Load Balancing Nginx Web Servers with OWASP Top 10 WAF in AWS

Quick Reference Guide

v1.0.2

ABOUT THIS GUIDE

This document provides a quick reference guide on how to load balance Nginx Web Servers and configure a WAF using the Enterprise AWS Loadbalancer.org Amazon cloud appliance.

- The WAF addresses the OWASP Top 10 vulnerabilities and is very quick and simple to deploy
- SSL offload is handled by STunnel, HAProxy handles back-end server re-encryption

RELATED DOCUMENTATION

For additional information about the Loadbalancer.org AWS Appliance, please also refer to the following documents:

- Administration Manual
- AWS Quick Start Guide

LOAD BALANCED PORTS

Port	Use	Transport Layer Protocol
80	HTTP	TCP
443	HTTPS	TCP

VPC SECURITY GROUP INBOUND RULES

The following inbound rules must be configured in your Security Group:

- For Management: TCP 22 (SSH), TCP 9443 (Appliance WebUI), 7777 (HAProxy Stats page)
- For Nginx services: TCP 80 (HTTP), TCP 443 (HTTPS)

LOAD BALANCER CONFIGURATION

DEPLOY THE LOADBALANCER.ORG AWS APPLIANCE

1. Deploy an AWS Loadbalancer.org appliance as detailed in the Quick Start Guide

ACCESSING THE APPLIANCE WEBUI

Using a browser, navigate to the Public DNS name or Public IP address on port 9443 , i.e.

https://<Public DNS name>:9443
or
https://<Public IP address>:9443

You'll receive a warning about the certificate as it's a self signed cert not related to an Internet based CA. Confirm you want to continue and a login prompt will be displayed. Use the following default credentials:

Username: loadbalancer
Password: <EC2 Instance-ID>

> Note:
>
> To change the password for the 'loadbalancer' account, use the WebUI option: *Maintenance > Passwords*.

Once logged in, the WebUI is displayed:

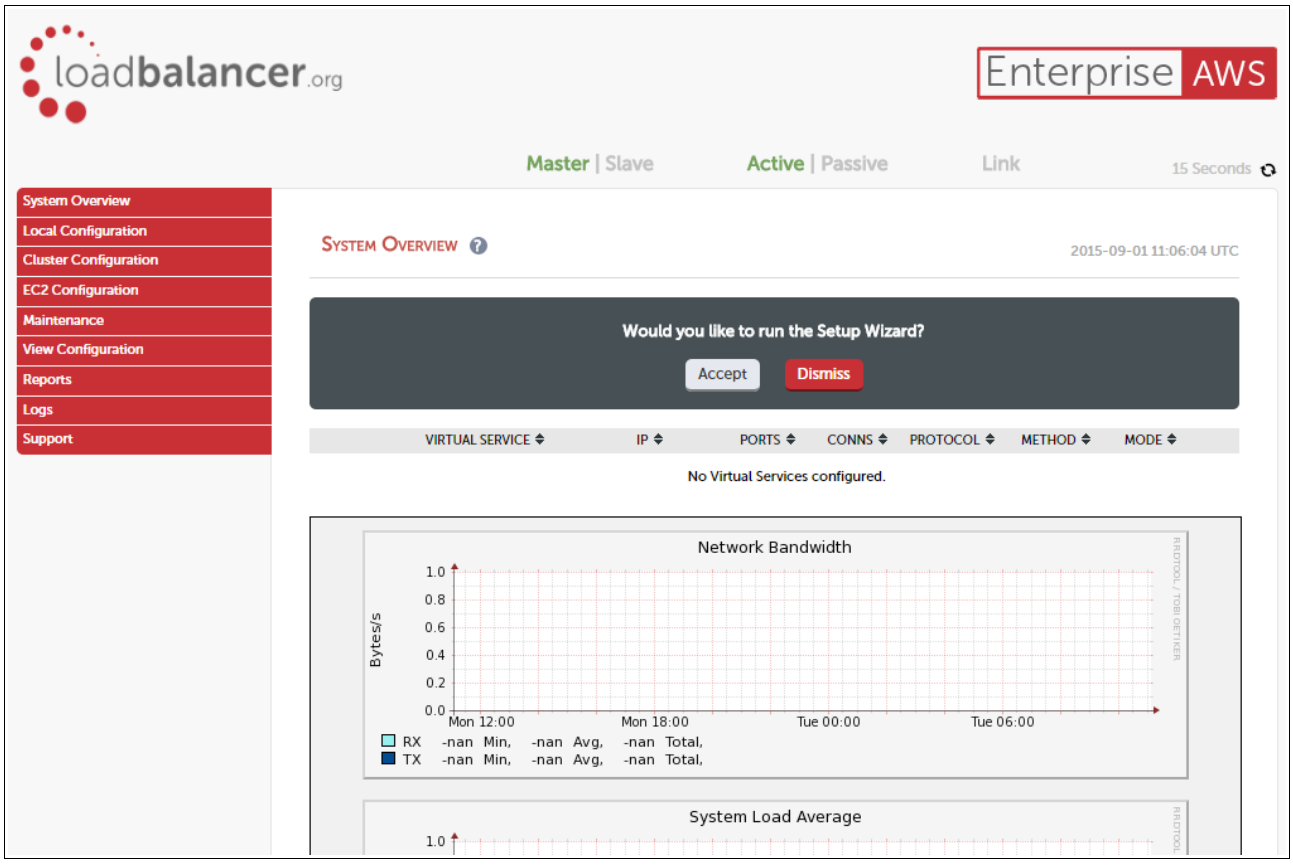

CONFIGURATION DIAGRAM

The diagram below shows how the system is configured.

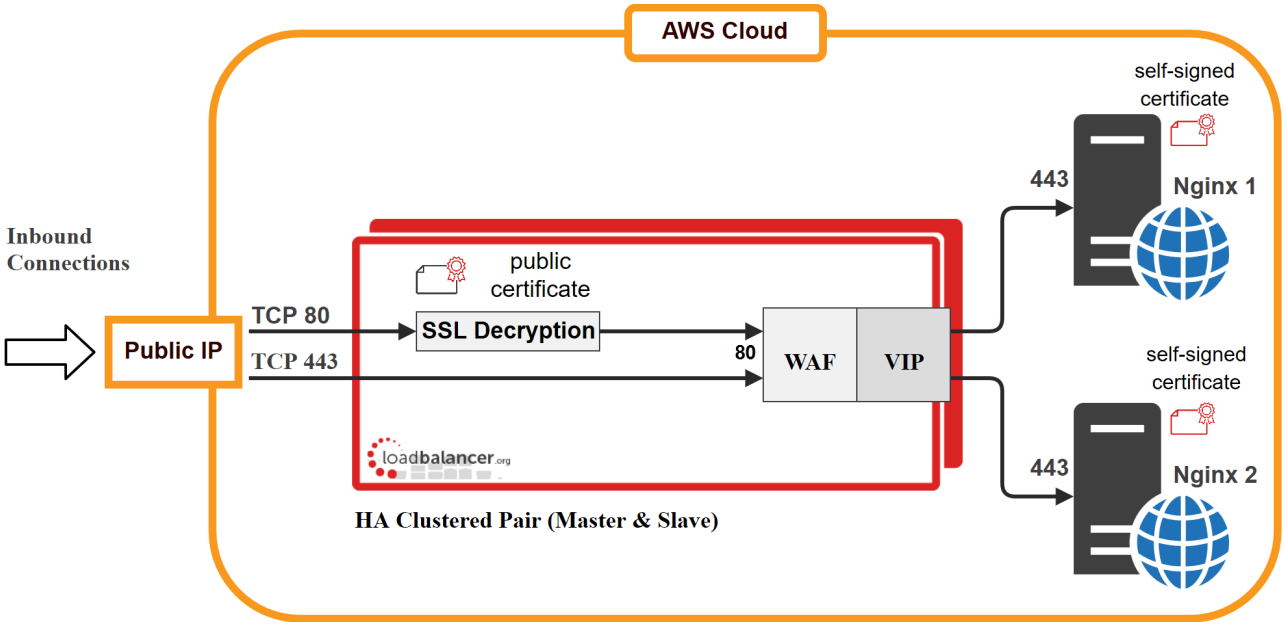

CONFIGURE THE VIRTUAL SERVICE (VIP)

