LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing Scality RING

Version 1.3.1

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced Scality RING environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any Scality RING 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 Scality RING. For full specifications of available models please refer to https://www.loadbalancer.org/products. Some features may not be supported in all cloud platforms due to platform specific limitations, please check with Loadbalancer.org support for further details.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.4.1 and later

> **Note** The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If using an older software version, note that the screenshots presented here may not match the WebUI exactly.

3.2. Scality RING

- 7.4.4 and later

4. Scality RING

Scality is a global company that develops software-defined object storage via commercial products such as RING. Scality RING software deploys on industry-standard x86 servers to store objects and files whilst providing compatibility with the Amazon S3 API.

Scality RING architecture supports High Availability (HA) clustering by putting a load balancer in front of it. Load balancers monitor and perform health checks on a node to ensure traffic is routed correctly to healthy nodes. Without the use of a load balancer, an off-line or failed node would still receive traffic, causing failures.

A variety of load balancing methods are currently supported by Scality RING, dependent on customer infrastructure, including layer 4, layer 7, and geo GSLB / location affinity. The RING service that should be load balanced is the S3 component.

5. Load Balancing Scality RING

> **Note** It's highly recommended that you have a working Scality RING environment first before

implementing the load balancer.

5.1. Load Balancing & HA Requirements

The function of the load balancer is to distribute inbound connections across a cluster of Scality RING nodes, to provide a highly available and scalable service. One virtual service is used to load balance the S3 aspect of RING.

5.2. Persistence (aka Server Affinity)

Client persistence is not required and should not be enabled.

5.3. Virtual Service (VIP) Requirements

To provide load balancing for Scality the following VIP is required:

- **S3**: handles requests from S3 client applications via HTTP and HTTPS

5.4. Port Requirements

The following table shows the ports that are load balanced:

Port	Protocols	Use
80	TCP/HTTP	Requests from S3 client applications
443	TCP/HTTPS	Requests from S3 client applications

5.5. SSL Termination

SSL termination on the load balancer is recommended for load balancing Scality RING.

5.6. Health Checks

The S3 service uses the "Negotiate HTTP (GET)" health check.

5.7. GSLB / Location Affinity

For multi-site RING deployments, it is possible to use the load balancer's GSLB functionality to provide high availability and location affinity across multiple sites. Using this optional, DNS based feature, in the event that a site's RING service and/or load balancers are offline then local clients are automatically directed to a functioning RING cluster at another site.

A full explanation and instructions on setting up this optional feature can be found in Configuring GSLB / Location Affinity.

5.8. Alternative Load Balancing Method for Read-Intensive Deployments (Direct Routing)

For deployments that are read-intensive, it is possible to use an alternative load balancing method known as ***Direct Routing***. This allows reply traffic to flow directly from the back end servers to the clients, thus removing the

load balancer as a potential bottleneck for reply traffic. Direct routing can benefit read-intensive deployments with a large reply traffic to request traffic ratio.

A more detailed explanation of this alternative load balancing method can be found in Alternative Load Balancing Method for Read-Intensive Deployments (Direct Routing).

6. Performance and Sizing for a Virtual Load Balancer Deployment with Scality RING

The Loadbalancer.org appliance can be deployed as a **Virtual Appliance**.

To achieve the best level of performance and throughput when load balancing a Scality RING deployment, the Loadbalancer.org appliance should be configured to actively use multiple CPU cores for the load balancing process. This must be considered when initially deploying and sizing virtual appliances.

A virtual host should be allocated a minimum of 4 vCPUs.

7. Deployment Concept

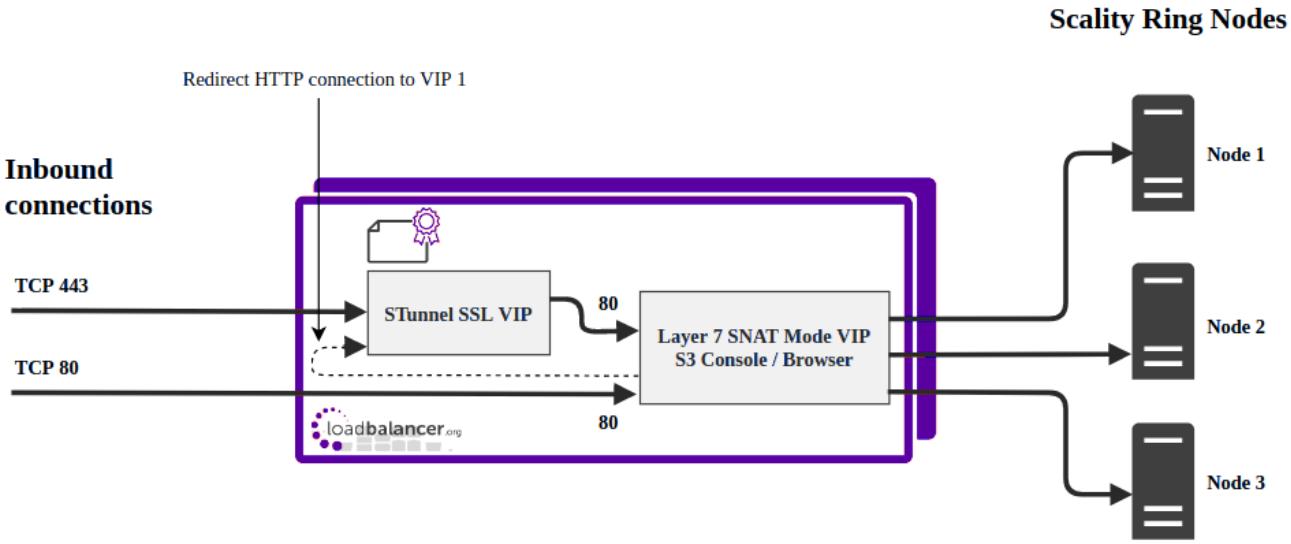

VIPs = **V**irtual **IP** Addresses

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 a clustered pair.

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 Server and other network 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.

> **Note** A number of compatibility issues have been found with various versions of Microsoft Internet Explorer and Edge. The WebUI has been tested and verified using both Chrome & Firefox.

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 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.

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

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

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

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

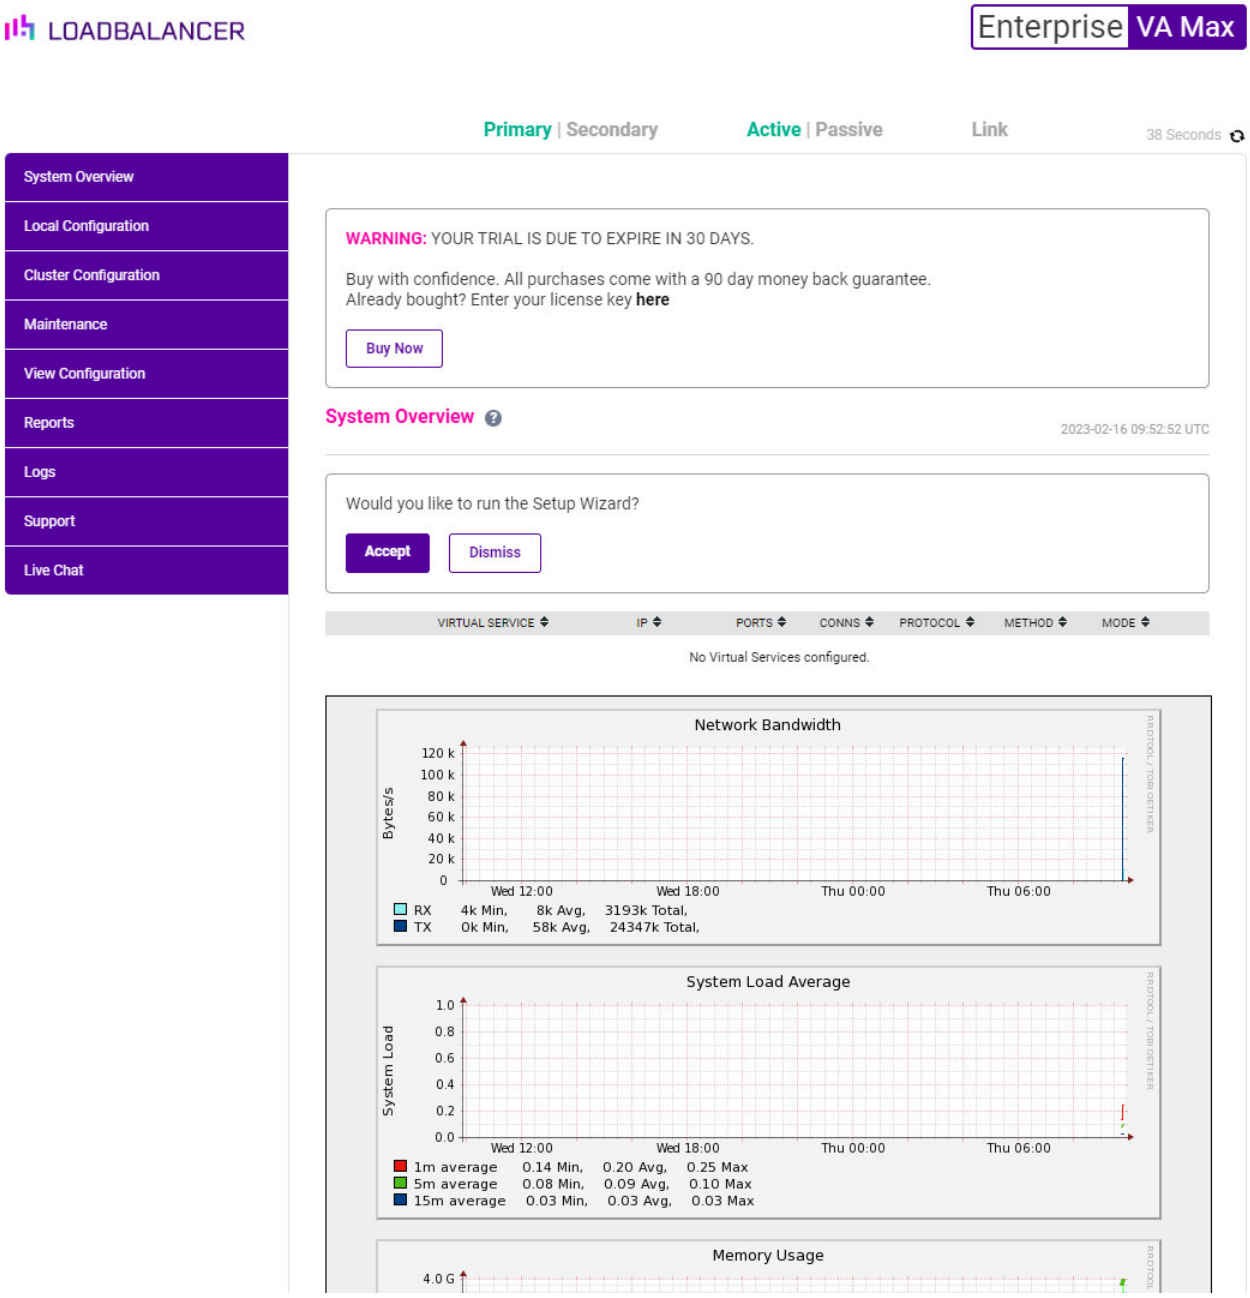

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

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 taking 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

To ensure that the appliance(s) are running the latest software version, we recommend a software update check is performed.

