LOADBALANCER

DEPLOYMENT GUIDE

Load Balancing NetApp StorageGRID

Version 1.3.0

Table of Contents

1. About this Guide

This guide details the steps required to configure a load balanced NetApp StorageGRID environment utilizing Loadbalancer.org appliances. It covers the configuration of the load balancers and also any NetApp StorageGRID 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 StorageGRID. For full specifications of available models please refer to: https://www.loadbalancer.org/products/enterprise.

Some features may not be available or fully supported in all cloud platforms due to platform specific limitations. For more details, please refer to the "Main Differences to our Standard (Non-Cloud) Product" section in the appropriate cloud platform Quick Start Guide or check with Loadbalancer.org support.

3. Software Versions Supported

3.1. Loadbalancer.org Appliance

- V8.9.1 and later

> **Note** The screenshots used throughout this document aim to track the latest Loadbalancer.org software version. If you're using an older version, or the very latest, the screenshots presented here may not match your WebUI exactly.

3.2. StorageGRID

- v11.3 and later

4. NetApp StorageGRID

NetApp StorageGRID is a software-defined, object-based storage solution that supports industry-standard object APIs such as Amazon S3 and Swift. It allows you to build a single name space across many sites, with multiple service levels for metadata-driven object life-cycle policies.

StorageGRID protects data via intelligent policy with options including replica, erasure coding and cloud tier.

StorageGRID can be deployed as optimized hardware appliances, virtual machines, Docker containers or a combination of all three.

4.1. StorageGRID Node Types

- **Admin Node** – Admin Nodes provide system administration services such as system configuration, monitoring, and logging. Each StorageGRID system includes one primary Admin Node. The primary Admin

Node hosts the Configuration Management Node (CMN) service which manages system-wide configurations and grid task. For redundancy, a StorageGRID system can have additional, Non-primary Admin Nodes.

- **Storage Node –** Storage Nodes manage the storage of objects to disk. This object management (both object data and object metadata) includes the evaluation of objects against ILM rules to determine how an object's data is stored over time and protected from loss.
- **Gateway Node (Optional) –** Gateway Nodes provide a load balancing interface to the StorageGRID system through which applications can connect. The Gateway Nodes host the Connection Load Balancer (CLB) service which acts as a switchboard for connecting clients to the most efficient Local Distribution Router (LDR) service.

> **Note**
>
> StorageGRID v11.3 introduced a new Load Balancer service which is included on Gateway Nodes and on all Admin Nodes. This service provides Layer 7 load balancing of S3 and Swift traffic from clients to Storage Nodes. The legacy Connection Load Balancer (CLB) service on Gateway Nodes is still supported; however, configuring endpoints for the new Load Balancer service is recommended.

5. Load Balancing NetApp StorageGRID

Load balancing StorageGRID maximizes speed and connection capacity by distributing the connections and workloads across multiple Storage Nodes.

5.1. Within a Single Site

Load balancing StorageGRID can be achieved in the following ways:

1. Using the native in-built Load Balancer service, which is installed on Admin Nodes and Gateway Nodes. The Load Balancer service provides Layer 7 load balancing and performs TLS termination of client requests, inspects the requests, and establishes new secure connections to the Storage Nodes. This is the recommended load balancing mechanism.
2. Using a third-party load balancer.

5.1.1. Using the Loadbalancer.org Appliance within a Single Site

For load balancing option 2 above, the Loadbalancer.org appliance is an ideal solution. The following diagram shows how the load balancer is deployed within a single site. The load balancer can be configured to load balance both the Storage Nodes and the Admin Nodes as shown in the following diagram.

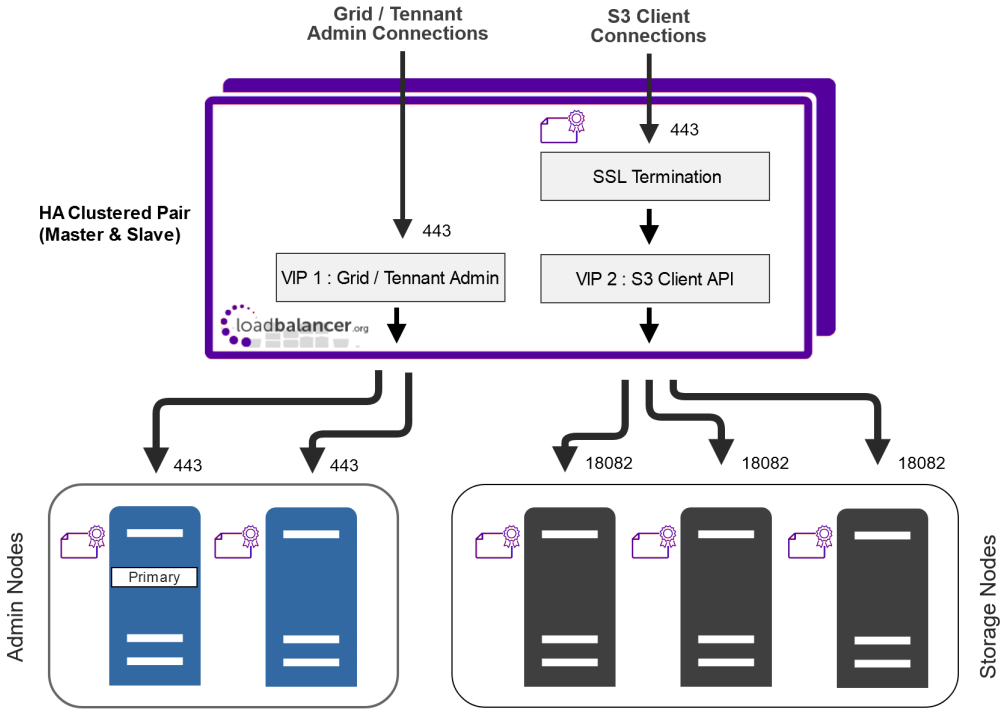

Configuration Notes

- SSL bridging is configured for S3 client connections, this requires an SSL certificate to be installed on the Loadbalancer.org appliance. For more information please refer to the sections Upload SSL Certificate & Configure SSL Termination.
- SSL pass-through is configured for HTTPS Grid / Tenant admin connections.

> **Note** If you're using the Swift API rather than the S3 API, VIP 2 would be modified to connect to the Storage Nodes on port 18083 rather than port 18082.

Load Balancing & HA Requirements

The function of the load balancer is to ensure that inbound connections to a NetApp StorageGRID cluster are distributed across healthy StorageGRID nodes. To provide HA for the load balancer, Loadbalancer.org recommends that 2 appliances are deployed as an HA clustered pair.

Virtual Service (VIP) Requirements

To load balance the client connections across the Storage Nodes and load balance the Grid / Tenant admin connections across the Admin Nodes, 2 Virtual Services (VIPs) are required:

- **VIP 1** – Grid / Tenant admin connections
- **VIP 2** – S3 Client connections

SSL Termination & Re-encryption

The S3 Inbound client connections are terminated on the load balancer. Connections from the load balancer to the load balanced Storage Nodes are re-encrypted (SSL bridging).

VIP Operating Mode

The VIPs are configured using Layer 7 SNAT mode. This mode offers high performance and requires no mode-specific configuration changes to the load balanced NetApp StorageGRID Nodes.

Timeouts

For NetApp StorageGRID, the load balancer's client and server timeouts are set to 10 minutes.

Health Checks

For the Admin Nodes, a HTTP GET method is used to perform the health check. The load balancer issues HTTP GET requests to each individual Admin Node and expects a 200 OK response.

For the Storage Nodes, NetApp recommends that the HTTP OPTIONS method is used to perform the health check. The load balancer issues HTTP OPTIONS requests to each individual storage node and expects a 200 OK response.

5.2. Across Multiple Sites

Where StorageGRID is deployed across multiple sites, Global Sever Load Balancing (GSLB) is used. All Loadbalancer.org appliances have this functionality built-in by default at no extra cost.

The GSLB functionality coupled with DNS delegation enables each load balancer to act as a smart DNS name server for the sub domains (in this guide **admin.company.com** & **s3.company.com**).

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

5.3. Load Balancing Scenarios

Since there are 2 ways to load balance the Storage Nodes within a single site (excluding the legacy CLB service), there are effectively 2 possible load balancing scenarios for a particular multi-site deployment as shown in the following table:

Scenario	Local Server Load Balancing Method	Inter-site Load Balancing Method
1	StorageGRID Load Balancer Endpoints configured on Admin Nodes & Gateway Nodes	Loadbalancer.org Appliance (GSLB functionality)
2	Loadbalancer.org Appliance (server load balancing Functionality)	Loadbalancer.org Appliance (GSLB functionality)

Both scenarios provide viable options when setting up load balancing for a multi-site StorageGRID deployment.

As explained here, when configuring Load Balancer Endpoints in StorageGRID for local site load balancing, it's important to remember that End Points configured on each Admin / Gateway node are independent from one another.

5.3.1. Scenario 1

GSLB: Handled by Loadbalancer.org Appliances

SLB: Handled by NetApp Admin & Gateway Nodes

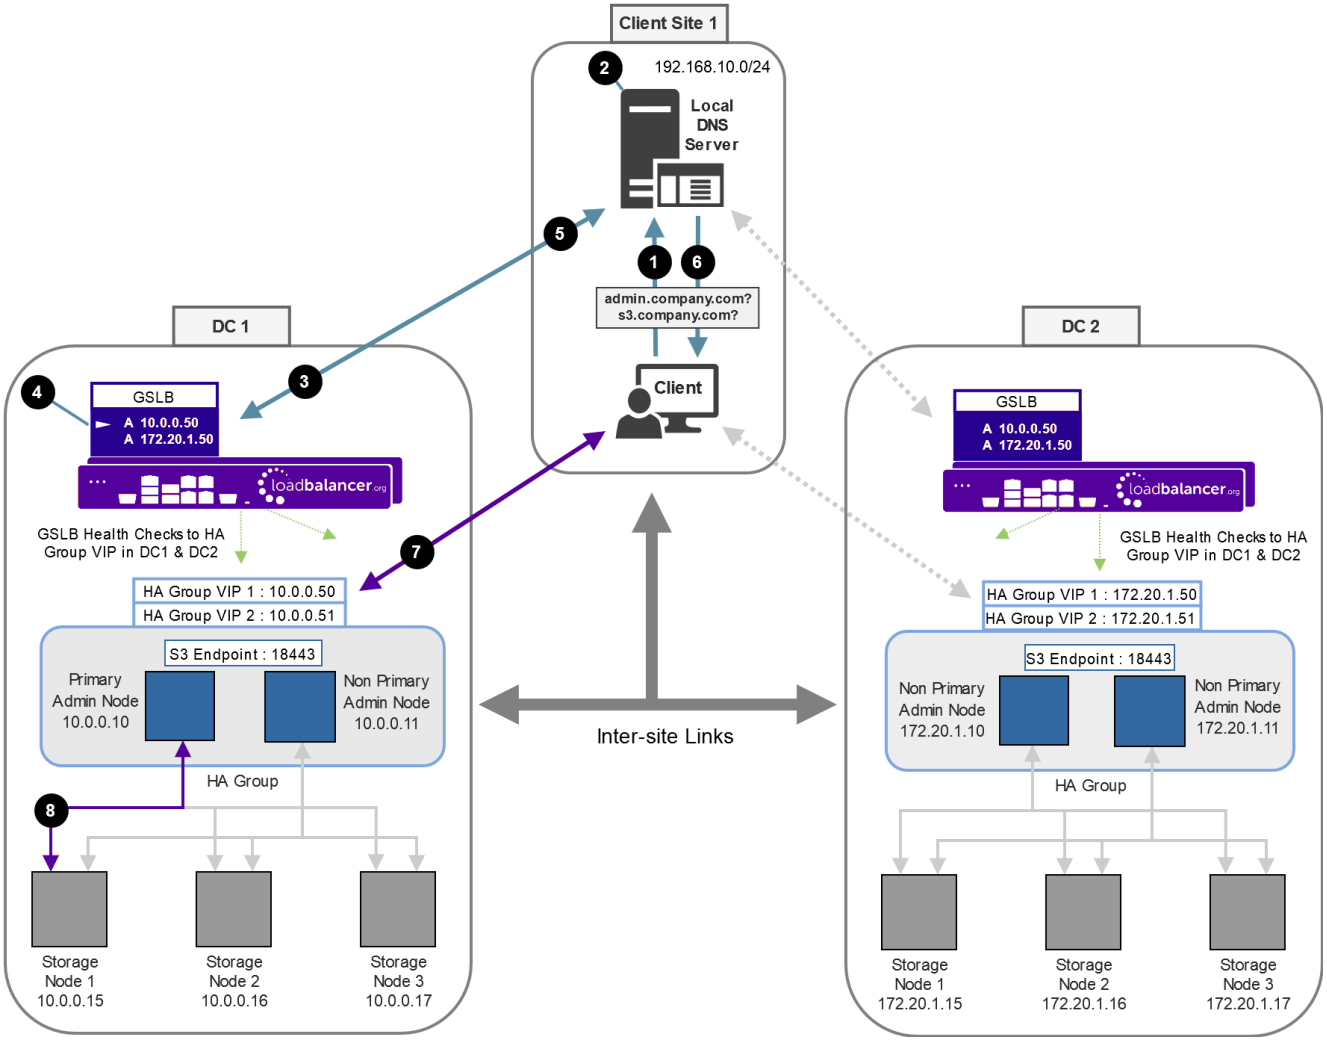

Explanation:

1. The client sends a DNS query for either **admin.company.com** or **s3.company.com** to the local DNS server.
2. The local DNS server has the sub domain delegated to all LB.org appliances (the LB.org appliances are configured as name servers for the sub domains).
3. One of the LB.org appliances receives the delegated DNS query.
4. If the query is for **s3.company.com** the LB.org appliance selects HA Group VIP 2 (10.0.0.51) in DC 1 based on the GSLB topology configuration and GSLB health checks.
5. The LB.org appliance returns the IP address to the DNS server.
6. The DNS server returns the IP address to the client.
7. The client connects to 10.0.0.51 on port 18443 for the S3 client connection.
8. The active Admin Node (10.0.0.10) then load balances the connection to Storage Node 1 (10.0.0.15) based on the StorageGRID load balancing algorithm.

