

Load Balancing Microsoft Always On VPN

Version 1.3.1



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. Microsoft Windows	3
4. Microsoft Always On VPN	3
4.1. Introduction	3
4.2. Always On VPN Components	3
4.3. How it Works	4
5. Always On VPN Prerequisites	4
6. Load Balancing Always On VPN	5
6.1. Basic Concepts	5
6.2. Load Balancer Deployment	5
6.3. Load Balancer Deployment Methods	6
6.4. Load Balanced Ports & Services	6
6.5. Persistence (Server Affinity)	6
6.6. Server Health Checking	6
6.7. SSL Offloading	6
7. Loadbalancer.org Appliance – the Basics	7
7.1. Virtual Appliance	7
7.2. Initial Network Configuration	7
7.3. Accessing the Appliance WebUI	7
7.3.1. Main Menu Options	9
7.4. Appliance Software Update	9
7.4.1. Online Update	9
7.4.2. Offline Update	10
7.5. Ports Used by the Appliance	10
7.6. HA Clustered Pair Configuration	11
8. Configuration for Always On VPN	11
8.1. Step 1 – Appliance Configuration	11
8.1.1. IKEv2 Virtual Service Configuration	11
8.1.2. SSTP Virtual Service Configuration	13
8.1.3. NPS Virtual Service Configuration	14
8.2. Step 2 – Always On VPN Server Configuration	16
8.2.1. NPS Server Configuration	16
8.2.2. Solving the ARP Problem For the VPN Servers	16
9. Testing & Verification	22
9.1. Using the System Overview	22
10. Technical Support	23
11. Further Documentation	23
12. Appendix	24
12.1. Configuring HA - Adding a Secondary Appliance	24
12.1.1. Non-Replicated Settings	24
12.1.2. Configuring the HA Clustered Pair	25
12.2. Useful Microsoft Resources & References	26
13. Document Revision History	27

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 [Administration Manual](#).

2. Loadbalancer.org Appliances Supported

All our products can be used with Always On VPN. 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

- V8.9.1 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. Microsoft Windows

- Windows 2016 and later

4. Microsoft Always On VPN

4.1. Introduction

Always On VPN provides a single, cohesive solution for remote access and supports domain-joined, non domain-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.

4.2. 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.

Note

For more information about GSLB, please refer to the [Administration Manual](#) and search for "Global Server Load Balancing".

4.3. 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.

5. 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.

6. 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.

6.1. 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.

6.2. Load Balancer Deployment

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



Note

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

6.3. Load Balancer Deployment Methods

For IKEv2, the load balancing method used must be transparent. This means that the client's source IP address is retained through to the Real Servers. Transparency is required for IKEv2 because Windows limits the number of IPSec Security Associations (SAs) coming from a single IP address. If a non transparent method was used, the source IP address for all traffic reaching the IKEv2 servers would either be the VIP address or the load balancer's own address, depending on the specific configuration.

Both layer 4 DR mode and layer 4 NAT mode are transparent and either can be used for IKEv2. When using DR mode, the "ARP problem" must be solved on all VPN Servers. For NAT mode, the default gateway for each VPN Server must be the load balancer.

For SSTP and NPS transparency is not required, although the load balancing method selected must support UDP. Therefore, whilst DR mode or NAT mode can be used, layer 4 SNAT mode is a simpler option since it requires no mode-specific configuration changes to the Real Servers.

In this guide layer 4 DR mode is used for IKEv2 and layer 4 SNAT mode is used for SSTP and NPS.

Note

For more information on the various load balancing methods supported, please refer to [Supported Load Balancing Methods](#).

Note

For more information on the ARP Problem, please refer to [DR Mode Considerations](#).

6.4. Load Balanced Ports & Services

The following ports/protocols must be load balanced:

Port	Protocol	Use
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

6.5. 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.

6.6. Server Health Checking

The load balancer performs regular checks to verify the health of each server / service. For the IKEv2 and NPS services an ICMP ping check is used, for SSTP a HTTPS negotiate check is used.

6.7. SSL Offloading

To provide scalability and effective load sharing we recommend that SSL is terminated on the VPN servers rather



than on the load balancer.

7. Loadbalancer.org Appliance – the Basics

7.1. 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.