Determining the Current Software Version

The software version is displayed at the bottom of the WebUI as shown in the example below:

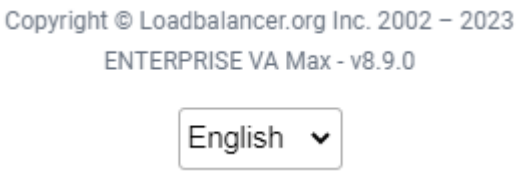

Checking for Updates using Online Update

> **Note** By default, the appliance periodically contacts the Loadbalancer.org update server and checks for updates. An update check can also be manually triggered as detailed below.

1. Using the WebUI, navigate to: *Maintenance > Software Update*.
2. Select **Online Update**.
3. If the latest version is already installed, a message similar to the following will be displayed:

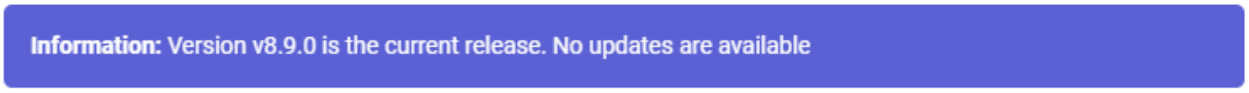

4. If an update is available, you'll be presented with a list of new features, improvements, bug fixes and security related updates.
5. Click **Online Update** to start the update process.

> **Note** Do not navigate away whilst the update is ongoing, this may cause the update to fail.

6. Once complete (the update can take several minutes depending on download speed and upgrade version) the following message will be displayed:

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

Using Offline Update

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

> **Note** Please contact support@loadbalancer.org to check if an update is available and obtain the latest offline update files.

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:

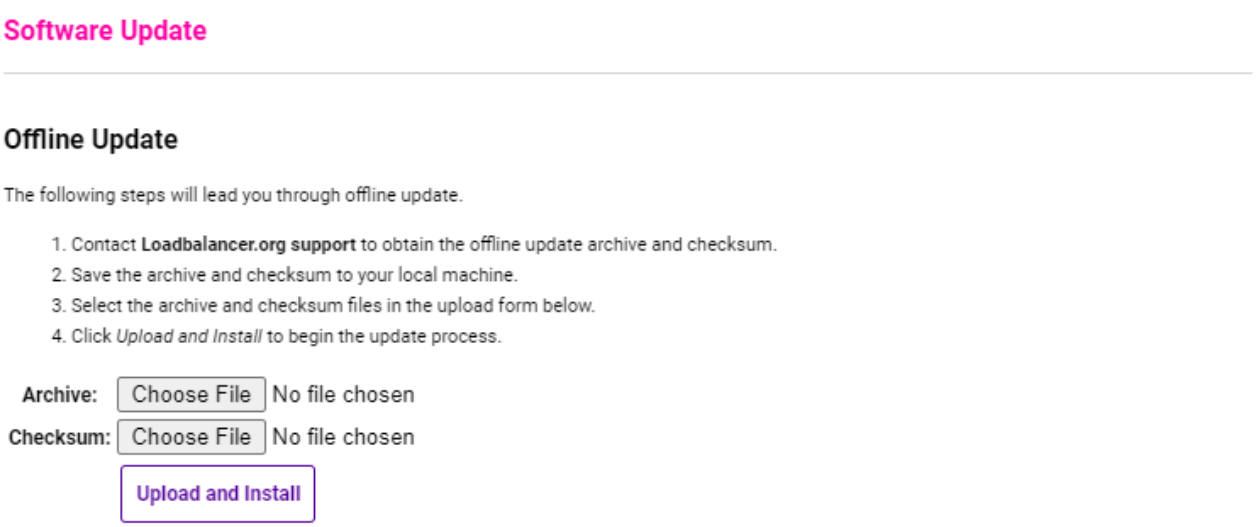

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

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	DNS
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	9080	WebUI - HTTP (disabled by default)
TCP	9081	Nginx fallback page
TCP	9443	WebUI - HTTPS

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 Configuring HA - Adding a Secondary Appliance.

9. Appliance Configuration for Scality RING

9.1. Enabling Multithreaded Load Balancing

> **Note**
>
> Multithreading is enabled by default for ***new*** load balancers starting from version 8.5.1 and does not require changing.
>
> ***If upgrading an older appliance*** then ensure that the multithreading configuration is set correctly, as described below.

The Loadbalancer.org appliance should be configured to actively use multiple CPU cores for the load balancing process. This is required to achieve the high level of performance and throughput required when load balancing a Scality RING deployment.

> **Note** A virtual host should be allocated a minimum of 4 vCPUs.

To enable multithreaded mode from the WebUI:

1. Navigate to *Cluster Configuration > Layer 7 - Advanced Configuration*.
2. Check the **Enable Multithreading** checkbox.
3. Check the **Default Number of Threads** checkbox.
4. Click **Update** to apply the changes.

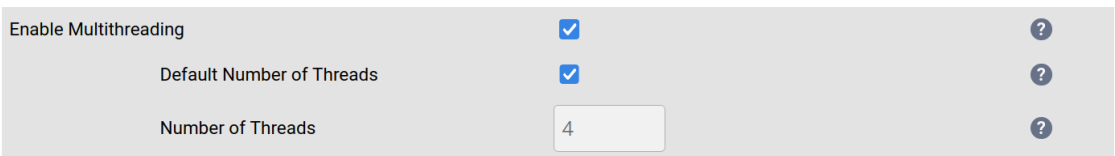

10. Appliance Configuration for Scality RING – Using Layer 7 SNAT

10.1. Configuring VIP 1 – S3

Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **S3**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **172.16.254.120**.
4. Set the *Ports* field to **80**.
5. Set the *Layer 7 Protocol* to **HTTP Mode**.
6. Click **Update** to create the virtual service.

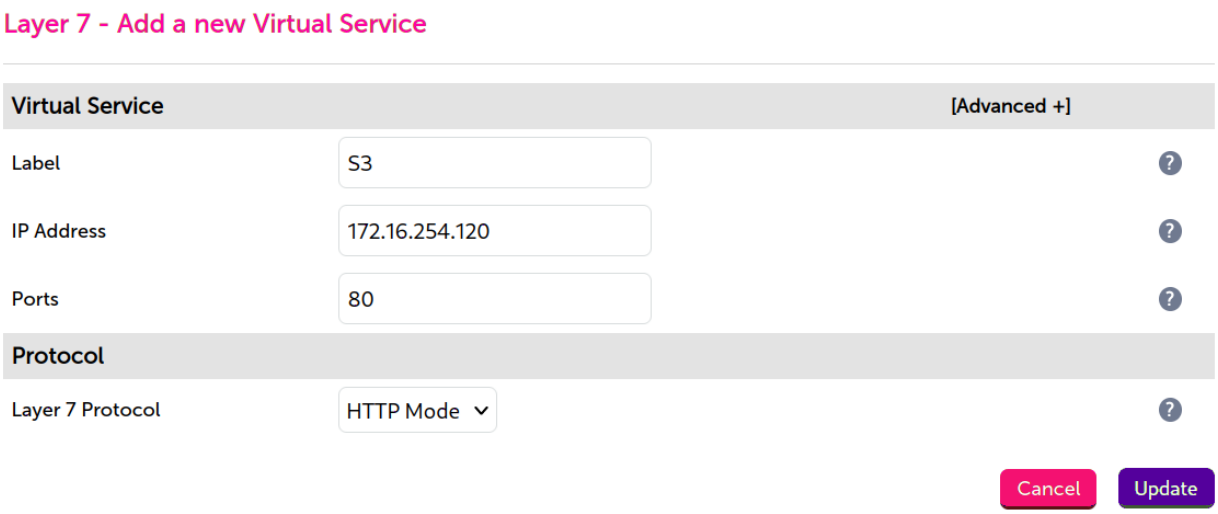

7. Click **Modify** next to the newly created VIP.
8. Set *Persistence Mode* to **None**.
9. Set *Health Checks* to **Negotiate HTTP (GET)**.
10. Set *Request to send* to **/_/healthcheck/deep/**.
11. Scroll to the *Other* section and click **Advanced**.
12. Enable *Force to HTTPS* by clicking the **Yes** radio button.

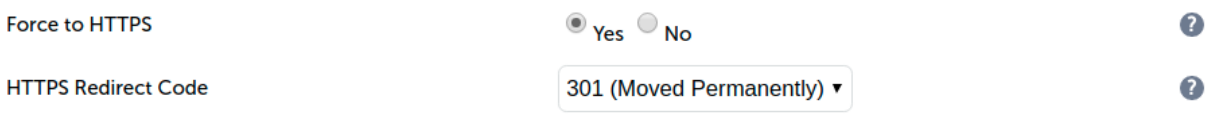

13. Click **Update**.

Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 7 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Define the *Label* for the real server as required, e.g. **ring-node1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **172.16.254.101**.
4. Click **Update**.
5. Repeat these steps to add additional RING nodes as real servers as required.

