


# Load Balancing Cribl Stream

Version 1.0.0



# Table of Contents

1. About this Brief .....	3
2. Loadbalancer.org Appliances Supported .....	3
3. Software Versions Supported .....	3
3.1. Loadbalancer.org Appliance .....	3
3.2. Cribl Stream .....	3
4. Cribl Stream .....	3
5. Load Balancing Cribl Stream .....	3
5.1. Virtual Service (VIP) Requirements .....	4
6. Deployment Concept .....	4
7. Load Balancer Deployment Methods .....	5
7.1. Layer 7 SNAT Mode .....	5
8. Configuring Cribl Stream for Load Balancing .....	6
8.1. Layer 7 SNAT Mode .....	6
9. Loadbalancer.org Appliance – the Basics .....	6
9.1. Virtual Appliance .....	6
9.2. Initial Network Configuration .....	6
9.3. Accessing the Appliance WebUI .....	6
Main Menu Options .....	8
9.4. Appliance Software Update .....	9
Determining the Current Software Version .....	9
Checking for Updates using Online Update .....	9
Using Offline Update .....	9
9.5. Ports Used by the Appliance .....	10
9.6. HA Clustered Pair Configuration .....	11
10. Appliance Configuration for Cribl Stream .....	11
10.1. Verify that the Gateway Service Does not Clash with the Virtual Services .....	11
10.2. VIP 1 - Cribl_Workers .....	11
Virtual Service (VIP) Configuration .....	11
Define the Associated Real Servers (RIPs) .....	12
10.3. VIP 2 - Cribl_Leaders .....	13
Virtual Service (VIP) Configuration .....	13
Define the Associated Real Servers (RIPs) .....	14
10.4. Finalizing the Configuration .....	14
11. Testing & Verification .....	15
11.1. Accessing Cribl Stream via the Load Balancer .....	15
11.2. Using System Overview .....	15
12. Technical Support .....	15
13. Further Documentation .....	15
14. Appendix .....	17
14.1. Configuring HA - Adding a Secondary Appliance .....	17
Non-Replicated Settings .....	17
Configuring the HA Clustered Pair .....	18
15. Document Revision History .....	20

# 1. About this Brief

This brief outlines the steps required to configure a load balanced Cribl Stream environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Cribl Stream 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 Cribl Stream. For full specifications of available models please refer to <https://www.loadbalancer.org/products>.

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

- V8.9.1 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. Cribl Stream

- All versions

## 4. Cribl Stream

Cribl Stream is a vendor-agnostic data pipelining engine for routing and processing IT and security data. With Stream, you can route, reduce, reformat, enrich, or shape data from any source to any destination.

## 5. Load Balancing Cribl Stream

#### Note

It's highly recommended that you have a working Cribl Stream environment first before implementing the load balancer.

Load balancing in Cribl Stream is essential for ensuring efficient, scalable, and resilient data processing. Depending on system architecture and the number of nodes, Load balancing can be configured for the following Cribl Stream components:



**Worker Nodes** : The load balancer distributes data across multiple worker nodes to optimize performance and maintain system availability. For more information, see [Distributed Worker Deployment](#).

**Leader Nodes** : In an HA deployment, 2 Leader nodes are deployed, one Primary and the other Standby. There is only ever one active Leader Node at one time. The standby Leader will become active only in the event of failover. For more information, see [Leader High Availability/Failover](#).

