

Load Balancing Cisco Duo Network Gateway

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. Duo Network Gateway	3
4. Duo Network Gateway	3
5. Load Balancing Duo Network Gateway	3
5.1. Load Balancing & HA Requirements	3
5.2. Virtual Service (VIP) Requirements	4
6. Deployment Concept	4
7. Load Balancer Deployment Methods	5
7.1. Layer 7 SNAT Mode	5
8. Loadbalancer.org Appliance – the Basics	6
8.1. Virtual Appliance	6
8.2. Initial Network Configuration	6
8.3. Accessing the Appliance WebUI	6
8.3.1. Main Menu Options	8
8.4. Appliance Software Update	9
8.4.1. Online Update	9
8.4.2. Offline Update	9
8.5. Ports Used by the Appliance	10
8.6. HA Clustered Pair Configuration	11
9. Configuring Duo Network Gateway for Load Balancing	11
9.1. Active/Active	11
9.1.1. Migrating from Standalone	11
10. Appliance Configuration for Duo Network Gateway	11
10.1. VIP 1 - DuoNetworkGateway	11
10.1.1. Virtual Service (VIP) Configuration	11
10.1.2. Configure the Associated Real Servers (RIPs)	12
10.2. Finalizing the Configuration	13
11. Testing & Verification	13
11.1. Using System Overview	13
11.2. Accessing Duo Network Gateway via the Load Balancer	14
12. Technical Support	14
13. Further Documentation	14
14. Appendix	15
14.1. Configuring HA - Adding a Secondary Appliance	15
14.1.1. Non-Replicated Settings	15
14.1.2. Configuring the HA Clustered Pair	16
15. Document Revision History	18

1. About this Brief

This brief outlines the steps required to configure a load balanced Cisco Duo Network Gateway environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Duo Network Gateway 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 Duo Network Gateway. 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. Duo Network Gateway

- All versions

4. Duo Network Gateway

Duo Network Gateway provides users with secure remote access to your on-premises private applications and internal servers without having to worry about managing VPN credentials. It also offers inline user enrollment, self-service device management, and support for a variety of authentication methods — such as passkeys and security keys, Duo Push, or Verified Duo Push - in the Universal Prompt.

5. Load Balancing Duo Network Gateway

Note


It's highly recommended that you have a working Duo Network Gateway environment first before implementing the load balancer.

5.1. Load Balancing & HA Requirements



Cisco Duo recommends that [high availability should be used](#) in all Duo Network Gateway deployments. Two configurations are supported:

- **Scenario 1 - Active/Passive** : A second Duo Network Gateway server is used as a standby/fallback server that becomes active if the active server fails
- **Scenario 2 - Active/Active** : Multiple Network Gateway servers are used concurrently, connections are distributed by the load balancer

 **Note** Active/Active is only supported in AWS.

5.2. Virtual Service (VIP) Requirements

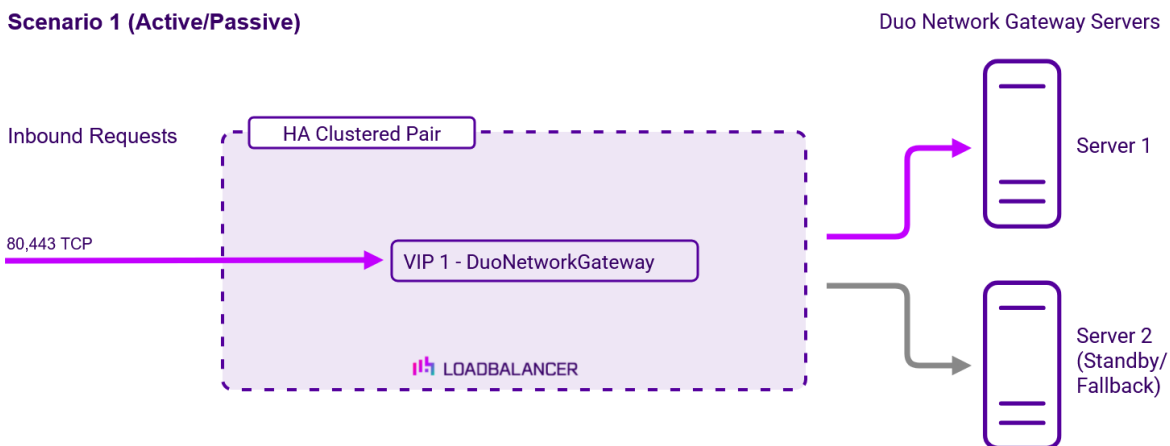
For both scenarios, the following VIP is required:

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	DuoNetworkGateway	L7 SNAT	80,443	Source IP	HTTPS (GET)

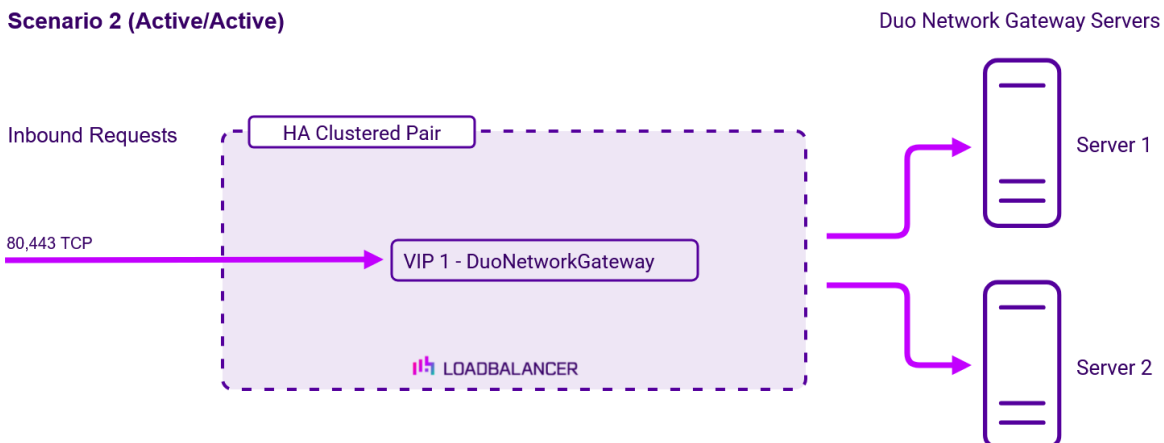
6. Deployment Concept

Once the load balancer is deployed, clients connect to the Virtual Service (VIP) rather than connecting directly to one of the Duo Network Gateway servers.

Scenario 1 (Active/Passive)



Scenario 2 (Active/Active)



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.