5.3.2. Scenario 2

GSLB: Handled by Loadbalancer.org Appliances

SLB: Handled by Loadbalancer.org Appliances

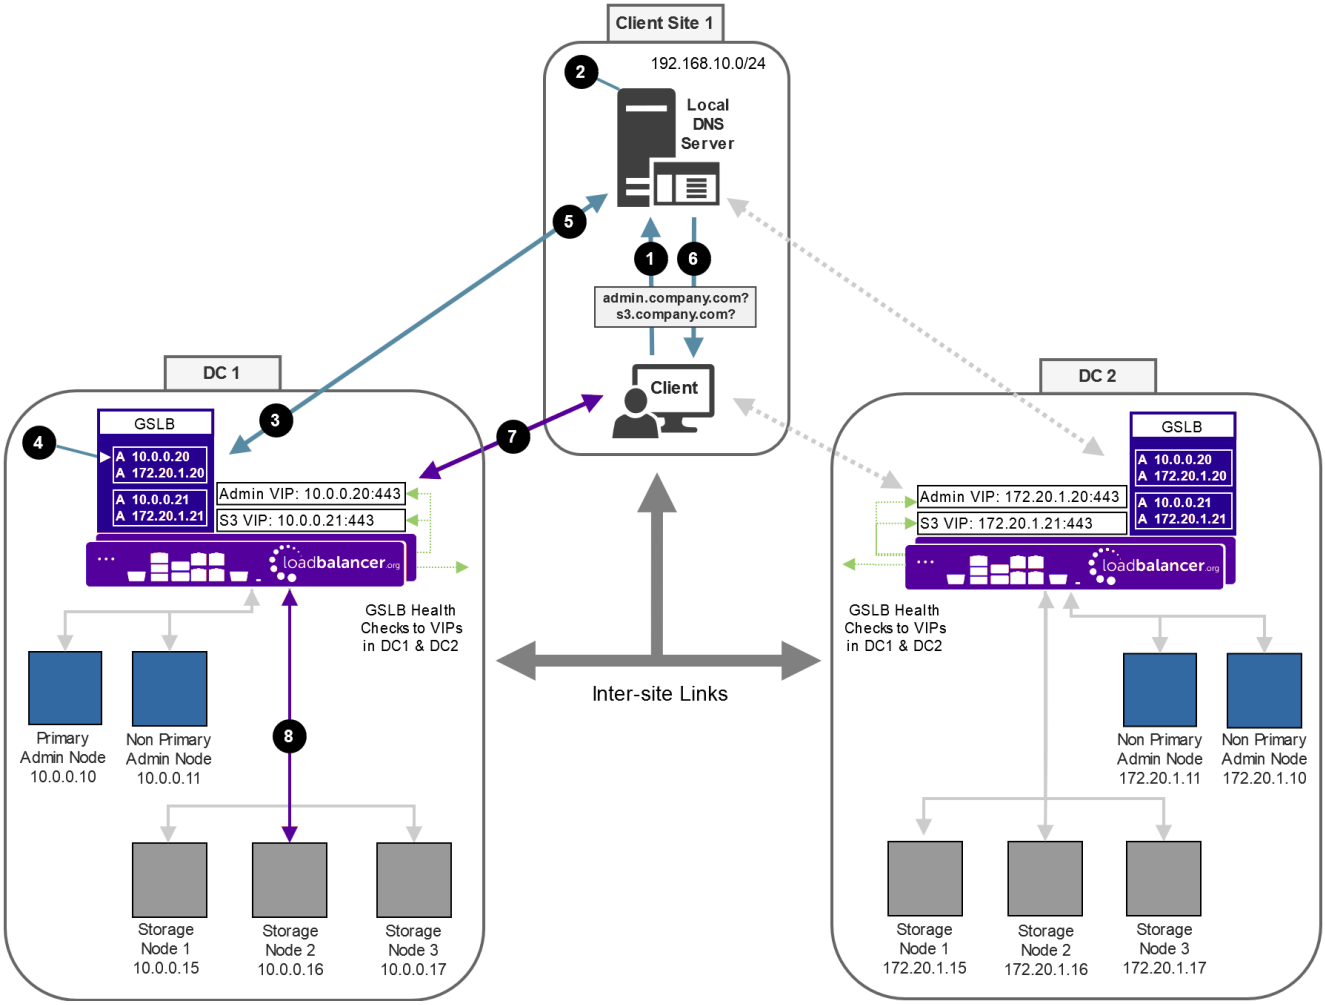

Explanation:

1. The client sends a DNS query for either **admin.company.com** or **s3.company.com** to the local DNS server.
2. The local DNS server has the sub domains delegated to all LB.org appliances (the LB.org appliances are configured as name servers for the sub domain).
3. One of the LB.org appliance receives the delegated DNS query.
4. If the query is for **s3.company.com** the LB.org appliance selects the S3 VIP (10.0.0.21) in DC 1 based on the GSLB topology configuration and GSLB health checks.
5. The LB.org appliance returns the IP address to the DNS server.
6. The DNS server returns the IP address to the client.
7. The client connects to 10.0.0.21 on port 443 for the S3 client connection.
8. The LB.org appliance then load balances the connection to Storage Node 2 (10.0.0.16) based on the LB.org appliance's load balancing algorithm.

6. Loadbalancer.org Appliance – the Basics

6.1. 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.

6.2. 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.

6.3. Accessing the Appliance WebUI

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

> **Note** There are certain differences when accessing the WebUI for the cloud appliances. For details, please refer to the relevant Quick Start / Configuration Guide.

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

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

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

   > **Note** If you need to change the port, IP address or protocol that the WebUI listens on, please refer to Service Socket Addresses.

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

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

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

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

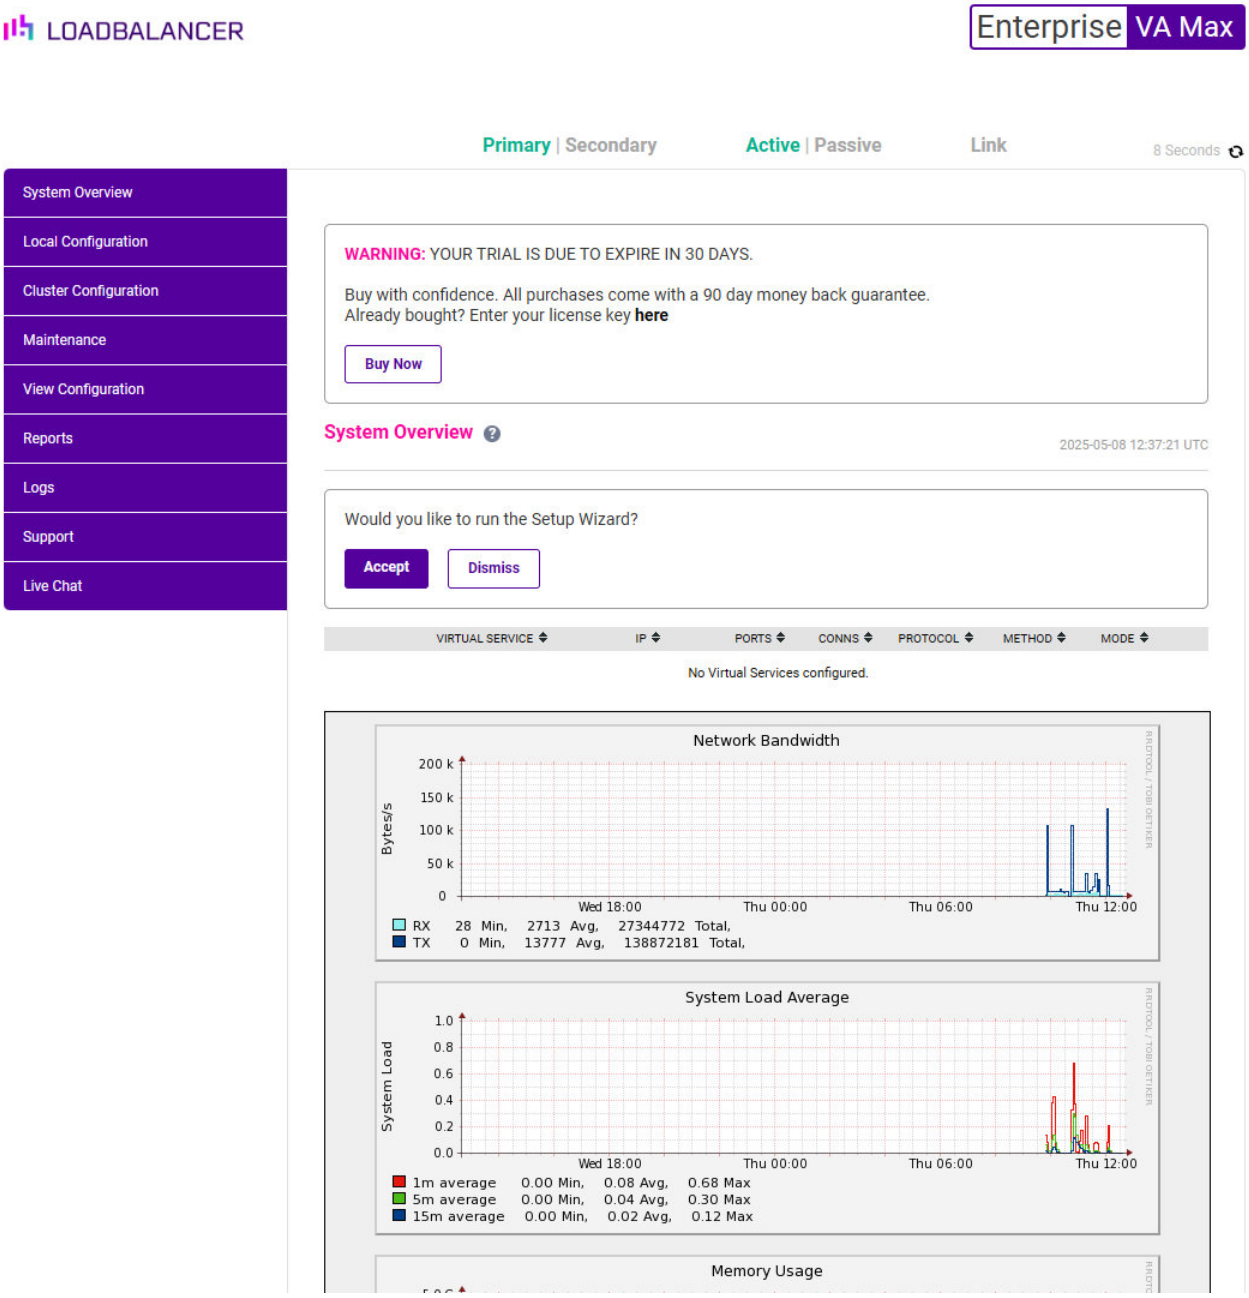

3. You'll be asked if you want to run the Setup Wizard. Click **Dismiss** if you're following a guide or want to configure the appliance manually. Click **Accept** to start the Setup Wizard.

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

6.3.1. 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

6.4. 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.

6.4.1. 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:

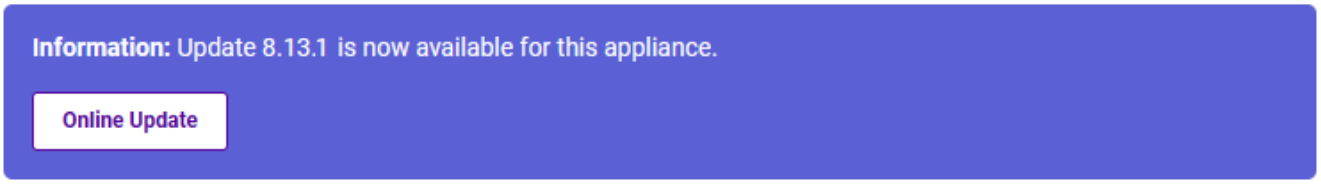

Click **Online Update**. A summary of all new features, improvements, bug fixes and security updates included in the update will be displayed. Click **Update** at the bottom of the page to start the update process.

> **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:

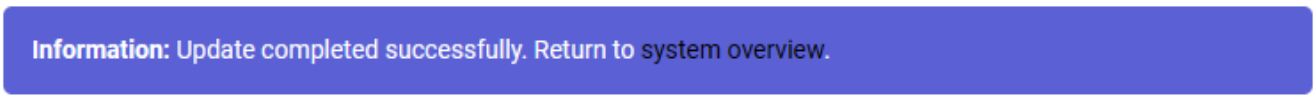

If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

6.4.2. 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:

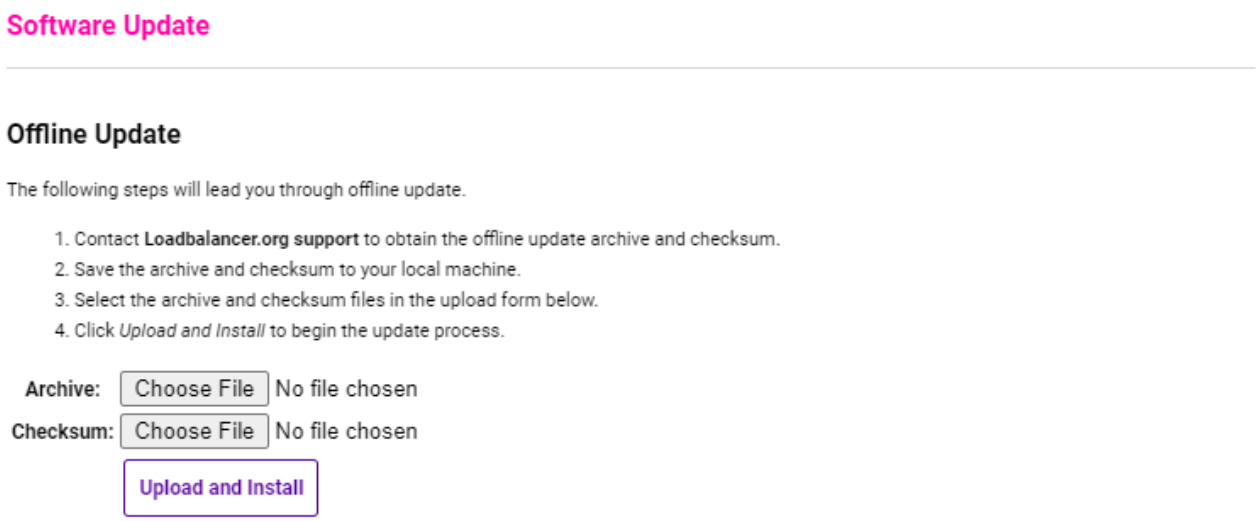

4. Select the *Archive* and *Checksum* files.
5. Click **Upload and Install**.
6. If services need to be reloaded/restarted or the appliance needs a full restart, you'll be prompted accordingly.

