LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing Ross Video Inception

Version 1.0.0

Table of Contents

1. About this Guide . . . 3
2. Loadbalancer.org Appliances Supported . . . 3
3. Software Versions Supported . . . 3
    3.1. Loadbalancer.org Appliance . . . 3
    3.2. Ross Video Inception . . . 3
4. Ross Video Inception . . . 3
5. Load Balancing Ross Video Inception . . . 3
    5.1. Load Balancing & HA Requirements . . . 3
    5.2. Persistence (aka Server Affinity) . . . 4
    5.3. Timeouts . . . 4
    5.4. Virtual Service (VIP) Requirements . . . 4
    5.5. SSL/TLS Termination (Optional) . . . 4
6. Deployment Concept . . . 5
7. Load Balancer Deployment Methods . . . 5
    7.1. Layer 7 SNAT Mode . . . 5
    7.2. A Note on Network Topology . . . 6
8. Loadbalancer.org Appliance – the Basics . . . 7
    8.1. Virtual Appliance . . . 7
    8.2. Initial Network Configuration . . . 7
    8.3. Accessing the Appliance WebUI . . . 7
        8.3.1. Main Menu Options . . . 9
    8.4. Appliance Software Update . . . 9
        8.4.1. Online Update . . . 9
        8.4.2. Offline Update . . . 10
    8.5. Ports Used by the Appliance . . . 10
    8.6. HA Clustered Pair Configuration . . . 11
9. Configuring Ross Video Inception for Load Balancing . . . 11
    9.1. Layer 7 SNAT Mode . . . 11
10. Appliance Configuration for Ross Video Inception . . . 11
    10.1. Configuring the Layer 7 Global Timeouts . . . 11
    10.2. VIP 1 - Inception . . . 12
        10.2.1. Virtual Service (VIP) Configuration . . . 12
        10.2.2. Define the Associated Real Servers (RIPs) . . . 13
    10.3. Finalizing the Configuration . . . 13
    10.4. Reference Configuration . . . 14
11. Testing & Verification . . . 14
    11.1. Using System Overview . . . 14
    11.2. Accessing Ross Video Inception via the Load Balancer . . . 15
12. Technical Support . . . 15
13. Further Documentation . . . 15
14. Appendix . . . 16
    14.1. Configuring HA - Adding a Secondary Appliance . . . 16
        14.1.1. Non-Replicated Settings . . . 16
        14.1.2. Configuring the HA Clustered Pair . . . 17
15. Document Revision History . . . 19

1. About this Guide

This guide details the steps required to configure a load balanced Ross Video Inception environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Ross Video Inception configuration changes that are required to enable load balancing.

For more information about initial appliance deployment, network configuration and using the Web User Interface (WebUI), please also refer to the Administration Manual.

2. Loadbalancer.org Appliances Supported

All our products can be used with Ross Video Inception. For full specifications of available models please refer to https://www.loadbalancer.org/products/enterprise.

Some features may not be available or fully supported in all cloud platforms due to platform specific limitations. For more details, please refer to the "Main Differences to our Standard (Non-Cloud) Product" section in the appropriate cloud platform Quick Start Guide or check with Loadbalancer.org support.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.13.2 and later

> **Note** The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If you're using an older version, or the very latest, the screenshots presented here may not match your WebUI exactly.

3.2. Ross Video Inception

- All versions

4. Ross Video Inception

Ross Video Inception is a web based newsroom computer system (NRCS) and social media management platform. It brings story planning, assignments, scripting and rundowns together with publishing to social media and the web, so that a news team can produce content for broadcast and digital channels from a single system. Users access Inception using a web browser.

5. Load Balancing Ross Video Inception

> **Note** It's highly recommended that you have a working Ross Video Inception environment first before implementing the load balancer.

5.1. Load Balancing & HA Requirements

Ross Video Inception is deployed on two (or more) Inception servers. Load balancing the web traffic across these servers provides load sharing, high availability and resilience - if one server fails or is taken out of service, users are automatically directed to the remaining server(s).

A single layer 7 Virtual Service in HTTP mode is used to load balance client connections to the Inception servers.

5.2. Persistence (aka Server Affinity)

Clients must continue to be served by the same Inception server for the duration of their session. By default, Ross Video use ***Source IP*** persistence, which is the method used in this guide. The persistence timeout is left at its default of 30 minutes.

Other persistence methods are also used on some Ross Video deployments, and may be selected instead where Source IP persistence is not suitable - for example, where many clients reach the load balancer from behind the same NAT device and so share one source IP address:

- **HTTP Cookie** - the load balancer inserts a cookie into the HTTP response that identifies the server to use for subsequent requests.
- **X-Forwarded-For and Source IP** - the client IP address held in the X-Forwarded-For header is used for persistence, falling back to the source IP address when the header is not present. This is useful where an upstream proxy is in the path.

5.3. Timeouts

For Ross Video Inception, the layer 7 ***Client Timeout*** and ***Real Server Timeout*** are both raised to **1800000** ms (30 minutes), from their default values of 42000 ms and 43000 ms respectively.

> **Note** These are global settings, so they apply to every layer 7 Virtual Service on the appliance. Check these values first when troubleshooting timeout related issues.

