

Load Balancing Everpure FlashBlade//S

Version 1.0.0



Table of Contents

1. About this Brief	4
2. Loadbalancer.org Appliances Supported	4
3. Software Versions Supported	4
3.1. Loadbalancer.org Appliance	4
3.2. Everpure FlashBlade//S	4
4. Everpure FlashBlade//S	4
5. Load Balancing Everpure FlashBlade//S	4
5.1. Load Balancing & HA Requirements	4
5.2. Virtual Service (VIP) Requirements	5
5.3. GSLB Requirements	5
5.4. SSL/TLS Termination	6
5.5. Health Checks	6
5.6. S3 Endpoint and DNS	6
6. Deployment Concept	6
7. Load Balancer Deployment Methods	7
7.1. Layer 7 SNAT Mode	7
8. Configuring Everpure FlashBlade//S for Load Balancing	8
9. Loadbalancer.org Appliance – the Basics	9
9.1. Virtual Appliance	9
9.2. Initial Network Configuration	9
9.3. Accessing the Appliance WebUI	9
9.3.1. Main Menu Options	11
9.4. Appliance Software Update	12
9.4.1. Online Update	12
9.4.2. Offline Update	12
9.5. Ports Used by the Appliance	13
9.6. HA Clustered Pair Configuration	14
10. Appliance Configuration for Everpure FlashBlade//S	14
10.1. VIP 1 - FlashBladeVIP	14
10.1.1. Virtual Service (VIP) Configuration	14
10.1.2. Configure the Associated Real Servers (RIPs)	15
10.2. Finalizing the Virtual Service Configuration	15
10.3. GSLB Configuration	15
10.3.1. Step 1 – Configure the Global Name	16
10.3.2. Step 2 – Configure the Members	16
10.3.3. Step 3 – Configure the Pool	18
10.3.4. Step 4 – Finalizing the GSLB Configuration	20
10.4. DNS Server Configuration	21
11. Testing & Verification	21
11.1. Using System Overview	21
11.2. Testing the GSLB Configuration	21
11.3. Accessing Everpure FlashBlade//S via the Load Balancer	22
12. Technical Support	22
13. Further Documentation	22
14. Appendix	23
14.1. Microsoft DNS Server Delegation	23
14.2. Configuring HA - Adding a Secondary Appliance	25
14.2.1. Non-Replicated Settings	25

14.2.2. Configuring the HA Clustered Pair	26
15. Document Revision History	29

1. About this Brief

This brief outlines the steps required to configure a load balanced Everpure FlashBlade//S environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Everpure FlashBlade//S 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 Everpure FlashBlade//S. 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

- V8.13.2 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. Everpure FlashBlade//S

- All versions

4. Everpure FlashBlade//S

Everpure (formerly known as Pure Storage) FlashBlade//S is unified, all-flash unstructured data storage that scales with you. It is a single, scale-out file and object storage array that scales effortlessly, with no performance drops, complex tuning, or downtime.

5. Load Balancing Everpure FlashBlade//S

Note

It's highly recommended that you have a working Everpure FlashBlade//S environment first before implementing the load balancer.

5.1. Load Balancing & HA Requirements

FlashBlade//S deployments typically consist of multiple FlashBlade//S clusters, often located in different data



centers. Two complementary load balancing functions are used:

- Within a site, a Virtual Service (VIP) on the load balancer distributes client connections across the data interfaces of the local FlashBlade//S cluster. This provides load sharing across the cluster and ensures that requests are only sent to healthy endpoints.
- Across sites, the Global Server Load Balancing (GSLB) functionality of the appliance is used to distribute requests between the FlashBlade//S clusters. This allows the storage solution to scale beyond a single cluster and provides resilience if a whole cluster or site becomes unavailable.

Clients use a single hostname to reach the storage service, irrespective of which cluster ultimately serves the request.

5.2. Virtual Service (VIP) Requirements

To provide load balancing and HA for Everpure FlashBlade//S, the following VIP is required at each site:

Ref.	VIP Name	Mode	Port(s)	Persistence Mode	Health Check
VIP 1	FlashBladeVIP	L7 SNAT (TCP)	443	None	Connect to Port

The Virtual Service operates in TCP mode, so HTTPS traffic is passed through to the FlashBlade//S cluster untouched. Persistence is not required: S3 is a stateless, request/response protocol, so each connection can be sent to whichever endpoint is least loaded.