7.2. 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.

7.3. 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.

7.3.1. 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

7.4. 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.

7.4.1. 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:

Information: Update 8.13.1 is now available for this appliance.

Online Update

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:



Information: Update completed successfully. Return to system overview.

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

7.4.2. 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:

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.

7.5. 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



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 (Centralized/Portal Management)
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
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](#).

7.6. 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 [Configuring HA - Adding a Secondary Appliance](#).

8. Configuration for Always On VPN

This is completed in 2 steps; step 1 covers the appliance configuration, step 2 covers the configuration changes required to the Always On VPN servers to enable load balancing.

8.1. Step 1 – Appliance Configuration

3 Virtual Services (VIPs) are required for Always On VPN. These are for IKEv2, SSTP and NPS. The following sections cover the configuration of each VIP.

8.1.1. IKEv2 Virtual Service Configuration

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:

Virtual Service		
Label	<input type="text" value="IKEv2_VIP"/>	?
IP Address	<input type="text" value="192.168.0.242"/>	?
Ports	<input type="text" value="500,4500"/>	?
Protocol		
Protocol	<input type="text" value="UDP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="Direct Routing"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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 **Direct Routing**.
8. Click **Update**.
9. Now click **Modify** next to the newly created Virtual Service.
10. Verify that the *Persistence Timeout* is set to **300**.
11. Under *Health Checks* ensure that *Check Type* is set to **ping server**.
12. Click **Update**.

Configuring the Associated 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:

Label	<input type="text" value="VPNSVR1"/>	?
Real Server IP Address	<input type="text" value="192.168.0.43"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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 the above steps to add your remaining VPN server(s).

8.1.2. SSTP Virtual Service Configuration

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:

Virtual Service		
Label	<input type="text" value="SSTP_VIP"/>	?
IP Address	<input type="text" value="92.168.0.242"/>	?
Ports	<input type="text" value="443"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="SNAT"/>	?

Cancel
Update

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. Set the *Forwarding Method* to **SNAT**.
8. Click **Update**.
9. Now click **Modify** next to the newly created Virtual Service.
10. Verify that the *Persistence Timeout* is set to **300**.
11. Under the *Health Checks* section set the *Check Type* to **Negotiate**.
12. Set the *Check Port* to **443**.
13. Set the *Protocol* to **HTTPS**.
14. Set the *Request to send* to **/sra_{BA195980-CD49-458b-9E23-C84EE0ADCD75}**.
15. Set the *Response expected* to **401**.

16. Click **Update**.

Configuring the Associated 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:

Label	<input type="text" value="VPNSVR1"/>	?
Real Server IP Address	<input type="text" value="192.168.0.43"/>	?
Real Server Port	<input type="text" value="443"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

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. Set the *Real Server Port* field to **443**.
6. Click **Update**.
7. Now repeat the above steps to add your remaining VPN server(s).

8.1.3. NPS Virtual Service Configuration

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:

Virtual Service		
Label	<input type="text" value="NPS_VIP"/>	?
IP Address	<input type="text" value="192.168.0.242"/>	?
Ports	<input type="text" value="1812,1813"/>	?
Protocol		
Protocol	<input type="text" value="UDP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="SNAT"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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. Verify that the *Persistence Timeout* is set to **300**.
11. Under *Health Checks* ensure that *Check Type* is set to **ping server**.
12. Click **Update**.

Configuring the Associated 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:

Label	<input type="text" value="NPS_SVR1"/>	?
Real Server IP Address	<input type="text" value="192.168.1.43"/>	?
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 name (Label) for the first Network Policy Server, e.g. **NPS_SVR1**.
4. Change the *Real Server IP Address* field to the required IP address, e.g. **192.168.1.43**.
5. Leave *Real Server Port* blank.
6. Click **Update**.
7. Now repeat the above steps to add your remaining NPS 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.

8.2. Step 2 – Always On VPN Server Configuration

8.2.1. NPS Server Configuration

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:

1. Expand **RADIUS Clients and Servers**.
2. Right-click *RADIUS Clients* and select **New**.
3. Enter a friendly name for the new RADIUS client.
4. Enter the IP address of the NPS Virtual Service in the **Address (IP or DNS)** field.
5. Enter and confirm the shared secret used between the NPS and VPN servers.
6. Click **OK**.
7. Repeat the above steps on all other NPS servers in the cluster.

