DEPLOYMENT GUIDE

Load Balancing Microsoft Always On VPN

v1.0.2

Quick Reference Guide

loadbalancer.org

Contents

1. About this Guide

This guide details the steps required to configure a load balanced Microsoft Always On VPN environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Microsoft Always On VPN 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 relevant Administration Manual:

- v7 Administration Manual
- v8 Administration Manual

2. Loadbalancer.org Appliances Supported

All our products can be used with Always On VPN. The complete list of models is shown below:

Discontinued Models	Current Models *
Enterprise R16	Enterprise R20
Enterprise VA R16	Enterprise MAX
Enterprise VA	Enterprise 10G
Enterprise R320	Enterprise 40G
	Enterprise Ultra
	Enterprise VA R20
	Enterprise VA MAX
	Enterprise AWS
	Enterprise AZURE **
	Enterprise GCP **

* For full specifications of these models please refer to: http://www.loadbalancer.org/products/hardware

** Some features may not be supported, please check with Loadbalancer.org support

3. Loadbalancer.org Software Versions Supported

- V8.4.1 and later

4. Microsoft Windows Versions Supported

- Windows 2016 and later

5. Microsoft Always On VPN

Introduction

Always On VPN provides a single, cohesive solution for remote access and supports domain-joined, nondomain-joined (workgroup), or Azure AD–joined devices, even personally owned devices. With Always On VPN, the connection type does not have to be exclusively user or device but can be a combination of both. For example, you could enable device authentication for remote device management and then enable user authentication for connectivity to internal company sites and services.

Always On VPN Components

Always On VPN is part of the Remote Access server role. The table below details the key components that must be available for Always On VPN to work.

These are the components that are made highly available using the load balancer:

Component	Purpose
Routing and Remote Access Servers (RRAS)	An Always On VPN deployment may require more than one RRAS server to provide redundancy or to increase capacity to service more VPN connections than a single server is capable of
Network Policy Servers (NPS)	To authenticate VPN connections, VPN servers are configured to forward authentication requests to an NPS server. Having more than one NPS server eliminates this single point of failure and may be required to support authentication for large scale deployments
Multisite redundancy	Unlike DirectAccess, Always On VPN has no concept of "multisite" configuration. To provide geographic redundancy, multiple VPN servers can be configured in various locations using a single, common public hostname. VPN client connections can then be routed to the most preferred location using the GSLB feature on the load balancer

How it Works

Using public DNS servers, the Windows 10 VPN client performs a name resolution query for the IP address of the VPN gateway.

Using the IP address returned by DNS, the VPN client sends a connection request to the VPN gateway.

The VPN gateway is also configured as a Remote Authentication Dial-In User Service (RADIUS) Client; the VPN RADIUS Client sends the connection request to the organization/corporate NPS server for connection request processing.

The NPS server processes the connection request, including performing authorization and authentication, and determines whether to allow or deny the connection request.

The NPS server forwards an Access-Accept or Access-Deny response to the VPN gateway.

The connection is initiated or terminated based on the response that the VPN server received from the NPS server.

6. Always On VPN Prerequisites

Several prerequisites must be in place before proceeding with this documentation. As such, it is assumed that the load balancer has been configured and that network connectivity to all networks has been validated. In addition, the following prerequisites must be in place before continuing:

- A public hostname for the VPN server which resolves to the IP address assigned to the VPN virtual service (or edge firewall if the load balancer is in a perimeter or DMZ network)
- An SSL certificate with a subject name that matches the VPN server's public hostname
- Each VPN server must be configured to assign unique IP addresses to its clients. Using DHCP for VPN client address assignment when there is more than one VPN server in a cluster is not supported
- An internal hostname for the NPS cluster which resolves to the IP address assigned to the NPS virtual service

7. Load Balancing Always On VPN

> Note: It's highly recommended that you have a working Always On VPN environment first before implementing the load balancer.

Basic Concepts

To provide resilience and high availability for your Always On VPN infrastructure, multiple Always On VPN servers should be deployed with a load balancer. This helps ensure that users can always connect to the corporate network by constantly checking the health of the Always On VPN servers and only forwarding connections to functional servers.

Load Balancer Deployment

The following diagram shows a typical load balanced Always On VPN deployment.

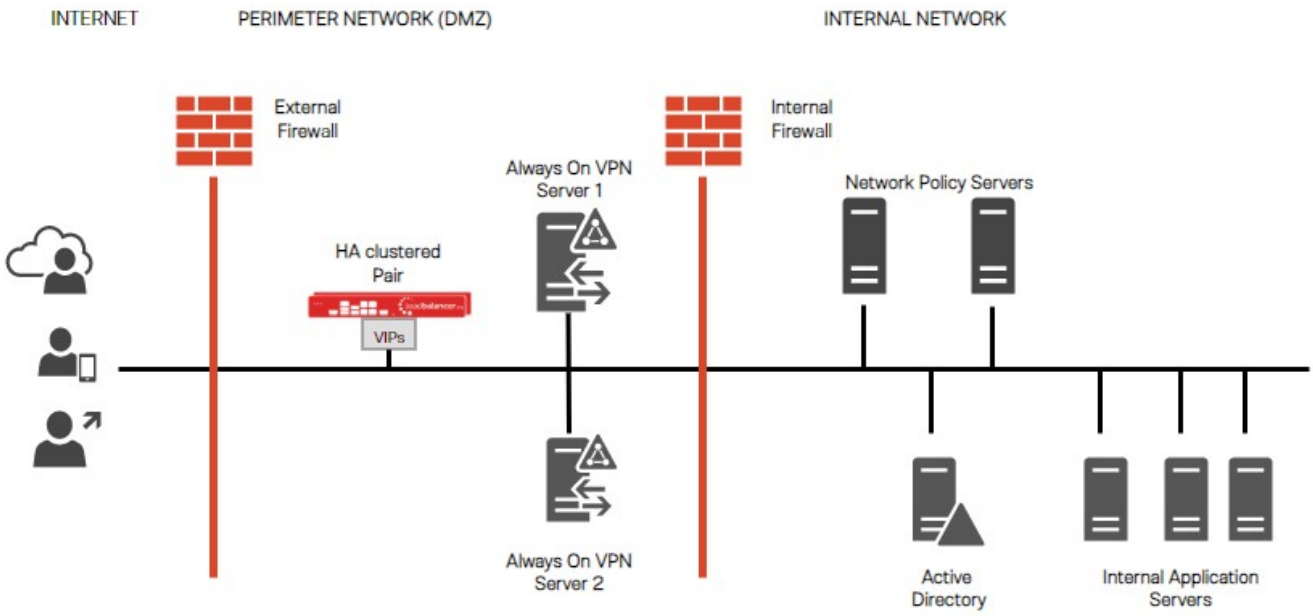

Notes:

- Load balancers can be deployed as single units or as a clustered pair. Loadbalancer.org recommends deploying a clustered pair for HA and resilience

Load Balancer Deployment Mode

Layer 4 SNAT mode is recommended for Always On VPN and is used for the configuration presented in this guide. This mode offers good performance and is simple to configure since it requires no configuration changes to the Always On VPN and Network Policy servers.

Layer 4 DR mode and NAT mode can also be used if preferred. To use DR mode it is *required* to solve the ARP problem on each Always On VPN server (please see the Administration Manual and search for "DR mode considerations"), while using NAT mode requires the default gateway on each real server to be configured to be the load balancer.

Load Balanced Ports & Services

The following ports/protocols must be load balanced:

Port	Protocol	Use
80	TCP/HTTP	All Always On VPN client to server communication when SSL Offloading
443	TCP/HTTPS	All Always On VPN client to server SSTP communication
500, 4500	UDP/IKEv2	IKEv2 communication
1812,1813	UDP	Network policy server communication

Persistence (Server Affinity)

Source IP address persistence is used for the Always On VPN servers. This ensures that a particular client will connect to the same Always On VPN server for the duration of the session and the Always On VPN server will connect to the same Network Policy server.

Server Health Checking

To constantly check and verify the health of the Always On VPN and Network Policy servers, the load balancer is configured to perform either a *Negotiate* check or a *Connect to port* check.

SSL Offloading

SSL pass-through is the recommend deployment method when load balancing Always On VPN servers (terminating SSL connections on the Real Servers). However, SSL offloading is possible if the need to reduce CPU utilization on the Real Servers is required.

8. Loadbalancer.org Appliance – the Basics

Virtual Appliance Download & Deployment

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 and XEN and has been optimized for each Hypervisor. By default, the VA is allocated 1 CPU, 2GB of RAM and has an 8GB 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 the Administration Manual and the ReadMe.txt text file included in the VA download for more detailed information on deploying the VA using various Hypervisors.

Initial Network Configuration

The IP address, subnet mask, default gateway and DNS settings can be configured in several ways as detailed below:

Method 1 - Using the Network Setup Wizard at the console

After boot up, follow the instructions on the console to configure the IP address, subnet mask, default gateway and DNS settings.

Method 2 - Using the WebUI

Using a browser, connect to the WebUI on the default IP address/port: **https://192.168.2.21:9443**

To set the IP address & subnet mask, use: *Local Configuration > Network Interface Configuration*

To set the default gateway, use: *Local Configuration > Routing*

To configure DNS settings, use: Local Configuration > Hostname & DNS

Accessing the Web User Interface (WebUI)

The WebUI can be accessed via HTTPS at the following URL: **https://192.168.2.21:9443/lbadmin**

** Note the port number → **9443***

(replace 192.168.2.21 with the IP address of your load balancer if it's been changed from the default)

Login using the following credentials:

Username: loadbalancer

Password: loadbalancer

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

Once logged in, the WebUI will be displayed as shown on the following page:

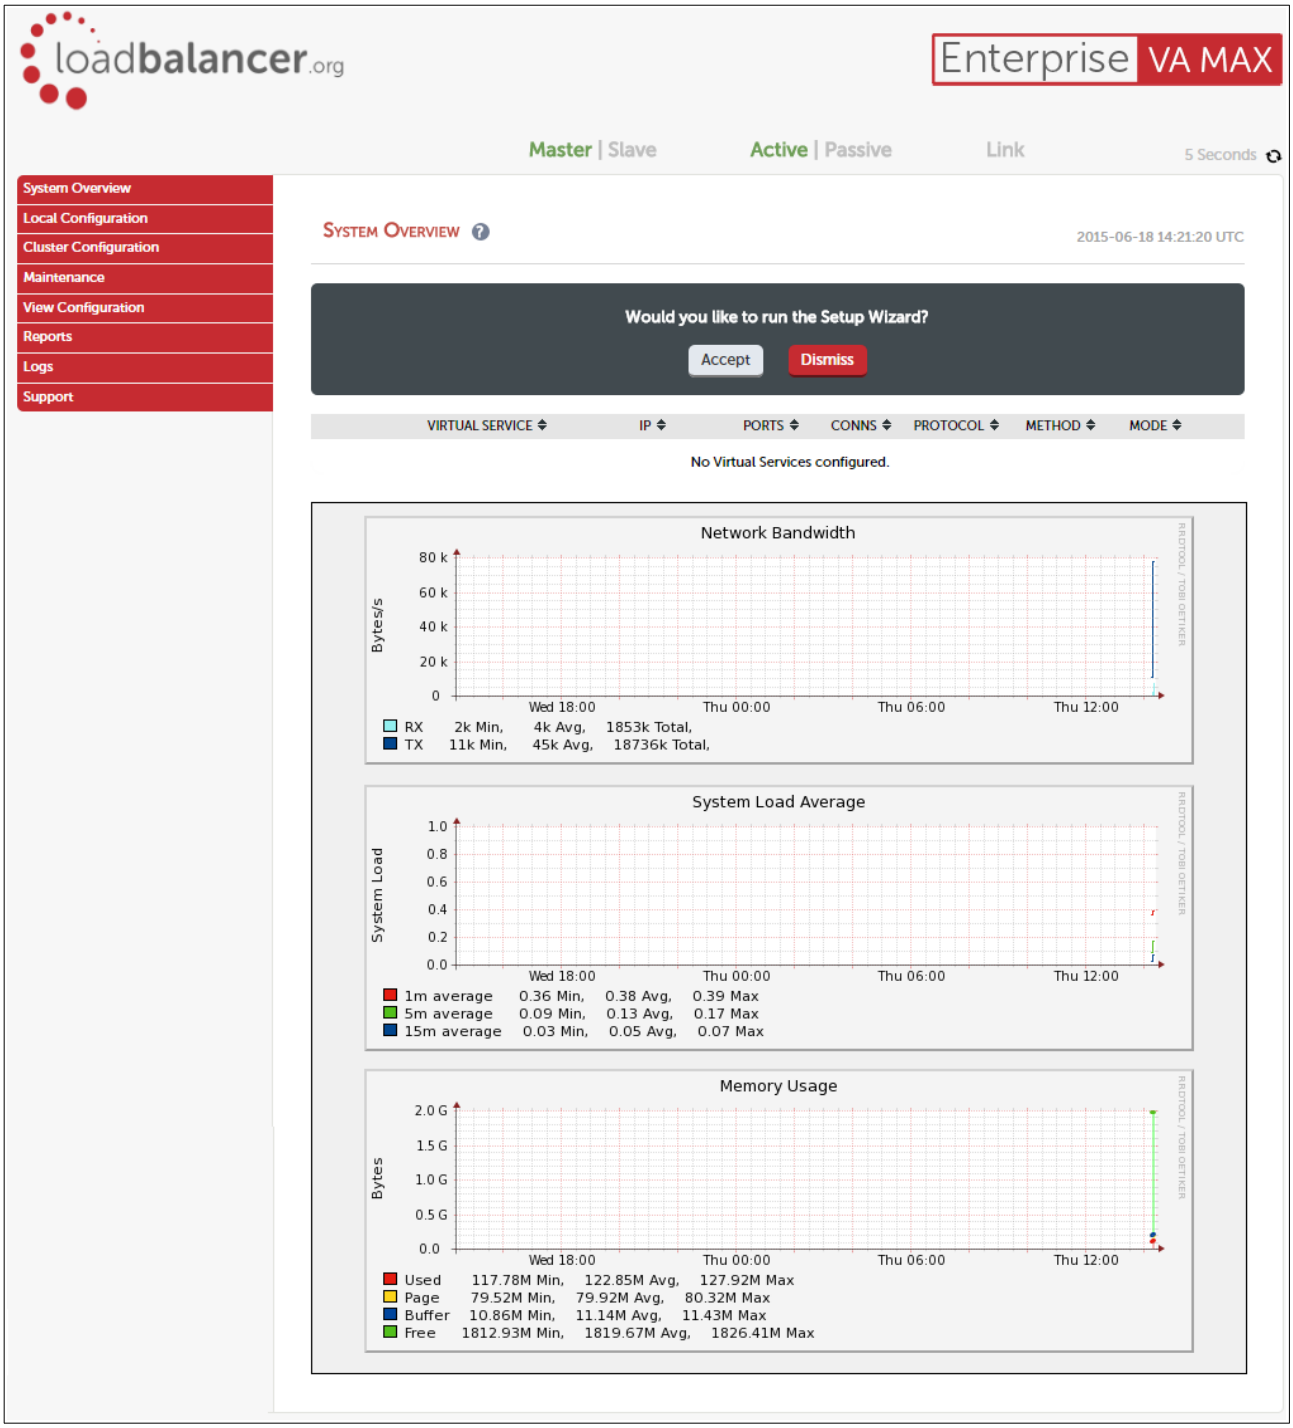

HA Clustered Pair Configuration

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 slave unit is covered in section 1 of the Appendix on page 19.

9. Appliance Configuration for Always On VPN

Configuring load balancing for Always On VPN is done in 3 steps:

A) Configuring load balancing for IKEv2

B) Configuring load balancing for SSTP (Secure Socket Tunneling Protocol)

C) Configuring load balancing for NPS (Network Policy Server)

A – Configuring load balancing for IKEv2

Setting up the Virtual Service (VIP)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Services* and click **Add a New Virtual Service**
2. Enter the following details:

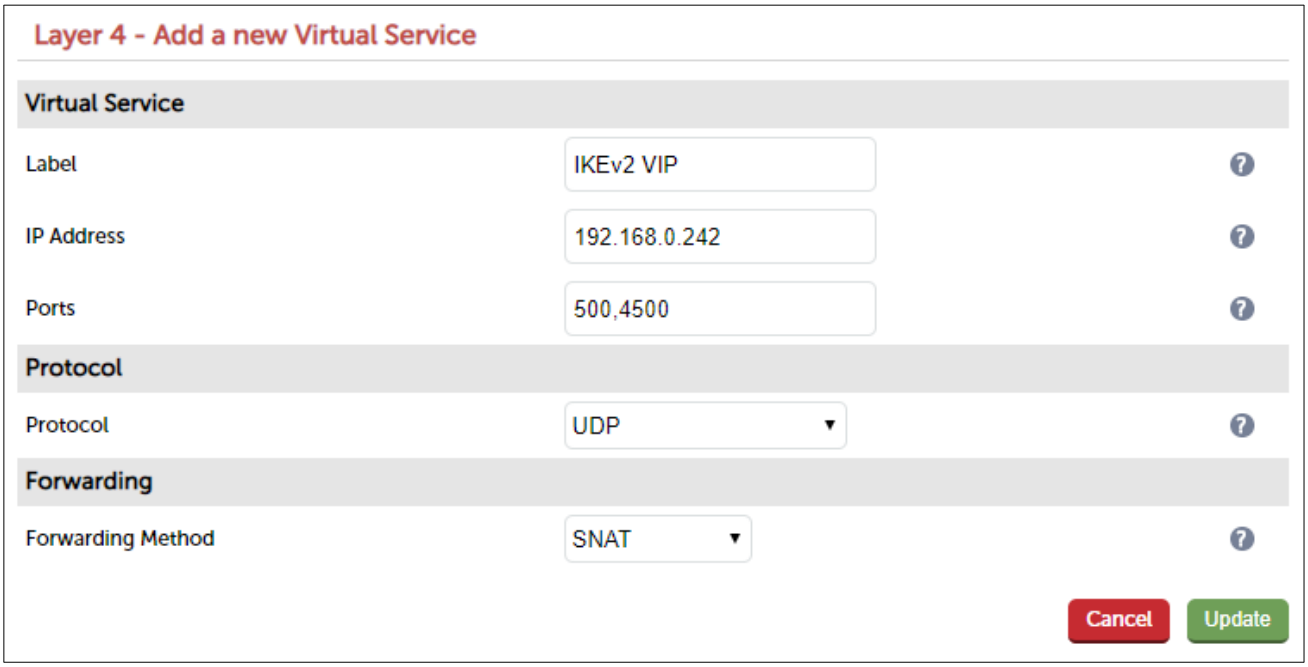

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **IKEv2_VIP**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.0.242**
5. Set the *Virtual Service Ports* field to **500,4500**
6. Set the *Protocol* to **UDP**
7. Set the *Forwarding Method* to SNAT
8. Click **Update**
9. Now click **Modify** next to the newly created Virtual Service
10. Set *Persistence Timeout* to **28800**
11. Under *Health Checks* set the *Check Type* to **Negotiate**

12. Set the *Check port* to **500**
13. Set the *Protocol* to **Radius (ipv4 only)**
14. Enter the appropriate **Radius Secret**
15. Enter the appropriate **Login**
16. Enter the appropriate **Password**
17. Click **Update**

Setting up the Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created Virtual Service
2. Enter the following details:

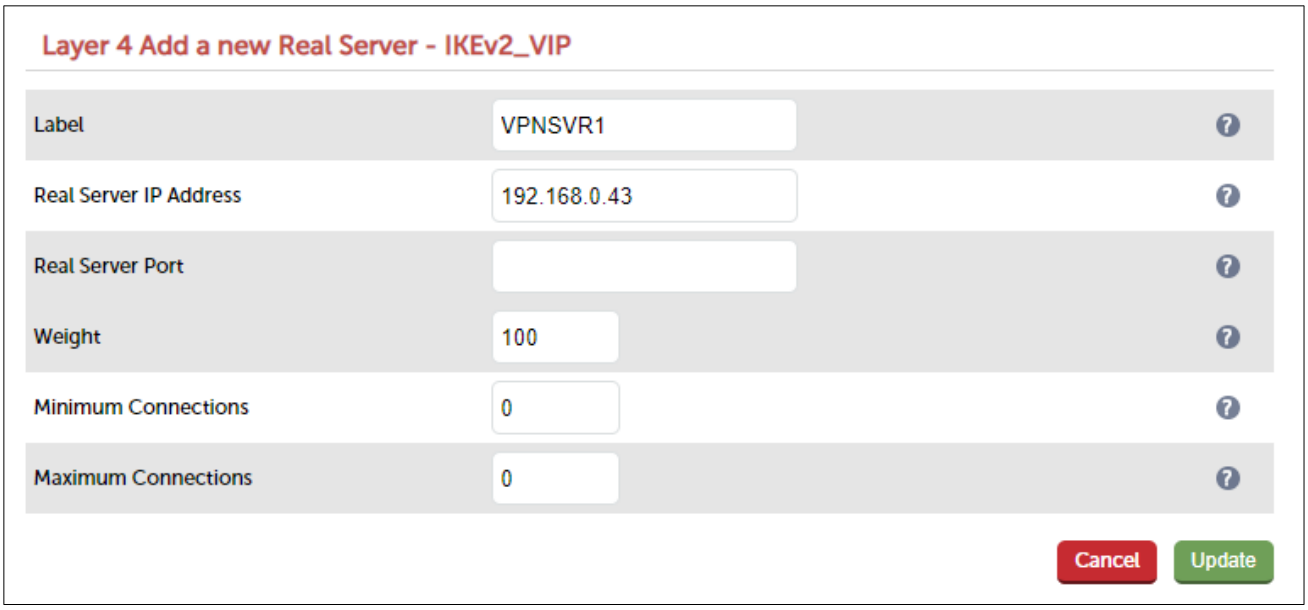

3. Enter an appropriate name (Label) for the first VPN server, e.g. **VPNSVR1**
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.0.43**
5. Click **Update**
6. Now repeat for your remaining VPN server(s)

B - Configuring load balancing for SSTP

Setting up the Virtual Service (VIP)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Services* and click **Add a New Virtual Service**
2. Enter the following details:

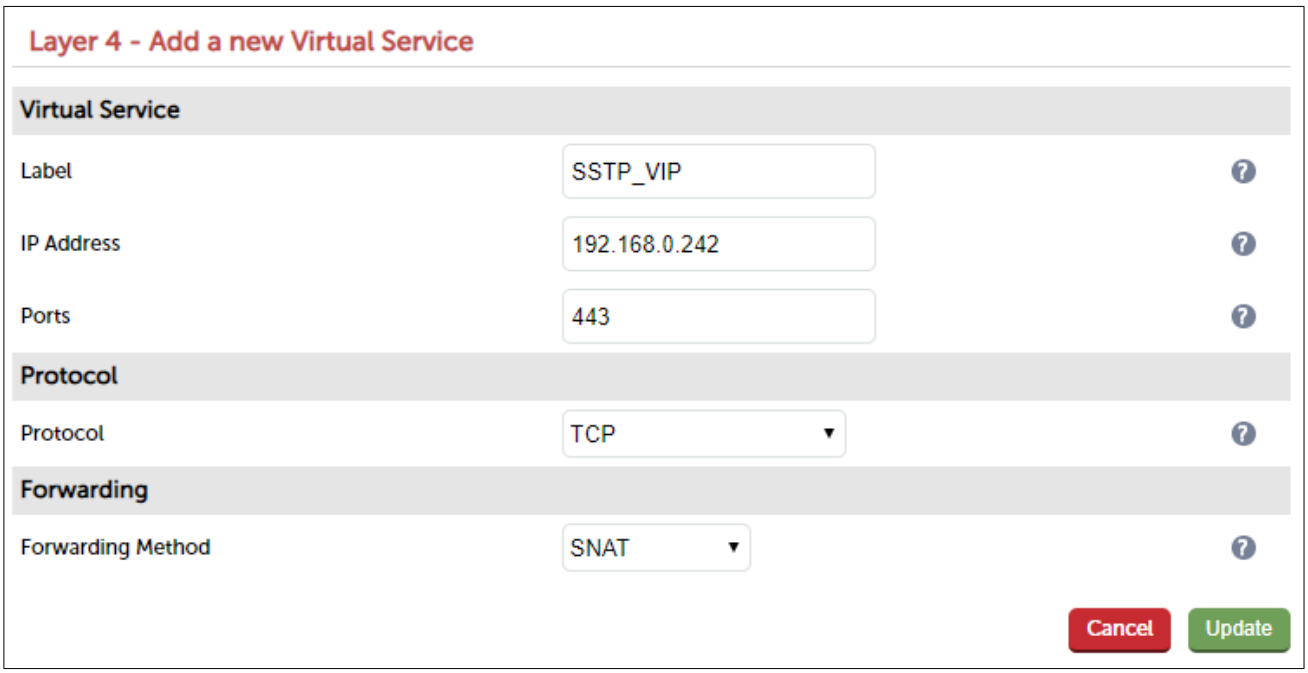

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **SSTP_VIP**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.0.242**
5. Set the *Virtual Service Ports* field to **443**
6. Set the *Protocol* to **TCP Mode**
7. Click **Update**
8. Now click **Modify** next to the newly created Virtual Service
9. Set *Persistence Timeout* **to 28800**
10. Under the *Health Checks* section set the *Check Type* to **Negotiate**
11. Set the *Check Port* to **443**
12. Set the *Protocol* to **HTTPS**
13. Set the *Request to send* to **/sra_{BA195980-CD49-458b-9E23- C84EE0ADCD75}/**
14. Set the *Response expected* to **401**
15. Click **Update**

Setting up the Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created Virtual Service
2. Enter the following details:

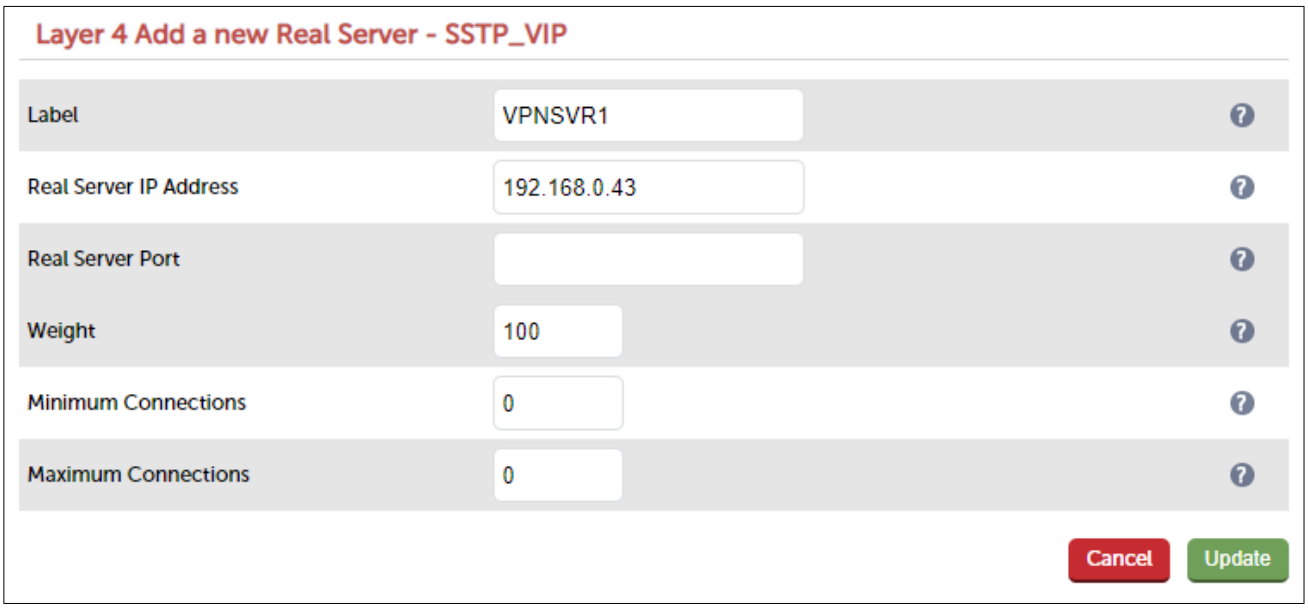

3. Enter an appropriate name (Label) for the first VPN server, e.g. **VPNSVR1**
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.0.43**
5. Click **Update**
6. Now repeat for your remaining VPN server(s)

C - Configuring load balancing for NPS

Setting up the Virtual Service (VIP)

1. Using the WebUI, navigate to: *Cluster Configuration* > *Layer 4 – Virtual Services* and click **Add a New Virtual Service**
2. Enter the following details:

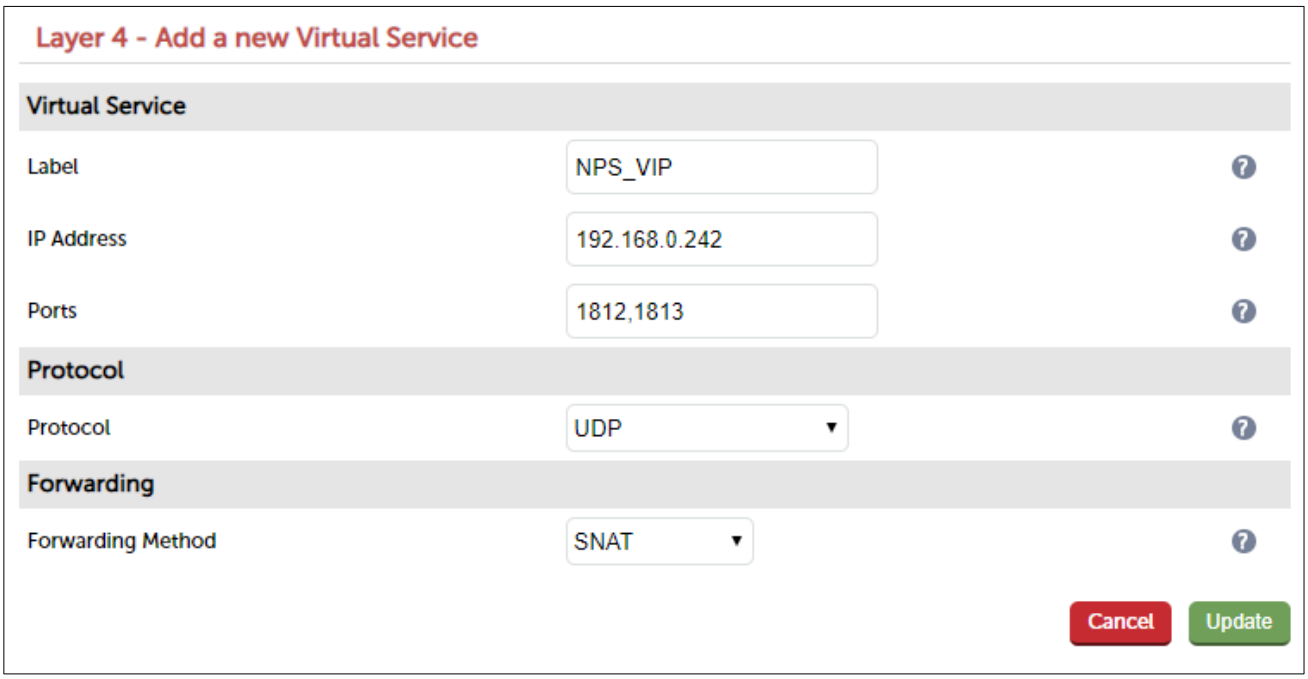

3. Enter an appropriate name (Label) for the Virtual Service, e.g. **NPS_VIP**
4. Set the *Virtual Service IP address* field to the required IP address, e.g. **192.168.0.242**
5. Set the *Virtual Service Ports* field to **1812,1813**
6. Set the *Protocol* to **UDP**
7. Set the *Forwarding Method* to SNAT
8. Click **Update**
9. Now click **Modify** next to the newly created Virtual Service
10. Set *Persistence Timeout* to **28800**
11. Click **Update**

Setting up the Real Servers (RIPs)

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created Virtual Service
2. Enter the following details:

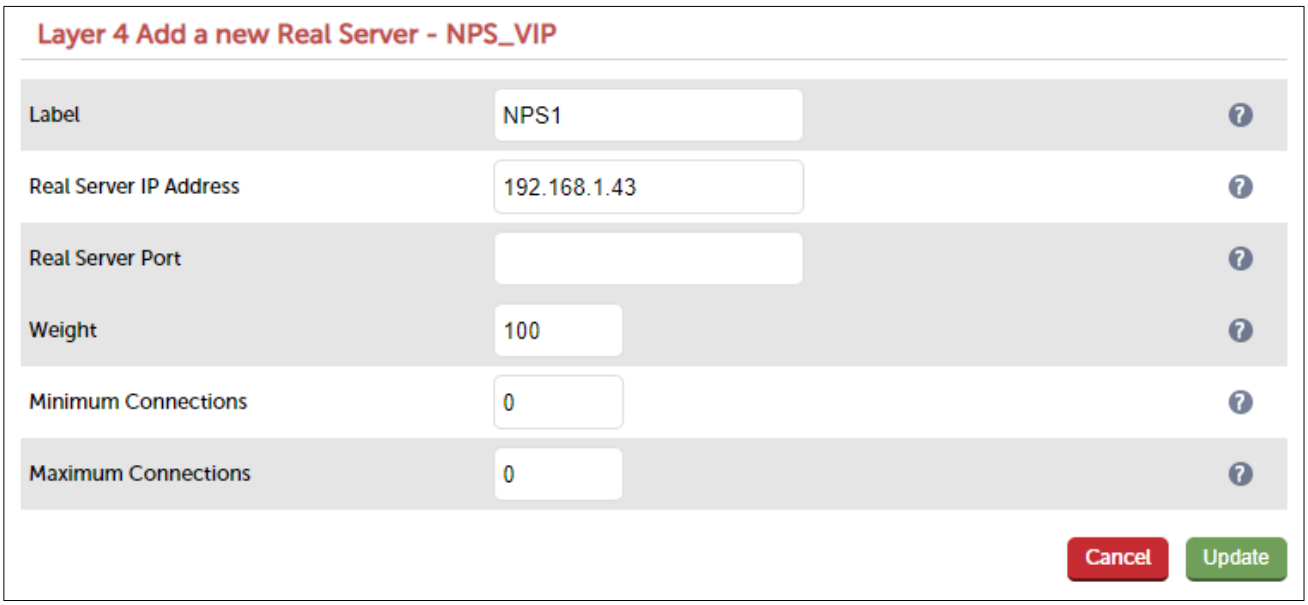

3. Enter an appropriate name (Label) for the first Network Policy Server, e.g. **NPS1**
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.1.43**
5. Click **Update**
6. Now repeat for your remaining Network Policy server(s)

Note:

- The certificate installed on the NPS server must be configured to use the cluster Fully Qualified Domain Name (FQDN) as the subject name on the certificate, with the Subject Alternative Name fields including the FQDNs of both the cluster and server names
- The source IP address of the RADIUS authentication and accounting requests is the Virtual IP Address (VIP) assigned to the virtual service. A RADIUS client must be configured in NPS to allow authentication and accounting requests to be processed. Open the NPS management console and perform the following steps

7. Expand **RADIUS Clients and Servers**
8. Right-click *RADIUS Clients* and select **New**
9. Enter a friendly name for the new RADIUS client
10. Enter the VIP of the NPS virtual service in the **Address (IP or DNS)** field
11. Enter and confirm the shared secret used between the NPS and VPN servers
12. Click **Ok**
13. Repeat the steps above on each NPS server in the cluster

10. Testing & Verification

Note: Make sure that the firewall on the clients and servers is enabled. This is a requirement for Always On VPN to work successfully.