6.5. 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
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)

Protocol	Port	Purpose
TCP	9080 *	WebUI - HTTP (disabled by default)
TCP	9081 *	Nginx fallback page
TCP	9443 *	WebUI - HTTPS
TCP	25565 *	Shuttle service (Centralized/Portal Management)

> **Note** The ports used for SSH, GSLB, SNMP, the WebUI, the fallback page, the gateway service and the shuttle service can be changed if required. For more information, please refer to Service Socket Addresses.

7. StorageGRID and Loadbalancer.org Appliance Configuration – Scenario 1

This configuration steps presented in this section relate to the deployment scenario presented here. With this scenario, local site server load balancing is handled by the NetApp Admin Nodes, global load balancing is handled by the Loadbalancer.org appliances.

7.1. StorageGRID Configuration

In this guide, an HA Group is configured which contain 2 Admin nodes and has 2 VIPs; one for the Grid Admin / Tennant connections and the other for the S3 client connections. A Load Balancer Endpoint is also configured to load balance the S3 client connections.

7.1.1. High Availability Groups

To configure HA Groups, use the StorageGRID menu option: ***Configuration > High Availability Groups***.

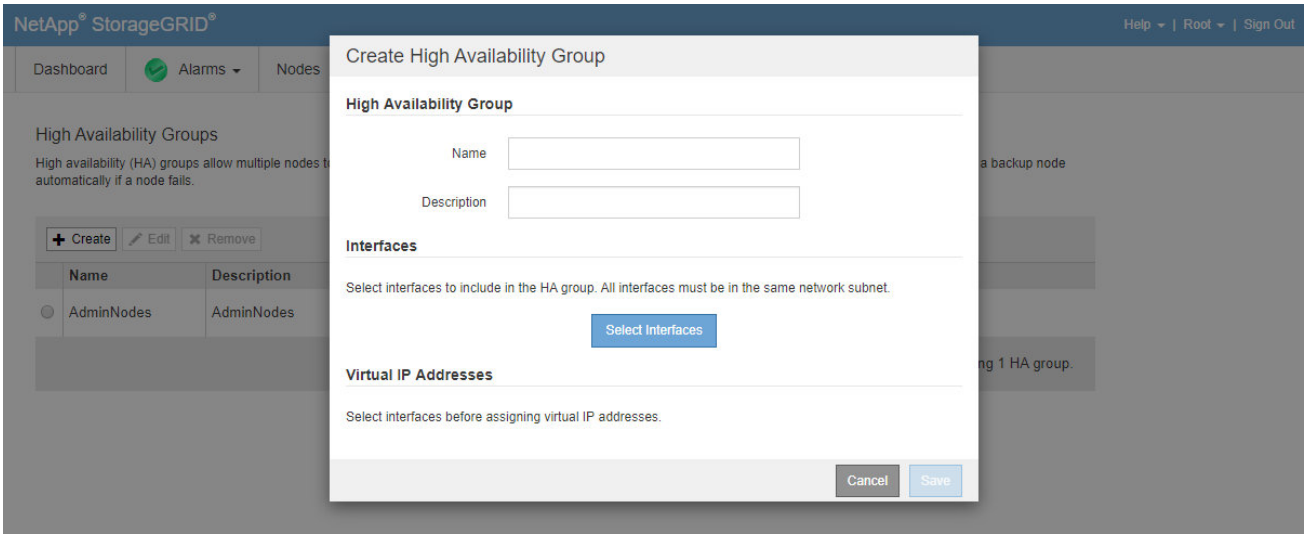

As mentioned, an HA Group with 2 VIPs is configured as shown below:

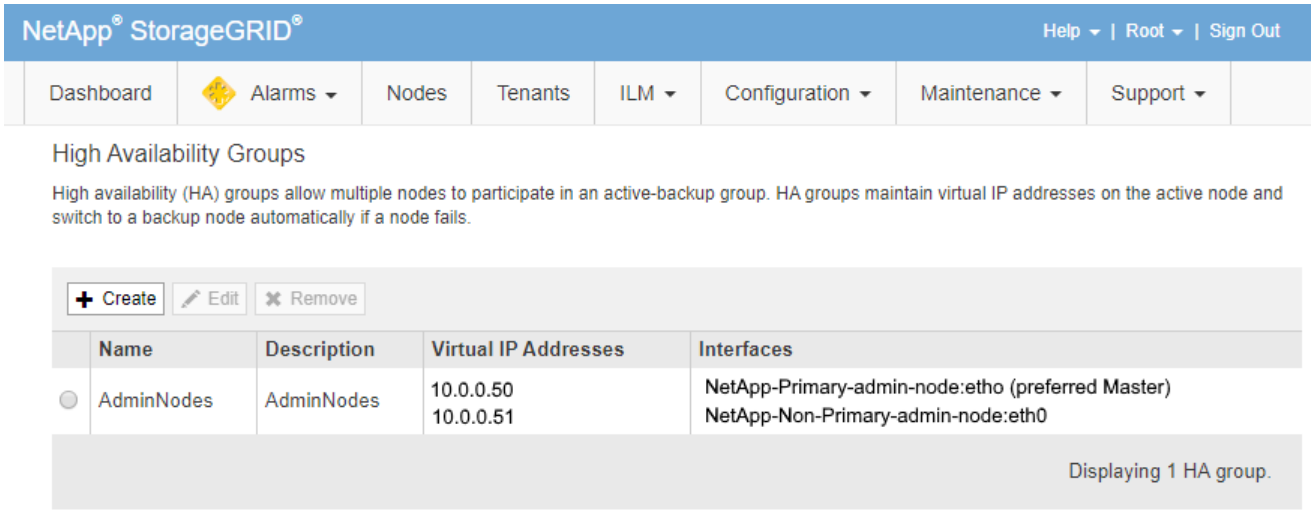

> **Note** For more information on configuring High Availability Groups, please refer to the NetApp Documentation Center.

7.1.2. Load Balancer Endpoints

To configure Endpoints, use the StorageGRID menu option: ***Configuration > Load Balancer Endpoints***.

> **Note** As explained in the second message in the screenshot below, endpoints configured on ports 80 and 443 only function on Gateway Nodes. Therefore If endpoints are configured on Admin Nodes, the standard HTTP & HTTPS ports cannot be used. In this guide, S3 Endpoints are configured on Admin Nodes, therefore the S3 HTTPS endpoint is configured on the alternative arbitrary port 18443. This can of course be changed as required.

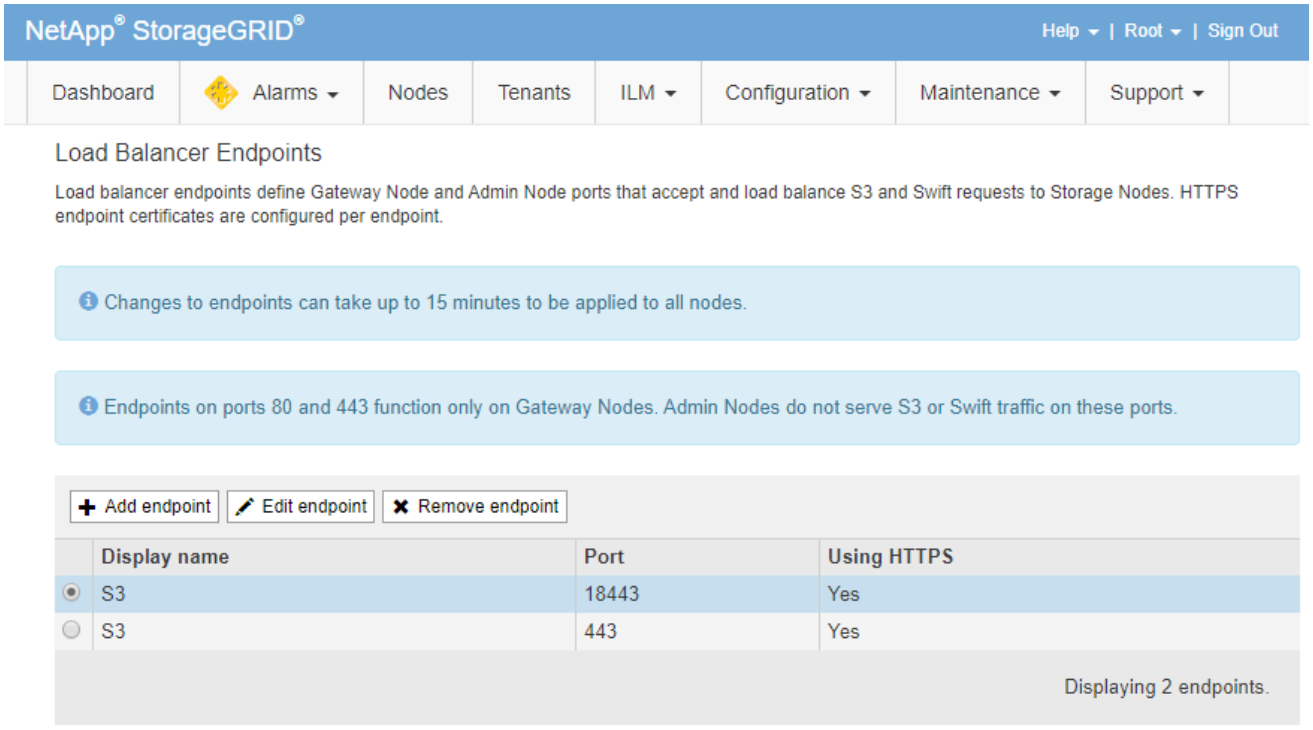

> **Note** For more information on configuring Load Balancer Endpoints, please refer to the NetApp

Documentation Center.

7.2. Loadbalancer.org Appliance Configuration

7.2.1. Step 1 – Configure the HA Pair

If you intend to deploy 2 LB.org load balancers at each site in order to configure an HA clustered pair (our recommended configuration) then the HA pair should be configured first before other 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 the section DNS Server Configuration.

Once the HA pair is configured, all configuration steps should take place on the Primary unit, the Secondary unit will then be kept in sync automatically. Please refer to the section Configuring HA - Adding a Secondary Appliance for details on configuring an HA pair.

7.2.2. step 2 – Configure GSLB

For Scenario 1, the LB.org appliance requires 2 sub domain definitions; one for the admin connections and the other for the S3 client connections. Configuration takes place in the WebUI under ***Cluster Configuration > GSLB Configuration***. The GSLB configuration must be identical across all appliances and sites to ensure that DNS replies are consistent.

a) Configure the Global Names

This is where the sub domains that are handled by GSLB are defined.

To add Global Names:

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.

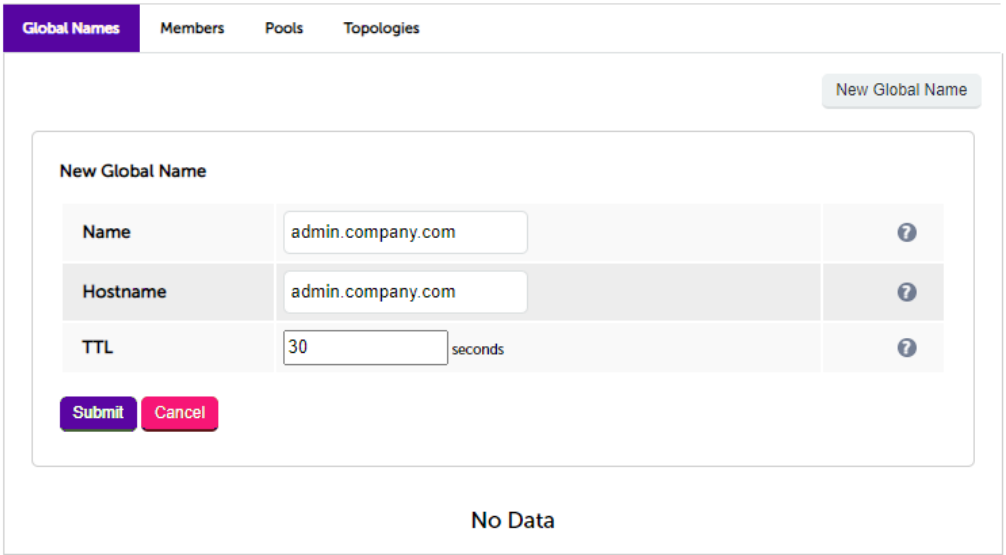

4. Define a friendly *Name* for the new hostname, which can just be the subdomain itself, e.g. **admin.company.com**
5. Define the *Hostname* of what will be the delegated subdomain, e.g. **admin.company.com**
6. Set the required *TTL*, the default is **30s**.
7. Click Submit.

>>> Now repeat steps 1 – 7 to define the **s3.company.com** sub domain. Once complete, both sub domains will be displayed:

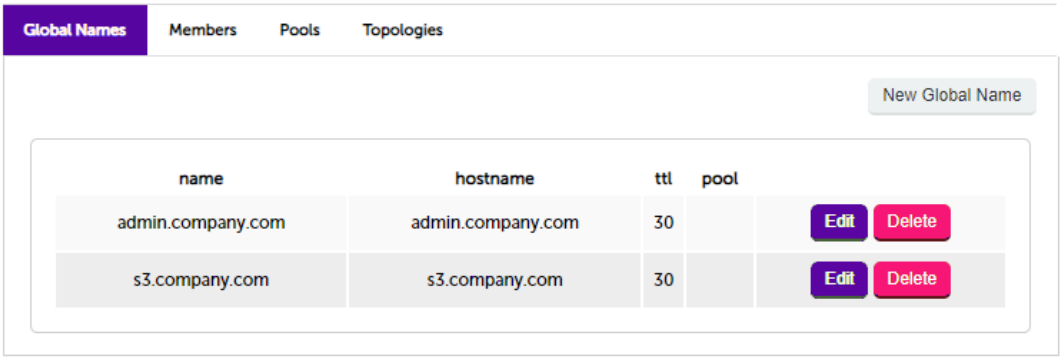

b) Configure the Members

In this Scenario the members are the HA Group VIPs defined in StorageGRID.