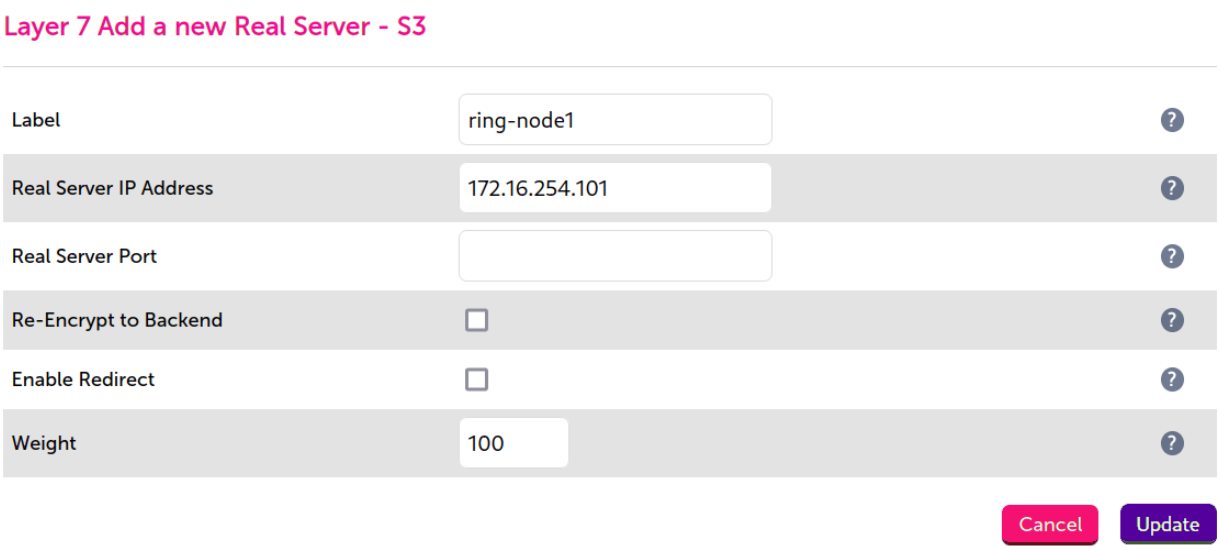

11. Additional Configuration Options & Settings

11.1. SSL Termination

SSL termination can be handled in the following ways:

1. On the Real Servers - aka **SSL Pass-through**.
2. On the load balancer – aka **SSL Offloading** (*recommend for Scality RING*).
3. On the load balancer with re-encryption to the backend servers – aka **SSL Bridging**.

In the case of Scality RING, it is recommended that SSL be terminated on the load balancer ***(SSL offloading)*** with **Force to HTTPS** enabled.

Notes

1. SSL termination on the load balancer can be very CPU intensive.
2. By default, a self-signed certificate is used for the new SSL VIP. Certificates can be requested on the load balancer or uploaded as described in the section below. The default self-signed certificate can be regenerated if needed using the WebUI menu option: SSL Certificate and clicking the **Regenerate Default Self Signed Certificate** button.

3. The backend for the SSL VIP can be either a Layer 7 SNAT mode VIP or a Layer 4 NAT or SNAT mode VIP. Layer 4 DR mode cannot be used since stunnel acts as a proxy, and the RING servers see requests with a source IP address of the VIP. However, since the RING servers believe that they own the VIP (due to the loopback adapter configured to handle to ARP problem) they are unable to reply to stunnel.
4. Finally, ensure that the Scality RING S3 Console and S3 Browser URL are configured as HTTPS via the S3 Service as per the example image below:

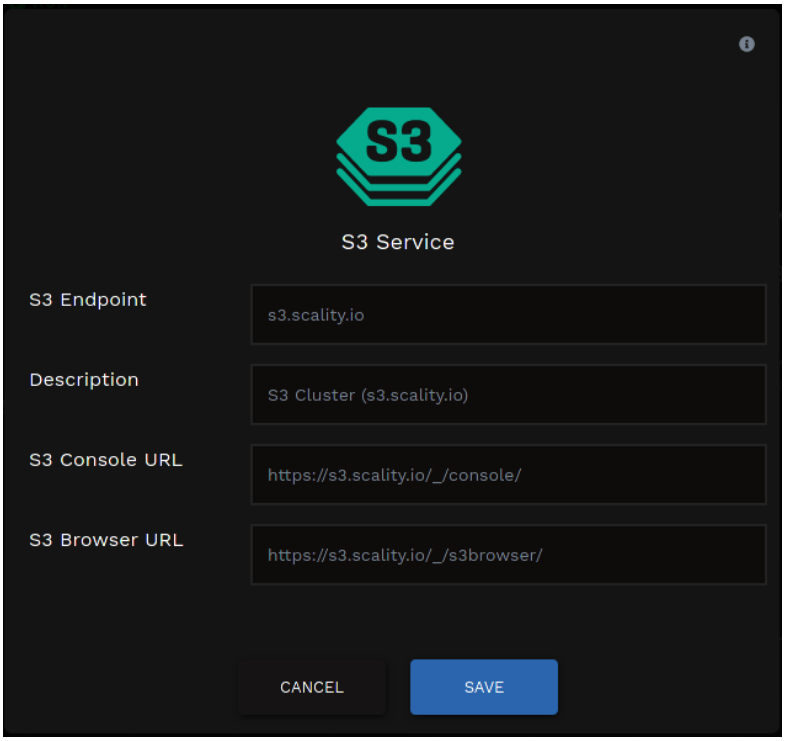

11.2. SSL Termination on the load balancer - SSL Offloading

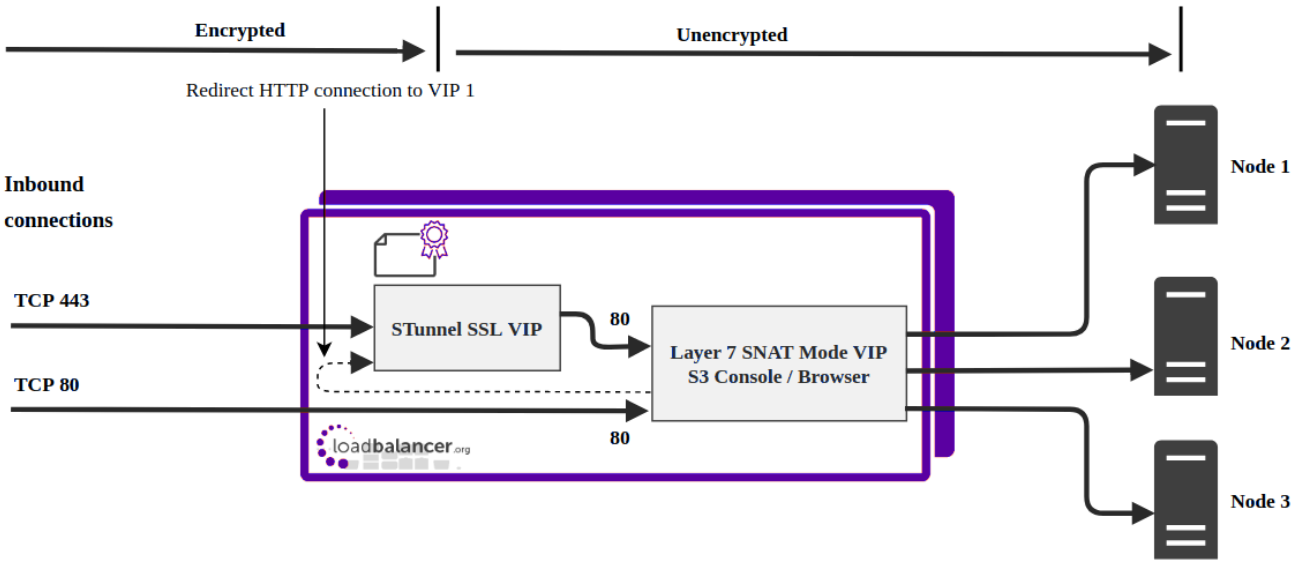

In this case, an SSL VIP utilizing STunnel is configured on the appliance and an SSL certificate is uploaded and associated to the Virtual Service. Data is encrypted from the client to the load balancer, but is unencrypted from the load balancer to the backend servers as shown above.

Certificates

If you already have an SSL certificate in either PFX or PEM file format, this can be uploaded to the Load balancer using the certificate upload option as explained in Uploading Certificates. Alternatively, you can create a Certificate Signing Request (CSR) on the load balancer and send this to your CA to create a new certificate. For more information please refer to Generating a CSR on the Load Balancer.

Uploading Certificates

If you already have a certificate in either PEM or PFX format, this can be uploaded to the load balancer.

To upload a Certificate:

1. Using the WebUI, navigate to: Cluster Configuration > SSL Certificates.
2. Click **Add a new SSL Certificate** & select Upload prepared PEM/PFX file.

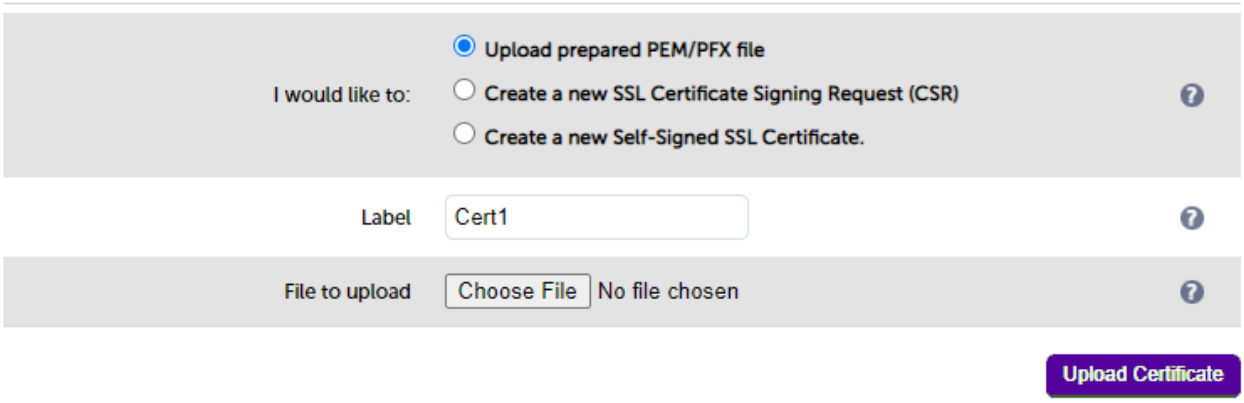

3. Enter a suitable Label (name) for the certificate, e.g. **Cert1**.
4. Browse to and select the certificate file to upload (PEM or PFX format).
5. Enter the password , if applicable.
6. Click **Upload Certificate**, if successful, a message similar to the following will be displayed:

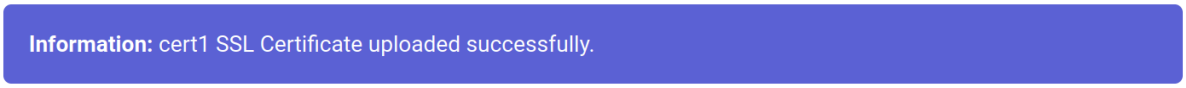

> **Note** It's important to backup all your certificates. This can be done via the WebUI from *Maintenance > Backup & Restore > Download SSL Certificates.*

