add new node**.
6. A warning will be displayed indicating that the pairing process will overwrite the new Secondary appliance's existing configuration, click **OK** to continue.
7. The pairing process will start as shown below:




### Create a Clustered Pair

 **LOADBALANCER** 192.168.2.20

Primary

Attempting to pair..

 **LOADBALANCER** 192.168.2.21

Secondary

Local IP address

192.168.2.120

IP address of new peer

192.168.2.121

Password for *loadbalancer* user on peer

••••••••••

8. Once completed successfully, the following message will be displayed:

Commit changes

The configuration of the following services has been changed. When reconfiguration is complete, restart/reload the services to commit the changes

Restart Heartbeat

### High Availability Configuration - primary

 **LOADBALANCER** 192.168.2.20

Primary

 **LOADBALANCER** 192.168.2.21

Secondary

Break Clustered Pair

Make Active

9. To finalize the configuration, click **Restart Heartbeat**.

## Checking the Status

A successfully configured clustered pair will display the following status:

1) On the Primary unit verify that the system status appears as follows:

**Primary** | Secondary      **Active** | Passive      **Link**

2) On the Secondary unit verify that the system status appears as follows:

Primary | **Secondary**      Active | **Passive**      **Link**

Note

Once the VIP has been defined, the Active text will be colored green on the Primary and Passive will be colored green on the Secondary.

## Verify Heartbeat Settings

- Using the WebUI on the Primary appliance , navigate to: *Cluster Configuration > Heartbeat Configuration*.
- The default Heartbeat settings are normally fine for most situations. For details of all heartbeat options please refer to **Configuring Heartbeat**.

## Verify the Secondary Configuration

To verify that the new VIP & RIP have been replicated correctly, open the WebUI on the Secondary and navigate to: *Cluster Configuration > Layer 4 - Virtual Services* and *Cluster Configuration > Layer 4 - Real Servers* and check that your configuration appears there also. For a correctly configured pair, the VIPs and RIPs are automatically replicated to the Secondary as they are defined on the Primary.

If not, double check that both units are configured correctly and that the IP address for the Secondary defined on the Primary is correct. Then on the Primary navigate to: *Maintenance > Backup & Restore* and click *Synchronize Configuration with peer*. This will force the VIPs & RIPs to be copied from the Primary to the Secondary, then check again.

## Basic Testing & Verification

A few quick checks can be performed to verify the configuration:

1. On the Primary, use *System Overview* to check that the VIP & RIPs are shown as active (green).
2. Using a browser, navigate to the VIP address, i.e. <http://192.168.2.150> to verify that you can reach the Real Servers via the Virtual Service.
3. On the Primary, check *Reports > Layer 4 Current Connections* to ensure that client connections are reported in state '**ESTABLISHED**'. If not and instead '**SYN-RCV'** is shown, double-check that you have set the default gateway on all Real Servers to be the floating IP address on the load balancer.

## Example 4 - One-Arm SNAT Mode & SSL Termination (Single Appliance)

This example uses HAProxy and STunnel at layer 7. STunnel is used to terminate SSL on the load balancer. STunnel then passes unencrypted HTTP traffic to the HAProxy VIP/RIP cluster. HAProxy does not offer the raw throughput of layer 4, but is still a high performance solution that is appropriate in many situations.

<b>Note</b>	Pound can also be used for SSL termination, although STunnel is the preferred and default method.
-------------	---

In this example it's assumed that the Real Server application has not been designed to track & share session details between Real Servers. Therefore, cookie based persistence will be enabled on the load balancer to ensure that clients connect to the same Real Server on each subsequent connection (within the persistence timeout window). If persistence is not configured then new connections may get distributed to a different Real Server which may result in failure of the application.

<b>Note</b>	Because HAProxy is a full proxy, any server in the cluster can be on any accessible subnet including across the Internet or WAN.
-------------	--

<b>Note</b>	In this mode, no changes are required to the Real Servers.
-------------	--

<b>Note</b>	We generally recommend that SSL is terminated on the Real Servers rather than on the load balancer. This ensures that the SSL load is distributed and also ensures scalability.
-------------	---

## Configuration Overview

- **Configure Network Settings** - A single Interface is needed, eth0 is normally used
- **Define the Virtual Service (VIP)** - All Real Servers are accessed via this IP address
- **Define the Real Servers (RIPs)** - Define the Real Servers that make up the cluster


- **Configure SSL Termination** - Configure STunnel for SSL termination

## Network Settings


Configure the various network settings as outlined below:

1. Using the WebUI, navigate to: *Local Configuration > Network Interface Configuration*.


**IP Address Assignment**




eth0



eth1



eth2



eth3

eth0

192.168.2.120/24

MTU 

1500

 bytes

2. Specify the IP address & mask for eth0 - normally eth0 is used for one-arm configurations although this is not mandatory, e.g. **192.168.2.120/24**.
3. Click **Configure Interfaces**.
4. Using the WebUI, navigate to: *Local Configuration > DNS & Hostname*.
5. Specify the DNS server(s).

Domain Name Server	Primary	<div style="border: 1px solid #ccc; padding: 2px 10px;">192.168.2.254</div>	?
	Secondary	<div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>	?
	Tertiary	<div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>	?

6. Click **Update**.
7. Using the WebUI, navigate to: *Local Configuration > Routing*.

**Default Gateway**

IP v4	<div style="border: 1px solid #ccc; padding: 2px 10px;">192.168.2.254</div>	via interface	<div style="border: 1px solid #ccc; padding: 2px 10px;">auto</div>	?
IP v6	<div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>	via interface	<div style="border: 1px solid #ccc; padding: 2px 10px;">auto</div>	?

8. Specify the Default Gateway, e.g. 192.168.2.254.
9. Click **Configure Routing**.

## Virtual Service (VIP)

Next, configure the Virtual Service. This is the IP address that is presented to clients. Any packet arriving at the load balancer with that IP address and port number will be handled by the Real Servers associated with the Virtual Service.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Virtual Services* and click **Add a new Virtual Service**.

Virtual Service		[Advanced +]
Label	<input type="text" value="ExVIP4"/>	?
IP Address	<input type="text" value="192.168.2.150"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="HTTP Mode"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

2. Enter a suitable Label (name) for the VIP, e.g. **ExVIP4**.
3. Enter a valid IP address, e.g. **192.168.2.150**.
4. Enter a valid port, e.g. **80**.
5. Leave *Layer 7 Protocol* set to **HTTP Mode**.
6. Click **Update**.

## Real Servers (RIPs)

Each Virtual Service requires a cluster of Real Servers (backend servers) to forward the traffic to.

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 - Real Servers* and click **Add a new Real Server**.

Label	<input type="text" value="RIP1"/>	?
Real Server IP Address	<input type="text" value="192.168.2.151"/>	?
Real Server Port	<input type="text" value="80"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Enable Redirect	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

2. Enter a suitable Label (name) for the RIP, e.g. **RIP1**.
3. Enter a valid IP address, e.g. **192.168.2.151**.

### Note

In this mode it's possible to have a different port for the RIP than was configured for the VIP, in this example both are the same.

4. Enter a valid port, e.g. **80**.
5. The *Weight* defaults to 100 making Real Servers active as soon as HAProxy is restarted.

- 6. Click **Update**.
- 7. Repeat for the remaining Real Servers.
- 8. Reload HAProxy to apply the new settings using the link provided in the "commit changes" message box at the top of the screen.

SSL Termination

An SSL Virtual Service is configured on port 443 using the same IP address as the Layer 7 VIP created previously. This allows a single IP address to be used for HTTP & HTTPS client connections.

- 1. Using the WebUI, navigate to: *Cluster Configuration > SSL Termination* and click **Add a new Virtual Service**.

Label	<input type="text" value="SSL_ExVIP4"/>	?
Associated Virtual Service	<div>ExVIP4 ▾</div>	?
Virtual Service Port	<input type="text" value="443"/>	?
SSL Operation Mode	<div>High Security ▾</div>	
SSL Certificate	<div>Default Self Signed Certificate ▾</div>	?
Source IP Address	<input type="text"/>	?
Enable Proxy Protocol	<input checked="" type="checkbox"/>	?
Bind Proxy Protocol to L7 VIP	<div>ExVIP4 ▾</div>	?
<div>Cancel Update</div>		

- 2. Enter an appropriate *Label* (name) for the new Virtual Service.

Note

The label will be auto configured based on the *Associated Virtual Service* selected. This can be edited if required.

- 3. Set *Associated Virtual Service* to the Layer 7 VIP created earlier, e.g. **ExVIP4**.
- 4. Leave *Virtual Service Port* set to **443**.
- 5. Leave the other settings at their default values.
- 6. Click **Update**.
- 7. Reload STunnel to apply the new settings using the button in "commit changes" message box.

Note

When creating the SSL Virtual Service, by default a self-signed certificate is used. This is ideal for testing but needs to be replaced for live deployments. Certificates can be added using the WUI option: *Cluster Configuration > SSL Certificate*. Once added, these will appear in the *SSL Certificate* drop-down when creating the SSL VIP.

Note

For more information on configuring SSL termination please refer to [SSL Termination](#).

Basic Testing & Verification

A few quick checks can be performed to verify the configuration:

1. Using *System Overview*, verify that the VIP & RIP are shown as active (green).
2. Using a browser, navigate to the VIP address, i.e. **http://192.168.2.150** to verify that you can reach the Real Servers via the Virtual Service using HTTP.
3. Using a browser, navigate to the STunnel SSL VIP address, i.e. **https://192.168.2.150** to verify that you can reach the Real Servers via the Virtual Service using HTTPS.

## Example 5 - Load Balancing FTP

This section provides guidance on how FTP can be load balanced using both layer 4 and layer 7 Virtual Services.

FTP is a multi-port service in both active and passive modes:

**active** 20,21

**passive** 21,high\_port

### Layer 4

To load balance FTP at layer 4, configure a layer 4 VIP in the normal way and set the Virtual Service/Real Server port field to port 21. The other ports that are required for both active and passive mode will be configured automatically using Firewall Marks. When using NAT mode, the `ip_vs_ftp` module is used to ensure that the client connects back via the load balancer rather than attempting to connect directly to the Real Server.

**Note** | For VIPs configured in this way, the checkport is automatically set to port 21.

### Layer 7

#### Active Mode

In active mode, the FTP server connects back to the client, so it must be aware of the clients IP address. To achieve this, TProxy must be enabled to make the load balancer transparent at layer 7. For this to work, two subnets must be used - the Virtual Server (VIP) in one subnet, the RIPs (i.e. the FTP servers) in the another. For more details on TProxy please refer to [Transparency at Layer 7](#).

Also, to ensure that the client receives a connection from the same address that it established the control connection to, an iptables SNAT rule must be defined in the firewall script for each FTP server. The format of the required rule is as follows:

```
iptables -t nat -A POSTROUTING -p tcp -s <FTP-Server-IP> -j SNAT --to-source <FTP-VIP>
```

e.g.

```
iptables -t nat -A POSTROUTING -p tcp -s 10.20.1.1 -j SNAT --to-source 192.168.2.180
```

(one rule must be added for each FTP server in the cluster)

**Note** | These rules can be added to the firewall script using the WebUI menu option: *Maintenance > Firewall Script*.

#### Active Mode - Key Points:

- Use separate subnets for the VIP & RIPs
- Enable TProxy for the layer 7 VIP
- Set the default gateway on the FTP servers to be an IP on the load balancer (ideally a floating IP to permit failover to the Secondary unit)
- Setup a layer 7 VIP listening on port 21 & configure the RIPs also to listen on port 21
- Ensure the Layer 7 Protocol is set to **TCP Mode**
- Increase the default client & server HAProxy timeouts to 5 minutes
- Add the SNAT firewall rules for each FTP server

#### Windows Example

1. Create a L7 VIP with the following settings, changing the name and IP address as required:

#### Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	FTP-ClusterACTV	?
IP Address	192.168.2.150	?
Ports	21	?
Protocol		
Layer 7 Protocol	TCP Mode ▾	?
		<span>Cancel</span> <span>Update</span>

2. Click **Update** to create the VIP.
3. Click **Modify** next to the newly created VIP.
4. Scroll down to the section *Other* and click **[Advanced]**.
5. Enable (check) the *Timeout* checkbox then set both *Client Timeout* and *Server Timeout* to **5m**.
6. Define the FTP servers as RIPs for the VIP just created as illustrated below (these must be on a different subnet to the VIP to enable TProxy to work correctly):

#### Layer 7 Add a new Real Server - FTP-ClusterACTV

Label	ftp1	?
Real Server IP Address	10.10.1.1	?
Real Server Port	21	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	100	?
		<span>Cancel</span> <span>Update</span>

7. Enable TProxy using the WebUI menu option: *Cluster Configuration > Layer 7 - Advanced Configuration*.
8. Now restart HAProxy using the WebUI menu option: *Maintenance > Restart Services*.
9. Define a SNAT rule for each FTP server using the WebUI menu option: *Maintenance > Firewall Script*.

e.g.

```
iptables -t nat -A POSTROUTING -p tcp -s 10.10.1.1 -j SNAT --to-source 192.168.2.150
iptables -t nat -A POSTROUTING -p tcp -s 10.10.1.2 -j SNAT --to-source 192.168.2.150
```

10. Configure the default gateway on each FTP server to be the load balancer. Ideally this should be a floating IP address to allow it to float (move) between the Primary & Secondary appliance. This can be added using the WebUI menu option: *Cluster Configuration > Floating IPs*.
11. Active FTP clients should now be able to connect to the VIP address (192.168.2.150) and view the directory listing successfully.

