

Load Balancing Qumulo Hybrid File System

Version 1.3.0



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. Qumulo Hybrid File System	3
4. Qumulo Hybrid File System	3
5. Load Balancing Qumulo Hybrid File System	3
5.1. Load Balancing & HA Requirements	4
5.2. Load Balancer Configuration	4
5.3. Connection Distribution	4
5.4. Health Checks	4
6. Deployment Concept	4
6.1. Explanation	5
7. Loadbalancer.org Appliance – the Basics	6
7.1. Virtual Appliance	6
7.2. Initial Network Configuration	6
7.3. Accessing the Appliance WebUI	6
7.3.1. Main Menu Options	8
7.4. Appliance Software Update	8
7.4.1. Online Update	8
7.4.2. Offline Update	9
7.5. Ports Used by the Appliance	9
8. Appliance Configuration for Qumulo Hybrid File System	10
8.1. Step 1 – Configure the HA Pair	10
8.2. Step 2 – Configure GSLB	10
8.2.1. Handling Multiple Subdomains, Including Wildcard Subdomains	10
8.2.2. Appliance Configuration	11
8.2.3. DNS Server Configuration	14
8.2.4. Testing the Configuration	14
9. Microsoft DNS Server Configuration	14
9.1. Microsoft DNS Server	15
10. Testing & Verification	17
10.1. Accessing the Service	17
11. Technical Support	17
12. Further Documentation	17
13. Appendix	18
13.1. Configuring HA - Adding a Secondary Appliance	18
13.1.1. Non-Replicated Settings	18
13.1.2. Configuring the HA Clustered Pair	19
14. Document Revision History	21

1. About this Guide

This guide details the steps required to configure a load balanced Qumulo Hybrid File System environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Qumulo 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 for load balancing Qumulo Hybrid File System. 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

- V8.4.1 or 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. Qumulo Hybrid File System

- V3.0.0 or later

4. Qumulo Hybrid File System

Qumulo's advanced hybrid file software allows you to symmetrically scale capacity and performance. Work in real-time, without the need for tree walks. Qumulo offers the most efficient file system in the industry, with the best raw-to-useable capacity that leverages 100 percent of useable storage.

Qumulo's file system is built for the hybrid cloud, providing a single file solution whether your data is in the cloud or on-prem or scaling across both. Users can burst compute in AWS or GCP and shift primary workloads to the cloud without application changes.

5. Load Balancing Qumulo Hybrid File System

 **Note**

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



5.1. Load Balancing & HA Requirements

The function of the load balancer is to ensure that connections to a Qumulo Hybrid file System cluster are distributed across healthy Qumulo nodes. This is done to provide a highly available and scalable service. To provide HA for the load balancer, Loadbalancer.org recommends that 2 appliances are deployed as an HA clustered pair.

5.2. Load Balancer Configuration

Load balancing a Qumulo deployment requires no Virtual Services (VIPs). Instead, both load balancers are configured as smart DNS name servers for the FQDN of the Qumulo cluster address in question (**qumulo.company.com** in this guide). This is achieved using the load balancer's built-in GSLB service and by using DNS delegation.

5.3. Connection Distribution

The GSLB service uses weighted round-robin load balancing to distribute inbound client connections across the healthy Qumulo nodes. The weights for the nodes can be set as required depending on relative node performance.

5.4. Health Checks

Each Qumulo node is regularly health-checked by each load balancer and this information is used when providing the smart DNS response to inbound DNS queries.


 **Note**

For more details on configuring DNS delegation please refer to [DNS Server Configuration for Qumulo Hybrid File System](#).

6. Deployment Concept

The diagram below illustrates how the load balancer is deployed.





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 [Configuring HA - Adding a Secondary Appliance](#) for more details on configuring an HA clustered pair.

6.1. Explanation

1. The client sends a DNS query for the service FQDN, e.g. **qumulo.company.com** to the local DNS server.
2. The local DNS server has the sub domain delegated to both load balancers (both load balancers are configured as name servers for the sub domain).
3. One of the load balancers receives the delegated DNS query.
4. The load balancer selects a healthy Qumulo node based on the GSLB health checks and the round robin algorithm used.
5. The load balancer returns the IP address of the selected Qumulo node to the DNS server.
6. The DNS server returns the IP address of the selected Qumulo node to the client.
7. The client connects directly to the Qumulo node.