11.3. Configuring SSL Termination on the Load Balancer

1. Using the WebUI, navigate to: *Cluster Configuration > SSL Termination* and click **Add a new Virtual Service**.

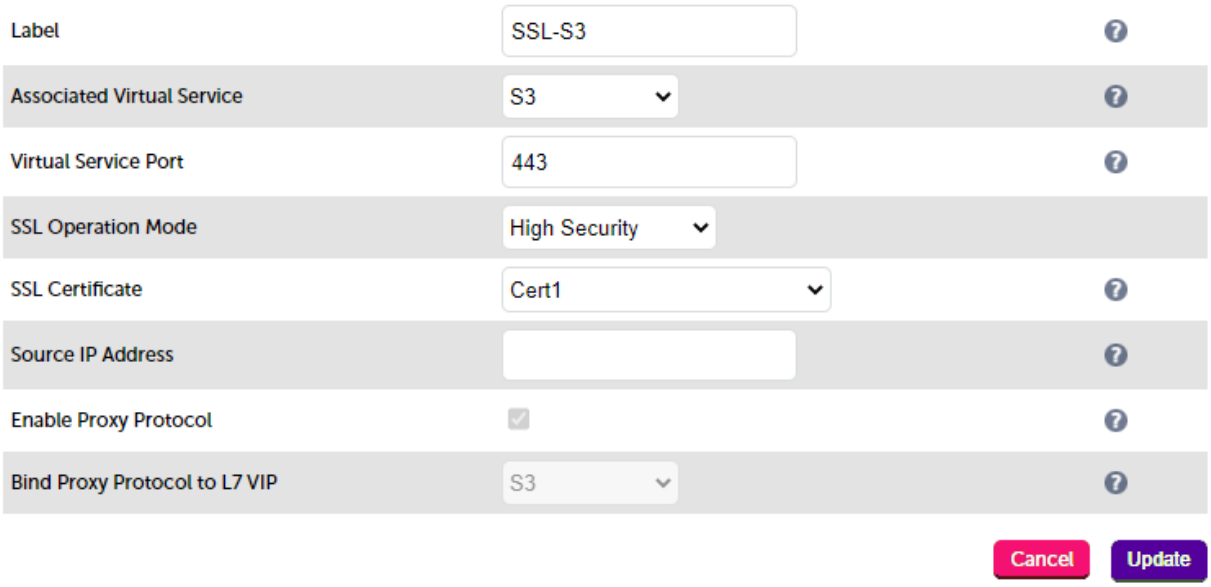

2. Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **S3**.

> **Note** Once the VIP is selected, the *Label* field will be auto-populated with **SSL-S3**. This can be changed if preferred.

> **Note** The Associated Virtual Service drop-down is populated with all single port, standard (i.e. non-manual) Layer 7 VIPs available on the load balancer. Using a Layer 7 VIP for the backend is the recommended method although as mentioned earlier, Layer 4 NAT mode and layer 4 SNAT mode VIPs can also be used if required. To forward traffic from the SSL VIP to these type of VIPs, you'll need to set Associated Virtual Service to **Custom**, then configure the IP address & port of the required VIP.

3. Leave *Virtual Service Port* set to **443**.
4. Leave *SSL Operation Mode* set to **High Security**.
5. Select the required certificate from the SSL Certificate drop-down.
6. Click **Update**.

Once configured, HTTP traffic will be load balanced by the Layer 7 SNAT mode VIP and HTTPS traffic will be terminated by the SSL VIP, then passed on to the Layer 7 SNAT mode VIP as unencrypted HTTP for load balancing.

11.4. Finalizing the Configuration

To apply the new settings, HAProxy and STunnel must both be reloaded. This can be done using the buttons 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**.

3. Click **Reload STunnel**.

12. Testing & Verification

> **Note** For additional guidance on diagnosing and resolving any issues you may have, please also refer to Diagnostics & Troubleshooting.

12.1. Using System Overview

The System Overview can be viewed in the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the RING Nodes) and shows the state/health of each server as well as the state of the each cluster as a whole. The example below shows that all RING nodes are healthy and available to accept connections.

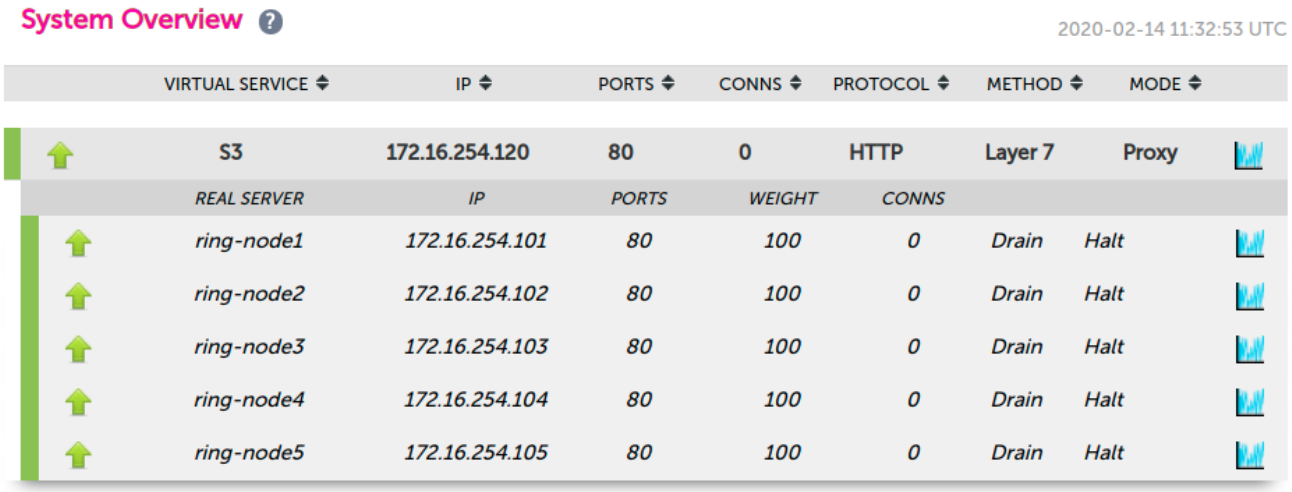

13. Technical Support

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.

14. Further Documentation

For additional information, please refer to the Administration Manual.

15. Appendix

15.1. Configuring GSLB / Location Affinity

Conceptual Overview

For **multi-site RING deployments**, it is possible to use the load balancer's global server load balancing (GSLB) functionality to provide both high availability and location affinity across multiple sites.

- Clients across multiple sites use the same fully qualified domain name to access RING services.
- **Under normal operation**: clients are directed to their local site's RING cluster.
- **In the event of a local service failure**: clients are automatically directed to a functioning RING cluster at another site. This would happen if the local site's RING service and/or load balancers were offline and unavailable.

For the sake of simplicity, the diagram presented below shows a two site setup. The principle can be extended to encompass as many sites as desired.

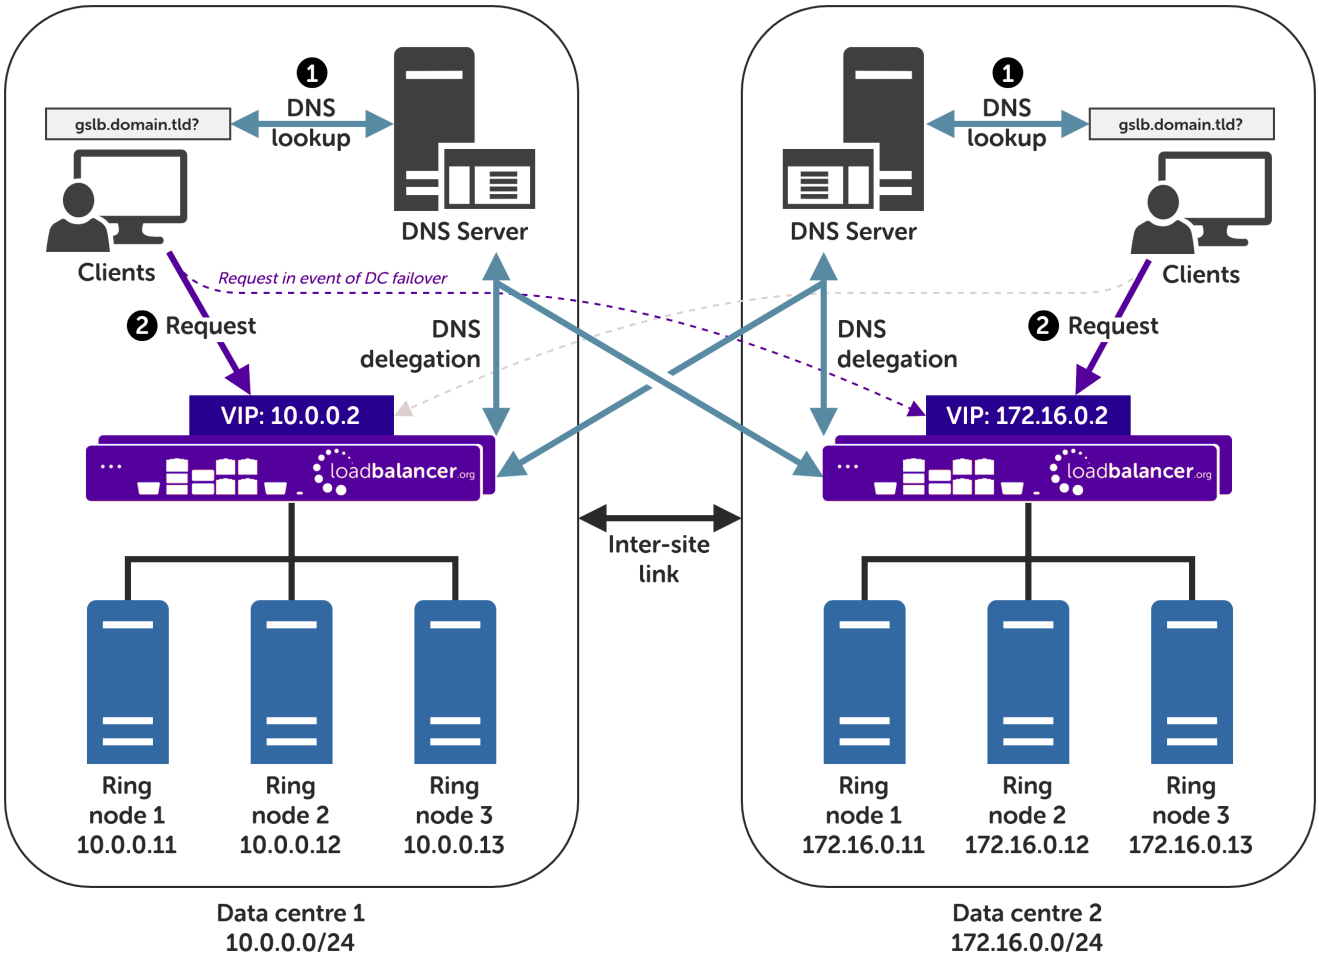

Explanation:

- **Start**: A client tries to access the RING service using the S3 protocol. To do this, the client uses the service's fully qualified domain name, in this example `gslb.domain.tld`

- The client sends a DNS query for `gslb.domain.tld` to its local DNS server.
- The DNS server has the domain `gslb.domain.tld` delegated to the load balancers.
- The DNS server sends a delegated DNS query for `gslb.domain.tld` to one of the load balancers.
- The load balancer that received the delegated DNS query replies to the DNS server. The load balancer answers with the IP address of the VIP (RING instance) that is **local to the DNS server making the query**, and hence local to the original client.
    - An example: if the delegated query from the DNS server originated from the 10.0.0.0/24 subnet then the VIP in that subnet is served up. Likewise, if the delegated query originated from the 172.16.0.0/24 subnet then the VIP in that subnet is served up. As such, clients are always directed to their local, on-site RING instance, provided that the local instance is online and available.
- The DNS server sends the delegated DNS answer to the client.
- **Finish**: The client connects to the S3 service at `gslb.domain.tld` by using the local VIP address.