## Passive Mode

In passive mode all connections are initiated by the client. The server passes the client a port to use for the inbound data connection. By default, FTP servers can use a wide range of ports for the inbound connection and it's often useful to limit this range. For more information on configuring passive mode for a range of OSs please refer to [Limiting Passive FTP Ports](#).

### Note

This method configures HAProxy to listen on port 21 (control channel) and all passive ports (data channel).

## Passive Mode - Key Points:

- It's sensible to use a controlled passive port range, this can be configured on the FTP server
- Configure the VIP to listen on port 21 and also the passive range selected, e.g. 50000-50100
- Configure the RIPv without specifying a port
- Ensure the Layer 7 Protocol is set to 'TCP Mode'
- If transparency is required (for passive mode this is optional), enable TProxy. This is done at the VIP level rather than globally as in previous versions. To enable TProxy at the VIP level, click **Modify** next to the VIP in question, scroll down to the *Other* section and click **[Advanced]**, then enable (check) *Transparent Proxy*.

### Note

If TProxy is enabled, certain topology requirements must be met. Also the default gateway on each FTP server must also be set to be an IP on the load balancer - preferably a floating IP which then allows failover to the Secondary unit. For More information on using TProxy please refer to [Transparency at Layer 7](#).

- Set the Client Timeout & Real Server Timeout to 5m (i.e. 5 minutes)
- Set the Persistence Mode to Source IP
- The Persistence Timeout can be left set to 30 (i.e. 30 minutes)
- To ensure the correct address is passed back to the client, on each FTP server specify the external address to be the VIP address.



e.g.

- for Windows 2008 R2 & later use the **External IP address of Firewall** field
- for Linux vsftpd use the directive: **pasv\_address=xxx.xxx.xxx.xxx**
- for Linux ProFTPD use the directive: **MasqueradeAddress=xxx.xxx.xxx.xxx**

#### Windows Example

1. Create a L7 VIP with the following settings changing the name, IP address & passive port range as required:

#### Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	<input type="text" value="FTP-ClusterPASV"/>	?
IP Address	<input type="text" value="192.168.2.150"/>	?
Ports	<input type="text" value="21,50000-50100"/>	?
Protocol		
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

2. Configure the VIP to listen on both the control port (21) and passive range (e.g. 50000-50100) as shown.
3. Click **Update** to create the VIP.
4. Click **Modify** next to the newly created VIP.
5. Scroll down to the section *Other* and click **[Advanced]**.
6. Enable (check) the *Timeout* checkbox then set both *Client Timeout* and *Server Timeout* to **5m**.
7. Define the FTP servers as RIPS for the VIP just created leaving the port field blanks as illustrated below:

#### Layer 7 Add a new Real Server - FTP-ClusterPASV

Label	<input type="text" value="ftp1"/>	?
Real Server IP Address	<input type="text" value="10.10.1.1"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

8. Now restart HAProxy using the WebUI menu option: *Maintenance > Restart Services*.
9. On each FTP server using IIS Manager define the same passive port range and set the external IP address to be the Virtual Server (VIP) address as shown in the example below:



## FTP Firewall Support

The settings on this page let you configure your FTP server to accept passive connections from an external firewall.

Data Channel Port Range:

50000-50100

Example: 5000-6000

External IP Address of Firewall:

192.168.2.180

Example: 10.0.0.1

### Note

The external IP address must be set to be the VIP address, this ensure that this IP address is passed back to the client to use for the subsequent inbound connection.

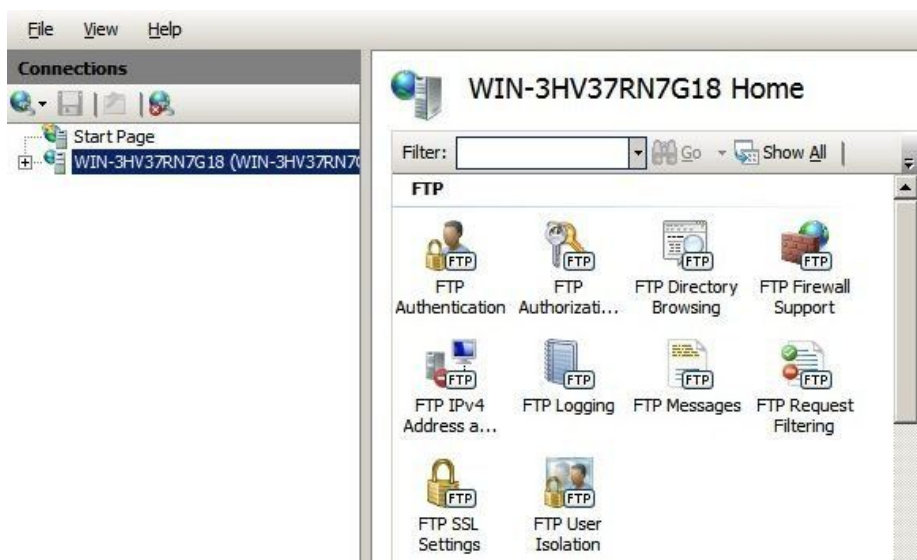
10. If TProxy is enabled, make sure the gateway of each FTP server is set to be an IP on the load balancer (preferably a floating IP to allow failover to the Secondary unit).
11. Now restart both IIS and the Microsoft FTP Service on each FTP server.
12. Passive FTP clients should now be able to connect to the VIP address (192.168.2.180) and view the directory listing successfully.

### Limiting Passive FTP Ports

Limiting passive ports allows your firewall to be more tightly locked down.

### For Windows 2008 R2 & Later

Open the IIS Management console, highlight the server node, then double-click the FTP Firewall Support icon.



The following screen will be displayed:



## FTP Firewall Support

The settings on this page let you configure your FTP server to accept passive connections from an external firewall.

Data Channel Port Range:

Example: 5000-6000

External IP Address of Firewall:

Example: 10.0.0.1

Specify a suitable range, in the example above this is 50000-50100

**Important** | Make sure you restart IIS and the Microsoft FTP Service to apply these settings.

### For Linux

- For **vsftpd**, the following line can be added to the vsftpd.conf file to limit the port range:

```
pasv_max_port - max is 65535
pasv_min_port - min is 1024
```

- For **proftpd**, the following line can be added to the proftpd.conf file to limit the port range:

```
PassivePorts 50000 - 50100
```

- For **pureftpd**, the following startup switch can be used:

```
-p --passiveportrange <min port:max port>
```

# Chapter 11 - Diagnostics & Troubleshooting

## Note

Don't hesitate to contact [support@loadbalancer.org](mailto:support@loadbalancer.org) if you're experiencing any issues or need assistance configuring your appliance.

## The Appliance

If you have are having problems with the appliance in general such as performance, administrative issues etc., the resource utilisation Graphs, administration log and the Apache logs may offer further insight.

### Resource Utilisation

The graphs in the System Overview can be used to give a quick view of how the load balancer is performing. By default, the system Overview includes graphs for **Network Bandwidth**, **System Load Average** and **Memory Usage**. For more information, please refer to [Resource Utilisation Graphs](#).

### Administration Log

The appliance maintains a log of all administration tasks and all notifications. The log can be viewed using the WebUI menu option: *Logs > Load Balancer*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/lbadmin.log`.

### Apache Logs

Apache is used to run the WebUI. The following log files may be useful when diagnosing WebUI related issues:

#### Apache Access Log

This log can be viewed using the WebUI menu option: *Logs > Apache*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/httpd/user_access.log`.

#### Apache Error Log

This log can be viewed using the WebUI menu option: *Logs > Apache Error*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/httpd/error.log`.

## Layer 4 & Layer 7 Virtual Services

Virtual services may not function as expected for a variety of reasons. This section suggests various diagnostic methods that can help identify problems and also looks at potential Virtual Service configuration issues that could be the cause.

### Checking Service State using the System Overview

The system overview can be used to quickly check the state of layer 4 and layer 7 Virtual Services and their associated Real Servers.

## Note

for more information, please also refer to [Real Server Monitoring & Control using the System Overview](#).

The first screenshot shows that the *VIP Web-Cluster* is healthy and that all associated Real Servers are up.

	VIRTUAL SERVICE	IP	PORTS	CONN	PROTOCOL	METHOD	MODE	
	Web-Cluster	192.168.110.235	80	0	HTTP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONN			
	Web1	192.168.110.240	80	100	0	Drain	Halt	
	Web2	192.168.110.241	80	100	0	Drain	Halt	
	Web3	192.168.110.242	80	100	0	Drain	Halt	

The second screenshot shows that the VIP is colored yellow and marked with an exclamation mark indicating that one or more of the Real Servers is not available. The colored tab and the arrow on the left show the current status of each Real Server.

	VIRTUAL SERVICE	IP	PORTS	CONN	PROTOCOL	METHOD	MODE	
	Web-Cluster	192.168.110.235	80	0	HTTP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONN			
	Web1	192.168.110.240	80	100	0	Drain	Halt	
	Web2	192.168.110.241	80	100	0	Online (halt)		
	Web3	192.168.110.242	80	100	0	Drain	Halt	

In the example above:

- *Web-Cluster* is yellow indicating that one or more of the Real Servers in the cluster is down - either due to a failed health check or because it's been manually taken offline (either drained or halted)
- *Web1* is green, this indicates that it's passing it's health check
- *Web2* is blue, this indicates that it has been either Halted or Drained, in this example Halt has been used as indicated by Online (Halt) being displayed. If it had been drained it would show as Online (Drain)
- *Web3* is red, this indicates that it has failed it's health check

In addition to the above colours, Purple/Green is used to indicate that a particular VIP is used for HTTP to HTTPS redirection.

## VIP(s) Fail to appear in the System Overview

If you have configured new VIPs which have not automatically appeared in the System Overview:

- For layer 7 VIPs, have you restarted or reloaded HAProxy since adding the VIP? New VIPs are not displayed until a service reload or restart occurs as shown in the screen shot below:

**Commit changes**

The configuration of the following services has been changed. When reconfiguration is complete, restart/reload the services to commit the changes

[Reload HAProxy](#)
**System Overview**

2023-02-01 15:53:31 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
No Virtual Services configured.						

- Is the corresponding floating IP address active? When a new VIP is configured, a corresponding floating IP is automatically added and brought up.
- Configured Floating IPs can be viewed using the WebUI option: *Cluster Configuration > Floating IPs*.
- If all floating IPs are missing or are down, the System status bar will show both "Active" & "Passive" colored grey, also none of the VIPs will be displayed in the System Overview as shown below:

**System Overview**

2023-02-01 15:56:58 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
-----------------	----	-------	-------	----------	--------	------

- The actual running network configuration can be viewed using the WebUI option: *View Configuration > Network Configuration* as shown below:

```
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether 00:0c:29:62:d3:56 brd ff:ff:ff:ff:ff:ff
    inet 192.168.111.230/18 brd 192.168.127.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet 192.168.110.235/18 brd 192.168.127.255 scope global secondary eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::20c:29ff:fe62:d356/64 scope link
        valid_lft forever preferred_lft forever
```

The above example shows that the interface address (192.168.111.230) and the floating IP address (192.168.110.235) associated with the VIP are both up.

## VIPs & RIPv are Green but Users still Can't Connect

If you've configured your VIPs and RIPv and everything looks fine (green) in the System Overview but users still cannot connect, there are a number of causes for this depending on whether you've configured layer 4 or layer 7 VIPs.

### Layer 7 VIPs

- Have you configured the correct layer 7 protocol? The default protocol for new layer 7 VIPs is HTTP. This is fine for web based traffic typically on port 80, but if you've configured your layer 7 VIP to load balance

something else like RDP on port 3389 or HTTPS on port 443, then you'll need to change the *Layer 7 Protocol* drop-down to TCP.

- Is TProxy enabled? If you've enabled Transparent Proxy, there are certain topology and other requirements that must be met.

Note

For more information about enabling transparency at layer 7, please refer to [Transparency at Layer 7](#).

## Layer 4 VIPs

- Have you complied with the layer 4 network topology requirements? It's important to remember that the health checks performed by the load balancer verify that the *load balancer* can successfully access the server/service/application. This does not verify that each server has been configured correctly to enable *client* access. The sections below explain how the connection state can be used to determine if the Real Servers have been configured correctly, and also what are the configuration requirements for each mode.

### DR Mode:

**Connection State** - Use the WebUI option: *Reports > Layer 4 Current Connections* to view the current traffic in detail. Any packets in state SYN\_RECV can mean that the '**ARP Problem**' has not been correctly solved on the associated Real Server.