5.3. GSLB Requirements

GSLB is used to hand out the site VIP address of a healthy FlashBlade//S cluster in response to DNS queries for the storage service's hostname. To provide this, the following GSLB objects are required:

Object	Example Value	Purpose
Global Name	FlashBlade / www.myflashblade.com	The hostname that clients use and that GSLB answers for.
Member	FB1 - 192.168.1.111	The site VIP for the first FlashBlade//S cluster.
Member	FB2 - 192.168.2.111	The site VIP for the second FlashBlade//S cluster.
Pool	FlashBladePool	Associates the global name with both members, and defines the health check and the load balancing method.

Two pool configurations are presented in this brief. Both use the same global name and members, and differ only in the **LB Method** used:

Scenario 1 - Active/Active (wrr)

Weighted round robin. DNS answers alternate between the healthy cluster VIPs, so client requests are shared across both FlashBlade//S clusters. Member weights can be adjusted where one cluster is larger than the other.



Scenario 2 - Active/Passive (fogroup)

Failover group. The first healthy member in the list is handed out continuously; the second member is only used if the first becomes unhealthy. Use this where one cluster should be preferred and the other is a standby.

Note

When **LB Method** is set to **fogroup**, the order in which the members are listed in the pool is the order in which they are used. Make sure the preferred cluster is listed first.

Note

For multi-site deployments where clients should be directed to their nearest cluster, the **twrr** (topology weighted round robin) method can be used instead, together with a topology definition. This brief does not cover topology configuration - if this is a requirement, please [contact support](#) for assistance.

5.4. SSL/TLS Termination

For raw performance and simplicity, we recommend that SSL is terminated on the FlashBlade//S cluster. This is the method used in this brief: the Virtual Service is configured in TCP mode and passes HTTPS connections straight through to the cluster, so **Re-Encrypt to Backend** is not required on the Real Servers.

Terminating SSL on the load balancer should only be used when specifically required. For example, when centralizing certificate management or when applying Layer 7 features such as ACLs, redirects and headers. When SSL is terminated on the load balancer, it's recommended that backend encryption is enabled so that traffic between the load balancer and the FlashBlade//S cluster remains encrypted.

5.5. Health Checks

Two independent levels of health checking are used:

- The Virtual Service checks each FlashBlade//S endpoint by connecting to TCP port 443 (**Connect to port**).
- The GSLB pool checks each site VIP with an HTTPS GET to **/** on port 443 and expects an HTTP **200** response. The check is sent with a **Monitor Hostname** so that the correct SNI and **Host :** header are presented to the cluster.

5.6. S3 Endpoint and DNS

Clients should always use a hostname endpoint rather than an IP address. The hostname is served by GSLB, which means that DNS for that name must be delegated to the load balancers - see [DNS Server Configuration](#).


If you plan to use virtual-host-style S3 addressing (bucket as subdomain), appropriate DNS (wildcard or per-bucket) and matching certificates are also needed. If using path-style addressing, a single hostname is typically sufficient.

6. Deployment Concept

Once the load balancer is deployed, clients connect to the Virtual Services (VIPs) rather than connecting directly to one of the FlashBlade//S endpoints. These connections are then load balanced across the FlashBlade//S cluster to distribute the load according to the load balancing algorithm selected. GSLB determines which site's



VIP a client is given in the first place.



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 Everpure FlashBlade//S, layer 7 SNAT mode is recommended. This mode is described below and is used for the configuration presented in this brief.

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. Configuring Everpure FlashBlade//S for Load Balancing

Data Interfaces

Ensure that the FlashBlade//S data interfaces (or data VIPs) that will be load balanced are configured and reachable from the load balancer, and that the S3 service is listening on port 443.

Certificates

Install a valid certificate on the FlashBlade//S cluster that covers the hostname clients will use, e.g. www.myflashblade.com. The same certificate should be used on every cluster that serves the global

name.

DNS

Ensure that DNS is updated so that the client endpoint hostname is delegated to the load balancers rather than pointing at individual FlashBlade//S clusters - see [DNS Server Configuration](#).