Using the System Overview

Verify that all VIPs & associated RIPs are reported as up (green) as shown below:

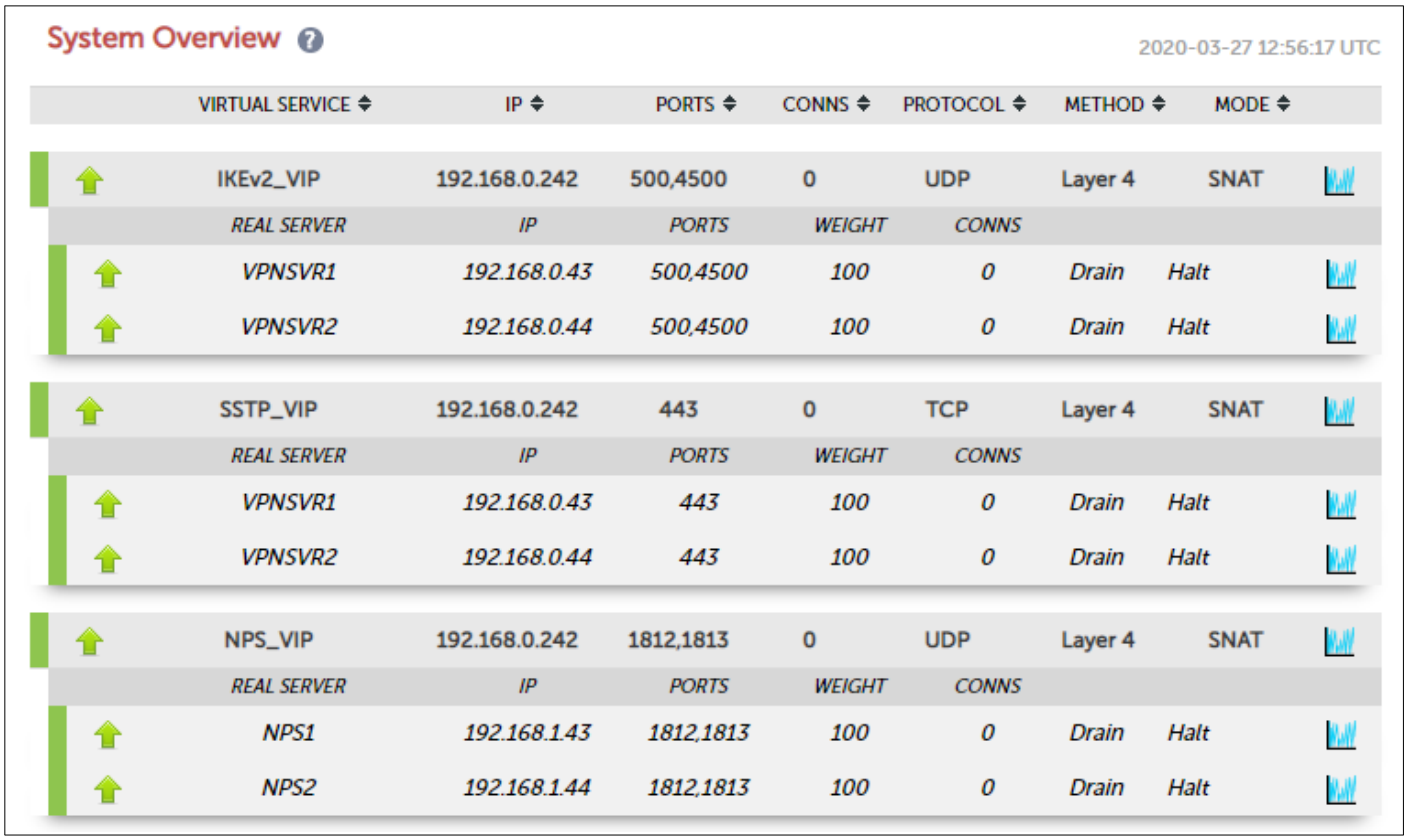

If certain servers are down, i.e. failing their health check, they will be highlighted red as shown below:

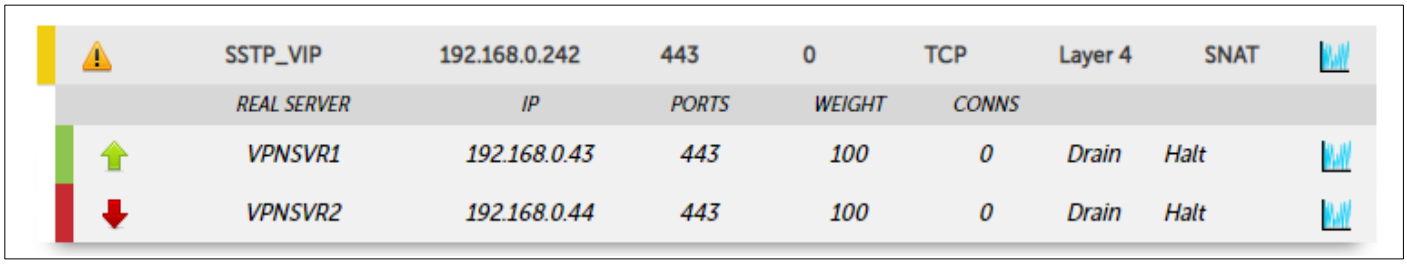

11. Technical Support

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.

12. Further Documentation

The Administration Manual contains much more information about configuring and deploying the appliance. It's available here: http://pdfs.loadbalancer.org/loadbalanceradministrationv8.pdf

13. Conclusion

Loadbalancer.org appliances provide a very cost effective and flexible solution for highly available load balanced Always On VPN Server environments.

14. Appendix

1 - Clustered Pair Configuration – Adding a Slave Unit

If you initially configured just the master unit and now need to add a slave - our recommended procedure, please refer to the relevant section below for more details:

Note: A number of settings are not replicated as part of the master/slave pairing process and therefore must be manually configured on the slave appliance. These are listed below:

- Hostname & DNS settings
- Network settings including IP addresses, bonding configuration and VLANs
- Routing configuration including default gateways and static routes
- Date & time settings
- Physical – Advanced Configuration settings including Internet Proxy IP address & port, Firewall table size, SMTP relay and Syslog server
- SNMP settings
- Graphing settings
- Firewall Script & Firewall Lockdown Script settings
- Software updates

Version 7:

Please refer to **Chapter 8 – Appliance Clustering for HA** in the v7 Administration Manual.