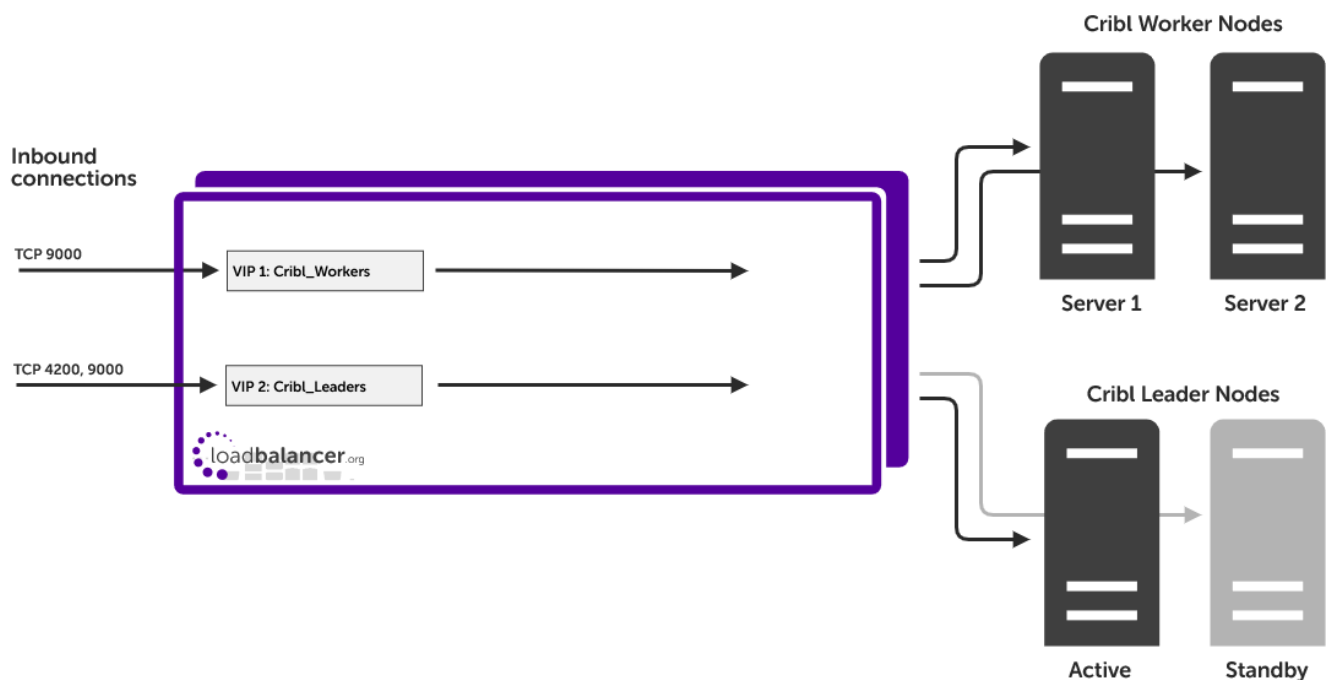
## 5.1. Virtual Service (VIP) Requirements

To provide load balancing and HA for Cribl Stream, the following VIPs are required:

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	Cribl_Workers	L7 SNAT	9000	Source IP	HTTPS (GET)
VIP 2	Cribl_Leaders	L7 SNAT	4200,9000	Source IP	HTTPS (GET)