**Real Server Configuration Requirements** - For layer 4 DR mode VIPs, the '**ARP Problem**' must be solved on all associated Real Servers - the exact steps required depend on the particular OS. For more information please refer to [DR Mode Considerations](#).

### NAT Mode:

**Connection State** - Use the WebUI option: *Reports > Layer 4 Current Connections* to view the current traffic in detail. Any packets in state SYN\_RECV can mean that the default gateway on the associated Real Server has not been set to be an IP address on the load balancer.

**Real Server Configuration Requirements** - For layer 4 NAT mode VIPs, the default gateway on all associated Real Server must be configured to be an IP address on the load balancer to ensure that client return traffic passes back via the load balancer, for an HA pair, this should be a floating IP address rather than the interface address to allow failover & failback.

## Diagnosing Real Server Issues

If Real Servers are down (red) in the System Overview, this means that the configured health check is failing which can be caused by a variety of reason:

- Verify that the application/service is actually running on each Real Server.
- Is the health check correctly configured and is it appropriate for the Real Servers? The default check for TCP services is a simple port connect - if this has been changed to a negotiate HTTP health check for example, has a valid Request & Response been configured?
- Check that you can ping the Real Server from the load balancer. This can be done using the WebUI option: *Local Configuration > Execute Shell Command* , at the console or via an SSH session, e.g.

```
ping -c 4 192.168.111.240
```

The **-c 4** causes 4 ping attempts, and the command then ends. This is important when running the command from the WebUI to ensure it terminates cleanly.

**Note** | The "Execute Shell Command" menu option is disabled by default. This can be enabled using the WebUI option: *Local Configuration > Security*. Set *Appliance Security Mode* to **Custom** then click **Update**.

**Note** | 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

- Verify that you can connect to the application port from the load balancer. This can be done using telnet at the console or via an SSH session:

```
telnet 192.168.111.240 <application Port>
```

e.g. for a web server listening on port 80:

This example shows that a telnet connection was successfully established:

```
[root@lbmaster ~]# telnet 192.168.110.240 80
Trying 192.168.110.240...
Connected to 192.168.110.240.
Escape character is '^['.
```

This example shows that the telnet connection failed:

```
[root@lbmaster ~]# telnet 192.168.110.240 80
Trying 192.168.110.240...
telnet: connect to address 192.168.110.240: Connection refused
```

**Note** | 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI menu option: *Local Configuration > Security*. You'll need to Set *Appliance Security Mode* to **Custom**, configure the required option(s) and click **Update**.

You can also use nmap on the load balancer to verify that the required Real Server port is open:

```
[root@lbmaster ~]# nmap 192.168.110.241

Starting Nmap 5.51 ( http://nmap.org ) at 2022-07-12 10:54 UTC
Nmap scan report for 192.168.110.240
Host is up (0.0012s latency).
Not shown: 996 filtered ports
PORT      STATE SERVICE
80/tcp    open  http
135/tcp   open  msrpc
443/tcp   open  https
3389/tcp  open  ms-term-serv
MAC Address: 00:50:56:82:0B:D3 (VMware)

Nmap done: 1 IP address (1 host up) scanned in 4.97 seconds
```

Another option is curl:



```
[root@lbmaster ~]# curl http://192.168.110.241
or
[root@lbmaster ~]# curl http://host.mydomain.com
```

For a working web server listening on port 80, the default page is returned.

#### Note

The ping, nmap and curl commands can also be run from the WebUI using the menu option: *Local Configuration > Execute Shell Command*.

- Check if there is a firewall preventing access to the Real Server.

## Requests are Not being Load Balanced as Expected

If the system Overview reports that there is an imbalance in the number of connections. This may be caused by a number of reasons as listed below:

- **Are clients connecting from behind a NAT device** - If this is the case, then all requests will appear to come from the same source IP address. This will be an issue if source IP address persistence is used because all client sessions would be load balanced to the same Real Server.
- **Has a real server been brought back online** - If *Balance Mode* has been set to **Weighed Least Connection**, when the server is brought online it will handle all new requests until balance is reached across all servers. This may occur quickly if lots of new requests are received. If not, the imbalance will continue for longer. If *Balance Mode* has been set to **Weighted Round Robin**, requests will continue to be distributed equally (assuming the weight is the same). If the other Real Servers in the cluster were already processing requests there will again be an imbalance.
- **Are you testing using a single test client** - If persistence is enabled for the VIP then you'll be load balanced to the same Real Server.

## Log File - Layer 4 Virtual Services

The layer 4 log contains all log entries from the layer 4 health checking process `Ldirectord` daemon. This is useful for monitoring when and why real services were down. The log can be viewed using the WebUI menu option: *Logs > Layer 4*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The logging here can be quite verbose but it clearly shows exactly what the health checking process is doing. The log file is located here:  
`/var/log/ldirectord.log`.

### Interpreting the log file

The following log snippet shows Real Server RIP1 failing its healthcheck and then subsequently passing the check. The latest entries are at the top:

```
2022-07-15T10:57:25.78+00:00 lbmaster ldirectord[18650]: Deleted fallback server (127.0.0.1:80), for virtual service VIP1 (192.168.111.152:80)
2022-07-15T10:57:25.77+00:00 lbmaster ldirectord[18650]: Added real server RIP1 (192.168.110.241:80), for virtual service VIP1 (192.168.111.152:80). Weight set to 100
2022-07-15T10:57:25.77+00:00 lbmaster ldirectord[18650]: Resetting soft failure count: 192.168.110.241:80 (tcp:192.168.111.152:80)
2022-07-15T10:57:12.76+00:00 lbmaster ldirectord[18650]: Deleted real server RIP1 (192.168.110.241:80), for virtual service VIP1 (192.168.111.152:80)
2022-07-15T10:57:12.76+00:00 lbmaster ldirectord[18650]: Added fallback server (127.0.0.1:80), for virtual service VIP1 (192.168.111.152:80). Weight set to 1
2022-07-15T10:57:04.75+00:00 lbmaster ldirectord[18650]: Soft failure real server: 192.168.110.241:80 (tcp:192.168.111.152:80) failure 1/2
```

This shows that:

1. Real Server RIP1 failed its first health check at 10:57:04.
2. By 10:57:12 RIP1 had failed its second health check and since this is the default failure threshold and since this was the failure of the last remaining working Real Server, the fallback server was added and RIP1 was

removed.

3. At 10:57:25 RIP1 started passing health checks so was re-added and the fallback server was removed.

## Log File - layer 7 Virtual Services

By default, logging for Layer 7 services is disabled.

*To enable Logging for Layer 7 Services:*

1. Using the WebUI navigate to: *Cluster Configuration > Layer 7 - Advanced Configuration*.
2. Set *Logging* to the required value, the options are:
  - **Off** - No logging enabled.
  - **Service State** - Logs service restarts, reloads and health checks.
  - **Errors** - Logs the service state and Errors.
  - **Everything** - Logs errors and all connections. Note that this logs every connection handled by HAProxy, which \*is likely to rapidly exhaust the log space on a busy Virtual Service.
3. Click **Update**.
4. Reload HAProxy using the Reload button in the "commit changes" message box at the top of the screen.

Once enabled, the log can be viewed using the WebUI menu option: *Logs > Layer 7*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/haproxy.log`.

## Interpreting the log file

The following log snippet shows Real Server RIP1 failing its healthcheck and then subsequently passing the check. The latest entries are at the top:

```
2022-07-15T13:12:38.00+00:00 lbmaster haproxy[30898]: Server VIP2/RIP1 is UP. 1 active and 1 backup servers online. 0 sessions requeued, 0 total in queue.
2022-07-15T13:12:30.99+00:00 lbmaster haproxy[30898]: Server VIP2/RIP1 is UP. 1 active and 1 backup servers online. 0 sessions requeued, 0 total in queue.
2022-07-15T13:12:30.99+00:00 lbmaster haproxy[30898]: Health check for server VIP2/RIP1 succeeded, reason: Layer4 check passed, check duration: 0ms, status: 2/2 UP.
2022-07-15T13:12:26.99+00:00 lbmaster haproxy[30898]: Health check for server VIP2/RIP1 succeeded, reason: Layer4 check passed, check duration: 0ms, status: 1/2 DOWN.
2022-07-15T13:11:18.96+00:00 lbmaster haproxy[30898]: Server VIP2/RIP1 is DOWN. 0 active and 1 backup servers left. Running on backup. 0 sessions active, 0 requeued, 0 remaining in queue.
2022-07-15T13:11:18.96+00:00 lbmaster haproxy[30898]: Health check for server VIP2/RIP1 failed, reason: Layer4 timeout, check duration: 4001ms, status: 0/2 DOWN.
2022-07-15T13:11:10.96+00:00 lbmaster haproxy[30898]: Health check for server VIP2/RIP1 failed, reason: Layer4 timeout, check duration: 4002ms, status: 1/2 UP.
```

This shows that:

1. Real Server RIP1 failed its first health check at 13:11:10.
2. By 13:11:18 RIP1 had failed its second health check and since this is the default failure threshold and since this was the failure of the last remaining working Real Server, the fallback server was activated and RIP1 was removed.
3. At 13:12:26 RIP1 started working and passed its first health check.
4. At 13:12:30 RIP1 passed its second health check (2 passes before being brought online is the default) so was brought back online and the fallback server was deactivated.

The following log snippet shows a client (192.168.65.194) making 3 HTTP GET requests to VIP1/RIP1. The latest entries are at the top:

```
2022-07-15T13:31:48.66+00:00 lbmaster haproxy[30898]: 192.168.65.194:53943 [15/Jul/2022:13:31:48.658] VIP2 VIP2/RIP1 0/0/0/1/1 404 1406 - - --NI 2/2/0/0/0 0/0 "GET /favicon.ico HTTP/1.1"
2022-07-15T13:31:48.64+00:00 lbmaster haproxy[30898]: 192.168.65.194:53943 [15/Jul/2022:13:31:48.623] VIP2 VIP2/RIP1 1/0/0/0/18 200 99947 - - --NI 2/2/0/0/0 0/0 "GET /iisstart.png HTTP/1.1"
2022-07-15T13:31:48.60+00:00 lbmaster haproxy[30898]: 192.168.65.194:53943 [15/Jul/2022:13:31:48.605] VIP2 VIP2/RIP1 0/0/0/1/1 200 943 - - --NI 1/1/0/0/0 0/0 "GET / HTTP/1.1"
```

HAProxy HTTP logs are divided into 16 fields. The following table breaks down the first line from the snippet above:

Field	Format	Extract from the example above
1	process_name '[' pid ']:'	haproxy[30898]:
2	client_ip ':' client_port	192.168.65.194:53943
3	'[' request_date ']'	[15/Jul/2022:13:31:48.658]
4	frontend_name	VIP2
5	backend_name '/' server_name	VIP2/RIP1
6	TR '/' Tw '/' Tc '/' Tr '/' Ta*	0/0/0/1/1
7	status_code	200
8	bytes_read*	943
9	captured_request_cookie	-
10	captured_response_cookie	-
11	termination_state	--NI
12	actconn '/' feconn '/' beconn '/' srv_conn '/' retries*	2/2/0/0/0
13	srv_queue '/' backend_queue	0/0
14	'{' captured_request_headers* '}'	n/a for this example
15	'{' captured_response_headers* '}'	n/a for this example
16	"" http_request ""	"GET /favicon.ico HTTP/1.1"

Note

For a detailed description of each field in a log entry for an HTTP request, please refer to [HTTP Log Format](#).

Note

TCP logs are similar but there are only 10 fields. For more information please refer to [TCP Log Format](#).

Note

For full details of all HAProxy logging features, please refer to [HAProxy Logging](#).

## SSL Termination - Pound

By default, Pound logging is disabled.

*To enable Logging for Pound:*

1. Using the WebUI navigate to: *Cluster Configuration > SSL - Advanced Configuration*.
2. Under the *Pound Global Settings* section set *Logging* to **Yes**.
3. Click **Update**.

Once enabled, the log can be viewed using the WebUI menu option: *Logs > SSL Termination (Pound)*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/pound.log`.

## SSL Termination - STunnel

By default, STunnel logging is set to **Emergency(0)**.

To change the Logging level for STunnel:

1. Using the WebUI navigate to: *Cluster Configuration > SSL - Advanced Configuration*.
2. Under the *STunnel Global Settings* section, set the *Debug Level* to the required value, the options are:

Value	Severity	Description
0	Emergency	System is unusable
1	Alert	Action must be taken immediately
2	Critical	Critical conditions
3	Error	Error conditions
4	Warning	Warning conditions
5	Notice	Normal but significant conditions
6	Informational	Informational messages
7	Debug	Debug-level messages

**Note**

Log Entries for the selected level plus all the levels below will be included. For example, if level 4 is selected, levels 0, 1, 2 and 3 will also be logged.

3. Click **Update**.

Once enabled, the log can be viewed using the WebUI menu option: *Logs > SSL Termination (STunnel)*. The log can be searched using the search box at the top of the screen or can be downloaded for external analysis. The log can be refreshed using the **Check Status** button. The log file is located here: `/var/log/stunnel.log`.