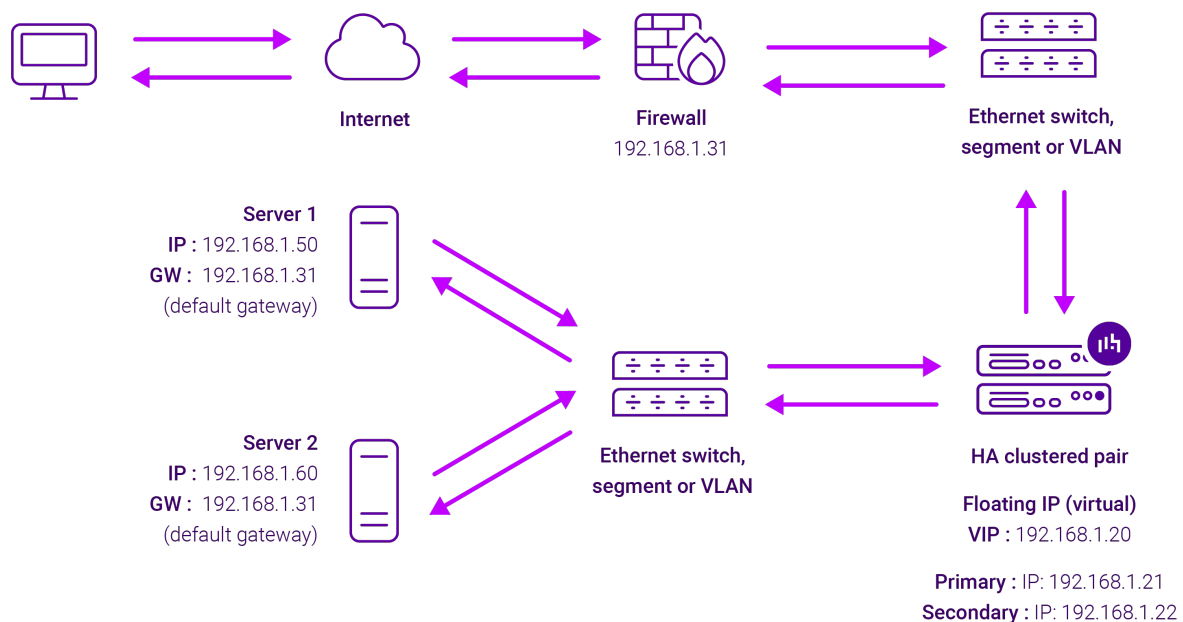
7. Load Balancer Deployment Methods

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 Duo Network Gateway, layer 7 SNAT mode is recommended. This mode is described below and is used for the configuration presented in this guide.

7.1. Layer 7 SNAT Mode

Layer 7 SNAT mode uses a proxy (HAProxy) at the application layer. Inbound requests are terminated on the load balancer and HAProxy generates a new corresponding request to the chosen Real Server. As a result, Layer 7 is typically not as fast as the Layer 4 methods. Layer 7 is typically chosen when either enhanced options such as SSL termination, cookie based persistence, URL rewriting, header insertion/deletion etc. are required, or when the network topology prohibits the use of the layer 4 methods. The image below shows an example network diagram for this mode.



- Because layer 7 SNAT mode is a full proxy, Real Servers in the cluster can be on any accessible network including across the Internet or WAN.
- Layer 7 SNAT mode is not transparent by default, i.e. the Real Servers will not see the source IP address of the client, they will see the load balancer's own IP address by default, or any other local appliance IP address if preferred (e.g. the VIP address). This can be configured per layer 7 VIP. If required, the load balancer can be configured to provide the actual client IP address to the Real Servers in 2 ways. Either by inserting a header that contains the client's source IP address, or by modifying the Source Address field of the IP packets and replacing the IP address of the load balancer with the IP address of the client. For more



information on these methods please refer to [Transparency at Layer 7](#).

- Layer 7 SNAT mode can be deployed using either a one-arm or two-arm configuration. For two-arm deployments, **eth1** is typically used for client side connections and **eth0** is used for Real Server connections, although this is not mandatory since any interface can be used for any purpose.
- Requires no mode-specific configuration changes to the load balanced Real Servers.
- Port translation is possible with Layer 7 SNAT mode, e.g. VIP:80 → RIP:8080 is supported.
- You should not use the same RIP:PORT combination for layer 7 SNAT mode VIPs and layer 4 SNAT mode VIPs because the required firewall rules conflict.

8. Loadbalancer.org Appliance – the Basics

8.1. Virtual Appliance

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

Note

The same download is used for the licensed product, the only difference is that a license key file (supplied by our sales team when the product is purchased) must be applied using the appliance's WebUI.

Note

Please refer to [Virtual Appliance Installation](#) and the ReadMe.txt text file included in the VA download for additional information on deploying the VA using the various Hypervisors.

Note

The VA has 4 network adapters. For VMware only the first adapter (**eth0**) is connected by default. For HyperV, KVM, XEN and Nutanix AHV all adapters are disconnected by default. Use the network configuration screen within the Hypervisor to connect the required adapters.

8.2. Initial Network Configuration

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

Important

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

8.3. Accessing the Appliance WebUI

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

Note

There are certain differences when accessing the WebUI for the cloud appliances. For details,



please refer to the relevant [Quick Start / Configuration Guide](#).

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

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

 **Note**

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

 **Note**

If you need to change the port, IP address or protocol that the WebUI listens on, please refer to [Service Socket Addresses](#).