> **Note**
>
> **In the event that the cluster of RING cluster and/or load balancers at one site should completely fail** then local clients will be directed to the RING cluster at the other site and the service will continue to be available.
>
> This style of multi-site failover is possible because the load balancers' GSLB functionality continuously health checks the service at each site. When the service at a site is observed to be unavailable then that site's IP address is no longer served when responding to DNS queries.

DNS Server Prerequisites

> **Important**
>
> Location affinity (ensuring clients 'stick' to their local site) **requires** a unique DNS server at each site.

For this setup to work and provide location affinity, a unique DNS server is required at each site, like the example deployment shown at the beginning of this section.

If multiple sites ***share*** a common DNS server then ***clients cannot be directed to their local, on-site RING instance***.

Example: Consider a two data centre deployment with a shared, common DNS server located at DC 1. From the perspective of a load balancer in this scenario, ***every*** delegated DNS request would be seen to come from the single, shared DNS server at DC 1. Specifically, the requests would all come from the DNS server's IP address, which would fall within DC 1's subnet.

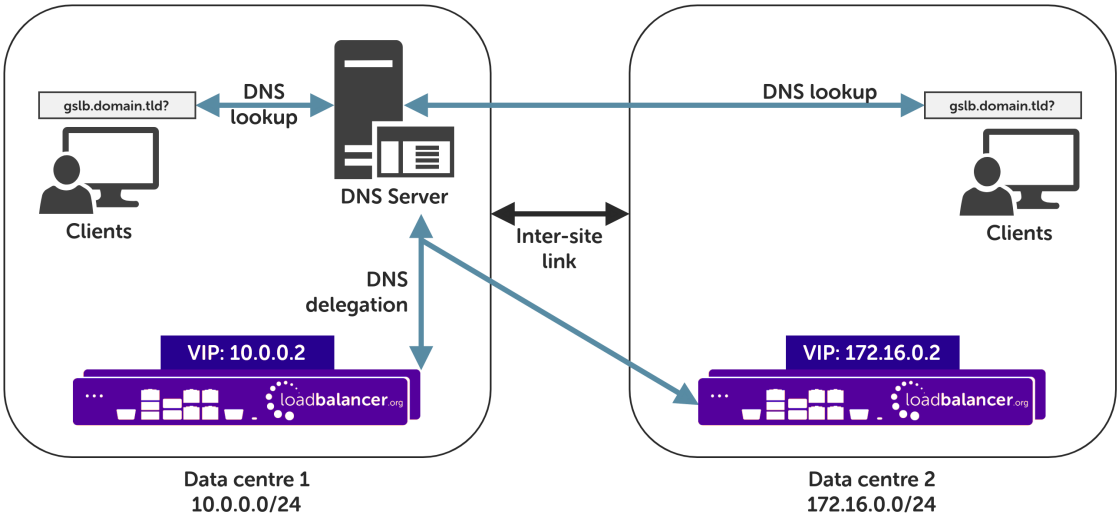

A load balancer would have *no way to distinguish between delegated requests for DC 1's clients and delegated requests for DC 2's clients*. All delegated requests would originate from within DC 1's subnet, therefore **all traffic would be directed to DC 1's RING instance**.

To resolve such a situation, a DNS server would need to be deployed at DC 2. The load balancers could then easily tell which site a given delegated DNS query has come from and, therefore, which site the client should be directed to.

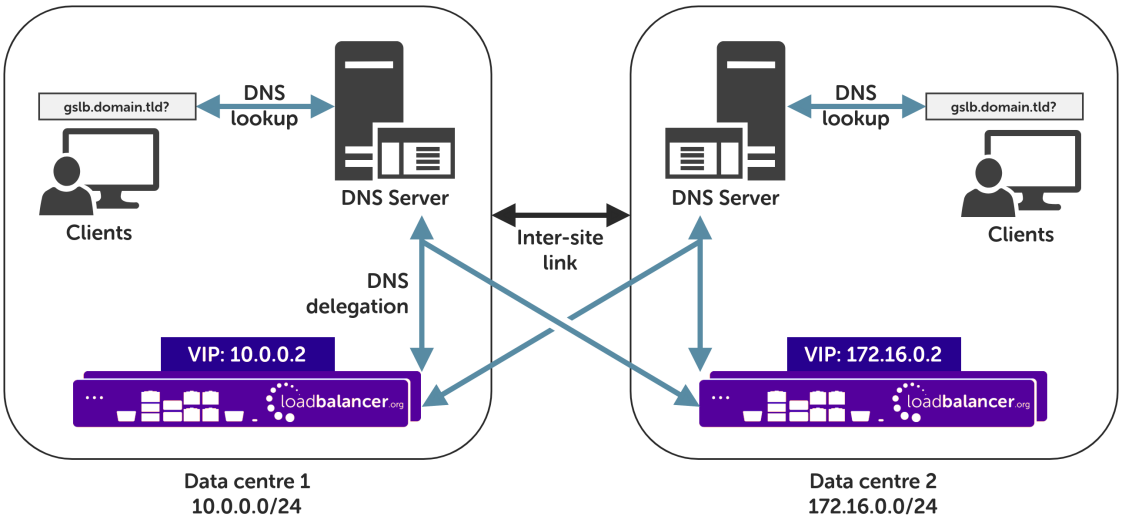

If having unique DNS servers per-site and splitting up sites using a topology configuration is *not* possible then clients **will** bounce between different VIPs (and hence bounce between sites) in a round-robin fashion. If this behaviour is acceptable then it can theoretically be used without significant issue.

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:

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

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

Appliance Configuration

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

Note that **the GSLB configuration must be identical across all sites**: inconsistent configurations will lead to unexpected behaviour.

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

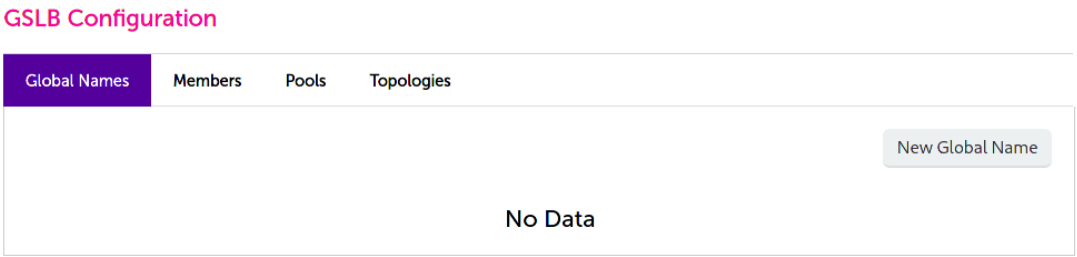

Step 1 – Configuring the Global Name

1. Using the WebUI on the primary appliance for the first site, 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**.

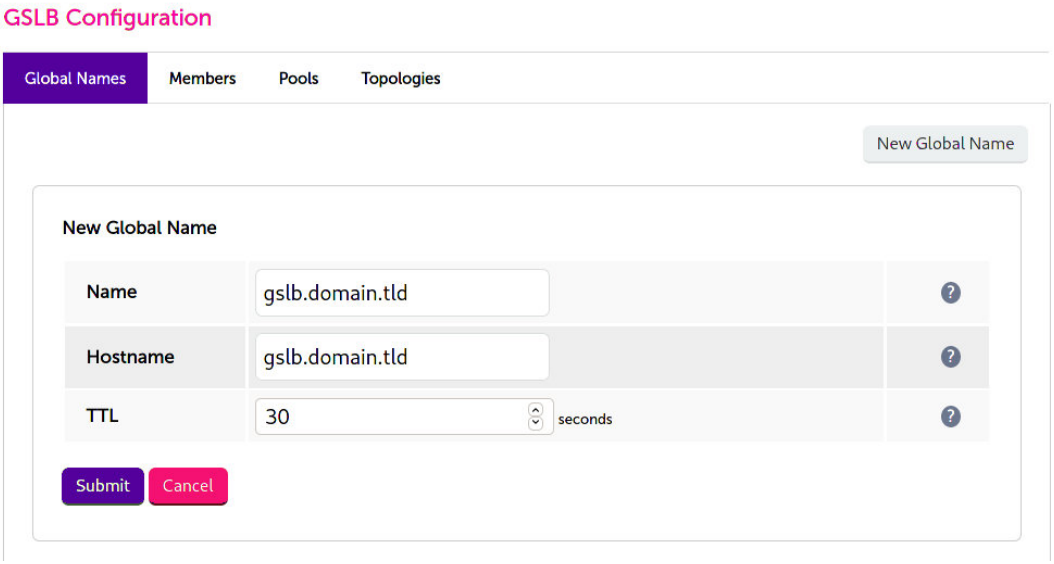

Step 2 – Configure the Members

Each *member* can be thought of as a single site.

1. Select the **Members** tab.
2. Click the **New Member** button.
3. Enter a friendly *Name* for the member, e.g. **DC1**.
4. Specify an *IP* address for the member: in this context, this should be the VIP address of the site's RING instance, e.g. **10.0.0.2**.
5. Ignore the example value in the *Monitor IP* field.
6. Click **Submit**.
7. Repeat these steps to add additional sites as members as required.