To add members:

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

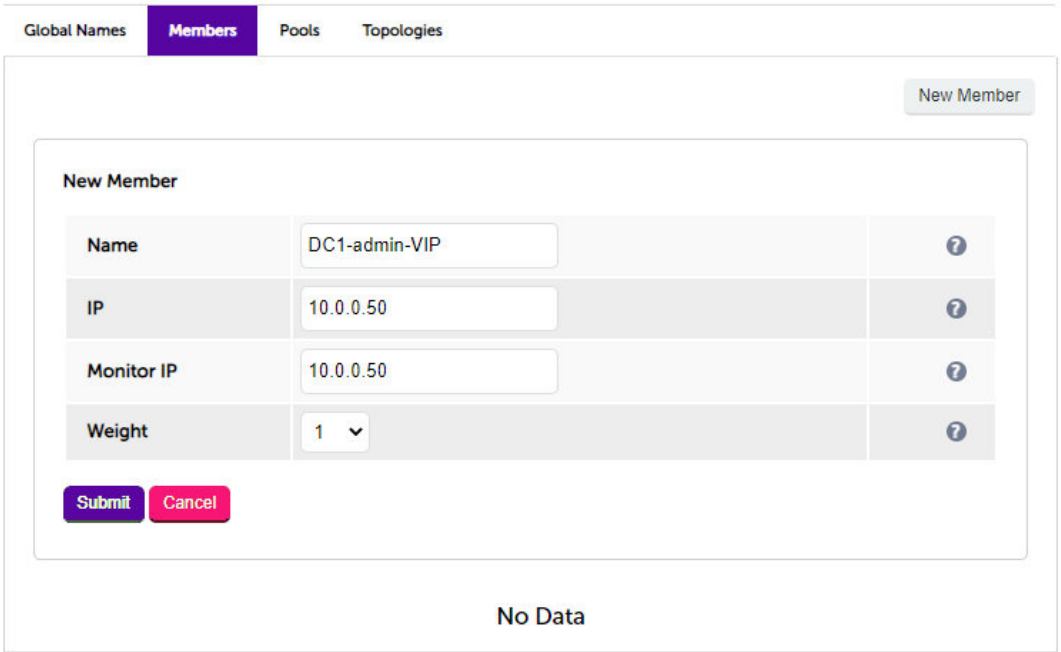

3. Enter a friendly *Name* for the member, e.g. **DC1-admin-VIP**.
4. Specify an *IP* address for the member, e.g. **10.0.0.50**.
5. Set *Monitor IP* to the same value, e.g. **10.0.0.50**.
6. Click **Submit**.

>>> Now repeat steps 1 – 6 to define the remaining members, in this guide **DC2-admin-VIP**, **DC1-S3-VIP** and **DC2-S3-VIP**. Once complete, all members will be displayed:

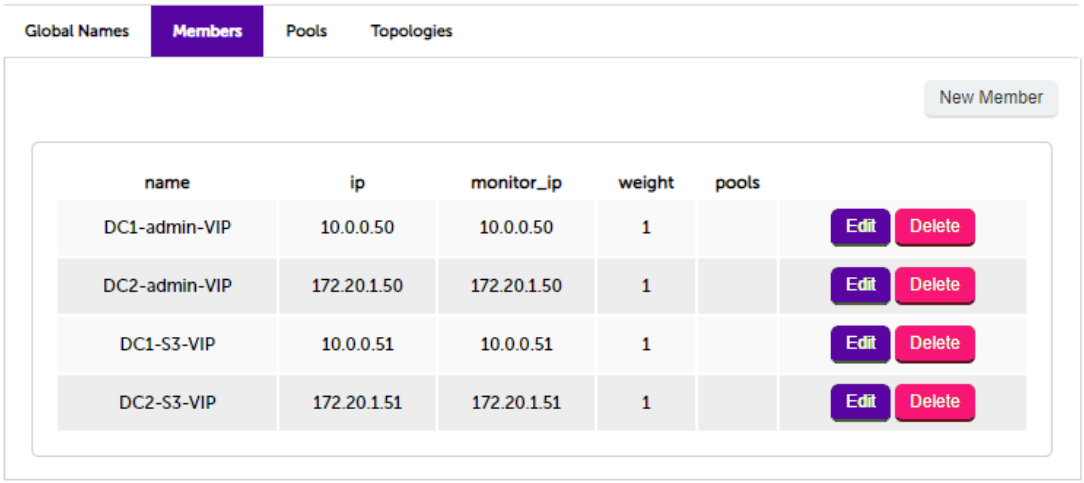

c) Configure the Pools

A pool must be created to link together each global name with the relevant members that must serve traffic for that global name.

To Add a Pool:

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

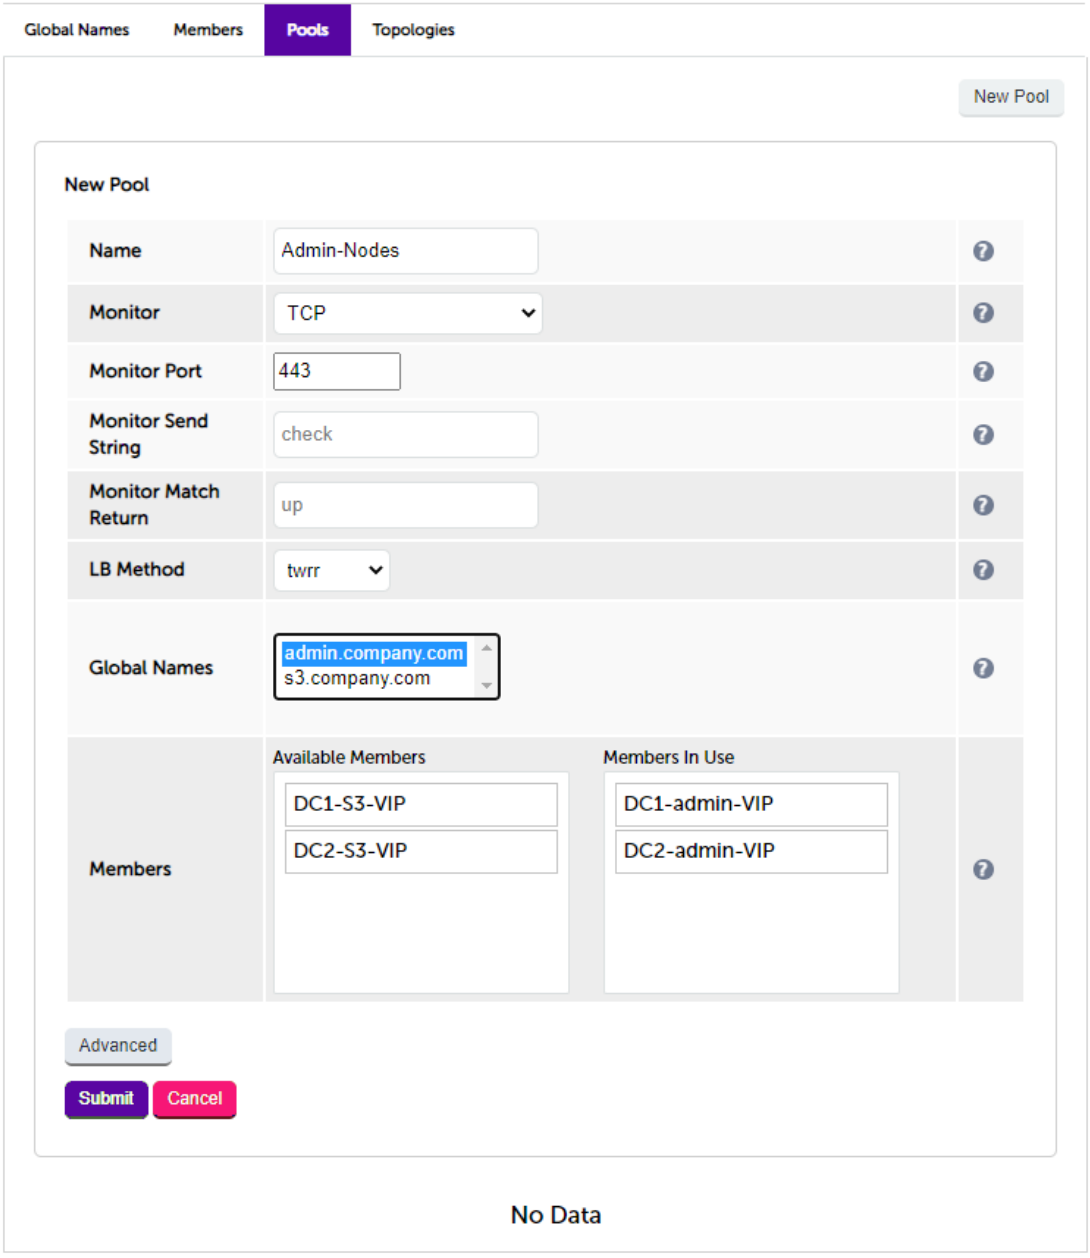

3. Enter a friendly *Name* for the pool, e.g. **Admin-Nodes**.
4. Set the *Monitor* to **TCP**.
5. Set *Monitor Port* to **443**.
6. Set *LB Method* to **twrr**.
7. From the *Global Names* list box, select the global name in question, e.g. **admin.company.com**
8. In the *Members* section, drag the appropriate members (i.e. the Admin nodes) from the *Available Members* box into the *Members In Use* box, e.g. **DC1-admin-VIP** & **DC2-admin-VIP**.
9. Click **Submit**.

>>> Now repeat steps 1 – 9 to define the remaining pool, in this guide **S3-Nodes** ensuring that the *Monitor Port* is set to **18443** (this is the value used in this guide). Once complete, both pools will be displayed:

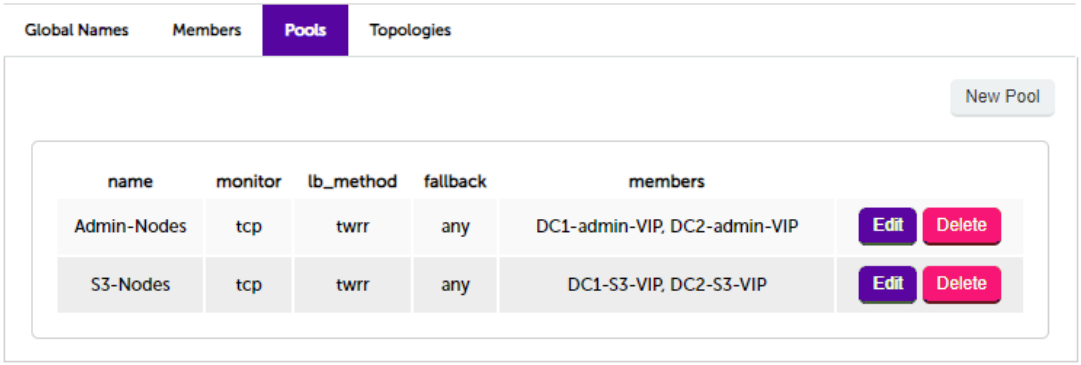

d) Configure the Topology

The example below relates to Scenario 1 which has one client site and 2 DCs. This topology definition associates Client Site 1 with DC 1. This ensures that under normal circumstances client connections from Client Site 1 (subnet 192.168.10.0/24) will be handled by DC1.

To Add a Topology:

1. Select the *Topologies* tab.
2. Click the **New Topology** button.

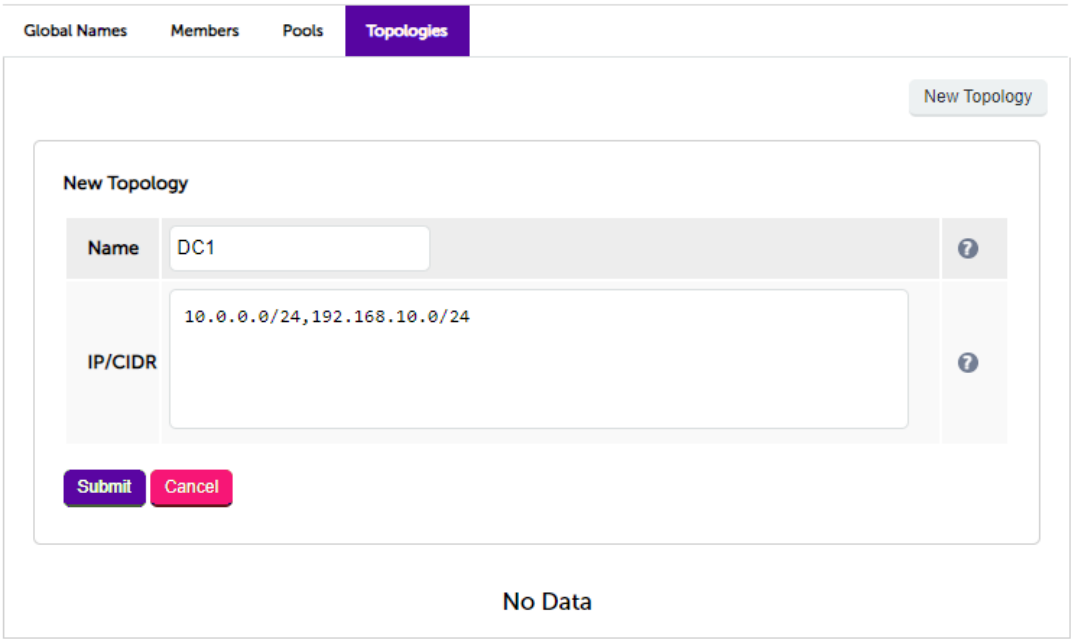

3. Enter a friendly *Name* for the topology, e.g. **DC1**.
4. In the *IP/CIDR* text box, define the subnet(s) that covers the site in question, e.g. **10.0.0.0/24, 192.168.10.0/24**.

> **Note** This can be a comma separated list of subnets as shown, or specific hosts. The key is that the site's local DNS server ***and*** the IP addresses of the StorageGRID nodes 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 nodes: the IP addresses of the local nodes are then served as DNS responses for clients at that site.

