LOADBALANCER

MANUALS

Enterprise GCP Configuration Guide

Version 8.13.2 Revision 1.0.0

Table of Contents

1. Introduction

Google Cloud Platform (GCP) is a broad suite of cloud-based services. It allows services to be deployed as and when required. Charges are made for what is used making it an extremely flexible and cost-effective solution. It offers scalable services deployed using a variety of management tools, such as GCP console, CLI and API. The Loadbalancer.org Enterprise GCP cloud based load balancer allows customers to rapidly deploy and configure a feature rich load balancing solution within the GCP cloud environment.

2. About Enterprise GCP

The core software is based on LBOS-7 which is a customized Linux build maintained by Loadbalancer.org, LVS, Ldirectord, Linux-HA, HAProxy & STunnel. At present, Enterprise GCP is available as a single appliance only due to the nature of the network design/constraints. HA (high-availability) clustering may be available in the future. Enterprise GCP is based on the same code base as our main hardware/virtual product. This means that Enterprise GCP supports many of the same features as the hardware & virtual based products. There are certain differences due to the way the GCP environment works. The main differences are listed below.

> **Note**
>
> Currently, Enterprise GCP can have one network interface. This interface has a single primary internal IP address. A single external public IP can be associated with the primary IP to expose load balanced services on the public Internet. Additional load balanced services can be presented on the same IP address using different ports.
>
> Multiple services on different internal IP addresses can be configured using Alias IPs. For a deployment example please refer to Example 2.
>
> If you want to present additional services on different External IPs, a GCP load balancer can be used. For more information and a deployment example, please refer to Example 4.

2.1. Main Differences to Our Standard (Non-Cloud) Product

1. Layer 4 DR mode is currently ***not*** supported.
2. HA (high availability) where a clustered pair of appliances is deployed is currently ***not*** supported.
3. Layer 4 NAT mode where the default gateway on the load balanced real servers is set to be the load balancer is ***not*** supported.
    - Instead, add a tagged route that only applies to traffic from the load balanced servers and set the next hop to be the load balancer – please refer to Example 3 for an example of how VPC routing is modified to route traffic back via the load balancer.
4. Layer 7 SNAT mode with TProxy enabled where the default gateway on the load balanced real servers is required to be the load balancer is ***not*** supported.
    - Instead, add a tagged route that only applies to traffic from the load balanced servers and set the next hop to be the load balancer – please refer to Example 3 for an example of how VPC routing is modified to route traffic back via the load balancer.
5. Enterprise GCP can have only one network interface.

2.2. Why use Enterprise GCP?

- **Comprehensive features** - Enterprise GCP supports a wide range of features:
  - Ability to load balance both GCP based and non-GCP based servers.
  - Customizable timeouts for custom applications.
  - Comprehensive back-end server health-check options.
  - Enables fallback servers to be configured and invoked when all load balanced servers/services fail.
  - Multiple persistence methods.
  - Full integration with Remote Desktop Services Connection Broker.
  - Multiple load balanced services running on multiple IP addresses (requires GCP load balancer).
  - GSLB for multisite load balancing.
  - Fully featured WAF (Web Application Firewall).
- **Ease of use** - The interface is virtually identical to our hardware/virtual product which is very simple and intuitive to use. This also makes migrations to GCP much easier for existing customers.
- **Freedom license** - Our freedom license enables customers to migrate from one environment to another (e.g. virtual to cloud) at no additional cost - and with free migration assistance.
- **Expert assistance is available 24 x 7** - Our highly experienced support team can assist when needed.

3. Accessing GCP

To start using GCP, you'll need a Google account. If you don't already have one you can create one at the following URL: https://cloud.google.com/

4. GCP Management

GCP resources can be managed in various ways:

- GCP Console
- Gcloud CLI
- Gcloud API

4.1. Accessing the GCP Console

The GCP Console can be accessed here.

4.2. GCP CLI & API

- For details on accessing and using the Gcloud CLI click here.
- For details on accessing and using the Gcloud API click here.

5. Deploying Enterprise GCP from the Marketplace

1. Login into the GCP Console.
2. Access the **Marketplace** and search for "Loadbalancer.org", you'll be presented with the following options:
   - **Loadbalancer Enterprise GCP BYOL –** for purchasing & applying your own license
   - **Loadbalancer Enterprise GCP R20 –** hourly billing with up to 5 VIPs / 4 RIPs
   - **Loadbalancer ADC Enterprise GCP –** hourly billing with unlimited VIPs / RIPs

   > **Note** The BYOL version will work completely unrestricted for 30 days without any license applied. During this period, only Google Compute usage charges will apply. After the 30 days, the trial will still function, but no configuration changes will be possible until the license is applied.