5.4. Virtual Service (VIP) Requirements

To provide load balancing and HA for Ross Video Inception, the following VIP is required:

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	Inception	L7 SNAT (HTTP)	80	Source IP	Connect to Port

5.5. SSL/TLS Termination (Optional)

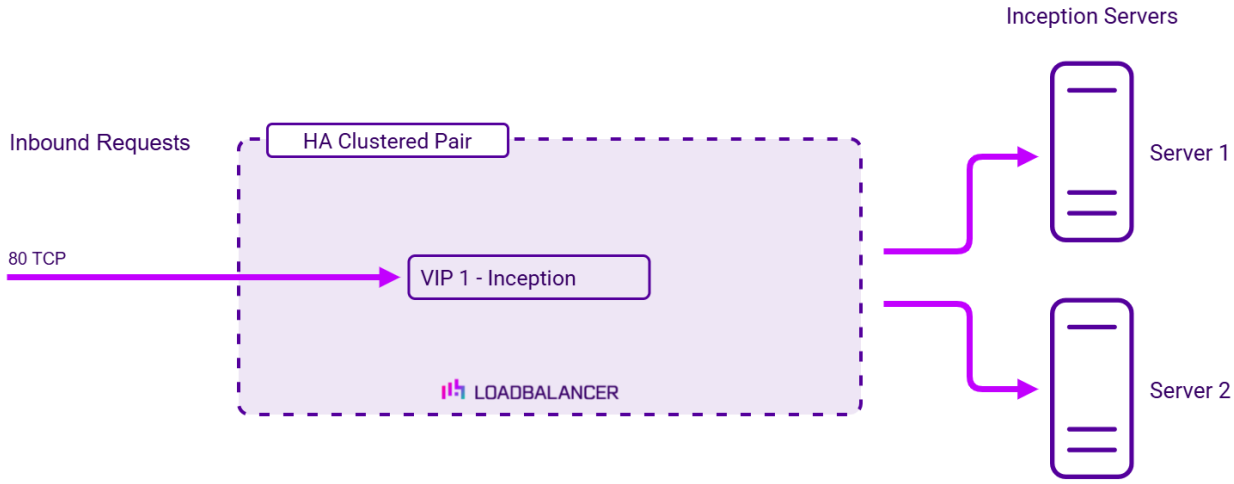
The configuration presented in this guide load balances HTTP on port 80. If clients must connect to Inception using HTTPS, SSL termination can be configured on the load balancer for the ***Inception*** VIP. Clients then connect to the VIP using HTTPS on port 443, while the connections from the load balancer to the Inception servers remain HTTP on port 80. Contact Ross Video to confirm whether any Inception configuration changes are required when clients connect using HTTPS.

For more information on uploading certificates and configuring SSL termination, please refer to SSL Termination

on the Load Balancer (SSL Offloading) in the Administration Manual.

6. Deployment Concept

Once the load balancer is deployed, clients connect to the Virtual Service (VIP) rather than connecting directly to one of the Inception servers. These connections are then load balanced across the Inception servers to distribute the load according to the load balancing algorithm selected.



> **Note** The load balancer can be deployed as a single unit, although Loadbalancer.org recommends a clustered pair for resilience & high availability. Please refer to the section Configuring HA - Adding a Secondary Appliance in the appendix for more details on configuring a clustered pair.

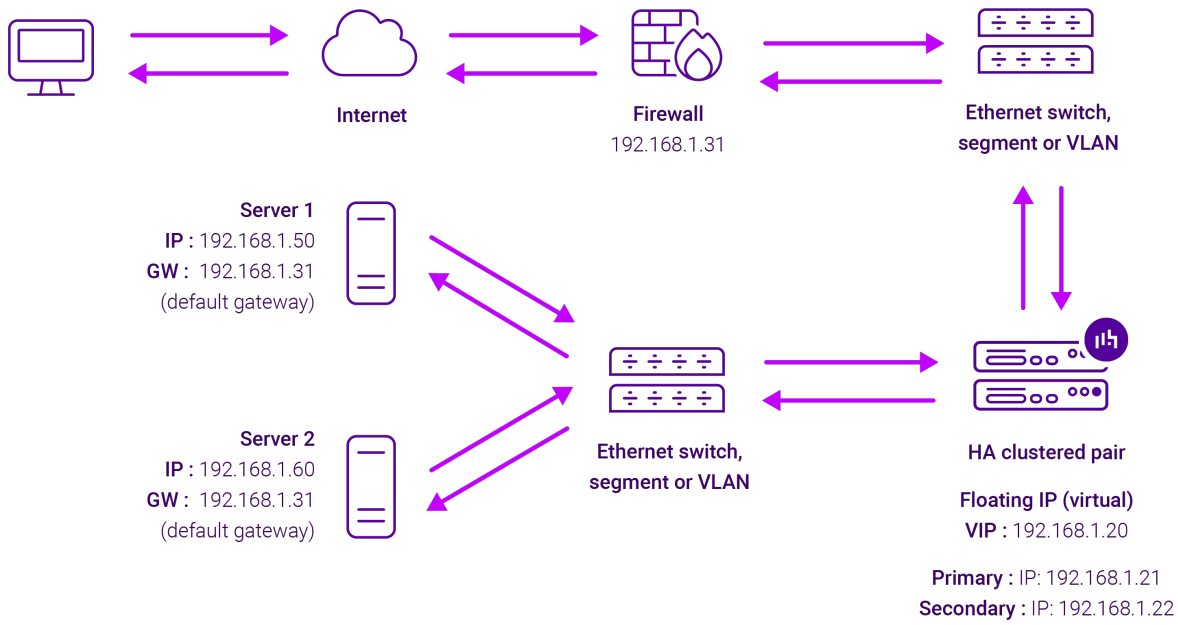
7. Load Balancer Deployment Methods

The load balancer can be deployed in 4 fundamental ways: ***Layer 4 DR mode***, ***Layer 4 NAT mode***, ***Layer 4 SNAT mode***, and ***Layer 7 SNAT mode***.