5. Click **Submit**.

>>> Now repeat steps 1 – 5 to define the other topologies, in this guide for **DC2**. Once complete, all Topologies will be displayed:

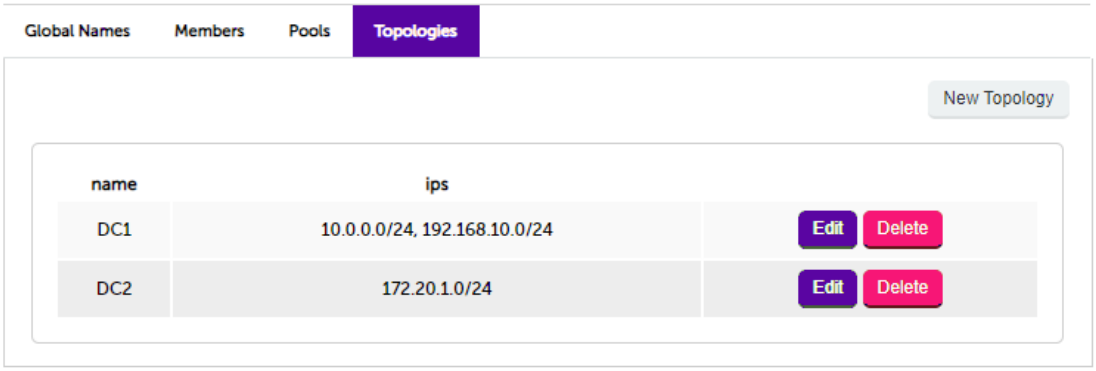

7.2.3. Step 3 – Finalising the Configuration

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**.

7.2.4. Step 4 – Verify the GSLB Configuration

The GSLB configuration should be tested to ensure it's working as expected and that both Primary and Secondary appliances are able to correctly respond to DNS queries for the sub domains. This must be operating correctly when configuring the DNS delegation in the following section.

From a Windows command prompt, the ***nslookup*** command can be used to send test DNS queries to the load balancers. 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 HA Group VIPs as shown below:

```
C:\Users\me>nslookup -norecurse admin.company.com 10.0.0.1
Server: UnKnown
Address: 10.0.0.1

Name: admin.company.com
Address: 10.0.0.50
```

This test should be repeated for the **s3.company.com** sub domain and then using the IP address of the Secondary unit to ensure that the Secondary is also able to correctly respond to the DNS queries.

7.3. DNS Server Configuration

Once the GSLB service has been configured on the Primary & Secondary load balancer at each site, the local DNS server at each client site must then be configured for GSLB.

The DNS server at each client site must be configured to delegate DNS requests for the subdomain in question to the load balancers. The load balancer's GSLB services will then serve the appropriate IP addresses to the local DNS server and back to the clients. For the example presented in this guide, the DNS server at the client site must be configured with a delegation for the sub domains **admin.company.com** and **s3.company.com**. The Primary and Secondary load balancers at DC1 and DC2 are configured as Name Servers for the sub domain which provides multi-site redundancy.

The exact steps for creating a DNS delegation vary between different DNS servers and are outside the scope of this document. For further information, a blog post that walks through creating a DNS delegation on a Microsoft DNS server in the context of setting up GSLB on our appliance can be found here (see the section titled "Delegating your subdomain to your GSLB's using Microsoft's DNS Server").

8. StorageGRID and Loadbalancer.org Appliance Configuration – Scenario 2

This configuration steps presented in this section relates to the deployment scenario presented here. With this scenario, local site server load balancing and global load balancing are both handled by the Loadbalancer.org appliances.

8.1. StorageGRID Configuration

High Availability Groups and Load Balancer Endpoints are not required for scenario 2. In this scenario the Loadbalancer.org appliance handles both local and global load balancing. To successfully use an external layer 7 load balancer to route requests to StorageGRID v11.4.0 and later, the load balancer must be authorized within StorageGRID. This process is described in the following section.

8.1.1. Configure the Loadbalancer.org Appliance as A Trusted Layer 7 Loadbalancer

If an external (third party) Layer 7 load balancer is used to route requests to the Storage Nodes, StorageGRID needs to determine the real sender's IP address. It does this by looking at the X-Forwarded-For (XFF) header, which is inserted into the request by the load balancer. As the X-Forwarded-For header can be easily spoofed in requests sent directly to the Storage Nodes, StorageGRID needs to ensure that each request is being routed by a trusted Layer 7 load balancer. If StorageGRID cannot trust the source of the request, it will ignore the X-Forwarded-For header.

In StorageGRID v11.4.0 and later, a new Grid Management API has been added to allow a list of trusted external Layer 7 load balancers to be configured.

To configure the Bucket policy and StorageGRID API, follow steps 1 – 4 below.

> **Note** More information for registered NetApp customers can be found here.

Step 1 – Configure the Bucket Policy

In the following example, everyone (including anonymous) is allowed to List the bucket and perform any Object operations on all objects in the bucket, provided that the requests come from a specified IP range (54.240.143.0 to 54.240.143.255, except 54.240.143.188). All other operations will be denied, and all requests outside of the IP range will be denied.

```
{
  "Statement": [
    {
      "Sid": "AllowEveryoneReadWriteAccessIfInSourceIpRange",
      "Effect": "Allow",
      "Principal": "*",
      "Action": [ "s3:*Object", "s3:ListBucket" ],
      "Resource": ["arn:aws:s3:::examplebucket","arn:aws:s3:::examplebucket/*"],
      "Condition": {
        "IpAddress": {"aws:SourceIp": "54.240.143.0/24"},
        "NotIpAddress": {"aws:SourceIp": "54.240.143.188"}
      }
    }
  ]
}
```

Step 2 – Generate an Authorization Token

- Navigate to the Authorization token section within the StorageGRID API:

  https://<Your-Admin-Node>/grid/apidocs.html?type=private#/auth/post_authorize

- Click **Try it out**
- Update the username and password in the **Post** body
- Click the blue **Execute** bar
- If valid credentials have been provided, the Authorization token will be generated and displayed at the top of the screen:

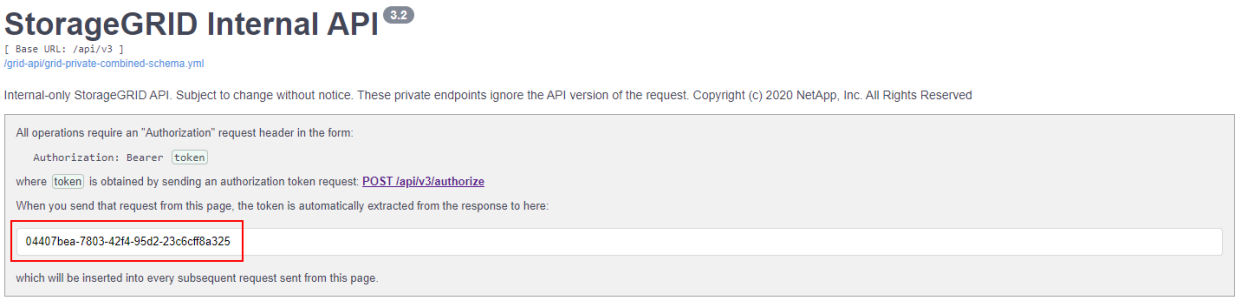

- The token is automatically inserted into every subsequent request sent from the page. If the page is closed, a new Authorization token must be generated.

Step 3 – Authorize the Load balancer

- Navigate to the External Load Balancer section within the StorageGRID API:

  https://<Your-Admin-Node>/grid/apidocs.html?type=private#/external-load-balancers/put_private_external_load_balancers

- Click **Try it out**
- Update the IP address in the **Put** body

> **Note** The IP address defined should be the address allocated to the load balancer's network interface when it was deployed. This can be viewed using the Loadbalancer.org WebUI by navigating to: *Local Configuration > Network Interface Configuration*.

- Click the blue **Execute** bar
- Scroll down to the **Server response** section and verify that the IP address has been set successfully:

```
{
  "responseTime": "2020-09-02T10:13:06.404Z",
  "status": "success",
  "apiVersion": "3.2",
  "data": [
    "192.168.10.100"
  ]
}
```

Step 4 – Verify the Configuration

> **Note** This step should be completed AFTER the VIPs have been configured on the Loadbalancer.org appliance and HAProxy has been restarted. Please refer to Step 2 – Configure Local Server Load Balancing below for details on configuring the VIPs.

- Send an anonymous request through the load balancer:

```
aws s3api --no-sign-request --endpoint-url https://<VIP address> list-objects --bucket
examplebucket
```

If you send the request from an IP address that is allowed in the policy, the request should succeed. If you send the request from any other IP address, the request should fail as *Access Denied*.

8.2. Loadbalancer.org Appliance Configuration

8.2.1. Step 1 – Configure the HA Pair

If you intend to deploy 2 LB.org load balancers at each site in order to configure an HA clustered pair (our recommended configuration) then the HA pair should be configured first before other 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. For more information please refer to the section DNS Server Configuration.

Once the HA pair is configured, all configuration steps 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 the section Configuring HA - Adding a Secondary Appliance.

8.2.2. Step 2 – Configure Local Server Load Balancing

Within each site 2 VIPs are required. One to load balance the Grid / Tenant Manager connections and the other to load balance the S3 client connections.

VIP 1 – Grid / Tennant Admin

a) Configure the VIP

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Enter the following details:

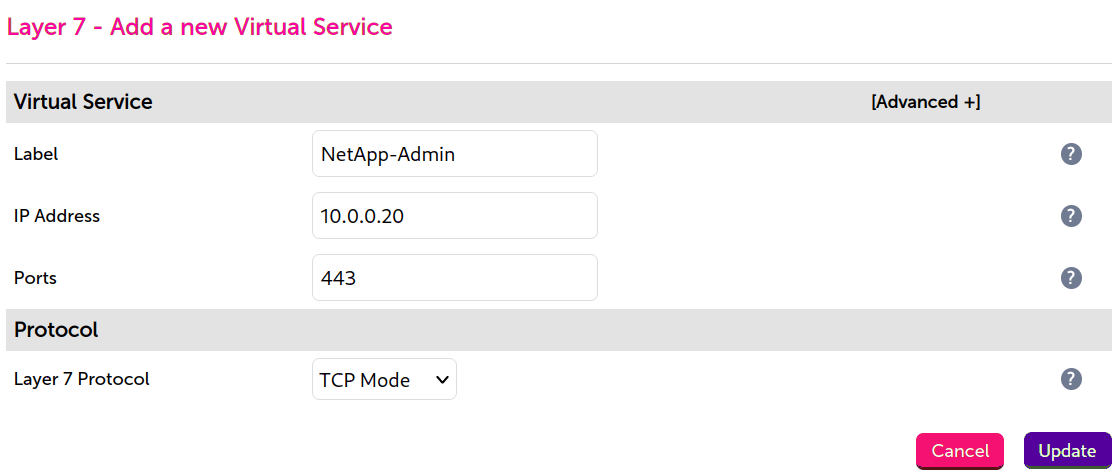

3. Enter an appropriate name for the VIP in the *Label* field, e.g. **NetApp-Admin**.
4. Set the *Virtual Service IP* address field to the required IP address, e.g. **10.0.0.20**.
5. Set the *Virtual Service Ports* field to **443**.
6. Set the *Layer 7 Protocol* to **TCP Mode**.
7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
10. Scroll down to the *Health Checks* section and set the *Health Check* to **Negotiate HTTPS (GET)**.
11. Leave *Response Expected* blank – this will cause the load balancer to look for a **HTTP 200 OK** response from each Storage Node.
12. Scroll down to the *Other* section and click **[Advanced]**.
13. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **10m** (i.e. 10 minutes).
14. Click **Update**.

b) Define the Associated Real Servers

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 NetApp-Admin VIP.

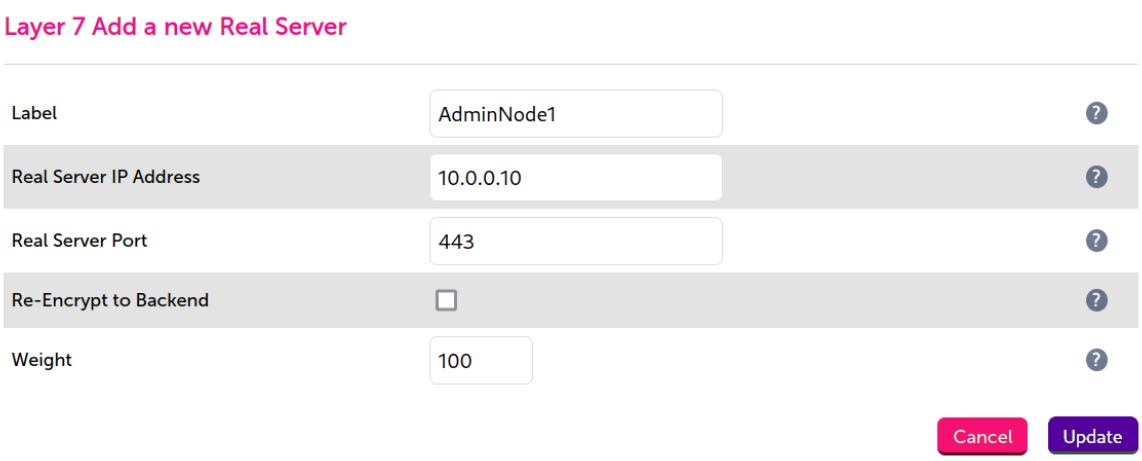

2. Enter an appropriate name for the server in the *Label* field, e.g. **AdminNode1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **10.0.0.10**.
4. Set the *Real Server Port* field to **443**.
5. Click **Update**.
6. Now repeat these steps to add the other Admin Node(s).

VIP 2 – S3 Client Access

a) Configure the VIP

1. Using the WebUI, navigate to *Cluster Configuration > Layer 7 – Virtual Services* and click on **Add a new Virtual Service**.
2. Enter the following details:

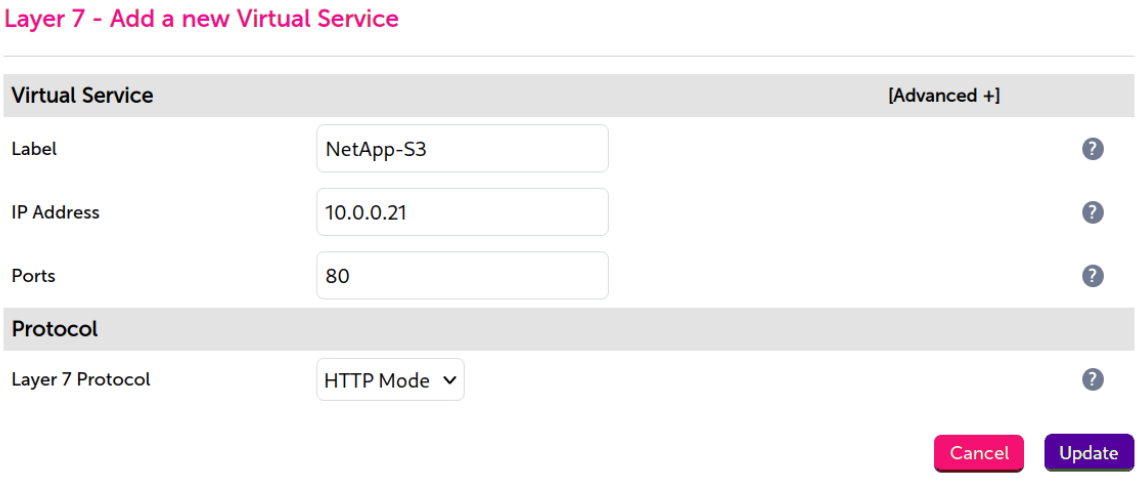

3. Enter an appropriate name for the VIP in the *Label* field, e.g. **NetApp-S3**.
4. Set the *Virtual Service IP* address field to the required IP address, e.g. **10.0.0.21**.
5. Set the *Virtual Service Ports* field to the required value, e.g. **80**.
6. Set the *Layer 7 Protocol* to **HTTP Mode**.