## WAF

For information on accessing and interpreting WAF logs, please refer to [WAF Gateway Error Logs](#).

## GSLB

For information on accessing and interpreting the GSLB reports, please refer to [GSLB Diagnostics](#).

## Heartbeat

For information on diagnosing Heartbeat issues and accessing the Heartbeat log, please refer to [Clustered Pair Diagnostics](#).

# Chapter 12 - Monitoring & Reporting

## Reports

All reports can be accessed using the *Reports* option in the WebUI.

### Layer 4 Status

This report shows the current weight and number of active & inactive connections for each Real Server. If a Real Server has failed a health check, it will not be listed. Use the *Logs > Layer 4* option to view the Ldirectord log file if expected servers are missing.

#### Layer 4 Status

Check Status

Virtual Service	Real Server	Forwarding Method	Weight	Active Connections	Inactive Connections
<b>HTTP-Cluster1</b> 192.168.110.120 port 192.168.110.120/tcp	<b>RIP1</b> 192.168.110.240	Route	100	0	0
	<b>RIP2</b> 192.168.110.241	Route	100	0	0
	<b>RIP3</b> 192.168.110.242				

```
IP Virtual Server version 1.2.1 (size=32768)
Prot LocalAddress:Port Scheduler Flags
  -> RemoteAddress:Port      Forward Weight ActiveConn InActConn
TCP  192.168.110.120:80 wlc persistent 300
  -> 192.168.110.240:80      Route  100    0        0
  -> 192.168.110.241:80      Route  100    0        0
```

In the example above, the details for RIP3 are not displayed because it's failing its health checks.

The **Check Status** button can be used to refresh the display.

### Layer 4 Traffic Rate

This report shows the current connections per second and bytes per second to each Real Server. If a Real Server has failed a health check, it will not be listed.

Check Status

Virtual Service	Real Server	Connections / s	Incoming Packets / s	Outgoing Packets / s	Incoming Bytes / s	Outgoing Bytes / s
<b>HTTP-Cluster1</b> 192.168.110.120 port 192.168.110.120/tcp		0	0	0	0	0
	<b>RIP1</b> 192.168.110.240	0	0	0	0	0
	<b>RIP2</b> 192.168.110.241	0	0	0	0	0
	<b>RIP3</b> 192.168.110.242					

IP Virtual Server version 1.2.1 (size=32768)

Prot	LocalAddress:Port	CPS	InPPS	OutPPS	InBPS	OutBPS
	-> RemoteAddress:Port					
TCP	192.168.110.120:80	0	0	0	0	0
	-> 192.168.110.240:80	0	0	0	0	0
	-> 192.168.110.241:80	0	0	0	0	0

In the example above, the details for RIP3 are not displayed because it's failing its health checks.

The **Check Status** button can be used to refresh the display.

## Layer 4 traffic Counters

This report shows the volume of traffic to each Real Server since the counters were last re-set. If a Real Server has failed a health check, it will not be listed.

[Check Status](#)[Reset Counters](#)

Virtual Service	Real Server	Connections	Incoming Packets	Outgoing Packets	Incoming Bytes	Outgoing Bytes
<b>HTTP-Cluster1</b> 192.168.110.120 port 192.168.110.120/tcp		0	0	0	0	0
	<b>RIP1</b> 192.168.110.240	0	0	0	0	0
	<b>RIP2</b> 192.168.110.241	0	0	0	0	0
	<b>RIP3</b> 192.168.110.242					

IP Virtual Server version 1.2.1 (size=32768)

Prot	LocalAddress:Port	Conns	InPkts	OutPkts	InBytes	OutBytes
	-> RemoteAddress:Port					
TCP	192.168.110.120:80	0	0	0	0	0
	-> 192.168.110.240:80	0	0	0	0	0
	-> 192.168.110.241:80	0	0	0	0	0

#### Note

Direct Routing is the default load balancing method and you will not see any stats for return packets as shown above (as they do not pass through the load balancer). They will be seen for NAT mode since return traffic does pass back via the load balancer.

In the example above, the details for RIP3 are not displayed because it's failing its health checks.

The **Check Status** button can be used to refresh the display.

## Layer 4 Current Connections

The current connections report is very useful for diagnosing issues with routing or ARP related problems. In the example below, the state is shown as **SYN\_RECV**. For layer 4 DR mode this is normally a good indication that the **ARP Problem** has not been solved. For NAT mode, this is a good indication that the Real Server's default gateway has not been configured to be the load balancer and therefore return traffic is not routed correctly.

## Check Status

### IPVS connection entries

pro	expire	state	source	virtual	destination
TCP	04:44	NONE	192.168.64.7:0	192.168.110.120:80	192.168.110.241:80
TCP	00:49	SYN_RECV	192.168.64.7:28808	192.168.110.120:80	192.168.110.241:80
TCP	00:49	SYN_RECV	192.168.64.7:28809	192.168.110.120:80	192.168.110.241:80

#### Note

The IPVS connection entries in state **NONE** represent the persistence related entries for client connections, and are not actual client connections. These only appear when persistence is enabled.

The **Check Status** button can be used to refresh the display.

## Layer 4 Current Connections (Resolve Hostnames)

This is the same as the current connections report but is slower as it looks up the DNS name of each IP address.

## Layer 7 Status

This report contains the current status of all configured layer 7 HAProxy Virtual Servers and Real Servers.

### HAProxy

#### Statistics Report for pid 7370

##### > General process information

pid = 7370 (process #1, nbproc = 1, nbthread = 1)  
uptime = 0d 0h00m34s  
system limits: memmax = unlimited; ulimit-n = 80036  
maxsock = 80036; maxconn = 40000; maxpipes = 0  
current conns = 1; current pipes = 0/0; conn rate = 1/sec; bit rate = 0.000 kbps  
Running tasks: 1/16; idle = 100 %

active UP  
active UP, going down  
active DOWN, going up  
active or backup DOWN  
active or backup DOWN for maintenance (MAINT)  
active or backup SOFT STOPPED for maintenance  
Note: "NOLB"/"DRAIN" = UP with load-balancing disabled.

##### Display option:

- Scope :
- Hide DOWN servers
- Disable refresh
- Refresh now
- CSV export

##### External resources:

- Primary site
- Updates (v2.0)
- Online manual

L7-HTTP																							
Queue				Session rate				Sessions				Bytes				Denied				Errors			
Cur	Max	Limit		Cur	Max	Limit		Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status
Frontend				0	0	-		0	0	40 000	0			0	0	0	0	0	0	0	0	0	OPEN
backup	0	0	-	0	0			0	0	-	0		?	0	0	0	0	0	0	0	0	0	no check
rip1	0	0	-	0	0			0	0	-	0		?	0	0	0	0	0	0	0	0	0	34s UP
Backend	0	0		0	0			0	0	4 000	0		?	0	0	0	0	0	0	0	0	0	34s UP

:prometheus-exporter																							
Queue				Session rate				Sessions				Bytes				Denied				Errors			
Cur	Max	Limit		Cur	Max	Limit		Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status
Frontend				0	0	-		0	0	40 000	0			0	0	0	0	0	0	0	0	0	OPEN
Backend	0	0		0	0			0	0	4 000	0		?	0	0	0	0	0	0	0	0	0	34s UP

:stats																							
Queue				Session rate				Sessions				Bytes				Denied				Errors			
Cur	Max	Limit		Cur	Max	Limit		Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status
Frontend				0	0	-		0	0	40 000	0			0	0	0	0	0	0	0	0	0	OPEN
Backend	0	0		0	0			0	0	4 000	0		?	0	0	0	0	0	0	0	0	0	34s UP

:admin_stats																							
Queue				Session rate				Sessions				Bytes				Denied				Errors			
Cur	Max	Limit		Cur	Max	Limit		Cur	Max	Limit	Total	LbTot	Last	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status
Frontend				1	1	-		1	1	40 000	2			974	29 143	0	0	0	0	0	0	0	OPEN
Backend	0	0		0	0			0	0	4 000	0		0s	974	29 143	0	0	0	0	0	0	0	34s UP

#### Note

From v8.8.1, all statistics are automatically refreshed every 30 seconds.

## Layer 7 Stick Table

Displays the layer 7 stick tables. For example, if a layer 7 VIP is created using RDP cookie persistence, a stick table will be used. The related VIP is then available in the drop-down as shown below:



## Reports > Stick Table (HAProxy)

L7-HTTP ? Refresh Clear Table

1 Entries Returned

ID	Key	Use	Expires	Server	Remove
0x1741310	192.168.85.1	use=0	1797362	rip1	

Page 1 of 1

Prev Next

### Notes

1. Stick tables are used when either source IP persistence or RDP cookie persistence is used with layer 7 Virtual Services.
2. Individual stick table entries can be removed by clicking the red 'X' in the remove column, the whole table can be cleared by clicking the **Clear Table** button.

### GSLB Reports

For details please refer to [GSLB Diagnostics](#).

## Graphs

Graphs are provided for appliance resource utilisation and for layer 4 and layer 7 Virtual Services / Real Servers. All graphs can be accessed using the *Reports > Graphing* option in the WebUI.

### Resource Utilisation Graphs

Resource utilisation graphs are available for Network Throughput, System Load Average, Memory usage and Disk Usage.

**Network Throughput** - This graph shows the Network throughput in bytes/second.

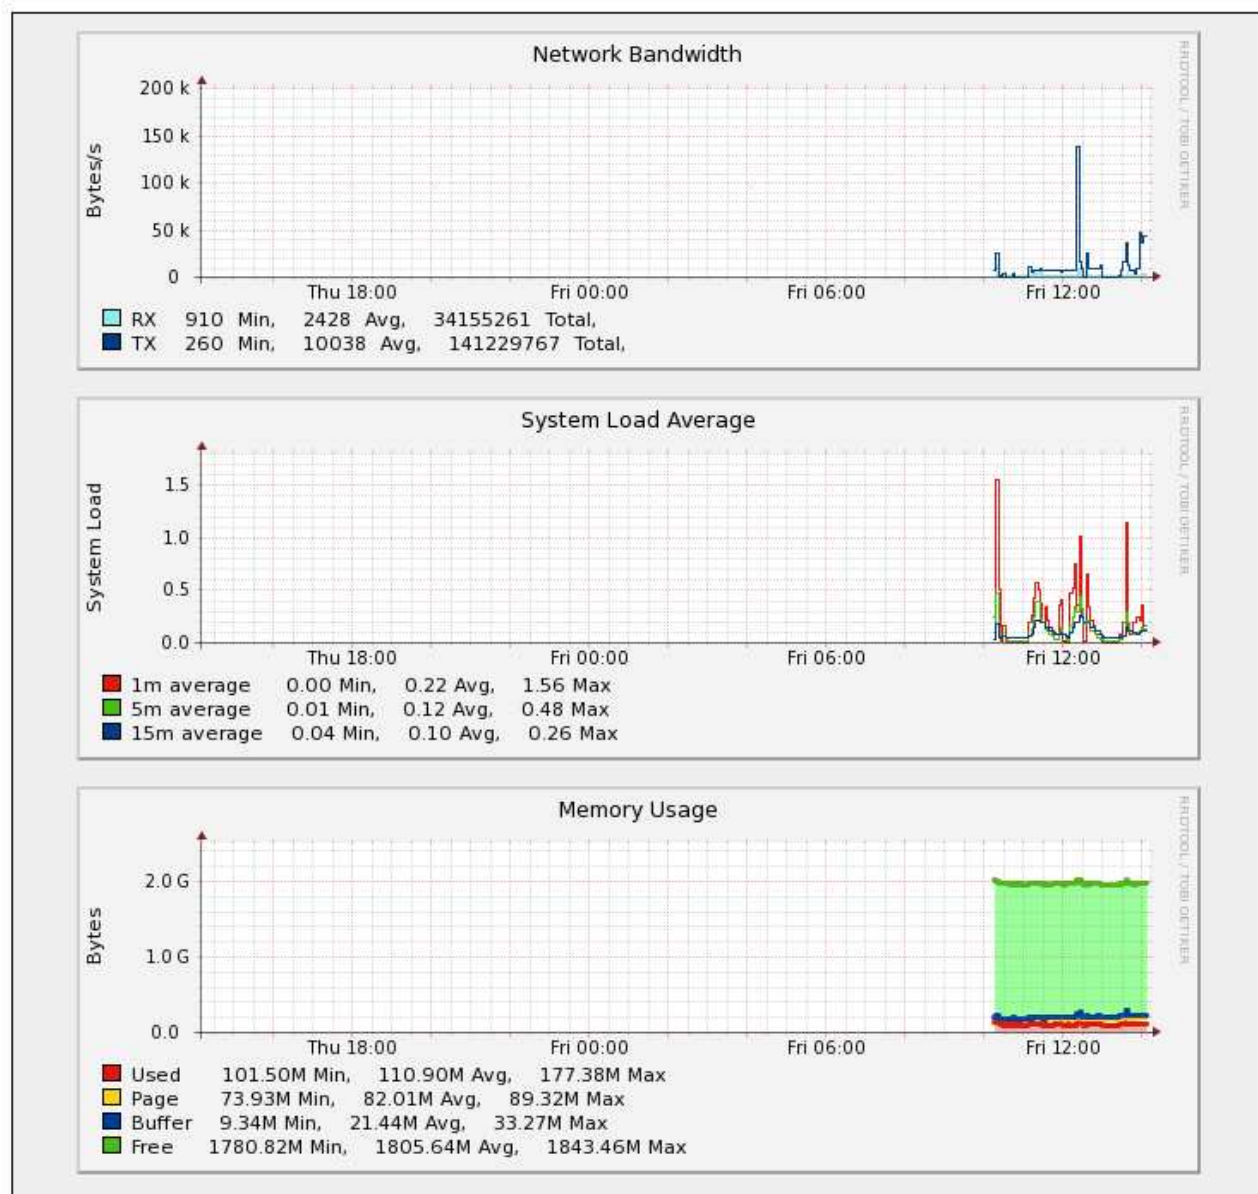
**System Load Average** - The system load average graph is a very useful way to determine if the appliance is overloaded. To understand how to interpret the value, a bridge analogy can be used:

- A value of 0 means there's no traffic on the bridge at all. In fact, between 0 and 1 means there's no congestion, and an arriving car will just drive straight onto the bridge.
- A value of 1 means the bridge is exactly at capacity. All is still good, but if traffic gets a little heavier, things are going to build up and getting accross the bridge will slow down.
- A value over 1 means there's congestion: A value of 2 means that there are two lanes of cars; one lane on the bridge and one lane waiting. A value of 3 means there are three lanes of cars; one lane on the bridge and two lanes waiting, etc.
- Using a single core with the load average somewhere around 0.75 - 1 is great. Using a quad core with the load average somewhere around 3.0 - 4.0 would also be fine.