## 6. Deployment Concept

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



VIP = **V**irtual **I**P Address

### 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.

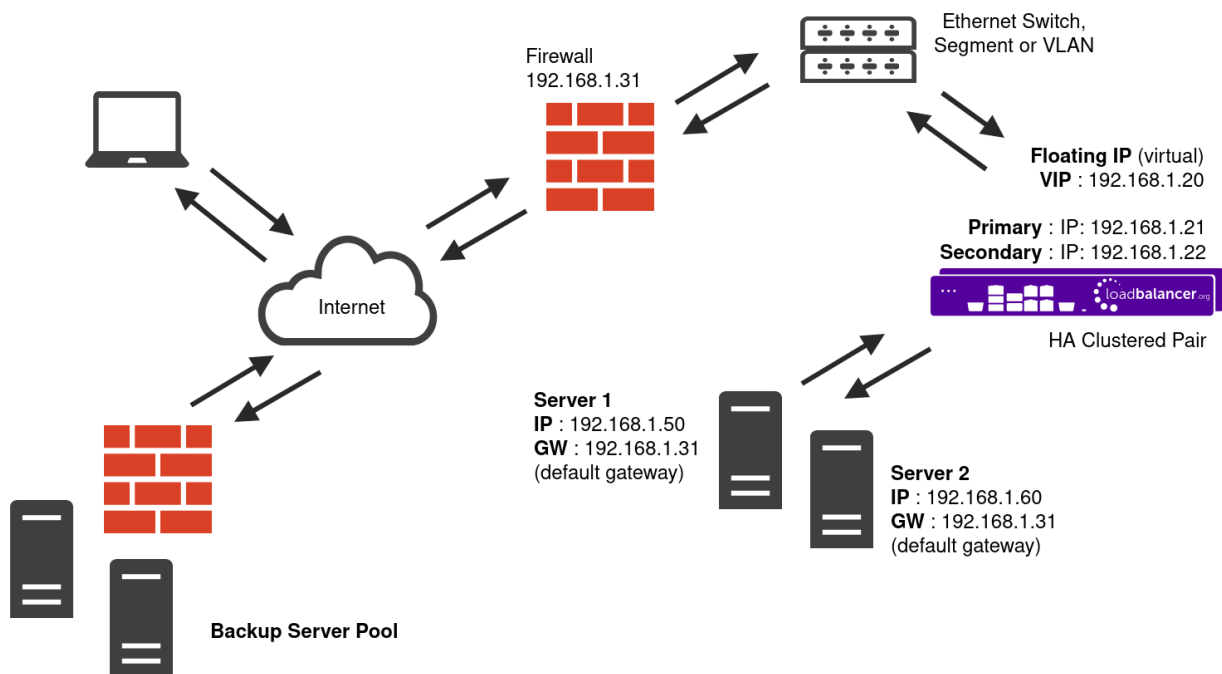
# 7. Load Balancer Deployment Methods

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 Cribl Stream, 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.



- 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.

## 8. Configuring Cribl Stream for Load Balancing

### 8.1. Layer 7 SNAT Mode

Layer 7 SNAT mode VIPs do not require any mode specific configuration changes to the load balanced Real Servers (the Cribl Worker and Leader nodes).

## 9. Loadbalancer.org Appliance – the Basics

### 9.1. 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.

### 9.2. 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.

### 9.3. 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:

Primary | Secondary    Active | Passive    Link    15 Seconds ↻

System Overview

Local Configuration

Cluster Configuration

Maintenance

View Configuration

Reports

Logs

Support

Live Chat

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

**System Overview** ? 2024-03-15 16:27:21 UTC

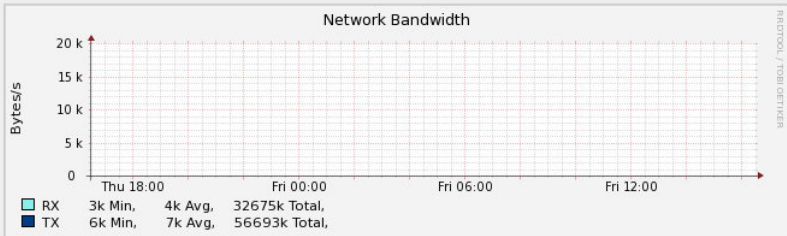
Would you like to run the Setup Wizard?

Accept
Dismiss

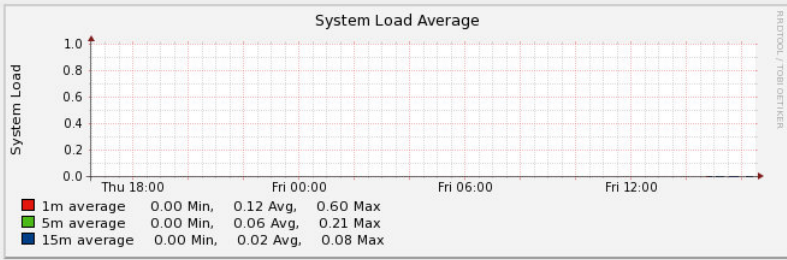
VIRTUAL SERVICE    IP    PORTS    CONNS    PROTOCOL    METHOD    MODE

No Virtual Services configured.

Network Bandwidth



System Load Average



Memory Usage

- You'll be asked if you want to run the Setup Wizard which can be used to configure layer 7 services. Click **Dismiss** if you're following a guide or want to configure the appliance manually or click **Accept** to start the wizard.

## Main Menu Options

**System Overview** - Displays a graphical summary of all VIPs, RIPv 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 & RIPv

**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





## 9.4. Appliance Software Update

To ensure that the appliance(s) are running the latest software version, we recommend a software update check is performed.

### 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 – 2024

ENTERPRISE VA Max - v8.11.4

English ▼

### Checking for Updates using Online Update

#### Note

By default, the appliance periodically contacts the Loadbalancer.org update server and checks for updates. An update check can also be manually triggered as detailed below.

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 v8.11.4 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.

### Using 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:

### Software Update

#### 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 chosen

Checksum:  No file chosen

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.

## 9.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22 *	SSH
TCP & UDP	53 *	DNS / GSLB
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	9000 *	Gateway service (Centralized/Portal Management)
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page



Protocol	Port	Purpose
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

### Note

The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to [Service Socket Addresses](#).

## 9.6. 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.

# 10. Appliance Configuration for Cribl Stream

## 10.1. Verify that the Gateway Service Does not Clash with the Virtual Services

By default, the Gateway Service on the load balancer is bound to the management IP address and listens on port 9000. If for some reason this has been changed from port 9000 to all ports (\*), this will clash with the Virtual Services. To check this setting:

1. Using the WebUI, navigate to: *Local Configuration > Physical - Advanced Configuration* and scroll to *Service Socket Addresses*.
2. Ensure that the Gateway Service is not listening on all IP addresses (\*). If it is, change this to an appropriate interface IP address.