7. Click **Update**.
8. Now click **Modify** next to the newly created VIP.
9. Scroll down to the *Persistence* section and set *Persistence Mode* to **None**.
10. Scroll down to the *Health Checks* section and set the *Health Check* to **Negotiate HTTPS (OPTIONS)**.
11. Leave *Response Expected* blank – this will cause the load balancer to look for a **HTTP 200 OK** response from each Storage Node.
12. Scroll down to the *Other* section and click **[Advanced]**.
13. Enable (check) the *Timeout* checkbox and set both *Client Timeout* & *Real Server Timeout* to **10m** (i.e. 10 minutes).
14. Click **Update**.

b) Define the Associated Real Servers

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 NetApp-S3 VIP.

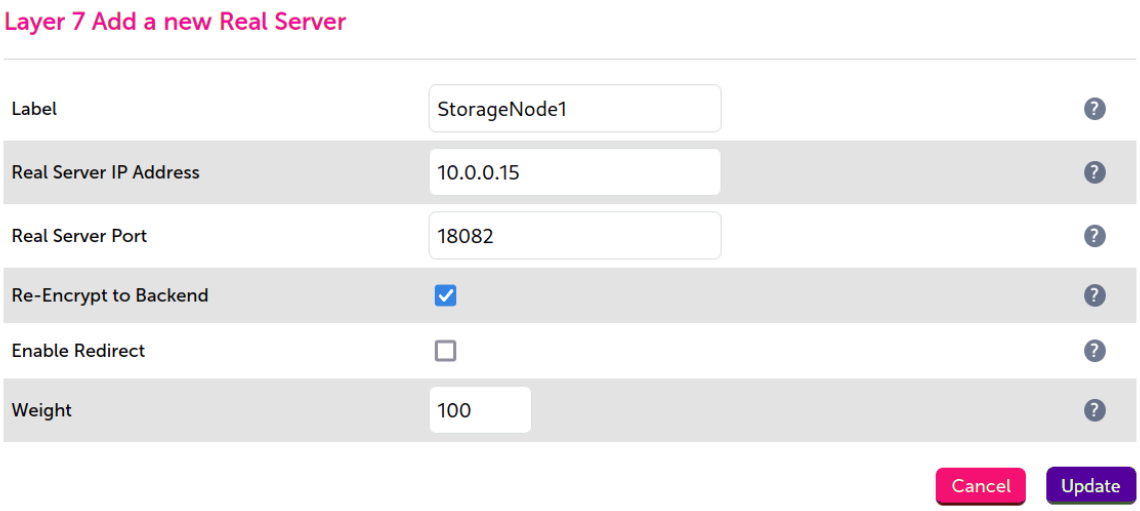

2. Enter an appropriate name for the server in the *Label* field, e.g. **StorageNode1**.
3. Change the *Real Server IP Address* field to the required IP address, e.g. **10.0.0.15**.
4. Set the *Real Server Port* field to **18082**.
5. Enable (check) *Re-Encrypt to Backend*.
6. Click **Update**.
7. Now repeat these steps to add the other Storage Node(s).

> **Note** As storage nodes are added or decommissioned due to capacity changes or hardware refresh, you will need to update the Real Server configuration accordingly.

c) Upload SSL Certificate

SSL certificates can be uploaded to the appliance and used with SSL termination VIPs. Both PFX and PEM format

certificates can be uploaded.

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*.

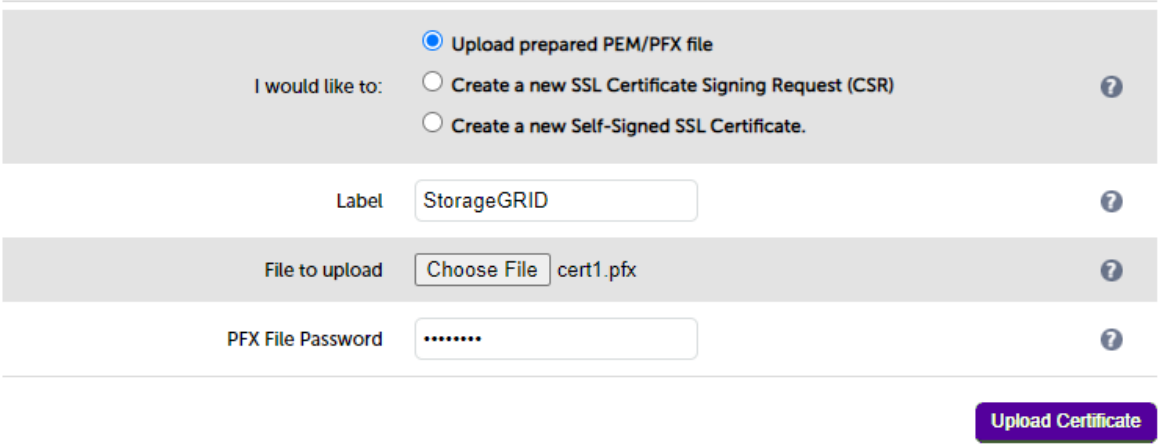

3. Enter a suitable *Label* (name) for the certificate, e.g. **StorageGRID**.
4. Browse to and select the certificate file to upload (PEM or PFX format).
5. Enter the password, if applicable.
6. Click **Upload Certificate**, if successful, a message similar to the following will be displayed:

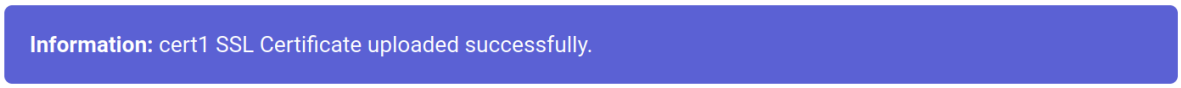

d) Configure SSL Termination

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

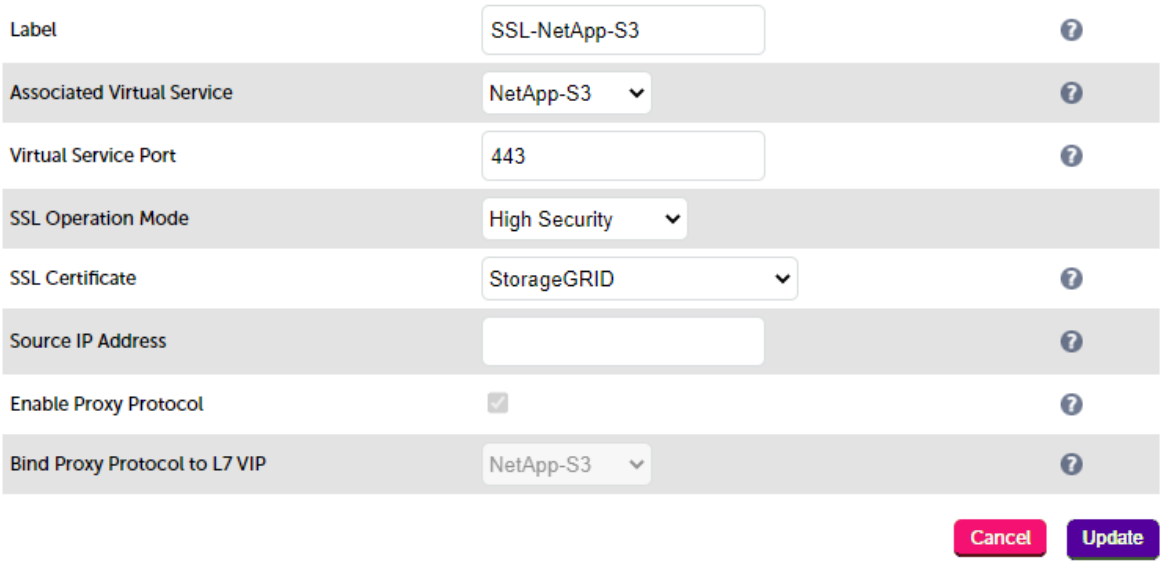

2. Using the *Associated Virtual Service* drop-down, select the Virtual Service created above, e.g. **NetApp-S3**.

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

3. Leave *Virtual Service Port* set to **443**.
4. Leave *SSL Operation Mode* set to **High Security**.
5. Select the *SSL Certificate* uploaded previously, e.g. **StorageGRID**.
6. Click **Update**.

e) Apply the Configuration – Reload Services

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**.

8.2.3. Step 3 – Configure GSLB

For Scenario 2, as with Scenario 1 the LB.org appliance requires 2 sub domain definitions; one for the admin connections and the other for the S3 client connections. Configuration takes place in the WebUI under *Cluster Configuration > GSLB Configuration*. The GSLB configuration must be identical across all appliances and sites to ensure that DNS replies are consistent.

a) Configure the Global Names

This is where the sub domains that are handled by GSLB are defined.

To add Global Names:

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.

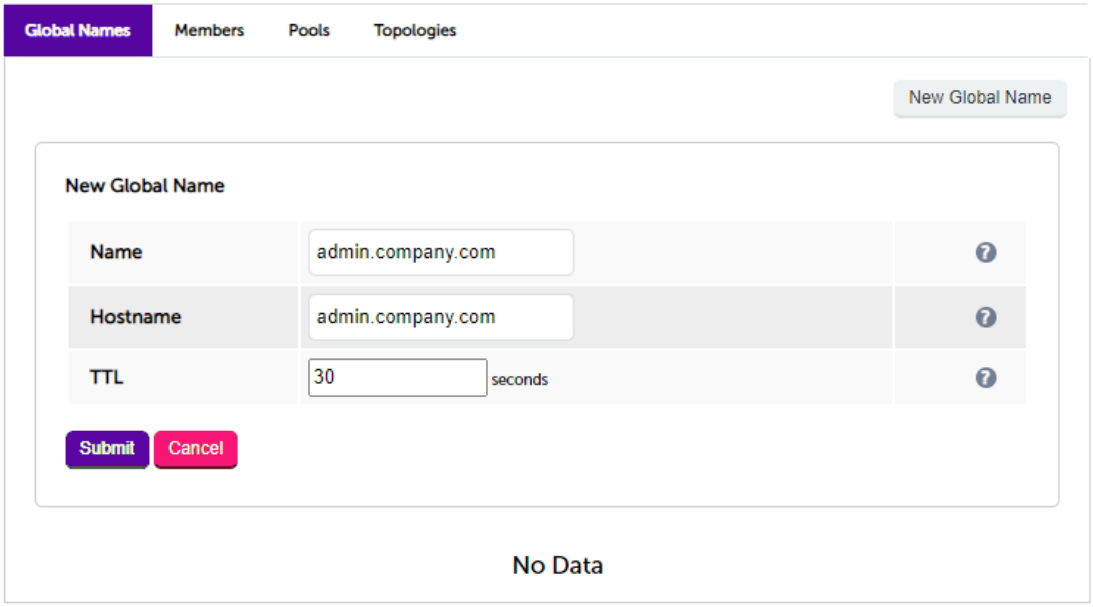

4. Define a friendly *Name* for the new hostname, which can just be the subdomain itself, e.g. **admin.company.com**
5. Define the *Hostname* of what will be the delegated subdomain, e.g. **admin.company.com**
6. Click Submit.

>>> Now repeat steps 1 – 6 to define the **s3.company.com** sub domain. Once complete, both sub domains will be displayed:

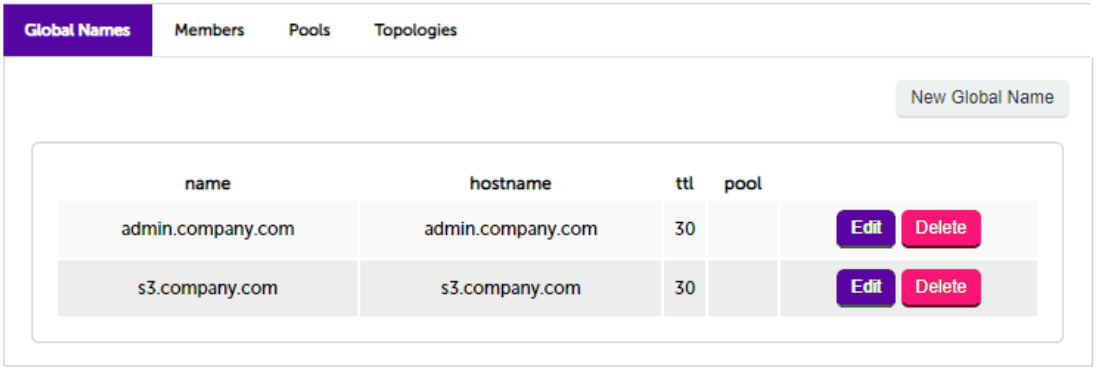

b) Configure the Members

In this Scenario the members are the LB.org appliance VIPs.