3. Click on the option you require, you will be presented with a more detailed overview of the product.
4. Click the **Get Started** button.
5. If you're happy to do so, agree to the **Terms and agreements**.

   > **Note** You only need to agree on the first deployment, subsequent deployments will skip this step.

6. Click the **Launch** button.
7. If prompted to do so, enable the required GCP APIs.
8. The new VM can be provisioned using the Deployment Manager as shown in the steps below, or by using the Command-Line by selecting the second tab.

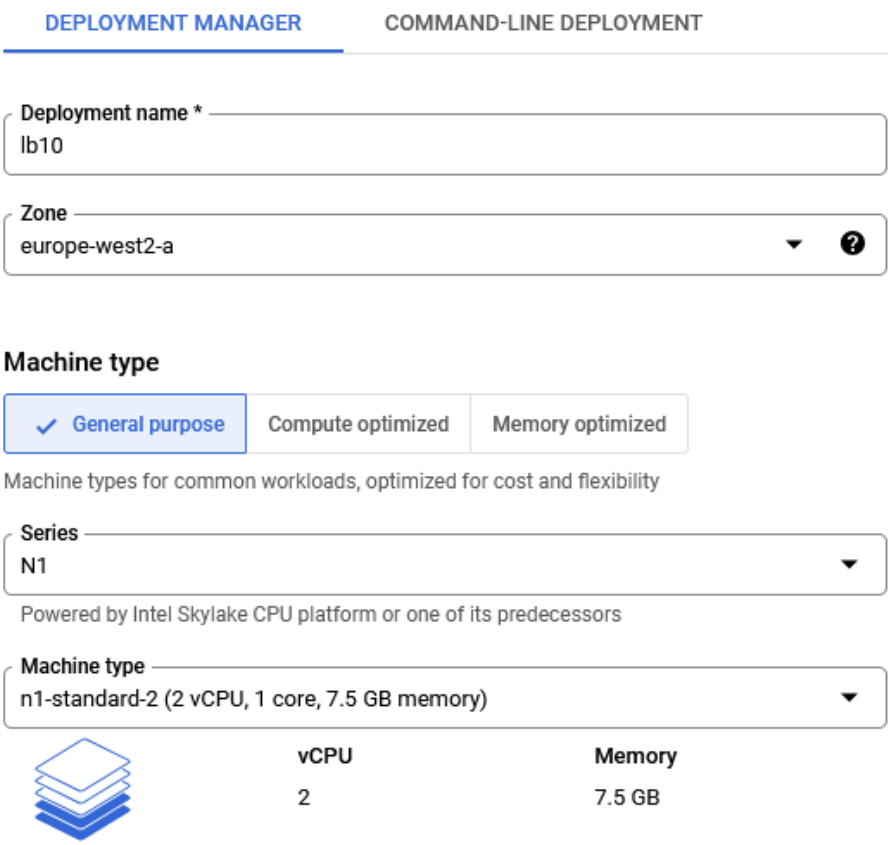

- Enter an appropriate *Deployment name*
- Select the required *Zone*
- Configure the *Machine Type* and associated settings according to your requirements

9. Scroll down to the *Boot disk* section.

- Configure the *Boot Disk* settings according to your requirements, these options can normally be left at their default values