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.

(①) 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>



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 8.13.2 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.

(1) 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](#).

8. Appliance Configuration for Qumulo Hybrid File System

8.1. Step 1 – Configure the HA Pair

If you intend to deploy 2 load balancers in order to configure an HA clustered pair (our recommended configuration) then the HA pair should be configured first before the GSLB configuration takes place. This simplifies the process since GSLB settings will then be automatically replicated to the paired appliance. This helps ensure that both appliances are correctly configured and ready for sub domain delegation (please refer to [\[DNS Server Configuration for Qumulo Hybrid File System\]](#)).

Once the HA pair is configured, the remaining configuration should take place on the Primary unit, the Secondary unit will then be kept in sync automatically. For details on configuring an HA pair, please refer to [Configuring HA - Adding a Secondary Appliance](#).

8.2. Step 2 – Configure GSLB

8.2.1. Handling Multiple Subdomains, Including Wildcard Subdomains

Scenario

Object storage-related DNS configurations may use various DNS subdomains, for example:

- s3-<region/location>.domain.tld (e.g. `s3-region1.domain.tld`)

Some scenarios also require the use of wildcard DNS entries, for example to cover bucket specific subdomains like `app-instance-f57ac0.s3-region1.domain.tld`.

Solution

Configuring DNS delegation can be complex. As such, the supported solution is to:



- Delegate a single subdomain to the load balancer, e.g. `gslb`.
- Use CNAME records to point everything else at the delegated subdomain

For example, the subdomain `gslb.domain.tld` would be delegated and everything else would point to it. This would look like so:

<code>gslb.</code>	Delegate to the load balancer
<code>s3-<region>.</code>	CNAME to <code>gslb.domain.tld</code>
<code>*.s3-<region>.</code>	CNAME to <code>gslb.domain.tld</code>
<code>s3-admin-console.</code>	CNAME to <code>gslb.domain.tld</code>

This approach simplifies DNS entry configuration, particularly when wildcard entries are involved.

8.2.2. Appliance Configuration

The GSLB service should be configured on the **primary** load balancer appliance and should be configured at each site if a multi-site deployment is being configured.

Configuration takes place in the WebUI under *Cluster Configuration > GSLB Configuration*:

Step 1 – Configuring the Global Name

1. Using the WebUI on the primary appliance, navigate to *Cluster Configuration > GSLB Configuration*.
2. Select the **Global Names** tab.
3. Click the **New Global Name** button.
4. Define a friendly **Name** for the new hostname, which can just be the subdomain itself, e.g. `gslb.domain.tld`


 **Note**

If only working with a *single* subdomain then it's perfectly acceptable to directly delegate the specific subdomain in question, e.g. `s3-region1.domain.tld`, rather than delegating a generic subdomain like `gslb.domain.tld`.

5. Define the **Hostname** of what will be the delegated subdomain, e.g. `gslb.domain.tld`
6. Click **Submit**.



GSLB Configuration



The screenshot shows the 'New Global Name' dialog box. It contains three input fields: 'Name' (gslb.domain.tld), 'Hostname' (gslb.domain.tld), and 'TTL' (30 seconds). Below the fields are 'Submit' and 'Cancel' buttons.

New Global Name	
Name	gslb.domain.tld
Hostname	gslb.domain.tld
TTL	30 seconds


Submit Cancel

Step 2 – Configure the Members

Each *member* is a single Qumulo node.

1. Select the **Members** tab.
2. Click the **New Member** button.
3. Enter a friendly *Name* for the member, e.g. **qumulo-node1**.
4. Specify an *IP* address for the member: in this context, this should be the IP address of the Qumulo node in question, e.g. **10.0.0.11**.
5. Ignore the example value in the *Monitor IP* field.
6. Click **Submit**.
7. Repeat these steps to add additional Qumulo nodes as members as required.