To add members:

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

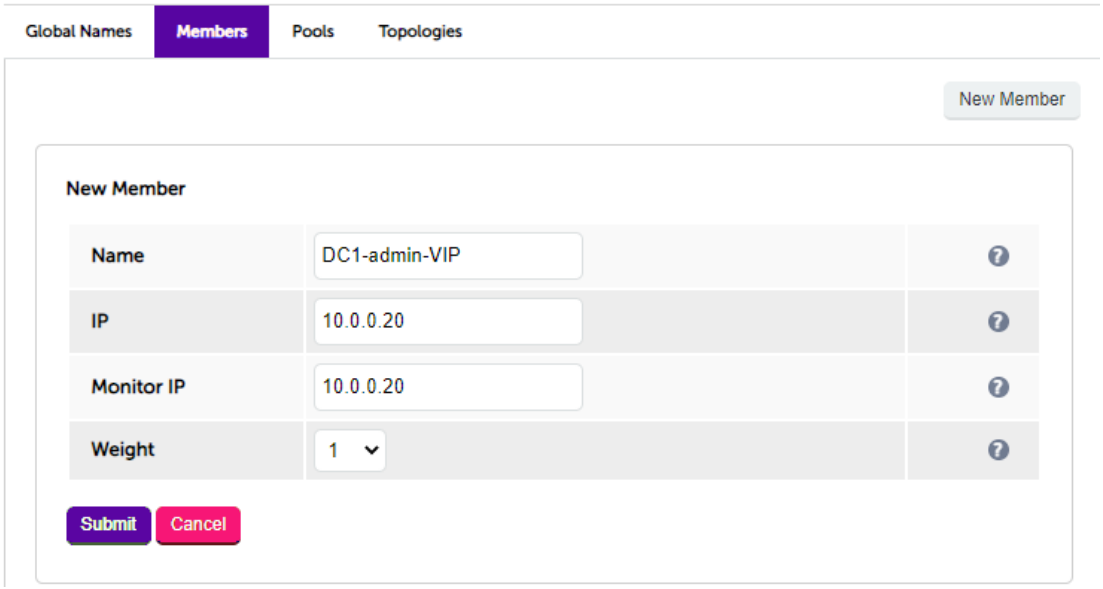

3. Enter a friendly *Name* for the member, e.g. **DC1-admin-VIP**.
4. Specify an *IP* address for the member, e.g. **10.0.0.20**.
5. Set *Monitor IP* to the same value, e.g. **10.0.0.20**.
6. Click **Submit**.

>>> Now repeat steps 1 – 6 to define the remaining members, in this guide **DC2-admin-VIP**, **DC1-S3-VIP** and **DC2-S3-VIP**. Once complete, all members will be displayed:

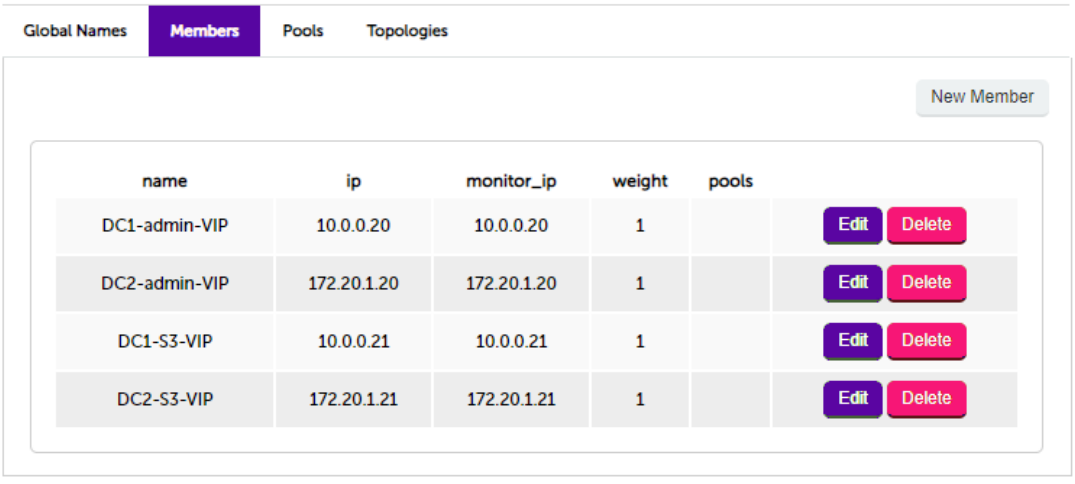

c) Configure The Pools

A pool must be created to link together each global name with the relevant members that must serve traffic for that global name.

To Add a Pool:

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

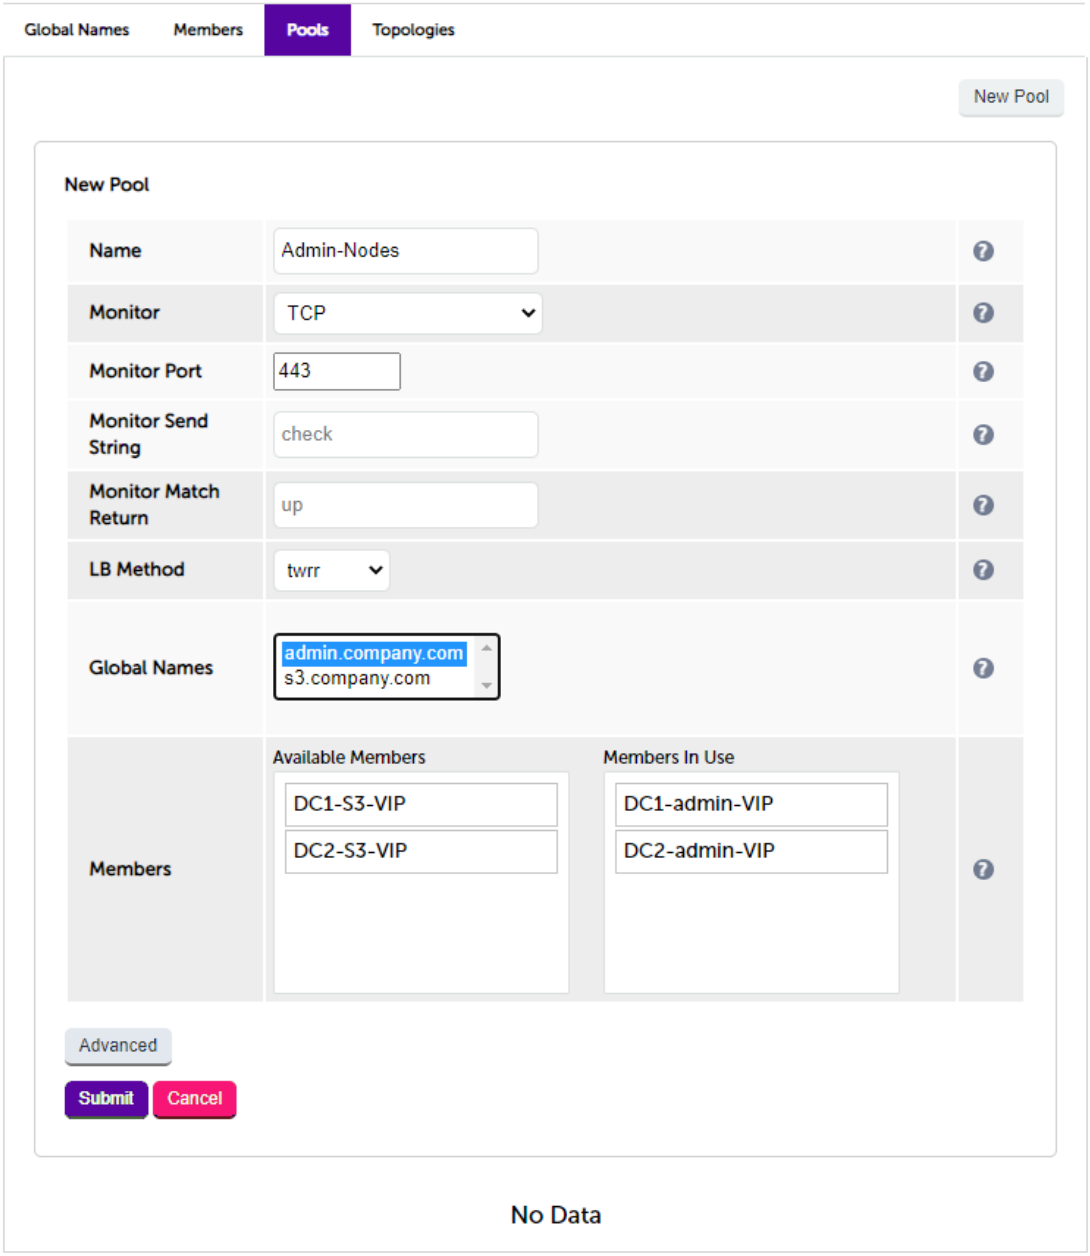

3. Enter a friendly *Name* for the pool, e.g. **Admin-Nodes**.
4. Set the *Monitor* to **TCP**.
5. Set *Monitor Port* to **443**.
6. Set *LB Method* to **twrr**.
7. From the *Global Names* list box, select the global name in question, e.g. **admin.company.com**
8. In the *Members* section, drag the appropriate members (i.e. the Admin nodes) from the *Available Members* box into the *Members In Use* box, e.g. **DC1-admin-VIP** & **DC2-admin-VIP**.
9. Click **Submit**.

>>> Now repeat steps 1 – 9 to define the remaining pool, in this guide **S3-Nodes**. Once complete, both pools will be displayed:

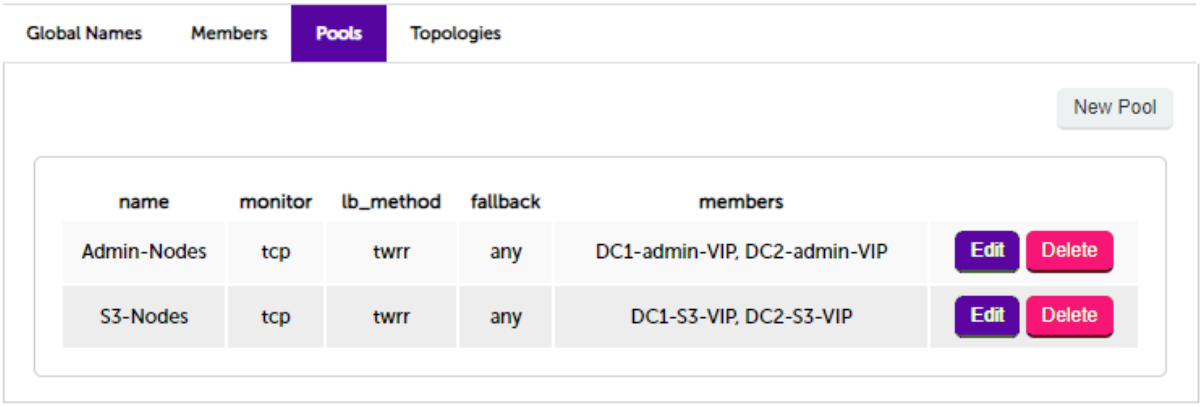

d) Configure The Topology

The example below relates to Scenario 2 which has one client site and 2 DCs. This topology definition associates Client Site 1 with DC 1. This ensures that under normal circumstances client connections from Client Site 1 (subnet 192.168.10.0/24) will be handled by DC 1.

To Add a Topology:

1. Select the *Topologies* tab.
2. Click the **New Topology** button.

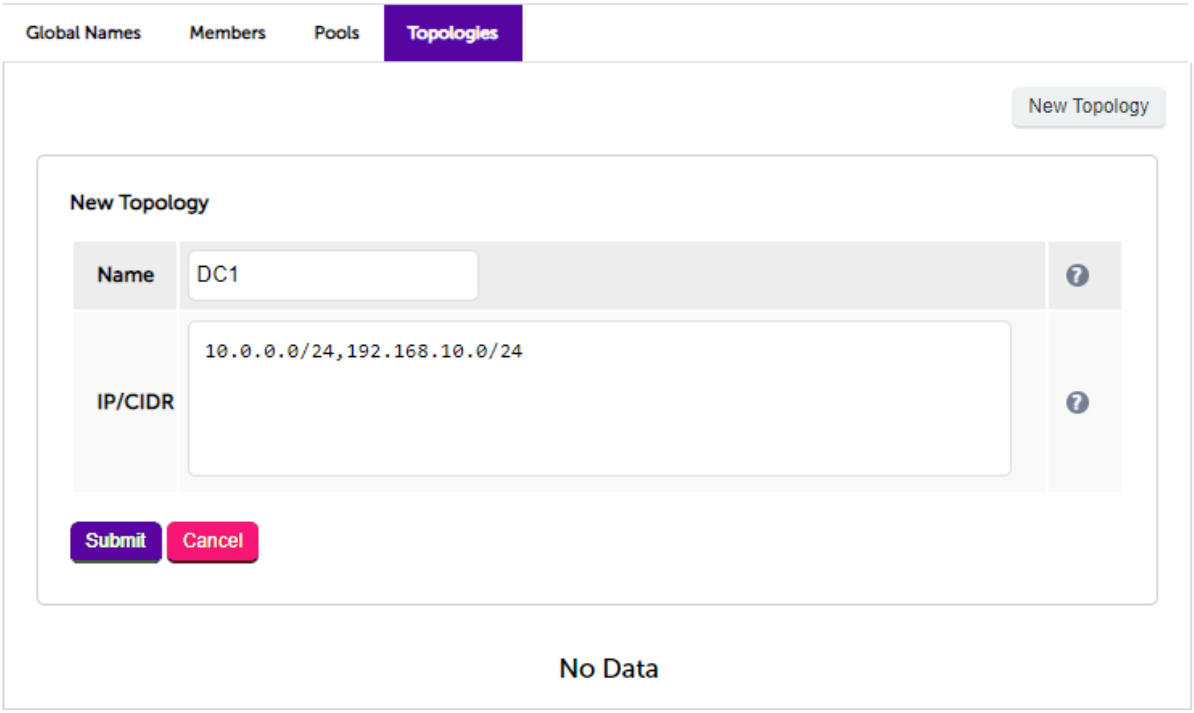

3. Enter a friendly *Name* for the topology, e.g. **DC1**.
4. In the *IP/CIDR* text box, define the subnet(s) that covers the site in question, e.g. **10.0.0.0/24, 192.168.10.0/24**.

> **Note** This can be a comma separated list of subnets as shown, or specific hosts. The key is that the site's local DNS server ***and*** the IP addresses of the StorageGRID nodes 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 nodes: the IP addresses of the local nodes are then served as DNS responses for clients at that site.