No other changes are required on the FlashBlade//S cluster: layer 7 SNAT mode VIPs do not require any mode specific configuration changes to the load balanced Real Servers.

9. Loadbalancer.org Appliance – the Basics

9.1. 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.

9.2. 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.

9.3. 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](#)

System Overview ? 2025-05-08 12:37:21 UTC

Would you like to run the Setup Wizard?

VIRTUALL SERVICE IP PORTS CONNS PROTOCOL METHOD MODE
No Virtual Services configured.

Network Bandwidth

	Min	Avg	Total
RX	28	2713	27344772
TX	0	13777	138872181

System Load Average

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

Memory Usage

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.

9.3.1. 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

9.4. 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.

9.4.1. 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 8.13.10 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.

9.4.2. 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.

9.5. 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](#)

9.6. HA Clustered Pair Configuration

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

10. Appliance Configuration for Everpure FlashBlade//S

The steps in this section should be followed on the load balancer (or HA pair) at each site. The example values used are those of the first site; at the second site the equivalent addresses on the 192.168.2.0/24 subnet would be used.

10.1. VIP 1 - FlashBladeVIP

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**.

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]
Label	FlashBladeVIP	?
IP Address	192.168.1.111	?
Ports	443	?
Protocol		[Advanced +]
Layer 7 Protocol	TCP Mode	?
		Cancel Update

2. Enter a suitable *Label* (name) for the Virtual Service, e.g. **FlashBladeVIP**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.1.111**.
4. Set *Ports* to **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 to the *Persistence* section.



- Set the *Persistence Mode* to **None**.
9. Scroll to the *Health Checks* section.
 - Set the *Health Checks* to **Connect to port**.
 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.

Layer 7 Add a new Real Server - FlashBladeVIP

Label	FBNode1	?
Real Server IP Address	192.168.1.112	?
Real Server Port		?
Re-Encrypt to Backend	☐	?
Weight	100	?

2. Enter a suitable *Label* (name) for the Real Server, e.g. **FBNode1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.1.112**.
4. Leave the *Real Server Port* field blank - the port defined on the Virtual Service (443) will be used.
5. Leave *Re-Encrypt to Backend* unchecked - SSL is terminated on the FlashBlade//S cluster.
6. Leave the *Weight* set to **100**.
7. Click **Update**.
8. Repeat these steps to add the remaining FlashBlade//S endpoints, e.g. **FBNode2** at **192.168.1.113**.

10.2. Finalizing the Virtual Service 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**.

10.3. GSLB Configuration

GSLB should be configured on the **primary** load balancer appliance at each site. The configuration must be



identical at every site so that consistent DNS responses are given irrespective of which appliance answers a query.

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

10.3.1. Step 1 – Configure 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.

GSLB - Configuration

The screenshot shows the 'GSLB - Configuration' web interface. At the top, there are four tabs: 'Global Names' (selected), 'Members', 'Pools', and 'Topologies'. In the top right corner, there is a button labeled 'New Global Name'. Below this, a form titled 'New Global Name' is displayed. The form contains three input fields: 'Name' with the value 'FlashBlade', 'Hostname' with the value 'www.myflashblade.com', and 'TTL' with the value '60' and a unit of 'seconds'. Each input field has a help icon (a question mark in a circle) to its right. At the bottom of the form, there are two buttons: 'Submit' (blue) and 'Cancel' (red).

4. Enter a friendly **Name** for the global name, e.g. **FlashBlade**.
5. Set the **Hostname** to the delegated hostname that clients will use, e.g. **www.myflashblade.com**.
6. Set the **TTL** to a suitable value, e.g. **60** seconds.
7. Click **Submit**.

Note

A short TTL allows clients to pick up a change of cluster quickly following a failure, at the cost of additional DNS traffic.

10.3.2. Step 2 – Configure the Members

Each **member** is the site VIP of one FlashBlade//S cluster.

1. Select the **Members** tab.
2. Click the **New Member** button.

GSLB - Configuration

Global Names

Members

Pools

Topologies

New Member

New Member

Name	FB1	?
IP	192.168.1.111	?
Monitor IP	192.168.1.111	?
Weight	1	?