GSLB Configuration



The screenshot shows the 'New Member' dialog box. It contains four input fields: 'Name' (qumulo-node1), 'IP' (10.0.0.11), 'Monitor IP' (10.2.0.1), and 'Weight' (1). Below the fields are 'Submit' and 'Cancel' buttons.

New Member	
Name	qumulo-node1
IP	10.0.0.11
Monitor IP	10.2.0.1
Weight	1

Submit Cancel

Step 3 – Configure the Pool

A pool must be created to link together a global name with the members that should serve traffic for that global



name.

Continuing with the example presented in this section, a pool would be created linking the global name `gslb.domain.tld` with the members (Qumulo nodes), all of which should serve Qumulo traffic.

1. Select the **Pools** tab.
2. Click the **New Pool** button.
3. Enter a friendly **Name** for the pool, e.g. **qumulo-nodes**.
4. Set the **Monitor** to **TCP**.
5. Set **Monitor Port** to **443**.
6. Set **LB Method** to **wrr**.
7. From the **Global Names** list box, select the global name in question, e.g. **gslb.domain.tld**
8. In the **Members** section, drag the appropriate members (Qumulo nodes) from the **Available Members** box into the **Members In Use** box.
9. Click **Submit**.

New Pool

Name	<input type="text" value="qumulo-nodes"/>	?								
Monitor	<input type="text" value="TCP"/>	?								
Monitor Port	<input type="text" value="443"/>	?								
Monitor Send String	<input type="text" value="check"/>	?								
Monitor Match Return	<input type="text" value="up"/>	?								
LB Method	<input type="text" value="wrr"/>	?								
Global Names	<input type="text" value="gslb.domain.tld"/>	?								
Members	<table><thead><tr><th>Available Members</th><th>Members In Use</th></tr></thead><tbody><tr><td></td><td><input type="text" value="qumulo-node1"/></td></tr><tr><td></td><td><input type="text" value="qumulo-node2"/></td></tr><tr><td></td><td><input type="text" value="qumulo-node3"/></td></tr></tbody></table>	Available Members	Members In Use		<input type="text" value="qumulo-node1"/>		<input type="text" value="qumulo-node2"/>		<input type="text" value="qumulo-node3"/>	?
Available Members	Members In Use									
	<input type="text" value="qumulo-node1"/>									
	<input type="text" value="qumulo-node2"/>									
	<input type="text" value="qumulo-node3"/>									

[Advanced](#) **Submit** **Cancel**

Step 4 – Finalising the Configuration



To apply the new settings, the GSLB service must be restarted as follows:

1. Using the WebUI, navigate to: **Maintenance > Restart Services** and click **Restart GSLB**.

8.2.3. DNS Server Configuration

Once the GSLB service has been configured on the primary load balancer at every site, the DNS server at each site must then be configured for GSLB.

The DNS server at each site must be configured to delegate DNS requests for the subdomain in question to the load balancers; the load balancers' GSLB services will serve the appropriate IP addresses to the DNS servers. Using the example presented throughout this section, the DNS server at each site would be configured with a delegation for the domain `gslb.domain.tld`. The domain would be delegated to every load balancer across every site, which provides multi-site redundancy.

Steps walking through creating a DNS delegation on a Microsoft DNS server in the context of setting up GSLB on our appliance can be found in the appendix, in the section [Microsoft DNS Server Configuration](#).

8.2.4. Testing the Configuration

The configuration can be tested to make sure it's working as expected.

From the command line on a Microsoft Windows machine, the `nslookup` program can be used to send test DNS queries to the load balancer(s). The primary load balancer is located at IP address 10.0.0.1 in the example presented here.

For the test, use the `-norecurse` option to instruct the load balancer **not** to attempt to query another server for the answer. A successful test would see the load balancer respond with the IP address of one of the online Qumulo nodes, like so:

```
C:\Users\me>nslookup -norecurse gslb.domain.tld 10.0.0.1
Server: UnKnown
Address: 10.0.0.1

Name: gslb.domain.tld
Address: 10.0.0.11
```

9. Microsoft DNS Server Configuration

Once the GSLB service has been fully configured on the primary load balancer at every site, as described in the previous sections, the DNS server at each site must be configured for GSLB.