For Ross Video Inception, layer 7 SNAT mode is recommended. This mode is described below and is used for the configuration presented in this guide.

7.1. 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. The image below shows an example network diagram for this mode.



- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network 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, **eth1** is typically used for client side connections and **eth0** is used for Real Server connections, although this is not mandatory since any interface can be used for any purpose.
- 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.

7.2. A Note on Network Topology

The Inception servers do not need to be on the same subnet as the load balancer. They can be on the local subnet, or on a remote subnet that is reached via a router - the load balancer must be able to reach the Inception servers on port 80, and the Inception servers must be able to route their replies back to the load balancer.

Because the Inception servers see connections arriving from the load balancer's own IP address, the load balancer adds an X-Forwarded-For header to every HTTP request by default. This header holds the IP address of the client, so the real client address is still available to the Inception servers.

8. Loadbalancer.org Appliance – the Basics

8.1. Virtual Appliance

A fully featured, fully supported 30 day trial is available if you are conducting a PoC (Proof of Concept) deployment. The VA is currently available for VMware, Virtual Box, Hyper-V, KVM, XEN and Nutanix AHV and has been optimized for each Hypervisor. By default, the VA is allocated 2 vCPUs, 4GB of RAM and has a 20GB virtual disk. The Virtual Appliance can be downloaded here.

> **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** Please refer to Virtual Appliance Installation and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.

> **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.

8.2. Initial Network Configuration

After boot up, follow the instructions on the appliance console to configure the management IP address, subnet mask, default gateway, DNS servers and other network and administrative settings.

> **Important** Be sure to set a secure password for the load balancer, when prompted during the setup routine.

8.3. Accessing the Appliance WebUI

The WebUI is accessed using a web browser. By default, users are authenticated using Apache authentication. Users can also be authenticated against LDAP, LDAPS, Active Directory or Radius - for more information, please refer to External Authentication.

> **Note** There are certain differences when accessing the WebUI for the cloud appliances. For details, please refer to the relevant Quick Start / Configuration Guide.

1. Using a browser, navigate to the following URL:

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

   > **Note** You'll receive a warning about the WebUI's SSL certificate. This is due to the default self signed certificate that is used. If preferred, you can upload your own certificate - for more information, please refer to Appliance Security Features.

   > **Note** If you need to change the port, IP address or protocol that the WebUI listens on, please

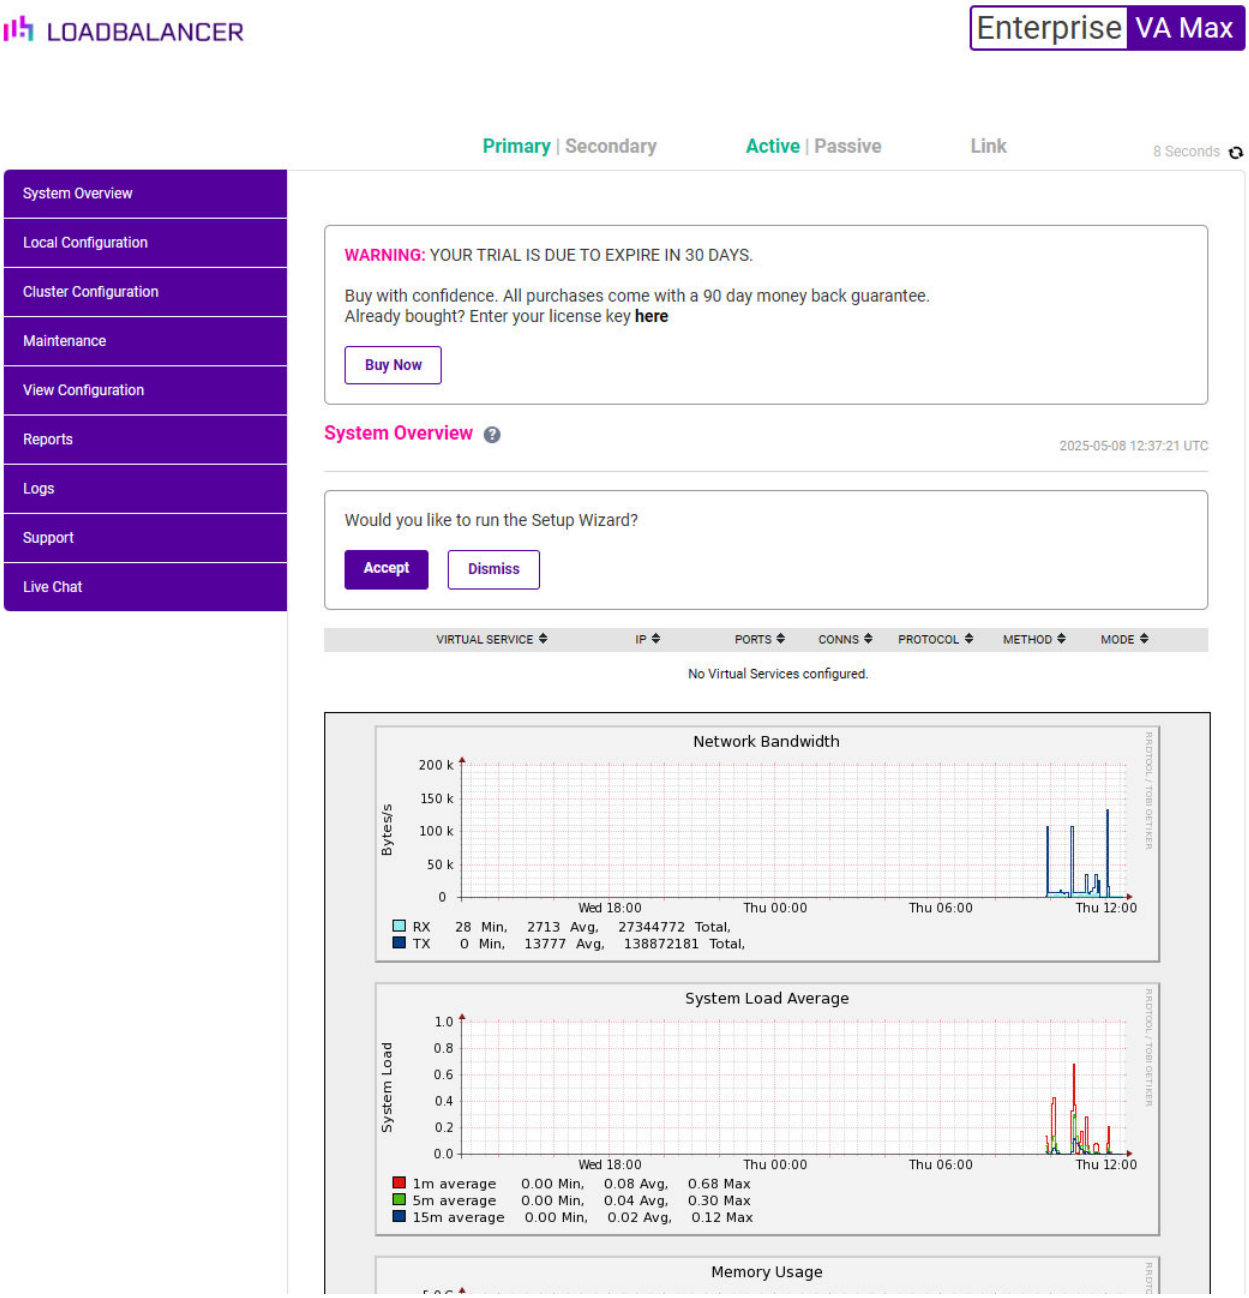
refer to Service Socket Addresses.