## 10.2. VIP 1 - Cribl\_Workers

### 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:

Virtual Service		[Advanced +]
Label	<input type="text" value="Cribl_Workers"/>	?
IP Address	<input type="text" value="192.168.2.150"/>	?
Ports	<input type="text" value="9000"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

3. Enter an appropriate *label* for the VIP, e.g. **Cribl\_Workers**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.2.150**.
5. Set the *Virtual Service Ports* field to **9000**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Click **Modify** next to the newly created VIP.
9. Scroll to the *Persistence* section.
  - Ensure that the *Persistence Mode* is set to **Source IP**.
10. Scroll to the *Health Checks* section.
  - Set the *Health Checks* to **Negotiate HTTPS (GET)**.
  - Set the *Request to send* to **/api/v1/health**.
  - Set the *Response Expected* to **Equals** and set the value to **healthy**.
11. Leave all other settings at their default value.
12. Click **Update**.

## Define the Associated Real Servers (RIPs)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	<input type="text" value="Worker1"/>	?
Real Server IP Address	<input type="text" value="192.168.2.160"/>	?
Real Server Port	<input type="text" value="9000"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Define the *Label* for the Real Server as required, e.g. **Worker1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.1.160**.
5. Set the *Real Server Port* field to **9000**.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining Worker node(s).

### 10.3. VIP 2 - Cribl\_Leaders

#### 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:

<b>Virtual Service</b>		[Advanced +]
Label	<input type="text" value="Cribl_Leaders"/>	?
IP Address	<input type="text" value="192.168.2.155"/>	?
Ports	<input type="text" value="4200,9000"/>	?
<b>Protocol</b>		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

3. Enter an appropriate *label* for the VIP, e.g. **Cribl\_Leaders**.
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.2.155**.
5. Set the *Virtual Service Ports* field to **4200,9000**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.

7. Click **Update**.
8. Click **Modify** next to the newly created VIP.
9. Scroll to the *Persistence* section.
  - Ensure that the *Persistence Mode* is set to **Source IP**.
10. Scroll to the *Health Checks* section and click **[Advanced]**.
  - Set the *Health Checks* to **Negotiate HTTPS (GET)**.
  - Set the *Request to send* to **/api/v1/health**.
  - Set the *Response Expected* to **Equals** and set the value to **healthy**.
11. Leave all other settings at their default value.
12. Click **Update**.

### Define the Associated Real Servers (RIPs)

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

Label	<input type="text" value="Leader1"/>	?
Real Server IP Address	<input type="text" value="192.168.2.165"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

3. Define the *Label* for the Real Server as required, e.g. **Leader1**.
4. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.1.165**.
5. Leave the *Real Server Port* field blank.
6. Leave all other settings at their default value.
7. Click **Update**.
8. Repeat these steps to add the remaining Leader node(s).

## 10.4. 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:

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

## 11. Testing & Verification

### Note

For additional guidance on diagnosing and resolving any issues you may have, please also refer to [Diagnostics & Troubleshooting](#).

### 11.1. Accessing Cribl Stream via the Load Balancer

Verify that you're able to successfully access all load balanced applications and services via the Virtual Services on the load balancer.

### Note

Make sure that DNS is updated so that any FQDNs used point to the VIPs rather than individual servers.

### 11.2. 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 Worker and Leader nodes) and shows the state/health of each server as well as the overall state of each cluster. The example below shows that all nodes are healthy (green) and available to accept connections:

#### System Overview ?

2024-09-26 15:13:18 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
	<b>Cribl_Workers</b>	192.168.2.150	9000	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	Worker1	192.168.2.160	9000	100	0	Drain	Halt	
	Worker2	192.168.2.161	9000	100	0	Drain	Halt	
	<b>Cribl_Leaders</b>	192.168.2.155	4200,9000	0	TCP	Layer 7	Proxy	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	Leader1	192.168.2.165	4200,9000	100	0	Drain	Halt	
	Leader2	192.168.2.166	4200,9000	100	0	Drain	Halt	

## 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](mailto: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.

### 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	Fallback Page	Fallback page configuration
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.