Submit

Cancel

3. Enter a friendly *Name* for the member, e.g. **FB1**.
4. Set the *IP* to the site VIP address of the first cluster, e.g. **192.168.1.111**.
5. Set the *Monitor IP* to the same address, e.g. **192.168.1.111**.
6. Leave the *Weight* set to **1**.
7. Click **Submit**.
8. Repeat these steps for the second cluster, e.g. **FB2** at **192.168.2.111**.

GSLB - Configuration

Global Names

Members

Pools

Topologies

New Member

New Member

Name	FB2	?
IP	192.168.2.111	?
Monitor IP	192.168.2.111	?
Weight	1	?

Submit

Cancel

10.3.3. Step 3 – Configure the Pool

A pool links the global name to the members that should serve traffic for it, and defines the health check and the load balancing method.

Scenario 1 – Active/Active (wrr)

1. Select the **Pools** tab.
2. Click the **New Pool** button.

GSLB - Configuration

Global NamesMembersPoolsTopologies

New Pool

Name	FlashBladePool	?
Monitor	HTTP	?
Monitor Use SSL	Yes	?
Monitor Hostname	www.myFB.com	?
Monitor URL Path	/	?
Monitor Port	443	?
Monitor Expected Codes	200	?
LB Method	wrr	?
Global Names	FlashBlade	?
Members	Available Members Members In Use FB1 FB2	?

Advanced

SubmitCancel

3. Enter a friendly *Name* for the pool, e.g. **FlashBladePool**.
4. Set the *Monitor* to **HTTP**.

5. Set **Monitor Use SSL** to **Yes**.
6. Set the **Monitor Hostname** to the hostname presented on the FlashBlade//S certificate, e.g. **www.myFB.com**.
7. Set the **Monitor URL Path** to **/**.
8. Set the **Monitor Port** to **443**.
9. Set the **Monitor Expected Codes** to **200**.
10. Set the **LB Method** to **wrr**.
11. From the **Global Names** list box, select the global name created in step 1, e.g. **FlashBlade**.
12. In the **Members** section, drag both members (**FB1** and **FB2**) from the **Available Members** box into the **Members In Use** box.
13. Click **Submit**.

Scenario 2 – Active/Passive (fogroup)

The configuration is identical to scenario 1, with the exception of the load balancing method and the order of the members:

1. Select the **Pools** tab.
2. Click the **New Pool** button.

Global Names
Members
Pools
Topologies

New Pool

Name	FlashBladePool	?
Monitor	HTTP	?
Monitor Use SSL	Yes	?
Monitor Hostname	www.myFB.com	?
Monitor URL Path	/	?
Monitor Port	443	?
Monitor Expected Codes	200	?
LB Method	fogroup	?
Global Names	FlashBlade	?
Members	Available Members Members In Use FB1 FB2	?

Advanced

Submit
Cancel

- Configure the pool as described in scenario 1, but set the **LB Method** to **fogroup**.
- In the **Members** section, drag the members into the **Members In Use** box in the order in which they should be used - the preferred cluster (**FB1**) first, the standby cluster (**FB2**) second.
- Click **Submit**.

10.3.4. Step 4 – Finalizing the GSLB 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**.

10.4. DNS Server Configuration

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

DNS must be configured to delegate requests for the hostname in question to the load balancers; the load balancers' GSLB services then serve up the appropriate site VIP address. Using the example presented in this brief, DNS would be configured with a delegation for www.myflashblade.com. The name should 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. Steps to configure DNS delegation on a Microsoft DNS server can be found in the appendix, in the section [Microsoft DNS Server Delegation](#).

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 FlashBlade//S endpoints) and shows the state/health of each endpoint as well as the overall state of each cluster. The example below shows that all endpoints are healthy (green) and available to accept connections:

System Overview ?

2026-06-08 14:03:29 UTC

	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE	
↑	FlashBladeVIP	192.168.1.111	443	0	TCP	Layer 7	Proxy	📊
	REAL SERVER	IP	PORTS	WEIGHT	CONNS			
↑	FBNode1	192.168.1.112	443	100	0	Drain	Halt	📊
↑	FBNode2	192.168.1.113	443	100	0	Drain	Halt	📊

11.2. Testing the GSLB Configuration

From the command line on a Microsoft Windows machine, the **nslookup** program can be used to send test DNS queries directly to the load balancer. The primary load balancer is located at IP address 192.168.1.1 in the example below.