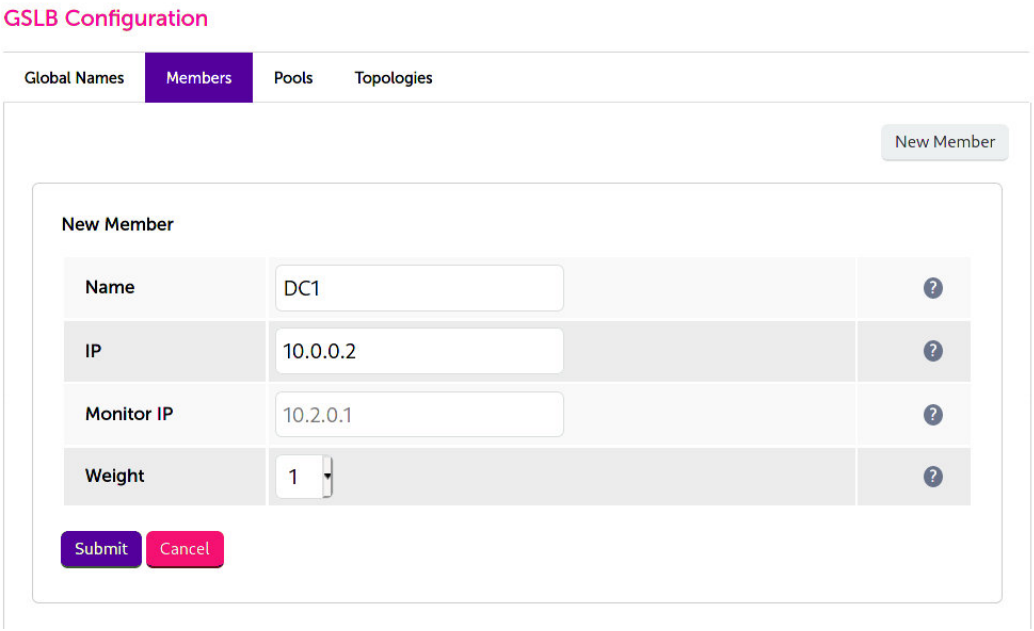

Step 3 – Configure the Pool

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, both sites have a functional RING instance ready for use. A pool would therefore be created linking the global name `gslb.domain.tld` with members (sites) DC1 and DC2, both of which should serve RING traffic.

1. Select the **Pools** tab.
2. Click the **New Pool** button.
3. Enter a friendly *Name* for the pool, e.g. **ring-sites**.
4. Set the *Monitor* to **HTTP**.
5. Set *Monitor Use SSL* to **Yes**.
6. Set *Monitor Hostname* to a hostname that should respond if the RING instance is online and healthy, e.g. **s3-region1.domain.tld**
7. Set *Monitor URL Path* to **/_/healthcheck/deep/**
8. Set *Monitor Port* to **443**.
9. Set *Monitor Expected Codes* to **200**.
10. Set *LB Method* to **twrr**.
11. From the *Global Names* list box, select the global name in question, e.g. **gslb.domain.tld**
12. In the *Members* section, drag the appropriate members (sites) from the *Available Members* box into the *Members In Use* box.
13. Click **Submit**.

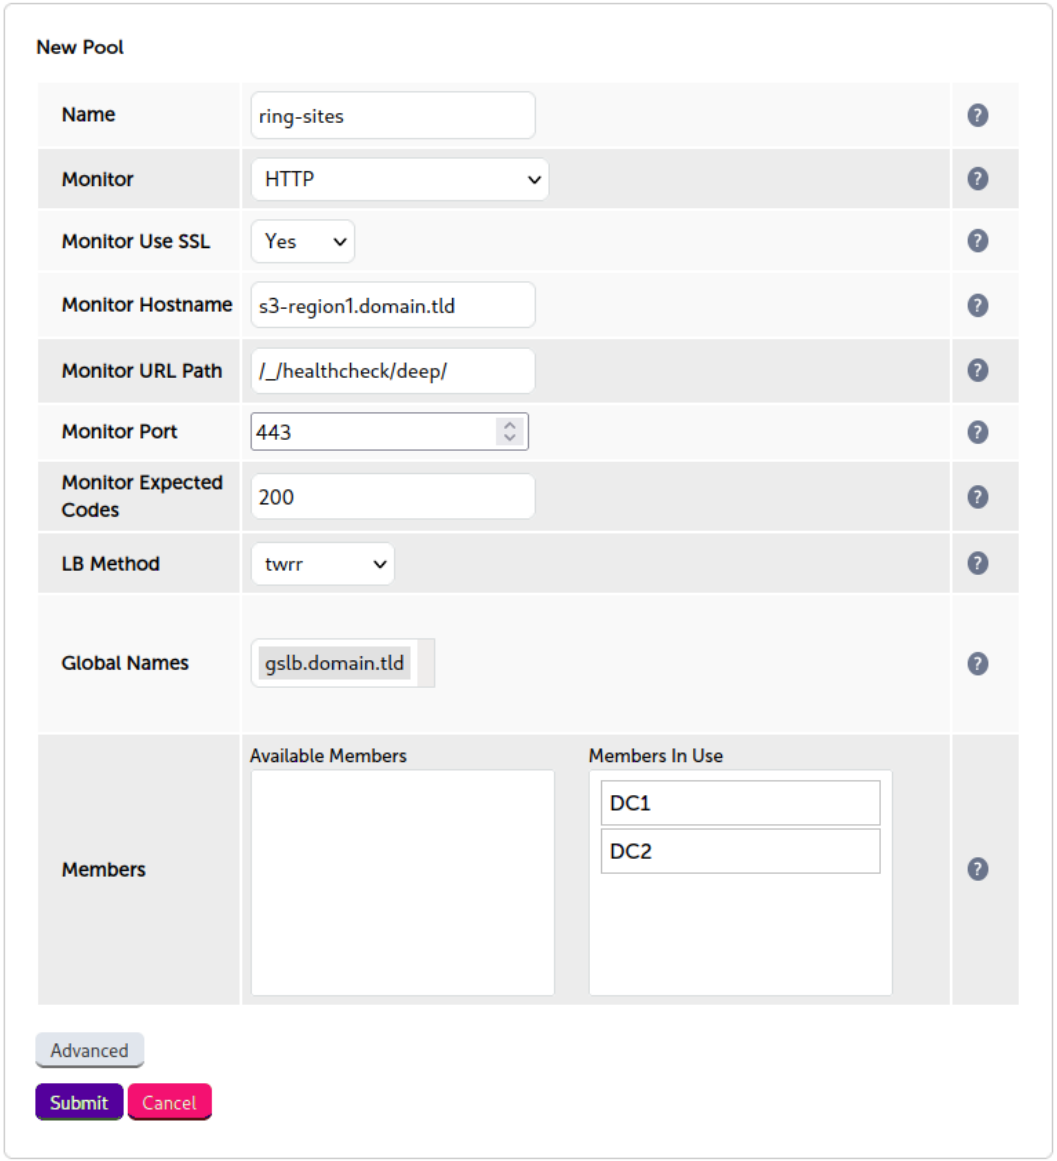

Step 4 – Configure the Topology

Topology configuration is used to map subnets to sites. This gives the solution its location awareness, allowing clients to be directed to their ***local*** RING instance instead of being bounced between every site which has been defined.

1. Select the **Topologies** tab.
2. Click the **New Topology** button.
3. Enter a friendly ***Name*** for the topology, e.g. **DC1**.
4. In the *IP/CIDR* text box, define the subnet(s) that covers the site in question, e.g. **10.0.0.0/24**.

   This can be a comma separated list of subnets and hosts, e.g. `10.0.0.0/24, 192.168.2.0/24, 192.168.17.57`. The key is that the site's DNS server ***and*** its RING VIP fall within the union of all subnets and hosts defined for the site. This is what allows DNS queries originating from the site to be matched up with that site's local VIP: the local VIP is then served as a DNS response for clients at that site.

5. Click **Submit**.
6. Repeat these steps to add additional topology configurations as required.

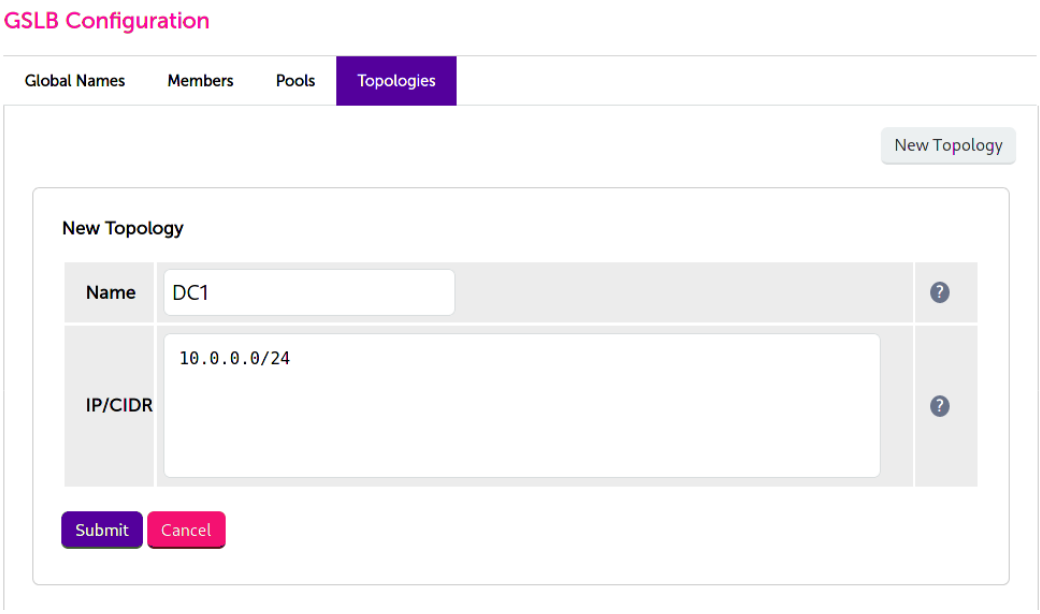

Step 5 – Finalising the Configuration

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**.

Optional: Defining a Default Site for External Traffic (Handling DNS Requests from Unpredictable Source Addresses)

It is plausible that a RING GSLB deployment may be required to answer DNS queries sourced from outside of the subnets defined in the topology configuration.

Consider a client on the public internet requesting a resource from the RING instance. The DNS query associated with the request may be sourced from a previously unseen, unpredictable public IP address. DNS queries from IP addresses that do not fall within the predefined network topology/subnets will be answered with DNS records pointing to ***any*** of the defined sites in a round-robin fashion.

An alternative is to define a ***default site***. All DNS queries from outside the predefined network topology will be answered with ***the same*** DNS record: a record pointing to the default site.

To configure this, add the widest possible subnet of 0.0.0.0/0 to the topology configuration of the site which is to be the 'default'. Any DNS query whose source IP address does not fall within one of the other, smaller subnets will be picked up by this new "catch all" subnet.