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

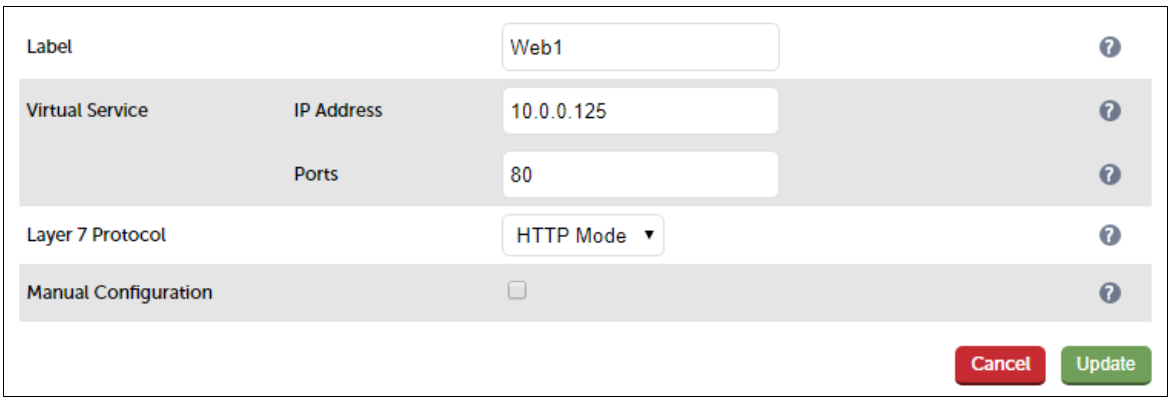

3. Enter the required *Label* (name) for the VIP, e.g. **Web1**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **10.0.0.125**
5. Set the *Virtual Service Ports* field to the required port, e.g. **80**
6. Leave *Layer 7 Protocol* set to **HTTP Mode**
7. Click **Update**

DEFINE THE REAL (NGINX) SERVERS

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:

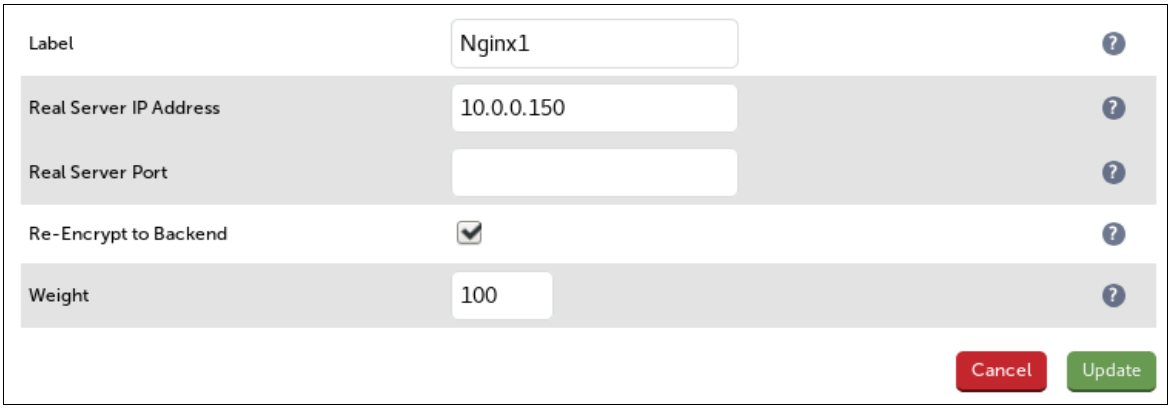

3. Enter an appropriate label for the Real Server , e.g. Nginx1
4. Set the *Real Server IP Address* field to the required address, e.g. **10.0.0.150**
5. Leave the *Real Server Port* field blank
6. Enable (check) **Re-Encrypt to Backend**
7. Click **Update**
8. Repeat the above steps to add your other Nginx server(s)