8.2.2. Solving the ARP Problem For the VPN Servers

When using Layer 4 DR mode, the ARP problem must be solved. This involves configuring each Real Server to be able to receive traffic destined for the VIP, and ensuring that each Real Server does not respond to ARP requests

for the VIP address – only the load balancer should do this.

The steps below are for Windows 2012 and later and must be completed on each VPN server.

Windows Server 2012 & Later

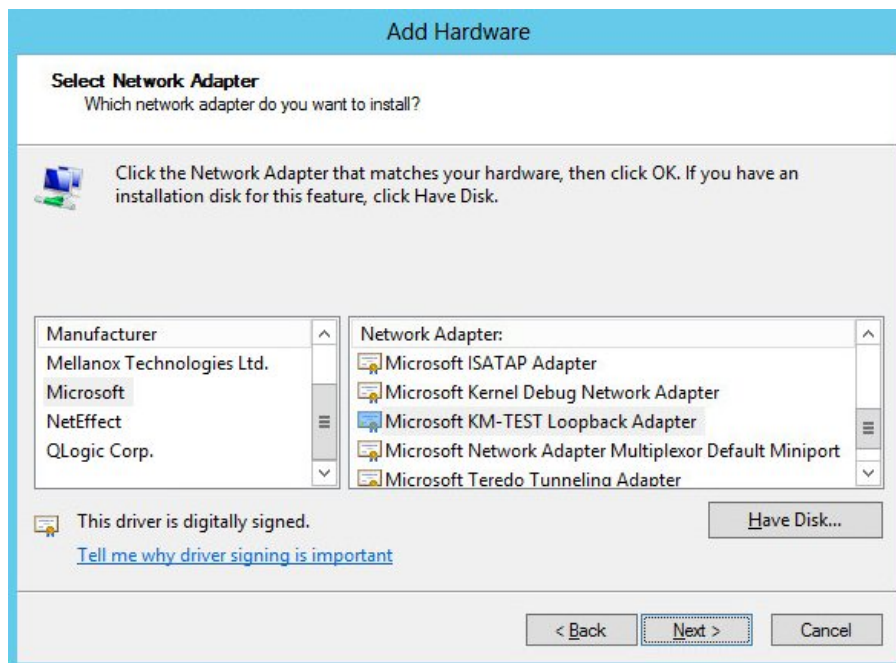
Windows Server 2012 and later support Direct Routing (DR) mode through the use of the Microsoft Loopback Adapter that must be installed and configured on each load balanced (Real) Server. The IP address configured on the Loopback Adapter must be the same as the Virtual Service (VIP) address. This enables the server to receive packets that have their destination set as the VIP address. If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be added to the Loopback Adapter.

In addition, the strong/weak host behavior must be configured on each Real Server. The weak host model allows packets with any IP to be sent or received via an interface. The strong host model only allows packets with an IP belonging to the interface to be sent or received.

(!) Important The following 3 steps must be completed on **all** Real Servers associated with the VIP.

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



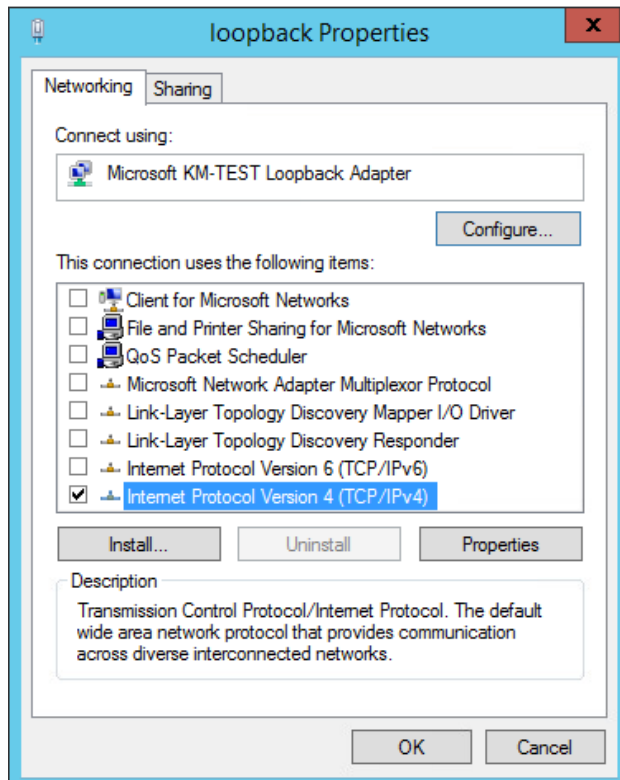
1. Open Control Panel and click **Network and Sharing Center**.
2. Click **Change adapter settings**.
3. Right-click the new Loopback Adapter and select **Properties**.