Version 8:

To add a slave node – i.e. create a highly available clustered pair:

- Deploy a second appliance that will be the slave and configure initial network settings
- Using the WebUI, navigate to: *Cluster Configuration > High-Availability Configuration*

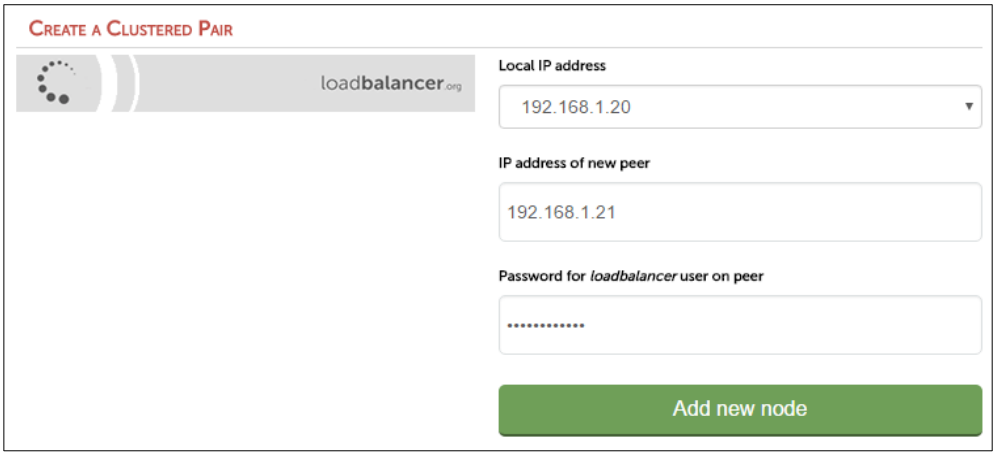

- Specify the IP address and the *loadbalancer* users password (the default is 'loadbalancer') for the slave (peer) appliance as shown above
- Click **Add new node**
- The pairing process now commences as shown below:

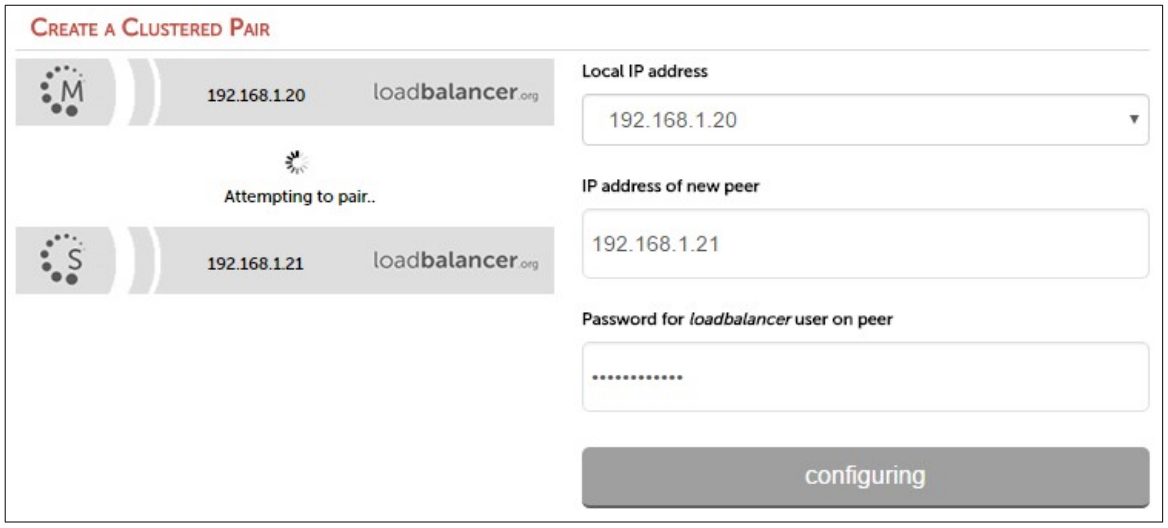

- Once complete, the following will be displayed:

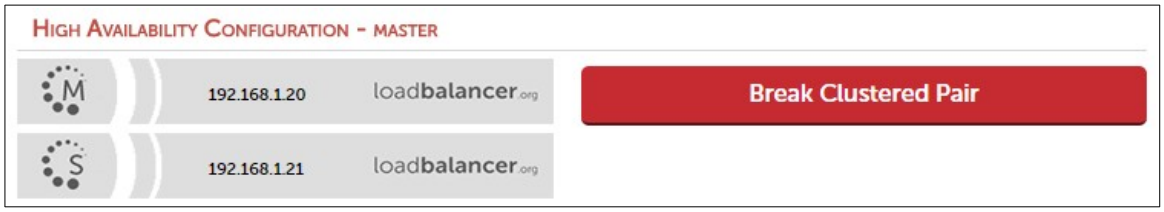

- To finalize the configuration, restart heartbeat and any other services as prompted in the blue message box at the top of the screen

Note: Clicking the **Restart Heartbeat** button on the master appliance will also automatically restart heartbeat on the slave appliance.

Note: Please refer to chapter 9 – Appliance Clustering for HA in the Administration Manual for more detailed information on configuring HA with 2 appliances.

2 – Useful Microsoft Resources & References

GSLB feature blog https://www.loadbalancer.org/blog/loadbalancer-org-releases-a-gslb/

Microsoft Windows 10 Always On VPN https://docs.microsoft.com/en-us/windows-server/remote/remote-access/vpn/always-on-vpn/

Microsoft Windows 10 Always On VPN Deployment Guide https://docs.microsoft.com/en-us/windows-server/remote/remote-access/vpn/always-onvpn/deploy/always-on-vpn-deploy

15. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	27 March 2020	Initial creation		IBG
1.0.1	3 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.0.2	24 September 2020	Health check update		IBG

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

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

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

Germany

Loadbalancer.org GmbH
Tengstraße 2780798,
München, Germany
TEL: +49 (0)89 2000 2179
sales@loadbalancer.org
support@loadbalancer.org