**Memory Usage** - This graph shows the appliance memory usage in GB.

**Disk Usage** - This graph shows the appliance disk usage in GB.

By default the Network Throughput, System Load Average and Memory Usage graphs are displayed in the System Overview. These graphs can also be accessed using the WebUI menu option: *Reports > Graphing* then selecting from the 4 options at the bottom of the drop-down.



Note

These graphs can be disabled/hidden if preferred using the WebUI menu option: *Local Configuration > Graphing*.

Note

The disk usage graph is not displayed in the System Overview, it can be accessed using the WebUI option: *Reports > Graphing* and selecting **Disk Usage** in the drop-down.




## Layer 4 and Layer 7 Virtual Services/Real Servers

Graphs are automatically configured when new Virtual and Real Servers are defined.

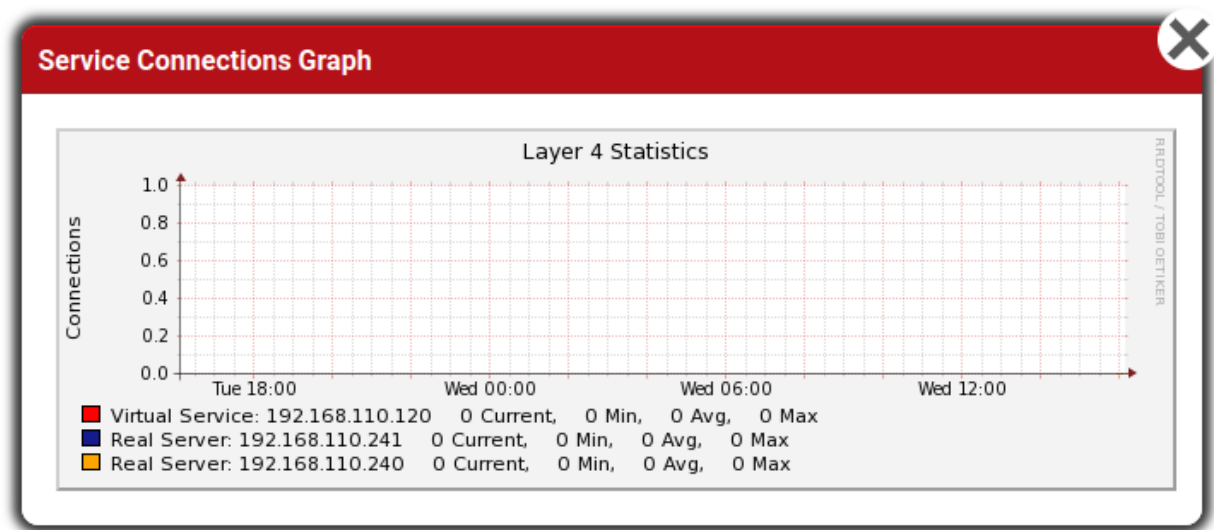
Graphs for the configured Virtual & Real Servers can be accessed either from the System Overview using the appropriate graph icon that appears next to each VIP and RIP or from the drop-down available in the WebUI under *Reports > Graphing*.

Using the System Overview:

The graph is displayed by clicking the relevant graph icon that's displayed next to each VIP/RIP:

↑	HTTP-Cluster1	192.168.110.120	80	0	TCP	Layer 4	DR	
	REAL SERVER	IP	PORTS	WEIGHT	CONN			
↑	RIP1	192.168.110.240	80	100	0	Drain	Halt	
↑	RIP2	192.168.110.241	80	100	0	Drain	Halt	

When this method is used, the daily Service Connections Graph (i.e. the last 24 hrs) is displayed for the particular VIP or RIP:



Clicking anywhere within this graph opens the complete list of graphs for the VIP/RIP in question. This is the same as selecting the VIP/RIP in the *Reports > Graphing* menu options as explained below.

Using the WebUI menu option: *Reports > Graphing*:

When selected, a drop-down similar to the following is displayed:

**Graphing**

Select an Option ▼

Select an Option

HTTP-Cluster1

RIP1

RIP2

Network Throughput

Load Average

Memory Usage

Disk Usage

Note

As VIPs & RIPs are added or removed, these are also automatically added/removed from the

drop-down list.






**Note**

For more information on using the System Overview please refer to [Chapter 8 - Real Server Health Monitoring & Control](#).

When selected in this way, a complete list of graphs is displayed for the VIP/RIP selected as shown below:

## Graphing

HTTP-Cluster1



The following graphs are displayed for each VIP or RIP selected:

- 5 x **Connection graphs** : Hourly, daily, weekly, monthly and yearly
- 5 x **Throughput graphs** : Hourly, daily, weekly, monthly and yearly

## Graphing Options

The graphing sub-system can be customized to suit your requirements.

*To change the Graphing Options:*

1. Using the WebUI, navigate to: *Local Configuration > Graphing*.

Layer 4	On ▼	?
Layer 7	On ▼	?
Interfaces	On ▼	?
Load Average	On ▼	?
Memory	On ▼	?
Disk Usage	On ▼	?

Advanced Configuration

Interval	10	?
Timeout	2	?
Threads	6	?
Logging	Off ▼	?

[Update](#)

2. Set the required option for each graphing category (either enabled (default), disabled or deleted).
3. Set the data collector *Interval* for the data collector in seconds. Change

**Note** | This is a global value and will effect all collectors. Do not change this value unless advised to do so by support.

**Warning** | Changing this value will reset the graphing databases. All the current data will be lost.

4. Set the data collector *Timeout* for the data collector in seconds.

**Note** | Do not change unless advised to do so by support.

5. Set the number of *Threads* for the data collector process.

**Note** | Do not change unless advised to do so by support.

6. Configure data collector logging for collectd.

**Note** | This is incredibly verbose and should only be enabled for debugging purposes.

## Layer 7 (HAProxy) Prometheus Exporter

v8.7.0 adds Prometheus support for layer 7 services. The Prometheus Exporter consumes the HAProxy stats page and converts the data to the [Prometheus time series](#).

*To Enable the Prometheus Exporter:*

1. Using the WebUI navigate to: *Cluster Configuration > Layer 7 - Advanced Configuration*.
2. Scroll to the bottom of the page and enable (check) the *Enable Prometheus Exporter* checkbox.
3. Reload HAProxy using the button in the "commit changes" message box at the top of the screen.

Once enabled, the Prometheus endpoint will be available here:

**`https://<appliance-ip-address>:9443/lbadmin/stats/l7prometheus/`**

You can authenticate using the credentials of any valid WebUI user account.

**Note** | For more information, please refer to our [Prometheus Blog](#).

## Grafana

Grafana is a tool that allows you to visualise data from sources such as Prometheus. For information on installing and configuring Grafana and setting up a Prometheus data source on the appliance, please refer to our [Grafana Blog](#).

## SNMP Reporting

By default, SNMP is disabled on the appliance. For information on enabling SNMP and configuring basic SNMP settings please refer to [SNMP Configuration](#).

### MIB Files

The appliance has a number of Management Information Base files (MIBs) available for use with your monitoring system. The appliance includes all the standard Linux MIBs and also a number of custom MIB that enable layer 4 and layer 7 services to be monitored. All MIB files are located in `/usr/share/snmp/mibs` on the appliance. The custom MIBs are also available for download using the links shown below.

**General** - For network stats, memory usage, CPU load etc, standard Linux MIBs can be used.

**Layer 4 Services** - For monitoring layer 4 services, [OC-MIB.txt](#) & [LVS-MIB.txt](#) are provided.

**Layer 7 Services** - for monitoring layer 7 services, [LBO-MIB.txt](#), [L7STAT-EXPERIMENTAL.txt](#) & [L7INFO-EXPERIMENTAL-MIB.txt](#) are provided.

**Note** | `/etc/snmp/snmp.conf` includes the directive `mibs +ALL` which means that all available MIBs are loaded.

The following 2 sections include various examples to help demonstrate how SNMP values for Layer 4 and Layer 7



services can be retrieved at the console or via an SSH session.

## SNMP for Layer 4 Services

The root OID for Layer 4 services is: **1.3.6.1.4.1.8225.4711**.

### Note

Lots of information is available, the layer 4 MIB file includes comprehensive descriptions of all OIDs.

View information for layer 4 services starting at the root OID:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.8225.4711

LVS-MIB::lvsVersion.0 = STRING: "1.2.1"
LVS-MIB::lvsNumServices.0 = INTEGER: 0
LVS-MIB::lvsHashTableSize.0 = INTEGER: 32768
...
etc.
```

The same can be achieved using the MIB object name *lvs*:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost lvs
```

## Monitoring Layer 4 VIPs & RIPs using SNMP

### Accessing VIP Information

List all VIPs using the OID for layer 4 VIPs:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.8225.4711.17.1.4

LVS-MIB::lvsServiceAddr.1 = IPAddress: 192.168.111.223
LVS-MIB::lvsServiceAddr.2 = IPAddress: 192.168.111.222
```

The same can be achieved using the MIB object name *lvsServiceAddr*:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost lvsServiceAddr
```

Display the *lvsService* table:

```
[root@lbmaster ~]# snmptable -c public -v 2c localhost lvsService
```

lvsServiceNum...	lvsServiceSch...	lvsServiceProto	lvsServiceAddr	lvsServicePort	lvsServiceFW...	lvsServicePers...	lvsSe
1	wlc	tcp	192.168.111.223	3389	undefined	300	255.0
2	wlc	tcp	192.168.111.222	80	undefined	300	255.0

### Accessing RIP Information

List the Real Servers that are passing their health check using the OID for layer 4 RIPs:



```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.8225.4711.18.1.3
LVS-MIB::lvsRealServerAddr.1.1 = IPAddress: 192.168.110.240
LVS-MIB::lvsRealServerAddr.1.2 = IPAddress: 192.168.110.243
LVS-MIB::lvsRealServerAddr.2.1 = IPAddress: 192.168.110.243
LVS-MIB::lvsRealServerAddr.2.2 = IPAddress: 192.168.110.240
```

**Note** | For layer 4 services, if the health check fails, the failed server will be omitted from the list.

The same can be achieved using the MIB object name *lvsRealServerAddr*:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost lvsRealServerAddr
```

Display the lvsReal table:

```
[root@lbmaster ~]# snmptable -c public -v 2c localhost lvsReal
```

lvsRealService...	lvsRealServer...	lvsRealServer...	lvsRealServer...	lvsRealServer...	lvsRealServer...	lvsRealStats...
1	1	192.168.110.240	3389	route	100	0
1	2	192.168.110.243	3389	route	100	0
2	1	192.168.110.243	80	route	100	0
2	2	192.168.110.240	80	route	100	0

## SNMP for Layer 7 Services

The root OID for Layer 7 services is: **1.3.6.1.4.1.54849**.

**Note** | Lots of information is available, the layer 7 MIB files include comprehensive descriptions of all OIDs.

View information for layer 7 services starting at the root OID:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.54849
L7INFO-EXPERIMENTAL-MIB::l7InfoLastPolled.0 = Timeticks: (141790592) 16 days, 9:51:45.92
L7INFO-EXPERIMENTAL-MIB::l7InfoUpdateInterval.0 = INTEGER: 4000
L7INFO-EXPERIMENTAL-MIB::l7InfoLastError.0 = Gauge32: 0
L7INFO-EXPERIMENTAL-MIB::l7InfoName.0 = STRING: HAProxy
L7INFO-EXPERIMENTAL-MIB::l7InfoVersion.0 = STRING: 1.8.25
...
etc.
```

The same can be achieved using the MIB object name:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost loadBalancerOrg
```

Display layer 7 HAProxy process information:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost L7info
```

## Note

This gives a similar output to the familiar HAProxy 'show info' command. for more information on using the Linux socat command to run the 'show info' command please refer to [Using Linux socket commands to configure Layer 7 Services](#).