2. Log in to the WebUI using the following credentials:

   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:



3. 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.

8.3.1. 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 creating 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

8.4. Appliance Software Update

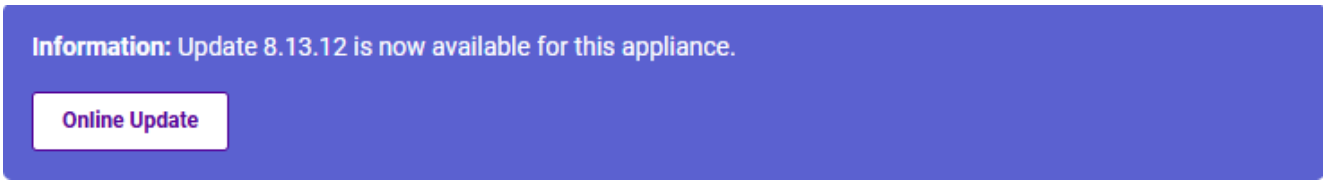
We recommend that the appliance is kept up to date to ensure that you benefit from the latest bug fixes, security updates and feature improvements. Both online and offline update are supported.

> **Note** For full details, please refer to Appliance Software Update in the Administration Manual.

> **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.

8.4.1. Online Update

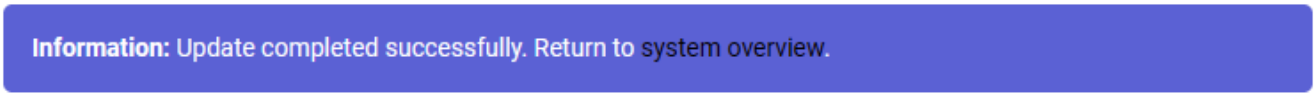
The appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. This is the default behavior and can be disabled if preferred. If an update is found, a notification similar to the example below will be displayed at the top of the WebUI:



Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

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

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



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

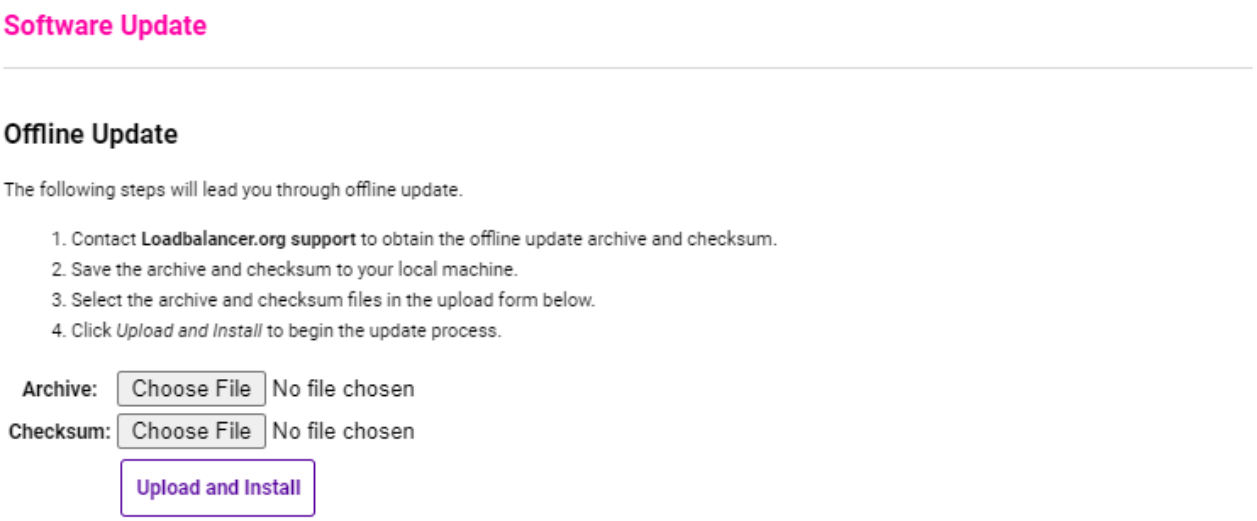
8.4.2. Offline Update

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

To check for the latest version, please refer to our product roadmap page available here. To obtain the latest offline update files contact support@loadbalancer.org.

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:



4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

8.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22	SSH
TCP & UDP	53	GSLB
TCP & UDP	123	NTP

Protocol	Port	Purpose
TCP & UDP	161	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000	Gateway service for ADC Portal comms
TCP	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback server
TCP	9443	WebUI - HTTPS
TCP	25565	Shuttle service for ADC Portal comms

> **Note** All ports listed above except port 123 (NTP) can be changed if required.

- To change the port used for heartbeat, refer to Configuring High Availability
- To change the port used for HAProxy replication, refer to Layer 7 - Advanced Configuration
- To change other ports, refer to Service Socket Addresses

8.6. HA Clustered Pair Configuration

Loadbalancer.org recommend that load balancer appliances are deployed in pairs for high availability. In this guide a single unit is deployed first, adding a secondary unit is covered in the section Configuring HA - Adding a Secondary Appliance of the appendix.

9. Configuring Ross Video Inception for Load Balancing

9.1. Layer 7 SNAT Mode

Layer 7 SNAT mode VIPs do not require any mode specific configuration changes to the load balanced Real Servers (the Inception servers).

Make sure that any firewall on the Inception servers allows inbound HTTP (TCP port 80) from the load balancer. In a clustered pair, allow this from both appliances.

10. Appliance Configuration for Ross Video Inception

10.1. Configuring the Layer 7 Global Timeouts

1. Using the WebUI, navigate to ***Cluster Configuration > Layer 7 – Advanced Configuration***.
2. Set the ***Client Timeout*** to **1800000** (ms).
3. Set the ***Real Server Timeout*** to **1800000** (ms).

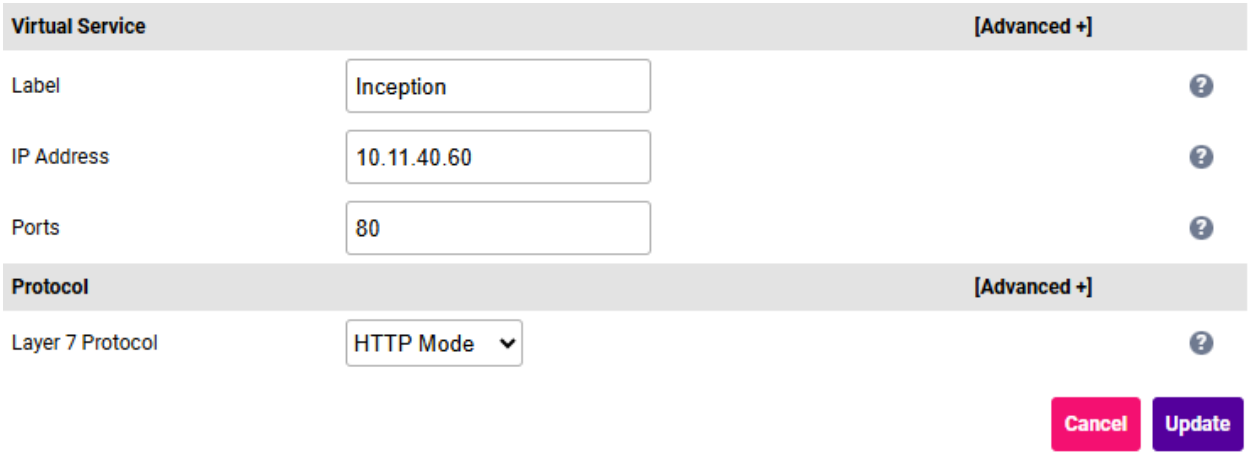


4. Leave all other settings at their default value.
5. Scroll to the bottom of the page and click **Update**.

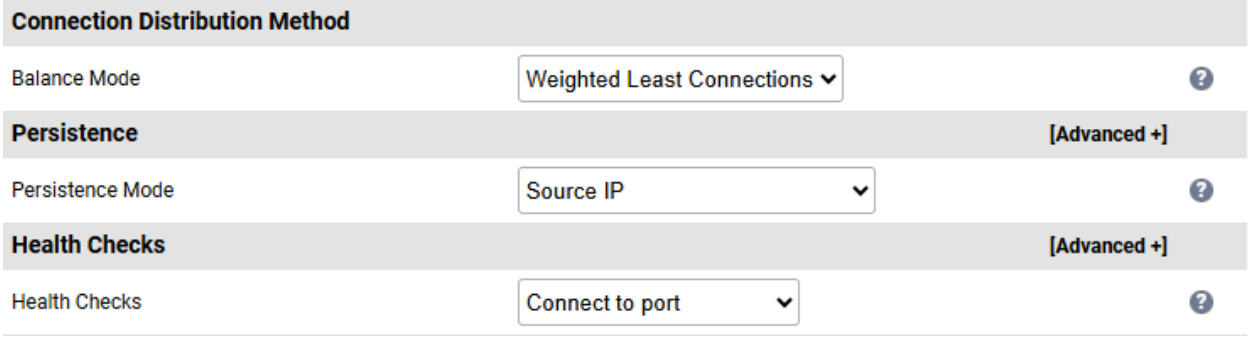
10.2. VIP 1 - Inception

10.2.1. Virtual Service (VIP) Configuration

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Virtual Services* and click **Add a new Virtual Service**.
2. Enter the following details:



3. Enter an appropriate *Label* for the VIP, e.g. **Inception**.
4. Set the *IP Address* field to the required IP address, e.g. **10.11.40.60**.
5. Set the *Ports* field to **80**.
6. Leave the *Layer 7 Protocol* set to **HTTP Mode**.
7. Click **Update**.
8. Click **Modify** next to the newly created VIP.

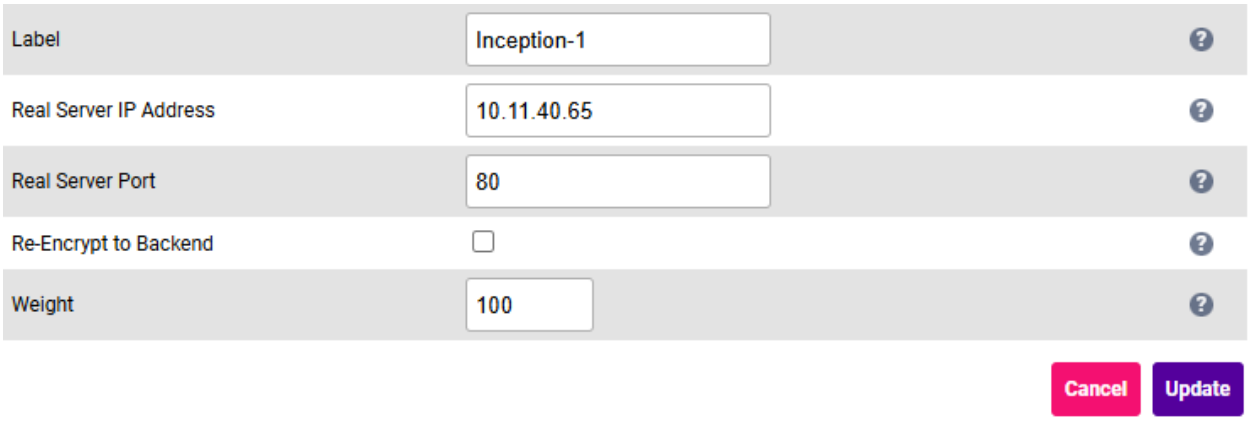


9. Scroll to the *Connection Distribution Method* section.
    - Ensure that the *Balance Mode* is set to **Weighted Least Connections**.
10. Scroll to the *Persistence* section.
    - Set the *Persistence Mode* to **Source IP**.
11. Scroll to the *Health Checks* section.
    - Ensure that *Health Checks* is set to **Connect to port**.
12. Leave all other settings at their default value.
13. Click **Update**.

> **Note** If a different persistence method is required, set the *Persistence Mode* to **HTTP Cookie** or **X-Forwarded-For and Source IP** instead - see Persistence (aka Server Affinity).

10.2.2. 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 newly created VIP.
2. Enter the following details:



3. Define the *Label* for the Real Server as required, e.g. **Inception-1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **10.11.40.65**.
5. Set the *Real Server Port* field to **80**.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining Inception server(s), e.g. **Inception-2** at **10.11.40.66**.

10.3. Finalizing the Configuration

To apply the new settings, HAProxy must be reloaded. This can be done using the button in the "Commit changes" box at the top of the screen or by using the Restart Services menu option:

1. Using the WebUI, navigate to: *Maintenance > Restart Services*.
2. Click **Reload HAProxy**.

10.4. Reference Configuration

Once the configuration is complete, the resulting HAProxy configuration for the Inception VIP can be checked using the WebUI option *View Configuration > Layer 7*. The key lines are shown below:

```
listen Inception
    bind 10.11.40.60:80 transparent
    mode http
    balance leastconn
    stick on src
    stick-table type ip size 10240k expire 30m peers loadbalancer_replication
    server backup 127.0.0.1:9081 backup  non-stick
    option http-keep-alive
    timeout http-request 5s
    option forwardfor
    option redispatch
    option abortonclose
    maxconn 40000
    server Inception-1 10.11.40.65:80 id 2  weight 100  check port 80 inter 4000  rise 2  fall 2
    server Inception-2 10.11.40.66:80 id 3  weight 100  check port 80 inter 4000  rise 2  fall 2
```