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:




 **Note** **192.168.2.20** is an example, make sure you specify the correct VIP address.

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


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:



Note **2001:470:1f09:e72::15/64** is an example, make sure you specify the correct VIP address.

Note If a Real Server is included in multiple DR mode VIPs, an IP address for each VIP must be

added to the Loopback Adapter.

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

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

The strong/weak host behavior can be configured using either of the following 2 methods:

- Option 1 - Using network shell (netsh) commands
- Option 2 - Using PowerShell cmdlets

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
```

9. Testing & Verification

Note

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

Note

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

9.1. Using the System Overview

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







System Overview ?

2020-03-27 12:56:17 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
↑	IKEv2_VIP	192.168.0.242	500,4500	0	UDP	Layer 4	DR	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	VPNSVR1	192.168.0.43	500,4500	100	0	Drain	Halt	
↑	VPNSVR2	192.168.0.44	500,4500	100	0	Drain	Halt	
↑	SSTP_VIP	192.168.0.242	443	0	TCP	Layer 4	SNAT	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	VPNSVR1	192.168.0.43	443	100	0	Drain	Halt	
↑	VPNSVR2	192.168.0.44	443	100	0	Drain	Halt	
↑	NPS_VIP	192.168.0.242	1812,1813	0	UDP	Layer 4	SNAT	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	NPS1	192.168.1.43	1812,1813	100	0	Drain	Halt	
↑	NPS2	192.168.1.44	1812,1813	100	0	Drain	Halt	



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

	SSTP_VIP	192.168.0.242	443	0	TCP	Layer 4	SNAT	
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
	VPNSVR1	192.168.0.43	443	100	0	Drain	Halt	
	VPNSVR2	192.168.0.44	443	100	0	Drain	Halt	

10. 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.

11. Further Documentation

For additional information, please refer to the [Administration Manual](#).

12. Appendix

12.1. 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.

12.1.1. 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	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.


12.1.2. 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

 **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 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

 **LOADBALANCER**

Primary

IP: 192.168.110.40

Attempting to pair..

 **LOADBALANCER**

Secondary

IP: 192.168.110.41

Local IP address

192.168.110.40

IP address of new peer

192.168.110.41

Password for *loadbalancer* user on peer


••••••••••

configuring

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



High Availability Configuration - primary



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](#).

12.2. Useful Microsoft Resources & References

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-on-vpn/deploy/always-on-vpn-deploy>

Troubleshooting Always On VPN:

<https://docs.microsoft.com/en-us/windows-server/remote/remote-access/vpn/always-on-vpn/deploy/always-on-vpn-deploy-troubleshooting>



13. 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
1.1.0	11 th August 2021	Changed the health check for the IKEv2 VIP to an ICMP ping check Changed the persistence timeout to 300 seconds (5mins) for all VIPs Changed load balancing method for the IKEv2 VIP from SNAT mode to DR mode	Incorrectly specified a Radius check Previous setting was unnecessarily high IKEv2 client connections must be transparent	RJC
1.2.0	1 January 2022	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.2.1	5 January 2023	Combined software version information into one section Added one level of section numbering Added software update instructions Added table of ports used by the appliance Reworded 'Further Documentation' section	Housekeeping across all documentation	AH
1.2.2	2 February 2023	Updated screenshots	Branding update	AH
1.2.3	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.3.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH
1.3.1	20 February 2025	Updated the request to send for the SSTP VIP's health check	Technical requirement	RJC



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](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.