## Monitoring Layer 7 VIPs & RIPs using SNMP

### Accessing VIP Information

List all VIPs using the OID for layer 7 VIPs:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.54849.1.2.5.1.2
L7STAT-EXPERIMENTAL-MIB::l7FePxname.4 = STRING: stats
L7STAT-EXPERIMENTAL-MIB::l7FePxname.493413195 = STRING: VIP4
L7STAT-EXPERIMENTAL-MIB::l7FePxname.1593140907 = STRING: VIP3
```

The same can be achieved using the MIB object name 'l7FePxname':

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost l7FePxname
```

Display the L7Frontend table:

```
[root@lbmaster ~]# snmptable -c public -v 2c localhost l7Frontend
```

l7FeId	l7FePxname	l7FeStatus	l7FeMode	l7FeConnRate	l7FeConnRate...	l7FeConn
4	stats	no1b	http	0	2	3
493413195	VIP4	no1b	tcp	0	0	0
1593140907	VIP3	no1b	http	0	0	0

Display the L7Backend table:

```
[root@lbmaster ~]# snmptable -c public -v 2c localhost l7Backend
```

l7BeId	l7BePxname	l7BeStatus	l7BeWeight	l7BeMode	l7BeCookie	l7BeAlgo
4	stats	up	0	http		roundRobin
493413195	VIP4	up	100	tcp		leastConnection
1593140907	VIP3	up	100	http	SERVERID	leastConnection

### Accessing RIP Information

List all RIPs using the OID for layer 7 RIPs:

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost 1.3.6.1.4.1.54849.1.2.7.1.4
L7STAT-EXPERIMENTAL-MIB::l7SrvSvname.493413195.1 = STRING: backup
L7STAT-EXPERIMENTAL-MIB::l7SrvSvname.493413195.551991902 = STRING: RDP2
L7STAT-EXPERIMENTAL-MIB::l7SrvSvname.493413195.1899219725 = STRING: RDP1
L7STAT-EXPERIMENTAL-MIB::l7SrvSvname.1593140907.1 = STRING: backup
L7STAT-EXPERIMENTAL-MIB::l7SrvSvname.1593140907.1269887581 = STRING: Web2
```

```
L7STAT-EXPERIMENTAL-MIB::l7SrvSvname.1593140907.1749728569 = STRING: Web1
```

The same can be achieved using the MIB object name 'l7SrvSvname':

```
[root@lbmaster ~]# snmpwalk -v 2c -c public localhost l7SrvSvname
```

Display the L7Server table:

```
[root@lbmaster ~]# snmptable -c public -v 2c localhost L7Server
```

l7SrvId	l7SrvSid	l7SrvPxname	l7SrvSvname	l7SrvAddressT...	l7SrvAddress...	l7SrvAddressI
493413195	1	VIP4	backup	ipv4	127.0.0.1	9081
493413195	551991902	VIP4	RDP2	ipv4	-64.-88.110.-13	3389
493413195	1899219725	VIP4	RDP1	ipv4	-64.-88.110.-16	3389
1593140907	1	VIP3	backup	ipv4	127.0.0.1	9081
1593140907	1269887581	VIP3	Web2	ipv4	-64.-88.110.-13	80
1593140907	1749728569	VIP3	Web1	ipv4	-64.-88.110.-16	80

### SNMPv3

For SNMPv3, the SNMP walk command format is:

```
snmpwalk -v3 -u <USERNAME> -l authNoPriv -a MD5 -A <PASS-PHRASE> -m LVS-MIB localhost  
1.3.6.1.4.1.8225.4711
```

e.g.

```
snmpwalk -v3 -u snmpv3user -l authNoPriv -a MD5 -A sNmPv31864zX -m LVS-MIB localhost  
1.3.6.1.4.1.8225.4711
```

## Chapter 13 - Useful Tools & Utilities

### Useful Diagnostics Tools

Full root access to the appliance is supported which enables many useful commands to be run directly at the console or via an SSH session. Many commands can also be run using the WebUI menu option: *Local Configuration > Execute Shell Command*. Several commonly used examples are listed below.

**Note** The "Execute Shell Command" menu option is disabled by default. This can be enabled using the WebUI option: *Local Configuration > Security. Set Appliance Security Mode to Custom* then click **Update**.

**Note** 'root' user console and SSH password access are disabled by default. These options can be enabled using the WebUI option: *Local Configuration > Security*. You'll need to *Set Appliance Security Mode to Custom*, enable the required option(s) and click **Update**.

### Netstat

Print network connections, routing tables, interface statistics, masquerade connections, and multicast memberships. Useful to check that services are listening on the correct IP/port.

```
e.g. netstat -anp
```

Command Output:

```
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address           Foreign Address         State       PID/Program name
tcp        0      0 127.0.0.1:7778          0.0.0.0:*                 LISTEN      7218/haproxy
tcp        0      0 192.168.100.238:80      0.0.0.0:*                 LISTEN      7218/haproxy
tcp        0      0 0.0.0.0:22              0.0.0.0:*                 LISTEN      9638/sshd
tcp        0      0 0.0.0.0:9081            0.0.0.0:*                 LISTEN      5136/nginx
tcp        0  336 192.168.100.237:22      10.10.0.26:61430        ESTABLISHED 9801/sshd
tcp        0      0 :::9443                  :::*                     LISTEN      732/httpd
tcp        0      0 :::22                   :::*                     LISTEN      9638/sshd
tcp        0      0 :::9080                  :::*                     LISTEN      732/httpd
udp        0      0 192.168.100.238:123     0.0.0.0:*                 1650/ntpd
udp        0      0 192.168.100.237:123     0.0.0.0:*                 1650/ntpd
udp        0      0 127.0.0.1:123           0.0.0.0:*                 1650/ntpd
udp        0      0 0.0.0.0:123             0.0.0.0:*                 1650/ntpd
udp        0      0 0.0.0.0:161             0.0.0.0:*                 10089/snmpd
etc.
```

### Telnet

The telnet command is used to communicate with another host using the TELNET protocol. It's very useful for testing that a connection to a specific port can be made. Note that this command should be run from the console or a terminal session rather than via the WebUI.

```
e.g. telnet 192.168.100.10 80
```

In this example, 192.168.100.10 is a Real Server, the command is useful to ensure that the load balancer is able to successfully connect to this server on port 80.

```
[root@lbmaster ~]# telnet 192.168.100.10 80
Trying 192.168.100.10...
Connected to 192.168.100.10.
```

Escape character is '^['.

## Tcpdump

Tcpdump enables network traffic to be dumped to a file for analysis. Filters can also be applied if required to select which traffic is captured. Very useful tool when diagnosing network issues. Note that this command should be run from the console or a terminal session rather than via the WebUI.

```
e.g. tcpdump -i any -s 0 -w tcpdump-file.pcap
```

This command captures all network traffic on all interfaces using the maximum packet size of 65535 bytes and dumps it to a file called tcpdump-file.pcap. To end the capture use CTRL+C.

Our support department may ask you to run this command and send the resulting output file to help them diagnose certain network issues.

## Ethtool

Ethtool is used for querying settings of an Ethernet device and changing them.

```
e.g. ethtool eth0
```

Command output:

```
Settings for eth0:
  Supported ports: [ TP ]
  Supported link modes:   1000baseT/Full
                        10000baseT/Full
  Supports auto-negotiation: No
  Advertised link modes:  Not reported
  Advertised pause frame use: No
  Advertised auto-negotiation: No
  Speed: 10000Mb/s
  Duplex: Full
  Port: Twisted Pair
  PHYAD: 0
  Transceiver: internal
  Auto-negotiation: off
  MDI-X: Unknown
  Supports Wake-on: uag
  Wake-on: d
  Link detected: yes
```

## NMAP

Nmap (Network Mapper) can be used to scan a range of hosts or a single host to determine which ports are open and which services are listening on those ports.

```
e.g. nmap 192.168.110.241
```

Command output:

```
Starting Nmap 5.51 ( http://nmap.org ) at 2022-03-15 14:54 UTC
Nmap scan report for 192.168.110.241
Host is up (0.0010s latency).
Not shown: 996 filtered ports
```

```
PORT      STATE SERVICE
80/tcp    open  http
135/tcp   open  msrpc
443/tcp   open  https
3389/tcp  open  ms-term-serv
MAC Address: 00:50:56:82:0B:D3 (VMware)
```

```
Nmap done: 1 IP address (1 host up) scanned in 4.53 seconds
```

## Wireshark

Wireshark is an open source application that can be used to analyze tcpdump output files. It can be downloaded from [here](#).

## Windows Specific Tools

### Microsoft Network Monitor

Network Monitor is a simpler alternative to Wireshark that has some nice features. It can be downloaded from [here](#).

### WinSCP

WinSCP is an open source application that allows files to be uploaded/downloaded to/from the load balancer using Windows. It can be downloaded from [here](#).

### PuTTY

PuTTY is an open source SSH client for Windows. It can be downloaded from [here](#).

# Chapter 14 - Backup & Restore and Disaster Recovery

## Backup & Restore

The WebUI can be used to perform backup and restore functions. From v8.7.0 all required configuration files can be backed up and restored from a single archive file.

### Backup

When a backup is performed, a zipped archive is created that includes the following:

- XML Configuration File
- SSL Certificates
- Manual Configuration
- SSH Keys
- Firewall Scripts
- PBR Rules
- Authentication Settings
- LBCLI Password File
- Fallback Page
- User Names and Passwords
- License Key

*To perform a backup:*

1. Using the WebUI navigate to: *Maintenance > Backup & Restore*.
2. Under the **Backup** tab, enter a password to encrypt sensitive files.

#### Note

SSL certificates and SSH keys are only backed up if a password is set. If the password is not set, the backup will exclude these files.

3. Click **Backup**.
4. The zip file will be created and downloaded to your browser's default location or you'll be prompted to select a location depending on the settings configured.

### Local XML Backup

It's also possible to create a local XML backup. This can be useful for example during a PoC where settings are being changed and you'd like to keep a quick configuration backup before trying new settings.

#### Note

The **local XML backup** only includes the main system configuration file (lb\_config.xml). If you want to create backups that can be used for a full system restore, use the backup option described above.

*To perform a local XML backup:*

1. Using the WebUI navigate to: *Maintenance > Backup & Restore*.
2. In the **Create Local XML backup** section under the **Backup** tab, click **Create**.
3. The XML file will be backed up on the appliance.

## Restore Options

### Restore

The restore option allows the following filetypes to be restored:

- Backup archives (\*.zip) created using the backup feature
- Certificate archives (\*.tar.gz) created using v8.6.3 and earlier (included for backward compatibility)
- XML backup \*.xml) created using v8.6.3 and earlier (included for backward compatibility)

*To perform a restore:*

1. Using the WebUI navigate to: *Maintenance > Backup & Restore*.
2. Select the **Restore** Tab.
3. Click **Choose File** and browse to and select the relevant backup file.
4. If you want the floating IP's not to be brought up when the file is restored, enable (check) the *Disable Floating IPs on restore* checkbox.
5. If the file being restored is a backup archive (zip file), enter the password (if applicable).

#### Note

If a password was entered when the backup was created, but the password is omitted when performing the restore, you can continue with the restore but sensitive files (certificates and SSH keys) will be omitted.

#### Note

If a password was not entered when the backup was created, you'll be notified that the backup does not contain the sensitive files and asked if you'd like continue.

6. Click **Upload & Restore**.
7. Click **OK** to confirm that you'd like to continue.

The screen shot below shows an ongoing restore:



### Installing and upgrading XML...

### Restoring Configuration from uploaded file...

### Resetting Configuration.

### Disabling Peer Comms During Restore.

### Upgrading XML.

### Restoring Network Configuration.

Once complete, heartbeat and possibly other services (depending on the configuration) will need to be either restarted or reloaded. You'll be prompted accordingly in the "commit changes" message box at the top of the screen. If the box does not appear, click on System Overview or any other menu option - the box will then be displayed. Don't use F5 as this will cause the restore to run again.

#### Local XML Restore

This option is used to restore the last local XML backup created with the *Create local XML Backup* option.

#### Note

The XML restore feature is not backward compatible with previous major versions of the software, e.g. it's not possible to restore a V.7.6 XML file to a v8.7.x appliance.

*To Restore from the last local XML Backup:*

1. Using the WebUI navigate to: *Maintenance > Backup & Restore*.
2. Select the **Restore** Tab.
3. In the **Restore from last local XML backup** section, click **Restore**.
4. Click **OK** to confirm you'd like to continue.

The restore will now commence.

Once complete, heartbeat and possibly other services (depending on the configuration) will need to be either restarted or reloaded. You'll be prompted accordingly in the "commit changes" message box at the top of the screen. If the box does not appear, click on System Overview or any other menu option - the box will then be displayed. Don't use F5 as this will cause the restore to run again.

#### Restore System Defaults

This option can be used to reset all settings to their factory defaults.

*To restore the appliance to factory defaults:*

1. Using the WebUI navigate to: *Maintenance > Backup & Restore*.
2. In the **Restore to system defaults** section under the **Backup** tab, click **Restore**.

3. Click **OK** to confirm you'd like to continue.

## Template Deployment

This feature enables selected Virtual Services to be exported to a JSON configuration file. It also allows JSON configuration files to be imported. When importing, new IP addresses can be assigned to each VIP and multiple Backend (Real) Servers can be defined.

*To Access Template Deployment:*

1. Using the WebUI, navigate to: *Maintenance > Backup & Restore*.
2. Select the **Template Deployment** tab.

*To generate a Template Deployment file:*

1. Select the required VIPs to include in the deployment file. Use the CTRL key to select multiple VIPs or use the **Select All** button for all.
2. Click **Export Selected**.

*To Upload a Template Deployment file:*

1. Click **Choose File** and browse to and select the required JSON file.
2. Click **Upload Deployment File**.
3. For each VIP to be created, specify the required *IP address*.
4. Use the *Backend L4 Servers* section to specify the *Name* and *IP Address* of all Backend (Real) Servers that are to be associated with layer 4 VIPs.
  - Use the **Add Row** button to specify additional layer 4 Backend Servers.
5. Use the *Backend L7 Servers* section to specify the *Name* and *IP Address* of all Backend (Real) Servers that are to be associated with layer 7 VIPs.
  - Use the **Add Row** button to specify additional layer 7 Backend Servers.
6. Click **Apply Configuration**.

## Disaster Recovery

To recover a single appliance, you'll need to restore your configuration files. For virtual and cloud based appliances, other hypervisor / cloud recovery methods can also be used.

To recover from a hardware failure that requires the firmware to be re-installed, please refer to [Firmware Recovery using a USB Memory Stick](#).

To recover a failed member of a HA pair the *Peer Recovery* method can be used. For more information, please refer to [Disaster Recovery After Node \(Primary or Secondary\) Failure](#).

## Being Prepared - Creating Backups

To be able to quickly recover your appliance, it's important that you create regular configuration backups and keep them stored in a secure location. Ideally, you should keep a backup of both the Primary and Secondary appliances.

## Firmware Recovery using a USB Memory Stick

If a hardware appliance has suffered a hardware failure that requires the firmware to be re-installed, follow steps 1 to 5 below.

### Step 1 - Download the latest appliance disk image

The latest disk image can be downloaded from our website - please contact support for more information.

### Step 2 - Extract the image from the compressed archive

Extract the image using tar under Linux or something like WinRar or 7-Zip under Windows (not the built-in Windows extractor).

### Step 3 - Prepare the USB stick

#### *Using Linux*

After formatting the USB stick run the following command:

```
dd if=/imagefilename.img of=/dev/name-of-usb-disk
```

e.g.

```
dd if=/tmp/v7.5.0_r3368.img of=/dev/sda
```

#### Note

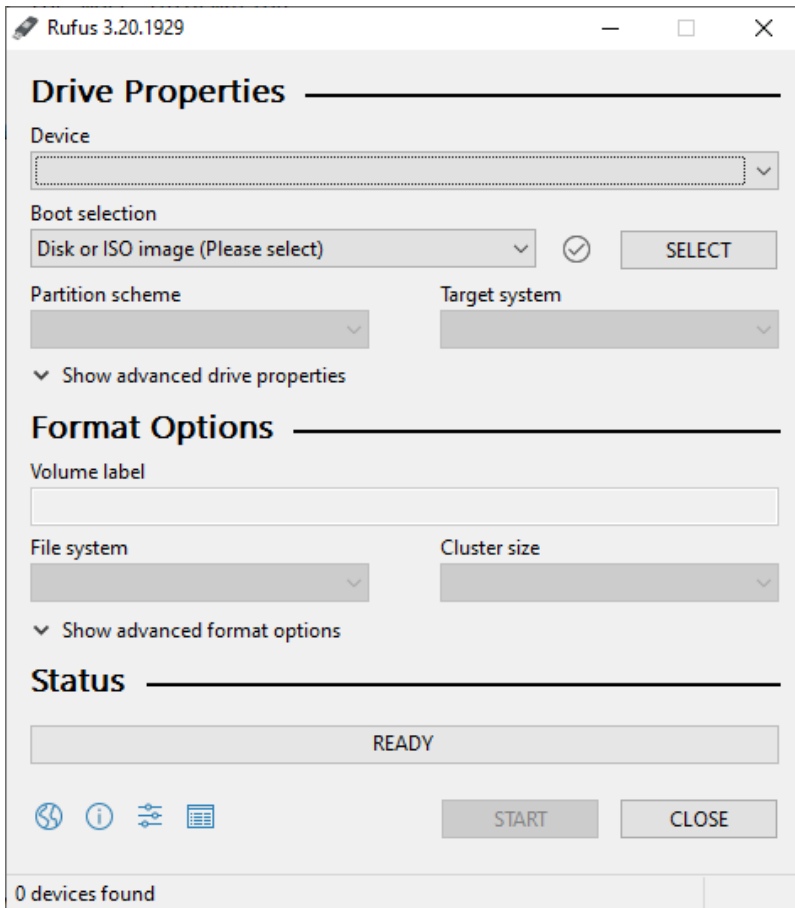
Do not use /dev/sdax where 'x' is a number, for example - /dev/sda1 as this will install to a partition on your usb stick. Use the whole disk **/dev/sda** Instead.

#### Important

**Be careful when using this command - make sure you specify the correct disk!**

#### *Using Windows*

For Windows, a third party image writer must be used. Several free options are available, the example below uses **Rufus** which can be downloaded [here](#). The download is an executable file.



Start Rufus, select the relevant USB device and the source image and click **START**.

**Important** | Be careful - make sure you specify the correct disk!

#### Step 4 - Restore the appliance using the USB stick

1. Boot the appliance into System Setup.
2. Ensure that UEFI boot is enabled and that the USB device is first in the boot order.

**Note** | From v8.7.2 all ISO images require UEFI to boot.

3. Boot the appliance, after the initial boot messages the following prompt will appear:

```
DO YOU WISH TO CONTINUE?  
Please enter yes or no  
Type "yes" and press <ENTER>
```

The installation should take around 2-3 minutes, once complete the following message will be displayed:

```
Installation Finished
```

4. As directed, press any key to shutdown the load balancer.
5. Once shutdown, remove the USB stick.
6. Power up the appliance.

7. Login at the console:

**Username:** root

**Password:** loadbalancer

8. Now run the model specific command to ensure that the hardware is configured correctly.

**Note** | Please contact support for the command for your specific model.

9. Now run the following additional commands:

```
# lbrestore <ENTER>
# lbcleanboot <ENTER>
```

10. Reboot the appliance once again.

11. Set the required IP address and other settings using the Network Setup Wizard as described [here](#).

**Note** | You'll need to reapply your license key file to ensure the newly restored appliance is correctly licensed. Please contact support if you have any issues.

## Disaster Recovery After Node (Primary or Secondary) Failure

For a Clustered Pair of load balancers, recovery of a failed node is quick and simple. The procedure is the same whether the Primary has failed or the Secondary has failed.

**Note** | Using the Peer Recovery feature means that the HA pair is re-established without disrupting any of the services that are currently running on the working appliance.

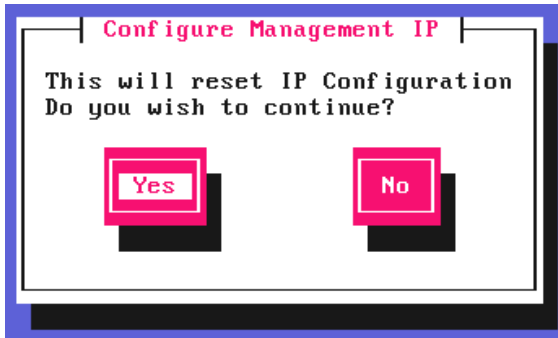
**Note** | Make sure that SSH Password access is enabled on the remaining operational node. This can be configured using the WebUI menu option: *Local Configuration > Security*, clearing the *Disable SSH Password Access* checkbox and clicking **Update**.

1. If the failed node is still on, power it down.
2. For a hardware appliance:
  - Disconnect all cables.
  - If the SSD/HD has failed and has been replaced and needs to be re-imaged, follow [these steps](#) to restore the appliance firmware.
3. Power up the new/repaired/re-imaged appliance.
4. Login to the console as:

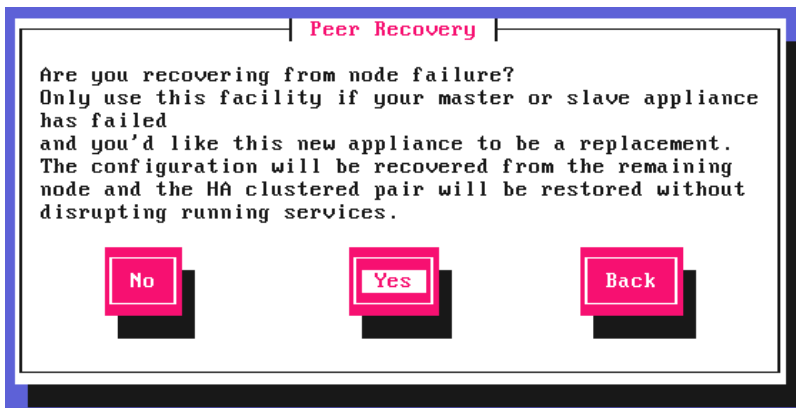
**username:** setup

**password:** setup

5. The Configure Management IP screen will be displayed:

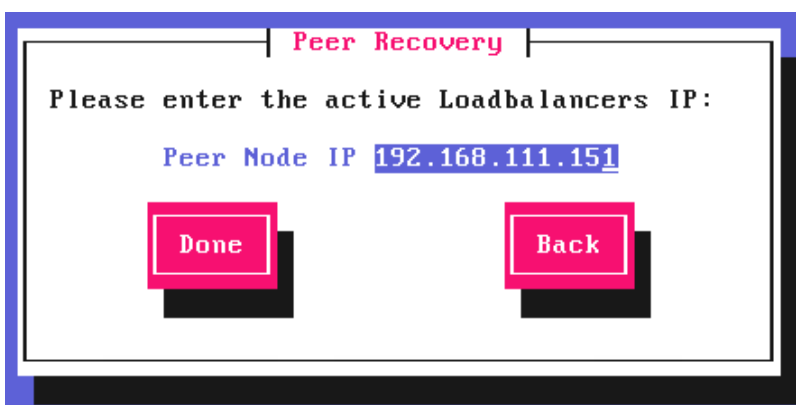


6. Select **Yes** and hit <ENTER> to continue.
7. On the Peer Recovery screen you'll be asked if you're recovering from node failure:



8. Select **Yes** and hit <ENTER> to continue.
9. Now continue through the Network Setup Wizard to configure the initial network settings - ensure these are the same as the failed appliance. Continue until you reach the Peer Recovery screen.

Note For more information about configuring network settings please refer to [Configuring Initial Network Settings](#).



10. Enter the IP address of the active appliance, select **Done** and hit <ENTER> to continue.

The screenshot shows a terminal window titled "Peer Recovery". The text inside says "Please enter the active Loadbalancers WUI Port:". Below this, it says "Peer Node Port (Default:9443) 9443" with the "9443" entered in a blue input field. At the bottom, there are two red buttons with white text: "Done" and "Back".

11. Enter the WebUI port of the active appliance, select **Done** and hit <ENTER> to continue.

The screenshot shows a terminal window titled "Peer Recovery". The text inside says "Please enter the password for the WUI loadbalancer user on the active peer". Below this, it says "Peer Node WUI Password loadbalancer" with "loadbalancer" entered in a blue input field. At the bottom, there are two red buttons with white text: "Done" and "Back".

12. Enter the WebUI password for the 'loadbalancer' user on the active appliance, select **Done** and hit <ENTER> to continue.

The screenshot shows a terminal window titled "Summary". It displays the following configuration details: "Interface: eth0", "IP Address: 192.168.111.150/18", "Default Gateway: 192.168.64.1", and "DNS Server/s: 8.8.8.8". Below this, it says "You will be able to access the web interface from: https://192.168.111.150:9443". At the bottom, there are three red buttons with white text: "Configure", "Start Over", and "Back".

13. A summary of the local settings is displayed, hit <ENTER> to continue. The settings from the peer appliance will then be applied. Once the restore completes, the Configuration Complete message will be displayed:

The screenshot shows a terminal window titled "Configuration complete". The text inside says "The configuration has been written successfully. Please reboot the appliance." At the bottom, there is a single red button with white text: "OK".

14. To complete the restore and re-pairing process, reboot the newly restored appliance as mentioned in the on-screen message. Once rebooted, the HA pair will be re-synchronized and fully recovered.

The Primary's WebUI should now display:

**Primary** | **Secondary**      **Active** | **Passive**      **Link**

The Secondary's WebUI should now display:

**Primary** | **Secondary**      **Active** | **Passive**      **Link**



# Chapter 15 - Technical Support

## Introduction

Loadbalancer.org have a team of very experienced support Engineers who are available to assist with your load balancer deployment.

Unlimited support is available as follows:

- During the cover period of any active support agreement

(to purchase a support package, please contact: [sales@loadbalancer.org](mailto:sales@loadbalancer.org))

- During the 30 day Virtual Appliance trial period

(to download the trial please go to: <https://www.loadbalancer.org/get-started/>)

## WebUI Support Options

The WebUI's Support menu option includes 3 sub-options that are provided to assist when support is required.

### Contact Us

This option provides details on how to contact the support team, the information we need to help us resolve the issue as quickly as possible and guidance on how to generate a technical support download which should also be included with your initial email.

Sending an email to [support@loadbalancer.org](mailto:support@loadbalancer.org) creates a ticket in our help desk system and enables all technical support staff to view the case. This is the most efficient way to contact support and guarantees that any reported issues will be acted upon and addressed as quickly and efficiently as possible.

### Technical Support Download

This option enables the Support Download to be created. The download is a compressed archive containing all log files and configuration files from the appliance and should be attached to your email.

#### Technical Support Download

When contacting **Loadbalancer.org Support**, you may be asked to supply the load balancer's configuration and log files. This page generates an archive of all the required files, which can then be downloaded to your PC.

Please click the button below to generate the archive.

The load balancer will collect the configuration files and logs into a compressed archive.

Download will begin shortly after clicking the button.

Send the archive by email to **Loadbalancer.org support**. If this is your first contact with support on this issue, please include your company name and details of the problem you are experiencing.

☒ **Do not include GZ files.**(This can decrease archive size in some circumstances)

**Note:** Generating the archive may take several minutes on a load balancer with extensive log files.

Generate Archive

Please click the button above to start the process.

- To minimize the size of the support download, **.GZ** files (archived logs) are excluded by default. If these must be included in the archive, clear the check-box shown above
- To generate the archive, click the **Generate Archive** button

Once downloaded, attach the file to your email when contacting support.

If the file is large, it can be uploaded using our [Upload Server](#).

#### Note

You'll first need to email support, then once you have a support ticket ID, you'll be able to upload the support archive. The Upload Sever link is also displayed in the Useful Links menu.

## Useful Links

This option presents a number of self explanatory web links.

## Remote Support

Where necessary, our support team may arrange a remote session to assist with trouble shooting. **Zoom** is the preferred tool and is used where possible.

## Live Chat

Online chat can be invoked directly from the WebUI.

*To enable Online Chat:*

1. Using the WebUI, navigate to: *Local Configuration > Live Chat*.
2. Ensure that *Enable Live Chat* is checked (enabled).
3. Click Update.


*To invoke Online Chat:*

1. Using the WebUI, navigate to: *Live Chat*.
2. Enter the question you have.



## Appendix

### Front & Rear Panel Layouts



#### Enterprise 1G



#### Enterprise 10G/25G/40G/50G



#### Enterprise 100G



\* The number of interfaces depends on your choice of interface cards

### IPMI (Remote Management) Configuration


The Enterprise 1G supports IPMI to allow remote control & management. This can either be accessed via the dedicated IPMI Ethernet interface or via one of the standard Ethernet interfaces in bridged mode.

To use the dedicated IPMI interface, ensure that a network cable is plugged into the interface before powering up the appliance.

#### Configuring the IP Address

By default the IP address is set using DHCP. The address allocated is displayed in the IPMI sub-menu in system setup. If preferred, a static IP address can also be set using the same menu. To access system setup, hit <DEL> as directed at boot time.

#### IPMI BIOS Menu:



To set the address:

Change **Update IPMI LAN configuration** to 'Yes'  
 Change **Configuration Address Source** to 'Static'

Now set the IP address, mask etc. as required.