Following on from the previous example, setting data centre 1 to be the 'default' site would look like so:

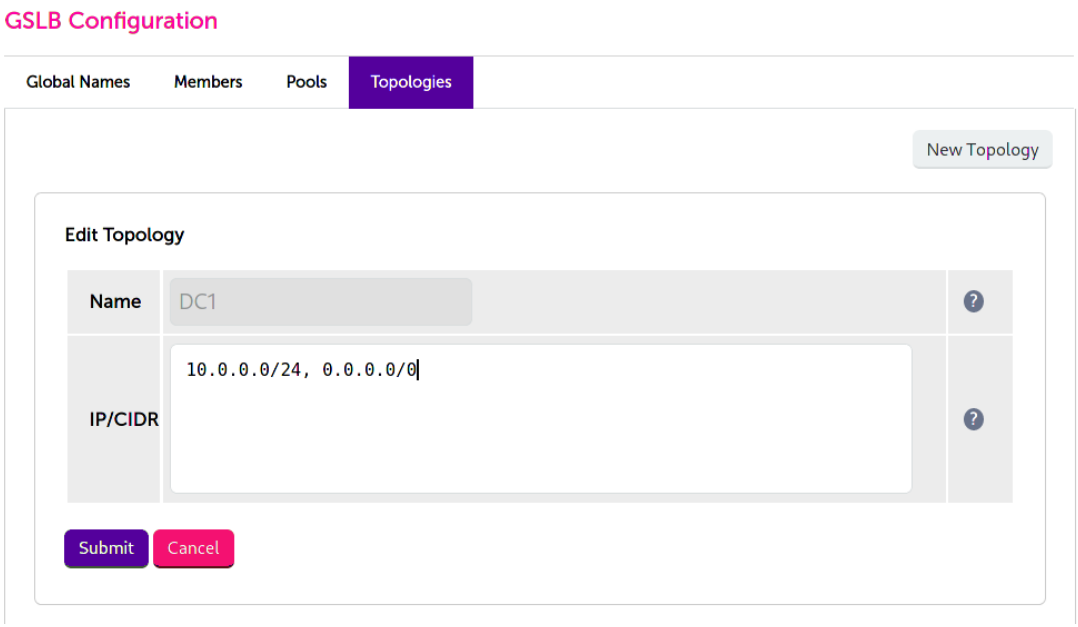

DNS Server Configuration

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.

15.2. 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.

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`).

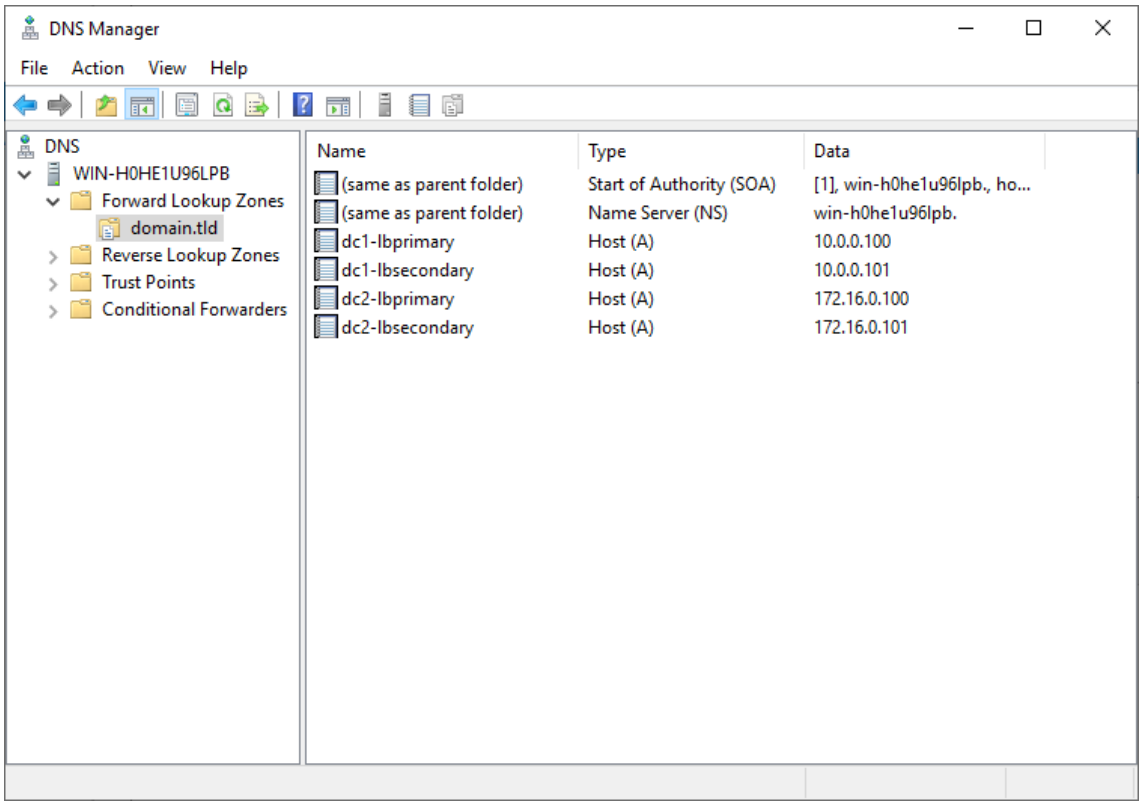

2. Provided that the load balancer part of the GSLB configuration has been completed and is working, the **New Delegation** wizard should now be used to delegate the subdomain to the load balancers. The delegation will use the new FQDNs for the load balancers, as defined in the previous step. The delegation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