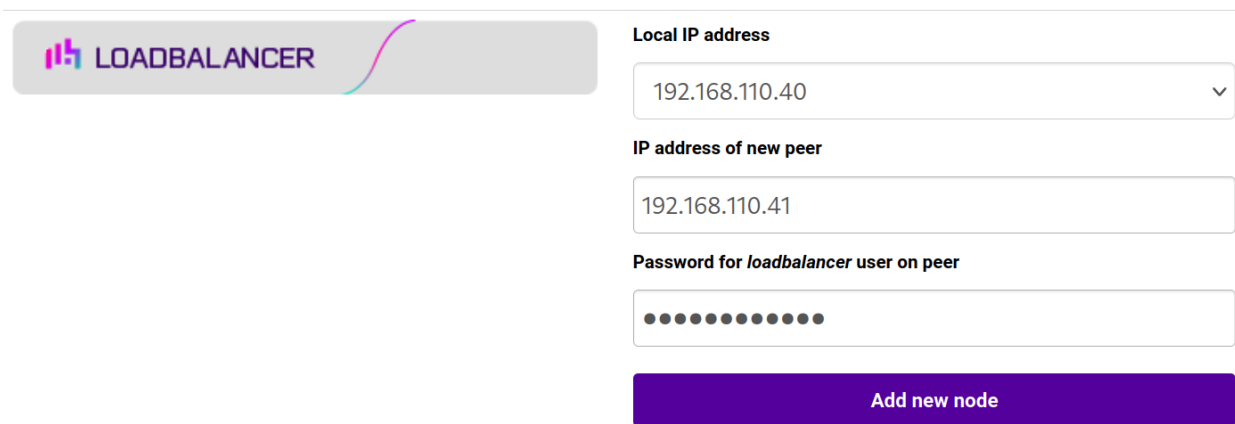
## Configuring the HA Clustered Pair

### 📌 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.

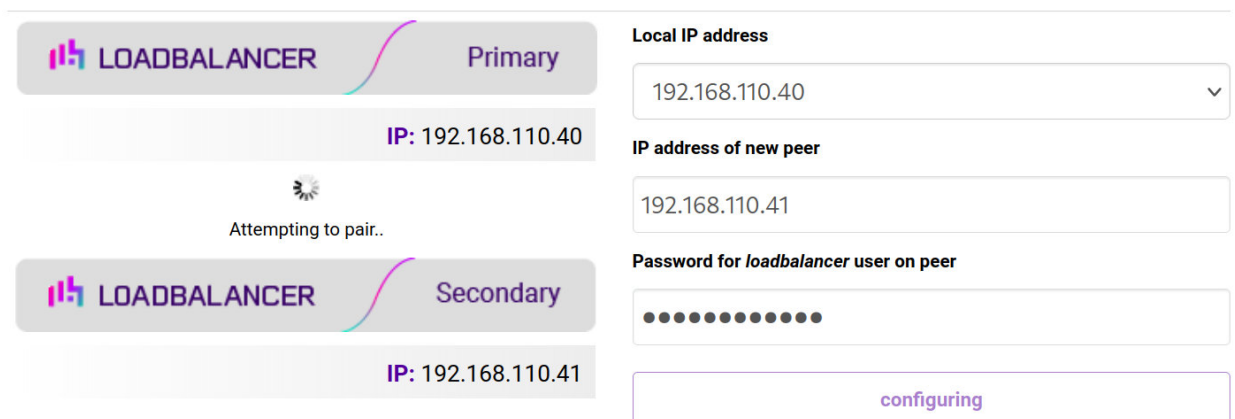
1. Deploy a second appliance that will be the Secondary and configure initial network settings.
2. Using the WebUI on the Primary appliance, navigate to: *Cluster Configuration > High-Availability Configuration*.

### Create a Clustered Pair



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:

### Create a Clustered Pair



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

## High Availability Configuration - primary

The screenshot displays a configuration interface for a High Availability (HA) setup. It features two load balancer appliances, each represented by a grey rounded rectangle with a purple icon and the text 'LOADBALANCER'. The top appliance is labeled 'Primary' and has the IP address '192.168.110.40'. The bottom appliance is labeled 'Secondary' and has the IP address '192.168.110.41'. To the right of these appliances is a prominent red button with the text 'Break Clustered Pair'.

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 [Appliance Clustering for HA](#).

### 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	03 October 2024	Initial version		RJC





**Visit us:** [www.loadbalancer.org](http://www.loadbalancer.org)

**Phone us:** +44 (0)330 380 1064

**Phone us:** +1 833 274 2566

**Email us:** [info@loadbalancer.org](mailto:info@loadbalancer.org)

**Follow us:** [@loadbalancer.org](https://twitter.com/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.