UPLOAD THE PUBLIC SSL CERTIFICATE

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Certificate* and click **Add a New SSL Certificate**
2. Select **Upload prepared PEM/PFX file**
3. Enter the following details:

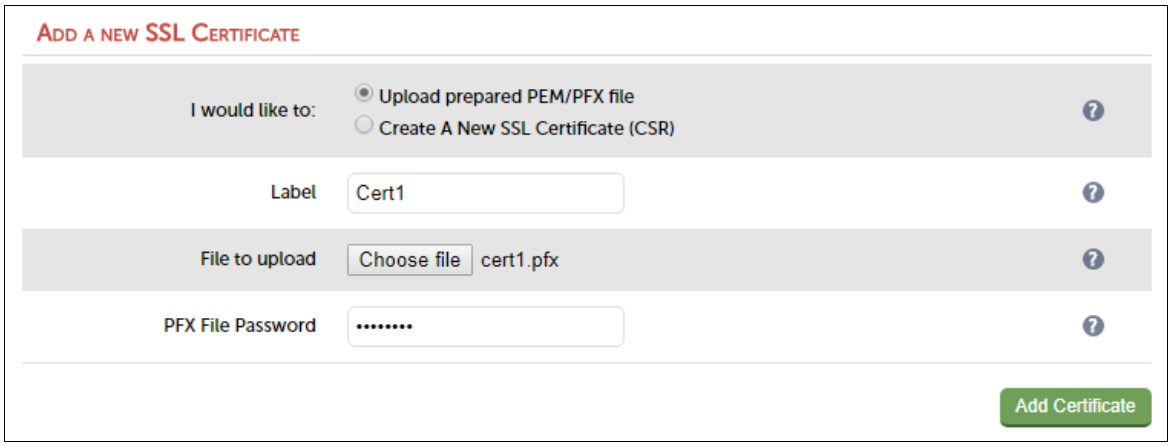

4. Specify and Label (name) for the certificate, e.g. **Cert1**
5. Click **Choose File** and browse to and select the relevant PFX or PEM file
6. Enter the *PFX file Password*
7. Click **Add Certificate**

CONFIGURE THE STUNNEL VIRTUAL SERVICE (VIP)

STunnel is used to terminate SSL on the load balancer.

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Termination* and click **Add a New Virtual Service**