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

Username: loadbalancer

Password: <configured-during-network-setup-wizard>

 **Note**

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

Once logged in, the WebUI will be displayed as shown below:

Primary | Secondary Active | Passive Link 8 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 ? 2025-05-08 12:37:21 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

RX	28 Min,	2713 Avg,	27344772 Total,
TX	0 Min,	13777 Avg,	138872181 Total,

System Load Average

1m average	0.00 Min,	0.08 Avg,	0.68 Max
5m average	0.00 Min,	0.04 Avg,	0.30 Max
15m average	0.00 Min,	0.02 Avg,	0.12 Max

Memory Usage

- 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.

i **Note** The Setup Wizard can only be used to configure Layer 7 services.

8.3.1. Main Menu Options

- System Overview** - Displays a graphical summary of all VIPs, RIPs and key appliance statistics
- Local Configuration** - Configure local host settings such as IP address, DNS, system time etc.
- Cluster Configuration** - Configure load balanced services such as VIPs & RIPs
- Maintenance** - Perform maintenance tasks such as service restarts and creating backups
- View Configuration** - Display the saved appliance configuration settings
- Reports** - View various appliance reports & graphs
- Logs** - View various appliance logs
- Support** - Create a support download, contact the support team & access useful links



Live Chat - Start a live chat session with one of our Support Engineers

8.4. Appliance Software Update

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

 **Note**

For full details, please refer to [Appliance Software Update](#) in the Administration Manual.

 **Note**

Services may need to be restarted/reloaded after the update process completes or in some cases a full appliance restart may be required. We therefore recommend performing the update during a maintenance window.

8.4.1. Online Update

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

Information: Update 8.13.5 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.

8.4.2. Offline Update

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

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

To perform an offline update:



1. Using the WebUI, navigate to: **Maintenance > Software Update**.
2. Select **Offline Update**.
3. The following screen will be displayed:

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.

8.5. Ports Used by the Appliance

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

Protocol	Port	Purpose
TCP	22	SSH
TCP & UDP	53	GSLB
TCP & UDP	123	NTP
TCP & UDP	161	SNMP
UDP	6694	Heartbeat between Primary & Secondary appliances in HA mode
TCP	7778	HAProxy persistence table replication
TCP	9000	Gateway service for ADC Portal comms
TCP	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback server
TCP	9443	WebUI - HTTPS
TCP	25565	Shuttle service for ADC Portal comms

Note

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



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

8.6. HA Clustered Pair Configuration

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

9. Configuring Duo Network Gateway for Load Balancing

The following setting should be configured on each Gateway server either during initial setup or by using the Duo Network Gateway Admin UI and navigation to: **Settings > Server Settings**.

Load Balancers This Duo Network Gateway is accessed through load balancers

Selecting this option and specifying the addresses of the load balancers will allow the Duo Network Gateway to use the true client IP address for logging, whitelisting IP restrictions, and passing upstream to web applications.

Load Balancer IP addresses

Separate multiple values with spaces or newlines. You can specify single IP addresses (192.168.99.105), a block (10.1.2.0/24) or a range of IP addresses (10.0.2.4-10.0.2.10). Example: '10.1.10.124 11.1.0.0/16 192.168.0.1-192.168.0.15'

The IP address(es) or range specified should be the base IP address of the Primary appliance and if HA is used, the Secondary appliance.

9.1. Active/Active

For more information on configuring active/active, please refer to: <https://duo.com/docs/dng#active-active>.

 **Note** Active/Active is only supported in AWS.

9.1.1. Migrating from Standalone

If you've already configured a standalone Duo Network Gateway and would like to migrate to Active/Active high availability, please refer to: <https://duo.com/docs/dng#migrating-from-standalone>.

10. Appliance Configuration for Duo Network Gateway

10.1. VIP 1 - DuoNetworkGateway

10.1.1. Virtual Service (VIP) Configuration

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



Virtual Service		[Advanced +]
Label	<input type="text" value="DuoNetworkGateway"/>	?
IP Address	<input type="text" value="10.10.90.136"/>	?
Ports	<input type="text" value="80,443"/>	?
Protocol		[Advanced +]
Layer 7 Protocol	<input type="text" value="TCP Mode"/>	?