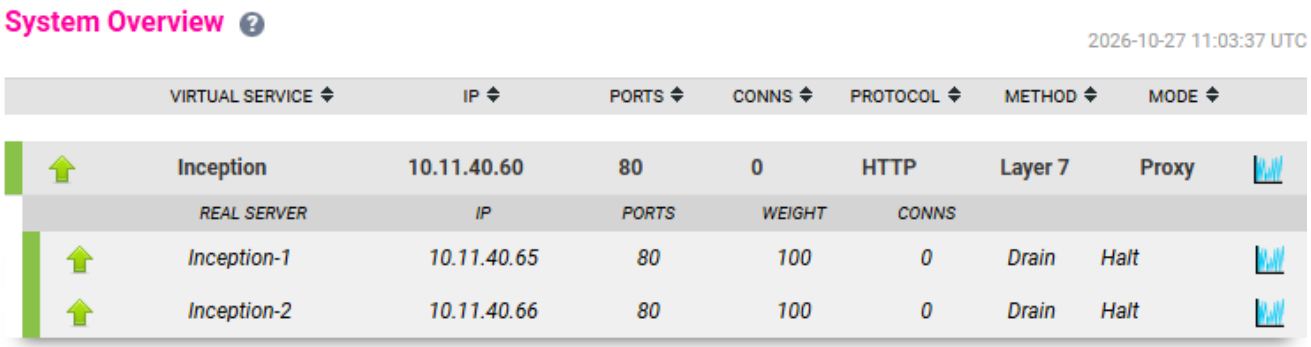
The remaining options at the end of each `server` line are omitted above for clarity. The raised timeouts appear in the global `defaults` section as `timeout client 1800000` and `timeout server 1800000`.

11. Testing & Verification

> **Note** For additional guidance on diagnosing and resolving any issues you may have, please also refer to Diagnostics & Troubleshooting.

11.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all Virtual Services & the associated Real Servers (i.e. the Inception servers) and shows the state/health of each server as well as the overall state of each cluster. The example below shows that both Inception servers are healthy (green) and available to accept connections:



11.2. Accessing Ross Video Inception via the Load Balancer

Verify that users are able to successfully access Inception via the Virtual Service on the load balancer.

Web access - using a web browser, browse to the VIP address, e.g. **http://10.11.40.60**, log in, and confirm that Inception works as expected.

Persistence - while logged in, refresh the browser several times. Using ***Reports > Layer 7 Stick Table*** in the WebUI, select the **Inception** VIP and confirm that there is an entry whose ***Key*** is the client's IP address and whose ***Server*** is one of the Inception servers. The web access logs on the Inception servers should also show that the requests continue to be handled by that same server.

Failover - stop the Inception service on one server, or take the server offline. The System Overview should mark that server as down and users should be directed to the remaining server(s) when they reconnect. Restart the service and confirm that the server returns to the up state.

> **Note** Make sure that DNS is updated so that any FQDNs used point to the VIP rather than individual servers.

12. Technical Support

For more details about configuring the appliance and assistance with designing your deployment please don't hesitate to contact the support team using the following email address: support@loadbalancer.org.

13. Further Documentation

For additional information, please refer to the Administration Manual.

14. Appendix

14.1. Configuring HA - Adding a Secondary Appliance

Our recommended configuration is to use a clustered HA pair of load balancers to provide a highly available and resilient load balancing solution. We recommend that the Primary appliance is fully configured first, then the Secondary appliance can be added to create an HA pair. Once the HA pair is configured, load balanced services must be configured and modified on the Primary appliance. The Secondary appliance will be automatically kept in sync.

> **Note** For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the documentation library

The clustered HA pair uses Heartbeat to determine the state of the other appliance. Should the active device (normally the Primary) suffer a failure, the passive device (normally the Secondary) will take over.

14.1.1. Non-Replicated Settings

A number of settings are not replicated as part of the Primary/Secondary pairing process and therefore must be manually configured on the Secondary appliance. These are listed by WebUI menu option in the table below:

WebUI Main Menu Option	Sub Menu Option	Description
Local Configuration	Hostname & DNS	Hostname and DNS settings
Local Configuration	Network Interface Configuration	Interface IP addresses, bonding configuration and VLANs
Local Configuration	Routing	Default gateways and static routes
Local Configuration	System Date & time	Time and date related settings
Local Configuration	Physical – Advanced Configuration	Various appliance settings
Local Configuration	Portal Management	Portal management settings
Local Configuration	Security	Security settings
Local Configuration	SNMP Configuration	SNMP settings
Local Configuration	Graphing	Graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Backup & Restore	Local XML backups
Maintenance	Software Updates	Appliance software updates
Maintenance	Firewall Script	Firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

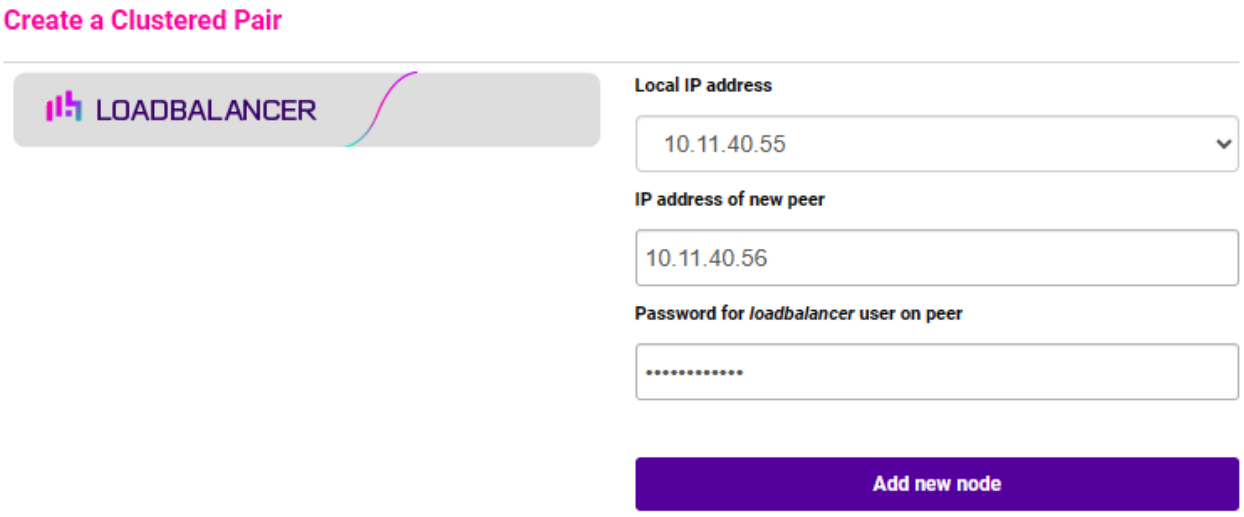
Important	Make sure that where any of the above have been configured on the Primary appliance, they're also configured on the Secondary.