5. Click **Submit**.

>>> Now repeat steps 1 – 5 to define the other topologies, in this guide for **DC2**. Once complete, all Topologies will be displayed:

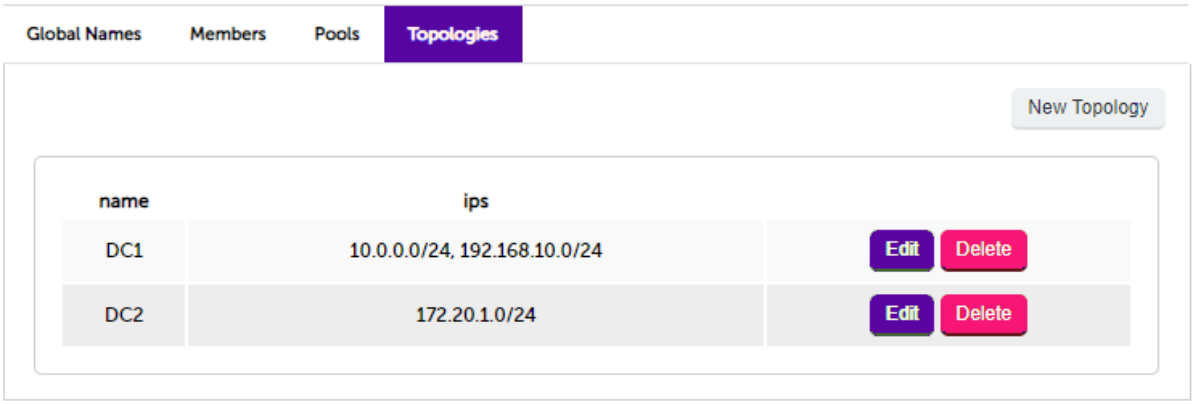

8.2.4. Step 4 – Finalising the Configuration

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**.

8.2.5. Step 5 – Verify the GSLB Configuration

The GSLB configuration should be tested to ensure it's working as expected and that both Primary and Secondary appliances are able to correctly respond to DNS queries for the sub domains. This must be operating correctly when configuring the DNS delegation in the following section.

From a Windows command prompt, the ***nslookup*** command can be used to send test DNS queries to the load balancers. 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 LB.org appliance VIPs as shown below:

```
C:\Users\me>nslookup -norecurse admin.company.com 10.0.0.1
Server: UnKnown
Address: 10.0.0.1

Name: admin.company.com
Address: 10.0.0.20
```

This test should be repeated for the **s3.company.com** sub domain and then using the IP address of the Secondary unit to ensure that the Secondary is also able to correctly respond to the DNS queries.

8.3. DNS Server Configuration

Once the GSLB service has been configured on the Primary & Secondary load balancer at each site, the local DNS server at each client site must then be configured for GSLB.

The DNS server at each client site must be configured to delegate DNS requests for the subdomain in question to the load balancers. The load balancer's GSLB services will then serve the appropriate IP addresses to the local DNS server and back to the clients. For the example presented in this guide, the DNS server at the client site must be configured with a delegation for the sub domains **admin.company.com** and **s3.company.com**. The Primary and Secondary load balancers at DC1 and DC2 are configured as Name Servers for the sub domain which provides multi-site redundancy.

The exact steps for creating a DNS delegation vary between different DNS servers and are outside the scope of this document. For further information, a blog post that walks through creating a DNS delegation on a Microsoft DNS server in the context of setting up GSLB on our appliance can be found here (see the section titled "Delegating your subdomain to your GSLB's using Microsoft's DNS Server").

9. Testing & Verification

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

Once the load balancer and storage nodes are configured, ensure that you can successfully connect to the StorageGRID deployment. You'll need to test & verify GSLB for deployment Scenario 1 and you'll need to test & verify GSLB and also the server load balancing configuration for deployment Scenario 2.

9.1. Testing GSLB

The complete GSLB & DNS configuration should be checked, ensuring that the client is able to resolve the sub domain FQDNs of the StorageGRID deployment via delegation and connect to a healthy node.

9.1.1. Verify the DNS Delegation

From a Windows command prompt, 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 this 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 other name servers, giving the other server's details as shown in the example below:

```
C:\Users\me>nslookup -norecurse admin.company.com 10.0.0.50
Server: UnKnown
Address: 10.0.0.50

Name: admin.company.com
Served by:
- lbmaster.company.com
        10.0.0.1
        admin.company.com
- lbslave.company.com
```

```
        10.0.0.2
        admin.company.com
```

9.1.2. Verify the Full DNS Request & Response

Now execute the same command **without** the *-norecurse* option. This should see the DNS server fetch the answer from one of the load balancers 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 HA Group VIPs (scenario 1) or one of the online LB.org appliance VIPs (scenario 2) as shown in the example below:

```
C:\Users\me>nslookup admin.company.com 10.0.0.50
Server: UnKnown
Address: 10.0.0.50

Non-authoritative answer:
Name: admin.company.com
Address: 10.0.0.20
```

9.1.3. Accessing the Service

A successful test will see the test connection passed from the test client to one of the online HA Group VIPs (scenario 1) or one of the online LB.org appliance VIPs (scenario 2).

9.2. Testing Local Server Load Balancing

For the server load balancing configuration of scenario 2, verify that you can connect to the StorageGRID deployment via the VIP addresses on the load balancer.

The System Overview can be viewed using the WebUI. It shows a graphical view of all VIPs & RIPs (i.e. the StorageGRID nodes) and shows the state/health of each node as well as the state of the cluster as a whole.

The example below shows that all NetApp StorageGRID nodes are healthy and available to accept connections.

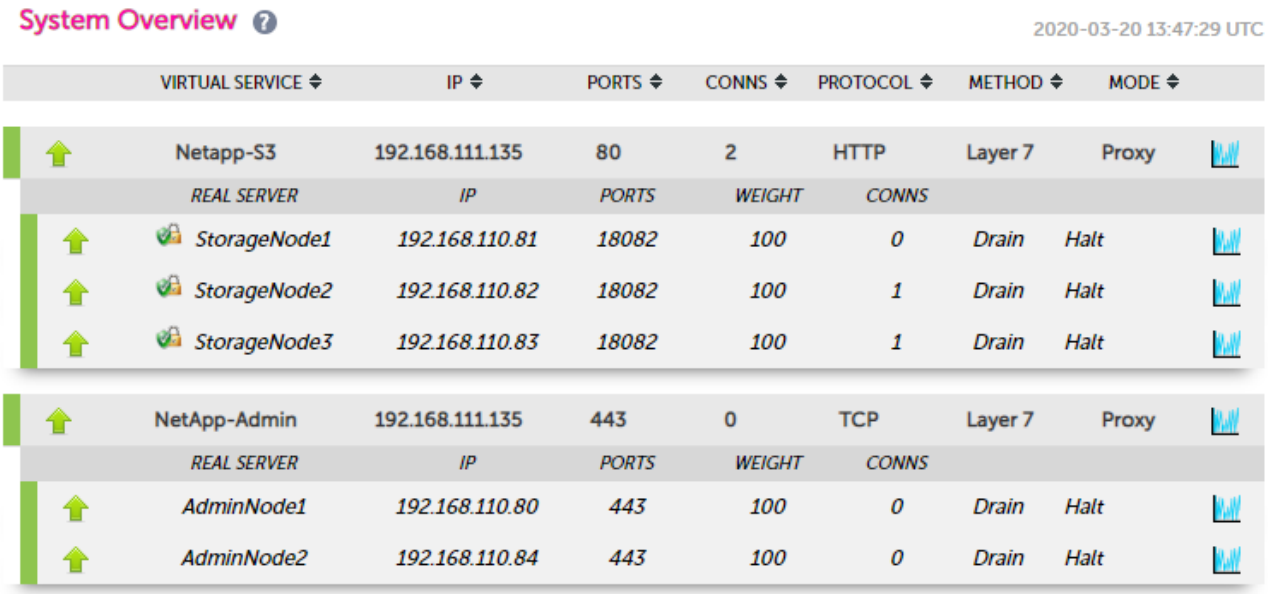

10. Technical Support

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.

11. Further Documentation

For additional information, please refer to the Administration Manual.

12. Appendix

12.1. Configuring HA - Adding a Secondary Appliance

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

> **Note** For Enterprise Azure, the HA pair should be configured first. For more information, please refer to the Azure Quick Start/Configuration Guide available in the documentation library

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

12.1.1. Non-Replicated Settings

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

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

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

12.1.2. Configuring the HA Clustered Pair

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

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

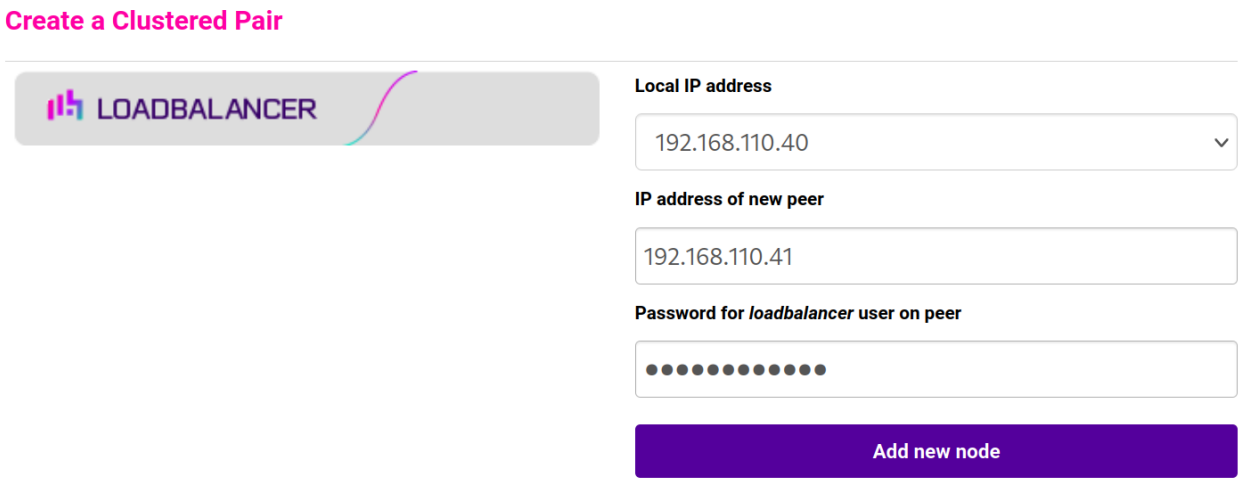

3. Specify the IP address and the *loadbalancer* user's password for the Secondary (peer) appliance as shown in the example above.
4. Click **Add new node**.
5. The pairing process now commences as shown below:

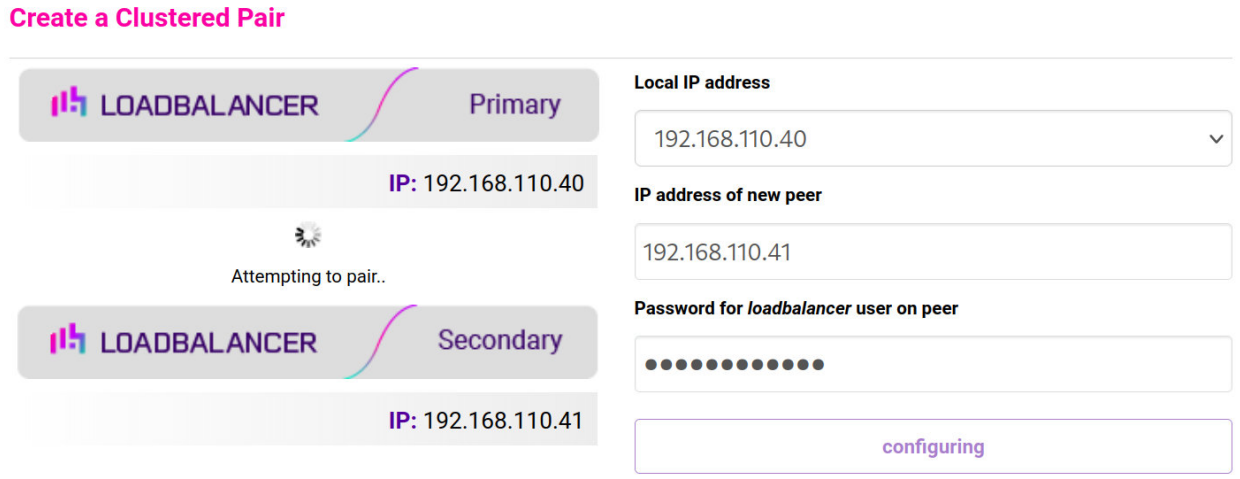

6. Once complete, the following will be displayed on the Primary appliance:

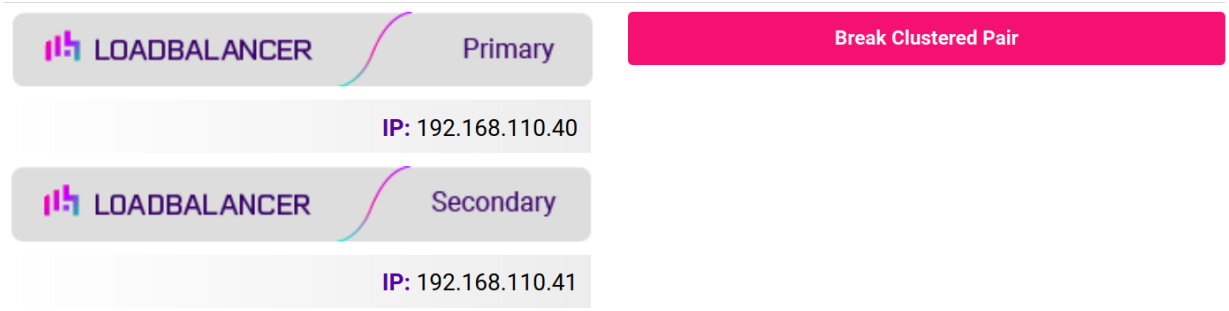

7. To finalize the configuration, restart heartbeat and any other services as prompted in the "Commit changes" message box at the top of the screen.

> **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.

13. Document Revision History

Version	Date	Change	Reason for Change	Changed By
1.0.0	5 November 2019	First draft		RJC
1.1.0	31 March 2010	Extensive document re-write	Revised load balancing methodology and addition of GSLB to handle multi-site scenarios	RJC
1.1.1	28 April 2020	Various minor updates	Required updates	RJC
1.1.2	2 September 2020	Added section describing how to configure 3rd party layer 7 load balancer authorization	New feature of StorageGRID v11.4	RJC
1.1.3	14 October 2020	New title page Updated Canadian contact details	Branding update Change to Canadian contact details	AH
1.1.4	10 February 2021	Updated GSLB configuration steps	GSLB is now configured graphically rather than by editing configuration files	RJC
1.2.0	18 August 2021	Converted the document to AsciiDoc	Move to new documentation system	RJC
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

Version	Date	Change	Reason for Change	Changed By
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
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

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.