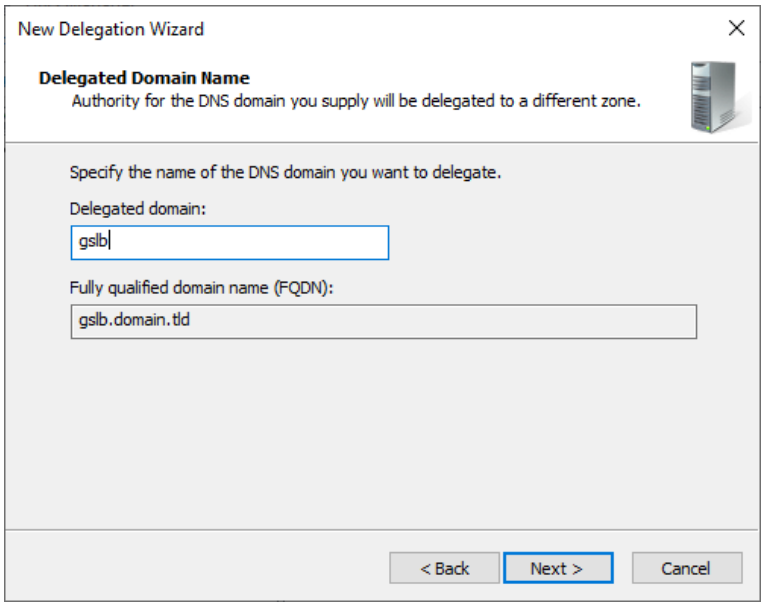

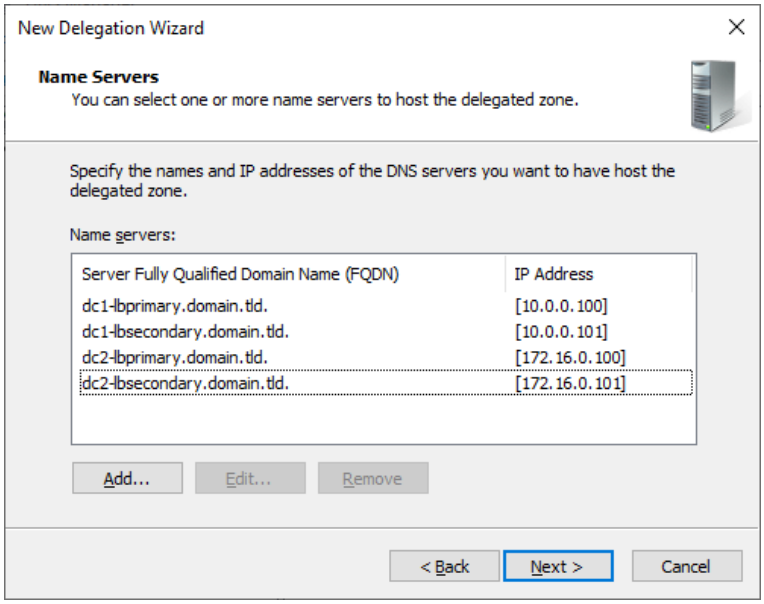

3. Test the delegation to make sure it is working as expected.

   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
```

15.3. Alternative Load Balancing Method for Read-Intensive Deployments (Direct Routing)

Direct routing, also known as direct server return or DSR, is a method of load balancing. With direct routing, reply traffic flows directly from the back end servers to the clients. In this way, the load balancer is completely bypassed on the return journey for a given connection, thus removing the load balancer as a potential bottleneck for traffic on the return path.

This alternative method of load balancing can benefit read-intensive deployments which feature a large reply traffic to request traffic ratio. For example, consider the scenario where a typical client request is 10 kB in size while a typical reply is 10 GB in size (perhaps file retrieval or video streaming). Direct routing benefits such scenarios: the much larger volume of reply traffic bypasses the load balancer and is *not* limited by the load balancer's network throughput. The reply traffic is instead limited by the total available network bandwidth between the servers and the clients, which is limited only by the underlying infrastructure.

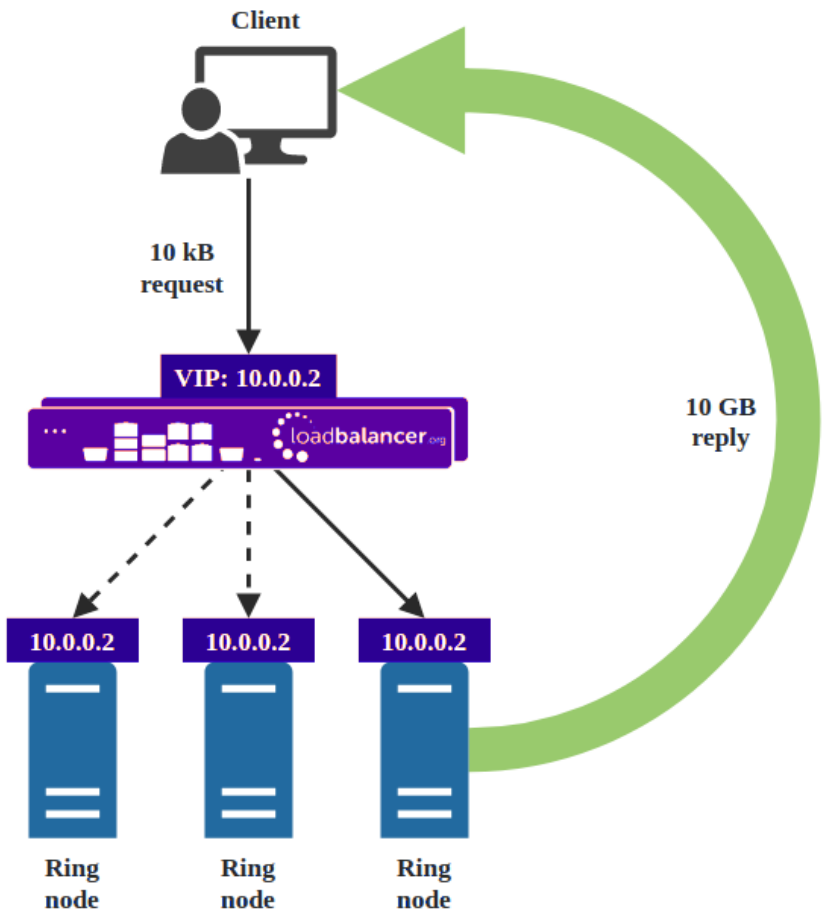

Caveats

There are caveats for using the direct routing load balancing method which should be considered:

- The load balancers must be on the same network segment / switching fabric as the RING nodes (due to the fact that this load balancing method works by rewriting MAC addresses, i.e. operates at layer 2 of the OSI model)
- Each RING node must own the VIP address so that they can all accept and reply to the load balanced traffic. This address should be assigned to a loopback network adaptor
- Each RING node must be configured to not reply to ARP requests for the VIP address or advertise that they own the address

For guidance on configuring the RING nodes for direct routing, in the context of the caveats described above, please consult with Scality Sales Engineering or Support.

Appliance Configuration for Scality RING – Using Layer 4 DR Mode (Direct Routing)

Configuring VIP 1 – S3

Configuring the Virtual Service (VIP)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Virtual Services* and click on **Add a new Virtual Service**.
2. Define the *Label* for the virtual service as required, e.g. **S3**.
3. Set the *Virtual Service IP Address* field to the required IP address, e.g. **192.168.87.67**.
4. Set the *Ports* field to **80**.
5. Leave the *Protocol* set to **TCP**.
6. Leave the *Forwarding Method* set to **Direct Routing**.
7. Click **Update** to create the virtual service.
8. Click **Modify** next to the newly created VIP.
9. Ensure that the *Persistence Enable* checkbox is unchecked.
10. Set the *Health Checks Check Type* to **Negotiate**.
11. Set the *Check Port* to **80**.
12. Set the Protocol to **HTTP**.
13. Set the Request to send to **/_/healthcheck/deep/**.
14. Click **Update**.

Defining the Real Servers (RIPs)

1. Using the web user interface, navigate to *Cluster Configuration > Layer 4 – Real Servers* and click on **Add a new Real Server** next to the newly created VIP.
2. Define the *Label* for the real server as required, e.g. **ring-node1**.
3. Set the *Real Server IP Address* field to the required IP address, e.g. **192.168.87.88**.
4. Click **Update**.

5. Repeat these steps to add additional RING nodes as real servers as required.

15.4. 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 configured first and then the Secondary should be added. Once the Primary and Secondary are paired, all load balanced services configured on the Primary are automatically replicated to the Secondary over the network using SSH/SCP.

> **Note** For Enterprise Azure, the HA pair should be configured first. In Azure, when creating a VIP using an HA pair, 2 private IPs must be specified – one for the VIP when it's active on the Primary and one for the VIP when it's active on the Secondary. Configuring the HA pair first, enables both IPs to be specified when the VIP is created.

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.

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	All network settings including IP address(es), bonding configuration and VLANs
Local Configuration	Routing	Routing configuration including default gateways and static routes
Local Configuration	System Date & time	All time and date related settings
Local Configuration	Physical – Advanced Configuration	Various settings including Internet Proxy, Management Gateway, Firewall connection tracking table size, NIC offloading, SMTP relay, logging and Syslog Server
Local Configuration	Security	Appliance security settings
Local Configuration	SNMP Configuration	Appliance SNMP settings
Local Configuration	Graphing	Appliance graphing settings
Local Configuration	License Key	Appliance licensing
Maintenance	Software Updates	Appliance software update management
Maintenance	Firewall Script	Appliance firewall (iptables) configuration
Maintenance	Firewall Lockdown Wizard	Appliance management lockdown settings

> **Important** Make sure that if these settings/updates have been configured on the Primary appliance, they're also configured on the Secondary appliance.

Adding a Secondary Appliance - Create an HA Clustered Pair

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

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

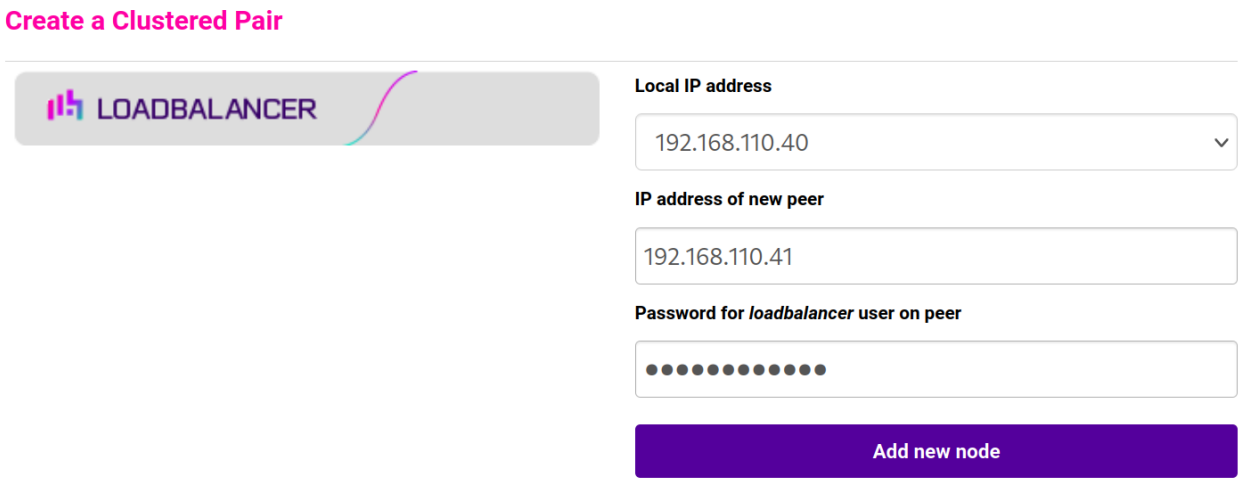

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

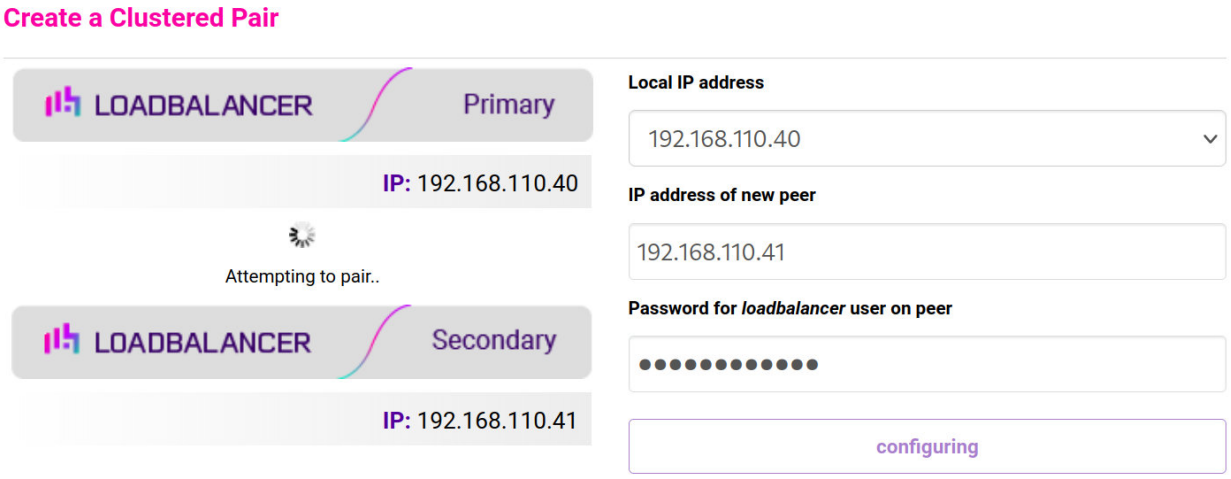

6. Once complete, the following will be displayed on the Primary appliance:

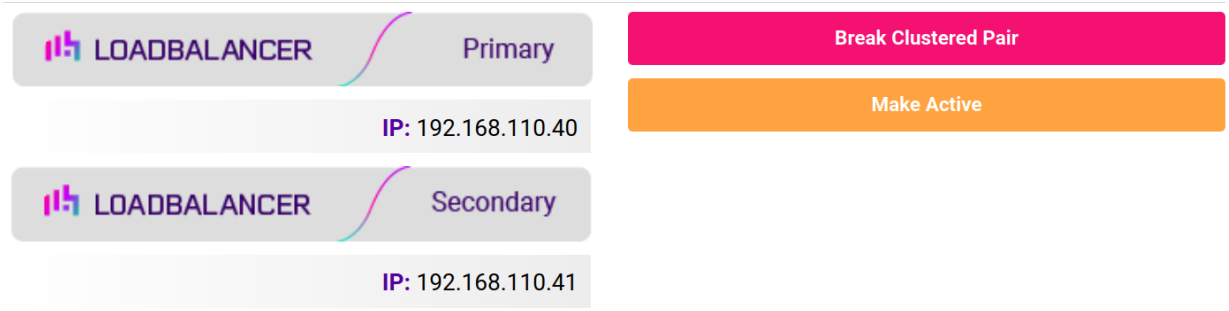

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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.

16. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	14 February 2020	Initial version		IBG
1.0.1	3 September 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.1.0	1 October 2021	Converted the document to AsciiDoc	Move to new documentation system	AH,RJC,ZAC
1.1.1	21 March 2022	Added new multithreading advice	Product change means multithreading is now enabled by default	AH
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	26 April 2022	Updated SSL related content to reflect latest software version	New software release	RJC
1.2.2	28 September 2022	Updated layer 7 VIP and RIP creation screenshots	Reflect changes in the web user interface	AH
1.2.3	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 Removed references to the colour of certain UI elements	Housekeeping across all documentation	AH

Version	Date	Change	Reason for Change	Changed By
1.2.4	2 February 2023	Updated screenshots	Branding update	AH
1.2.5	7 March 2023	Removed conclusion section	Updates across all documentation	AH
1.3.0	24 March 2023	New document theme Modified diagram colours	Branding update	AH
1.3.1	29 June 2023	Updated multithreading advice	New default option in the web user interface	AH

LOADBALANCER

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

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.