14.1.2. Configuring the HA Clustered Pair

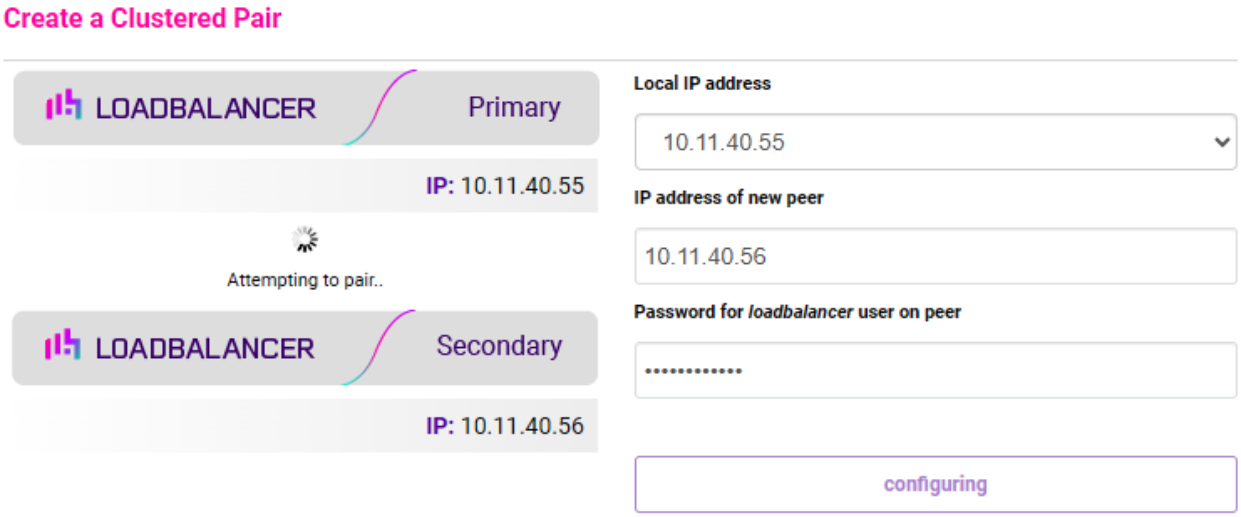
Important	During HA pairing, all WebUI users and passwords are synchronized from the Primary to the Secondary. After clustering completes (you will be logged out of the Secondary when this occurs), the Primary's credentials should be used to login to both nodes.

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.

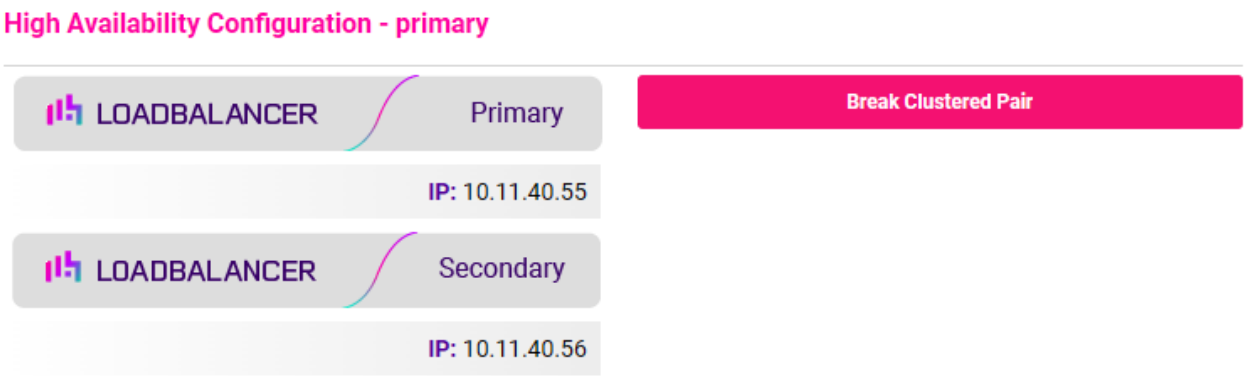
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: ***Cluster Configuration > High-Availability Configuration***.



3. Specify the IP address and the ***loadbalancer*** user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:



6. Once complete, the following will be displayed on the Primary appliance:



7. To finalize the configuration, restart heartbeat and any other services 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.

> **Note** For more details on configuring HA with 2 appliances, please refer to Configuring High Availability.

> **Note** For details on testing and verifying HA, please refer to Clustered Pair Diagnostics.

15. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	29 October 2026	Initial version		RJC

LOADBALANCER

Visit us: www.loadbalancer.org

Phone us: +44 (0)330 380 1064

Phone us: +1 833 274 2566

Email us: info@loadbalancer.org

Follow us: @loadbalancer.org

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.