Use the **-norecurse** option to instruct the load balancer **not** to attempt to query another server for the answer. A successful test sees the load balancer respond with the site VIP address of one of the online FlashBlade//S clusters, like so:



```
C:\Users\me>nslookup -norecurse www.myflashblade.com 192.168.1.1
Server: UnKnown
Address: 192.168.1.1

Name: www.myflashblade.com
Address: 192.168.1.111
```

Repeat the query several times. With **LB Method** set to **wrr**, the answer should alternate between the site VIPs of the two clusters. With **LB Method** set to **fogroup**, the same address should be returned each time.

Then halt the FlashBlade//S endpoints at the first site, or block the GSLB health check to that site's VIP, and repeat the query. Once the health check has failed, only the address of the second cluster should be returned.

11.3. Accessing Everpure FlashBlade//S via the Load Balancer

Connect to the hostname served by GSLB and use an S3 client (e.g. [aws s3api](#) or [s3cmd](#)) to perform various S3 operations (PUT, GET, LIST, DELETE). All operations should succeed when the clusters are healthy.

Repeat the tests with one FlashBlade//S endpoint halted using the **Halt** link in the System Overview, to confirm that the Virtual Service continues to serve requests from the remaining endpoints.

Note

Make sure that DNS is updated so that any FQDNs used point to the VIPs rather than individual FlashBlade//S endpoints.

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. Microsoft DNS Server Delegation

The Steps to create a DNS delegation on a Microsoft DNS server in the context of setting up GSLB on our appliance are presented below.

1. Using the **Action > New Host** menu option in **DNS Manager**, create A records for each load balancer at every site. For example:

- `dc1-lbprimary.domain.tld`
- `dc1-lbsecondary.domain.tld`
- `dc2-lbprimary.domain.tld`
- `dc2-lbsecondary.domain.tld`

Once created, the entries in **DNS Manager** would appear as follows:



The screenshot shows the DNS Manager window with the following structure:

- Left pane: DNS > WIN-H0HE1U96LPB > Forward Lookup Zones > domain.tld
- Right pane: A table of DNS records.

Name	Type	Data
(same as parent folder)	Start of Authority (SOA)	[1], win-h0he1u96lpb., ho...
(same as parent folder)	Name Server (NS)	win-h0he1u96lpb.
dc1-lbprimary	Host (A)	10.0.0.100
dc1-lbsecondary	Host (A)	10.0.0.101
dc2-lbprimary	Host (A)	172.16.0.100
dc2-lbsecondary	Host (A)	172.16.0.101

2. Ensure that the load balancer has been correctly configured for GSLB and is working, then use the **New Delegation** wizard to delegate the subdomain to the load balancers. The delegation must include the FQDNs for all the load balancer A records that were created in the previous step. The delegation wizard can be accessed using **Action > New Delegation**.

New Delegation Wizard

Delegated Domain Name
Authority for the DNS domain you supply will be delegated to a different zone.

Specify the name of the DNS domain you want to delegate.

Delegated domain:

Fully qualified domain name (FQDN):

< Back **Next >** Cancel

New Delegation Wizard

Name Servers
You can select one or more name servers to host the delegated zone.

Specify the names and IP addresses of the DNS servers you want to have host the delegated zone.

Name servers:

Server Fully Qualified Domain Name (FQDN)	IP Address
dc1-lbprimary.domain.tld.	[10.0.0.100]
dc1-lbsecondary.domain.tld.	[10.0.0.101]
dc2-lbprimary.domain.tld.	[172.16.0.100]
dc2-lbsecondary.domain.tld.	[172.16.0.101]

Add... Edit... Remove

< Back **Next >** Cancel

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

From the Windows command line, the **nslookup** command 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 a different server(s), giving the other server's details as shown in the following example:

```
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 nodes (in "direct to node" scenarios) or one of the Virtual Services (in GSLB to Virtual Service scenarios) as shown in the following example:

```

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

```

14.2. 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.

14.2.1. 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



WebUI Main Menu Option	Sub Menu Option	Description
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.

14.2.2. 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

- Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
- Click **Add new node**.
- 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

- 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	17 September 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.