The DNS server at each site must be configured to delegate DNS requests for the subdomain in question to the load balancers; the load balancers' GSLB services will serve the appropriate IP addresses to the DNS servers. Using the example presented throughout this document, the DNS server at each site would be configured with a delegation for the domain `gslb.domain.tld`. The domain would be delegated to every load balancer across every site, which provides multi-site redundancy.

The exact steps for creating a DNS delegation vary between different DNS servers. Presented below are steps




that walk through creating a DNS delegation on a Microsoft DNS server in the context of setting up GSLB on our appliance.



9.1. Microsoft DNS Server

Delegating a subdomain in Microsoft DNS Manager is a short process.

1. Open **DNS Manager** and create A records for every load balancer at every site, using *Action > New Host* (e.g. `dc1-lbprimary.domain.tld`, `dc1-lbsecondary.domain.tld`, `dc2-lbprimary.domain.tld`, and `dc2-lbsecondary`).



2. Provided that the load balancer part of the GSLB configuration has been completed and is working, the **New Delegation** wizard should now be used to delegate the subdomain to the load balancers. The delegation will use the new FQDNs for the load balancers, as defined in the previous step. The delegation wizard is located at *Action > New Delegation*.



3. Test the delegation to make sure it is working as expected.

From the Windows command line, the `nslookup` program can be used to send test DNS queries to the DNS server. The DNS server is located at IP address 10.0.0.50 in the example presented here.

For the first test, use the `-norecurse` option to instruct the DNS server **not** to query another server for the answer. A successful test would see the DNS server respond and indicate that the subdomain in question is served by another server(s), giving the other server's details, like so:

```
C:\Users\me>nslookup -norecurse gslb.domain.tld 10.0.0.50
Server: UnKnown
Address: 10.0.0.50

Name:   gslb.domain.tld
Served by:
- dc1-lbprimary.domain.tld
  10.0.0.100
  gslb.domain.tld
```



```
- dc1-lbsecondary.domain.tld
    10.0.0.101
    gslb.domain.tld
- dc2-lbprimary.domain.tld
    172.16.0.100
    gslb.domain.tld
- dc2-lbsecondary.domain.tld
    172.16.0.101
    gslb.domain.tld
```

For the second test, execute the same command **without** the `-norecurse` option. This should see the DNS server fetch the answer from the load balancer and then serve up the 'fetched' answer in its response. A successful test would see the server reply with the IP address of one of the online sites/services, like so:

```
C:\Users\me>nslookup gslb.domain.tld 10.0.0.50
Server: UnKnown
Address: 10.0.0.50

Non-authoritative answer:
Name:   gslb.domain.tld
Address: 10.0.0.2
```

10. Testing & Verification

Note

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

Since the load balancers have no Virtual Services, there is no graphical overview of 'healthy' services to check and verify on the System Overview page of the WebUI. Instead, the GSLB configuration should be checked, ensuring that the client is able to resolve the FQDN of the Qumulo deployment via delegation and connect to a healthy Qumulo node.

See the earlier section [Testing the Configuration](#) for instructions on how to test the GSLB setup.

10.1. Accessing the Service

A successful test will see the test connection passed directly from the test client to one of the online Qumulo nodes.

11. 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.

12. Further Documentation

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



13. Appendix

13.1. 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.

13.1.1. 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



① Important

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


13.1.2. Configuring the HA Clustered Pair

ℹ 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



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



Primary

IP: 192.168.110.40

Attempting to pair..

Secondary


IP: 192.168.110.41

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](#).

14. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	23 April 2020	First draft		RJC
1.0.1	29 April 2020	Various minor updates	Improved readability	RJC
1.0.2	3 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.1.0	1 December 2021	Converted the document to AsciiDoc	Move to new documentation system	AH, RJC, ZAC
1.2.0	6 April 2022	Updated GSLB set up instructions to use GUI-driven GSLB configuration Updated DNS server configuration instructions	GSLB updates across all documentation Changed to use new, consistent common component	AH
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





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://www.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.