2. Enter a suitable *Label* (name) for the Virtual Service, e.g. **DuoNetworkGateway**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **10.10.90.136**.
4. Set *Ports* to **80,443**.
5. Set the *Layer 7 Protocol* to **TCP Mode**.
6. Click **Update**.
7. Now click **Modify** next to the newly created VIP.
8. Scroll down to the *Health Checks* section.
 - Set *Health Checks* to **Negotiate HTTPS (GET)**.
 - Set *Request to Send* to **/health-check**.
9. If Scenario 1 (Active/Passive) is used, the second Gateway server must be added as a standby/fallback server. The fallback server only becomes active if the Active server fails. Scroll to the *Fallback Server* section.

Fallback Server		[Advanced +]
Disable Fallback	<input type="checkbox"/>	?
IP Address	<input type="text" value="10.10.90.131"/>	?
Port	<input type="text"/>	?


- Set the *IP Address* to the address of the second Gateway server, e.g. **10.10.90.131**.
 - Leave the *Port* field blank.
10. Leave all other settings at their default value.
 11. Click **Update**.

10.1.2. Configure 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.

Label	<input type="text" value="DuoGatewayServer1"/>	?
Real Server IP Address	<input type="text" value="10.10.90.130"/>	?
Real Server Port	<input type="text"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

2. Enter a suitable *Label* (name) for the Real Server, e.g. **DuoGatewayServer1**.
3. Set the *Real Server IP Address* to the IP address of the first Gateway server, e.g. **10.10.90.130**.
4. Leave the *Real Server Port* field blank.
5. Leave all other settings at their default value.
6. Click **Update**.
7. If Scenario 2 (Active/Active) is used, repeat these steps to add the second Gateway server. The load will then be shared between both Gateway servers.


 **Note** Active/Active is only supported in AWS.

10.2. Finalizing the Configuration

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

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

11. Testing & Verification

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

11.1. Using System Overview





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

Scenario 1 (Active/Passive)



System Overview







2026-05-06 07:57:43 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
 DuoNetworkGateway	10.10.90.136	80,443	0	TCP	Layer 7	Proxy	
<i>REAL SERVER</i>							
	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
 DuoGatewayServer1	10.10.90.130	80,443	100	0	Drain	Halt	

Scenario 2 (Active/Active)

System Overview

2026-05-06 07:59:05 UTC

VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
 DuoNetworkGateway	10.10.90.136	80,443	0	TCP	Layer 7	Proxy	
<i>REAL SERVER</i>							
	<i>IP</i>	<i>PORTS</i>	<i>WEIGHT</i>	<i>CONNS</i>			
 DuoGatewayServer1	10.10.90.130	80,443	100	0	Drain	Halt	
 DuoGatewayServer2	10.10.90.131	80,443	100	0	Drain	Halt	

11.2. Accessing Duo Network Gateway via the Load Balancer

Verify that you're able to successfully access Duo Network Gateway via the Virtual Service on the load balancer.

Note

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

12. Technical Support

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

13. Further Documentation

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



14. Appendix

14.1. Configuring HA - Adding a Secondary Appliance

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

 **Note**

For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the [documentation library](#)

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

14.1.1. Non-Replicated Settings

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

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

(!) Important

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

14.1.2. Configuring the HA Clustered Pair

(!) Important

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

Note

If you have already run the firewall lockdown wizard on either appliance, you'll need to ensure that it is temporarily disabled on both appliances whilst performing the pairing process.

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

Create a Clustered Pair

LOADBALANCER

Local IP address
10.11.40.55

IP address of new peer
10.11.40.56

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: 10.11.40.55

Attempting to pair..

LOADBALANCER Secondary
IP: 10.11.40.56

Local IP address
10.11.40.55

IP address of new peer
10.11.40.56

Password for loadbalancer user on peer
.....

configuring

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

High Availability Configuration - primary

LOADBALANCER Primary
IP: 10.11.40.55

LOADBALANCER Secondary
IP: 10.11.40.56

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 [Configuring High Availability](#).

Note For details on testing and verifying HA, please refer to [Clustered Pair Diagnostics](#).

15. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	7 May 2026	Initial version		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.