2. Enter the following details:

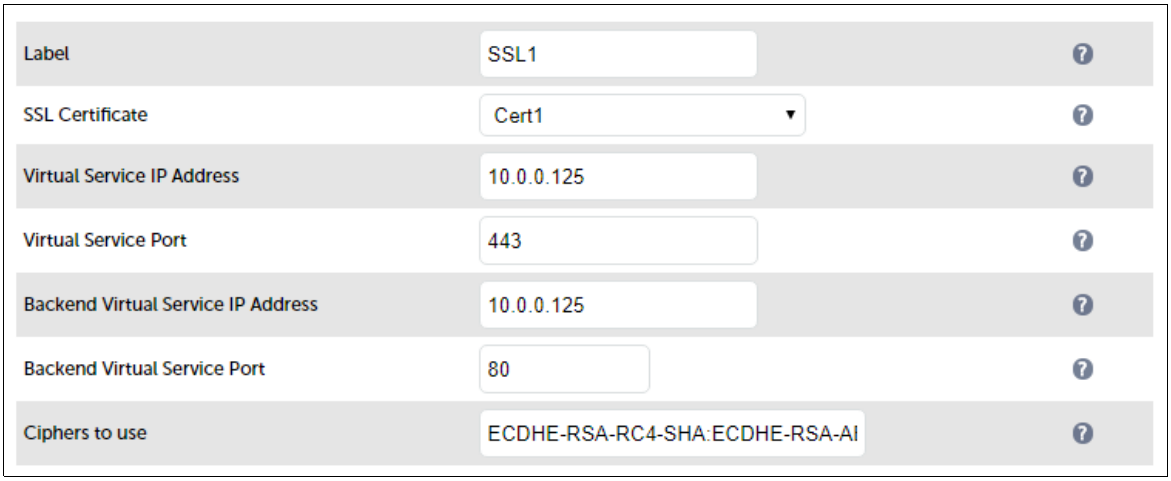

3. Enter the required *Label* (name) for the Virtual Service, e.g. **SSL1**
4. Select the required certificate in the *SSL Certificate* drop-down
5. Set the *Virtual Service IP address* to be the same as the VIP created previously, e.g. **10.0.0.125**
6. Set the *Virtual Service Port* field to **443**
7. Set the *Backend Virtual Service IP address* to be the same as the VIP created previously, e.g. **10.0.0.125**
8. Set the *Backend Virtual Service Port* field to **80**
9. The other settings can typically be left at their default values
10. Click **Update**

CONFIGURE THE WAF

1. Using the WebUI, navigate to: *Cluster Configuration > WAF – Gateway* and click **Add a New WAF Gateway**
2. Enter the following details:

3. Select the VIP created previously, e.g. **Web1**
4. Specify a suitable *WAF label* (name), e.g. **WAF1**

5. Leave *Rule Engine Traffic Blocking* unchecked for now

> Note:
>
> While disabled, this option ensures that the ModSecurity Rule Engine logs any critical errors. You should leave the WAF in this mode until you are confident that the error logs are not showing false positives. Once you are confident you can enable this mode and the WAF will start blocking any malicious requests with a 403 Forbidden response.

6. Click **Update**

APPLY THE NEW SETTINGS

1. Once the configuration is complete, use the **Reload HAProxy**, **Restart STunnel** and **Reload WAF** buttons at the top of the screen to commit the changes.

ASSOCIATE AN EIP WITH THE VIRTUAL SERVICES

An EIP is added and associated with the VIP to provide a public IP address for client connections.

1. Using the WebUI, navigate to: *EC2 > EC2 Network Configuration*
2. Click **Allocate New Elastic IP** , this will request an EIP from Amazon using API calls
3. Click **[Associate]** to associate the EIP to the Virtual Service private IP address