```

Lan channel 1
Update IPMI LAN configuration      [Yes]
Configuration Address source      [Static]
Station IP address                0.0.0.0
Subnet mask                      0.0.0.0
Station MAC address              00-00-00-00-00-00
Gateway IP address               0.0.0.0
  
```

IPMI BIOS Menu:

```

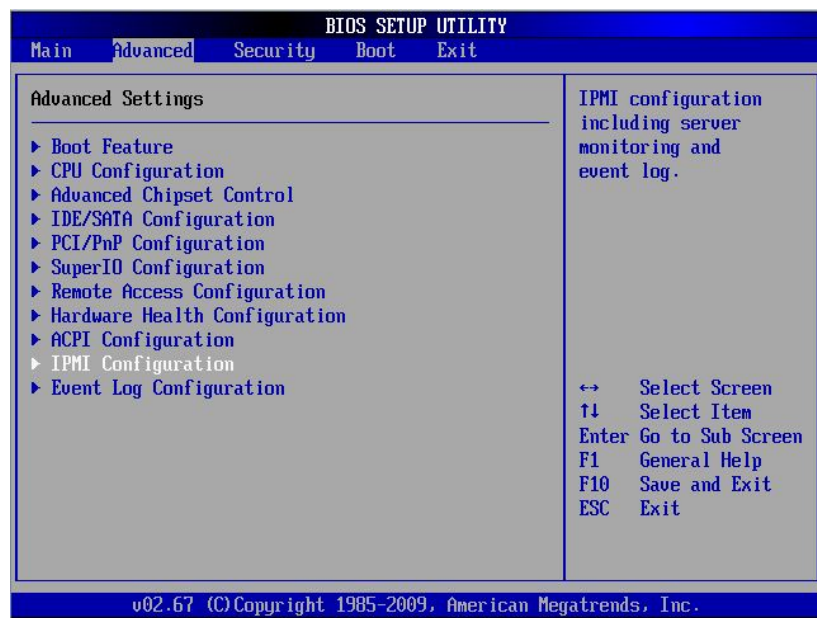
Channel Number                   [01]
Channel Number Status:Channel number is OK
IP Address Source                [Static]
IP Address                      [192.168.075.111]
Subnet Mask                     [255.255.192.000]
Gateway Address                 [192.168.064.001]
MAC Address                     [00.25.90.6F.39.DA]
  
```

To set the address:

Select **Set LAN Configuration**.

Change **IP Address Source** to 'Static'

Now set the IP address, mask etc. as required.



Accessing the login page:

Using a browser, connect to **http://<ip address>**

the following login prompt is displayed:

A login form with a light beige background. At the top, it says "Please Login" in bold. Below this are two input fields: "Username" and "Password". Each field has a white text box. Below the password field is a "login" button with a grey gradient and a black border.

username: **ADMIN**

default password: **ADMIN**

**Note**

For certain appliances the IPMI password may have been updated. If this is the case the updated password will be noted on a label applied on the top cover of the appliance. If you have any issues please contact [support@loadbalancer.org](mailto:support@loadbalancer.org).

Once logged in, the following screen is displayed:

System	Server Health	Configuration	Remote Control	Virtual Media	Maintenance	Miscellaneous
--------	---------------	---------------	----------------	---------------	-------------	---------------

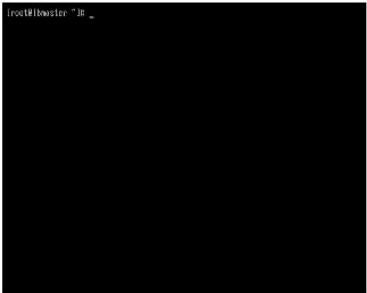
System
System Information
FRU Reading

### Summary

Firmware Revision : 01.64  
Firmware Build Time : 2012-05-08  
IP address : 192.168.075.111  
MAC address : 00:25:90:74:20:78

Remote Console Preview

Refresh Preview Image



Power Control via IPMI

Host is currently on
Power On
Power Down
Reset

Copyright © 2010 Super Micro Computer, Inc.

## IPMI Interface

As mentioned above IPMI can be accessed via the dedicated interface or via one of the standard on-board NICs. This can be configured in the IPMI interface using: *Configuration > Network > LAN Interface*.

- **Dedicate** - use the dedicated interface only
- **Share** - run in bridge mode using one of the standard NICs
- **Failover** - allows either connection method to be used (the default)

## Remote Control

To access the systems console, simple click on the Remote Console Preview image. A new window will open with access to the console of the appliance.

### Note

You cannot SSH into the module directly. You need to connect via the IPMI's web interface, then use the remote control option as mentioned above. This can also be accessed using the 'Remote Control' option in the top menu. From here you can use the Launch Console option to launch a virtual Java console which will allow you to use the device as if you stood in front of the device. Next the 'Power Control' options menu will give you several options such as Restart Server, Power off and Power Cycle server. these options will perform the same function as pressing the physical reset button on the unit (Reset Server) as well as being able to perform the same functions as the physical power switch as well.

Please do remember that the IPMI power control options are completely independent of the Load balancer software and that the reset option is the same as pressing reset on your PC.

## iDRAC (Remote Management) Configuration

iDRAC is supported on the Enterprise 10G, Enterprise 50G and Enterprise 100G appliances.

**Note** | For more information on iDRAC licensing options, please contact sales.

### iDRAC IP Address

The Default IP address is **192.168.0.120/255.255.255.0**. This can be changed in a number of ways. The recommended method is to use the **iDRAC Setting Utility**. All methods are described [here](#).

### Logging in to iDRAC

The default user credentials are:

```
Username: root
Password: calvin
```

Details on how to change the default password are available [here](#).

### iDRAC Password Reset

If you've forgotten your iDRAC password, it can be reset as described [here](#).

### iDRAC Licensing

If required, trial iDRAC licenses are available [here](#).

Details on importing your iDRAC license are available [here](#).

### More Information

The full iDRAC9 user guide is available [here](#).

## Appliance IPv4 Address Format (CIDR notation)

When specifying IP addresses on the appliance, CIDR format is used. The following table shows the various masks and the corresponding IPv4 IP/CIDR equivalents:

Mask	IP/CIDR Prefix
255.255.255.255	a.b.c.d/32
255.255.255.254	a.b.c.d/31
255.255.255.252	a.b.c.d/30
255.255.255.248	a.b.c.d/29
255.255.255.240	a.b.c.d/28
255.255.255.224	a.b.c.d/27
255.255.255.192	a.b.c.d/26
255.255.255.128	a.b.c.d/25
255.255.255.000	a.b.c.d/24
255.255.254.000	a.b.c.d/23

Mask	IP/CIDR Prefix
255.255.252.000	a.b.c.d/22
255.255.248.000	a.b.c.d/21
255.255.240.000	a.b.c.d/20
255.255.224.000	a.b.c.d/19
255.255.192.000	a.b.c.d/18
255.255.128.000	a.b.c.d/17
255.255.000.000	a.b.c.d/16
255.254.000.000	a.b.c.d/15
255.252.000.000	a.b.c.d/14
255.248.000.000	a.b.c.d/13
255.240.000.000	a.b.c.d/12
255.224.000.000	a.b.c.d/11
255.192.000.000	a.b.c.d/10
255.128.000.000	a.b.c.d/9
255.000.000.000	a.b.c.d/8
254.000.000.000	a.b.c.d/7
252.000.000.000	a.b.c.d/6
248.000.000.000	a.b.c.d/5
240.000.000.000	a.b.c.d/4
224.000.000.000	a.b.c.d/3
192.000.000.000	a.b.c.d/2
128.000.000.000	a.b.c.d/1



## About Loadbalancer.org

Loadbalancer.org's mission is to ensure that its clients' businesses are never interrupted. The load balancer experts ask the right questions to get to the heart of what matters, bringing a depth of understanding to each deployment. Experience enables Loadbalancer.org engineers to design less complex, unbreakable solutions - and to provide exceptional personalized support.



### United Kingdom

Loadbalancer.org Ltd.  
Compass House, North Harbour  
Business Park, Portsmouth, PO6 4PS  
UK: +44 (0) 330 380 1064  
sales@loadbalancer.org  
support@loadbalancer.org

### United States

Loadbalancer.org, Inc.  
4550 Linden Hill Road, Suite 201  
Wilmington, DE 19808, USA  
TEL: +1 833.274.2566  
sales@loadbalancer.org  
support@loadbalancer.org

### Canada

Loadbalancer.org Appliances Ltd.  
300-422 Richards Street, Vancouver,  
BC, V6B 2Z4, Canada  
TEL: +1 866 998 0508  
sales@loadbalancer.org  
support@loadbalancer.org

### Germany

Loadbalancer.org GmbH  
Tengstraße 2780798,  
München, Germany  
TEL: +49 (0)89 2000 2179  
sales@loadbalancer.org  
support@loadbalancer.org