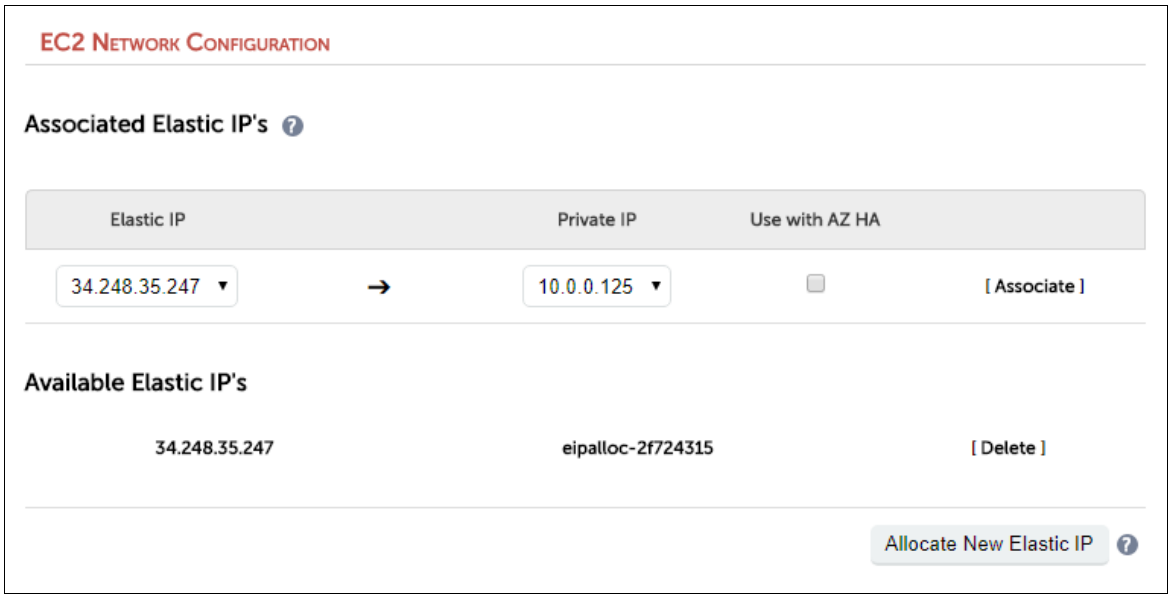

This association is then displayed as shown below:

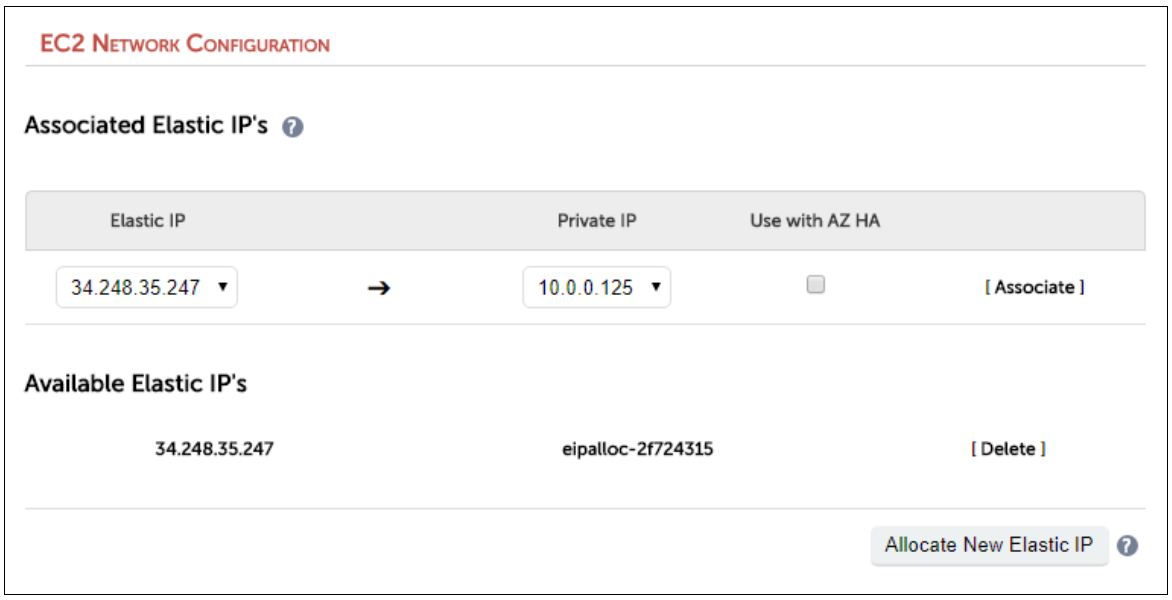

TESTING

The load balanced Nginx Web Servers should now be accessible on ports 80 & 443 using the EIP address or corresponding public DNS name.

LOGGING CLIENT SOURCE IP ADDRESSES IN NGINX

The Nginx service on a web server can be configured to store the value of X-Forwaded-For headers for incoming web traffic. These headers are added by default by the load balancer. This allows upstream servers and network devices to see the real source IP addresses of clients, even though the load balancer is acting as a proxy.

For full details on how to configure Nginx for this, see our blog post:
https://www.loadbalancer.org/blog/nginx-and-x-forwarded-for-header/

LOADBALANCER.ORG TECHNICAL SUPPORT

If you have any questions regarding the appliance or would like assistance designing your deployment, please don't hesitate to contact our support team: support@loadbalancer.org.