10. Scroll down to the *Networking* section.

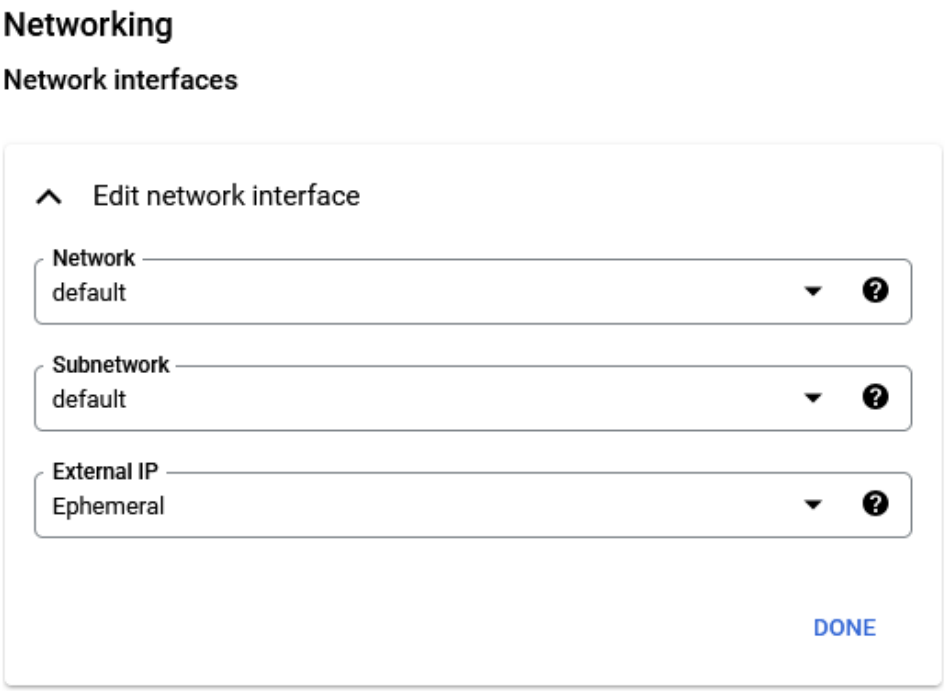

- Configure the ***Networking*** settings according to your requirements

> **Note** The internal IP and the external IP of the load balancer instance can be promoted to static addresses rather than ephemeral after deployment. This is explained in Configuring Static IP Addresses below.

- Click **DEPLOY**, a deployment summary will be displayed:

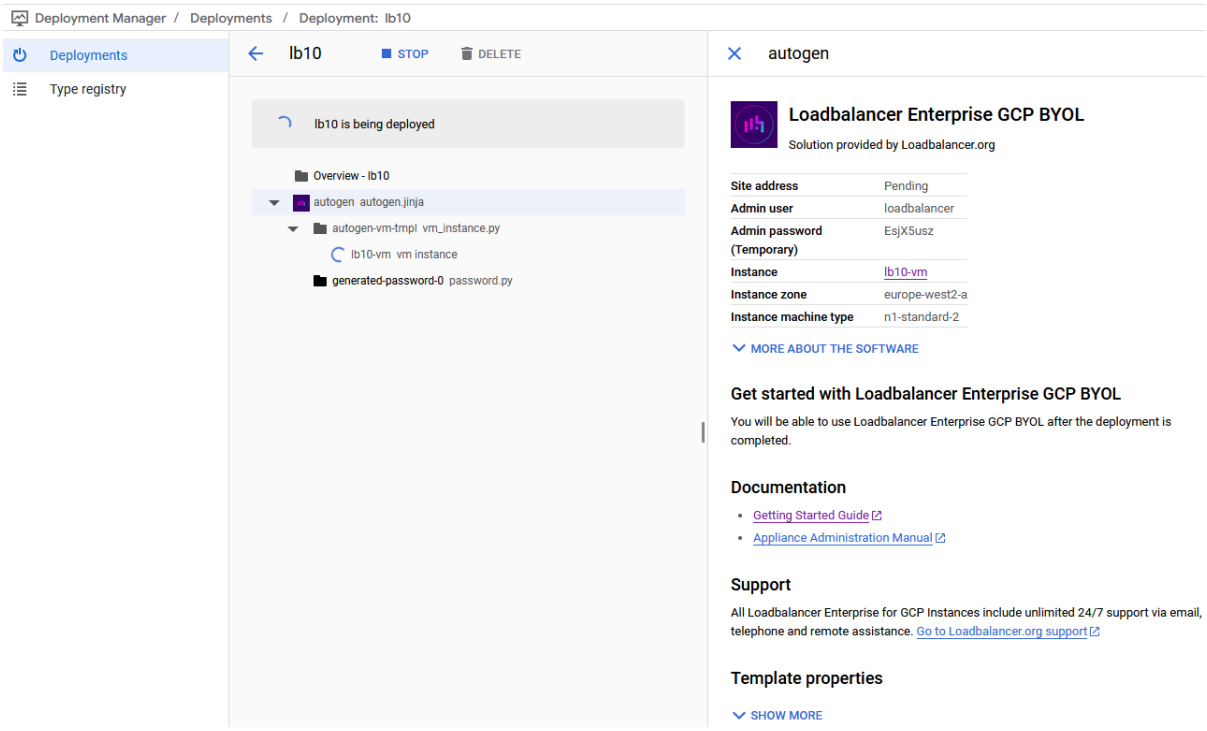

- Once the deployment is complete, the **VISIT THE SITE** button can be used to open the WebUI

- The **Admin User** "loadbalancer" and **Admin password (Temporary)** are used to login to the WebUI, the password should be changed once you logon, for details see Appliance Security below
- An ephemeral external IP is assigned to the appliance, for production deployments this should be changed to a static address, for details see Configuring Static IP Addresses below

> ⚠ Warning
>
> Do **NOT** change the private IP of the appliance using the appliance's WebUI. This will not update the GCP network stack and make the appliance completely unusable. There is currently no recovery from this and you'll need to deploy a new instance.

5.1. Configuring Static IP Addresses

5.1.1. To Reserve a Static Internal IP Address for the Appliance

1. Using the GCP console, edit the instance.
2. Expand the ***Network Interfaces*** section and using the dropdown for the internal IP address, select the promote internal address option.

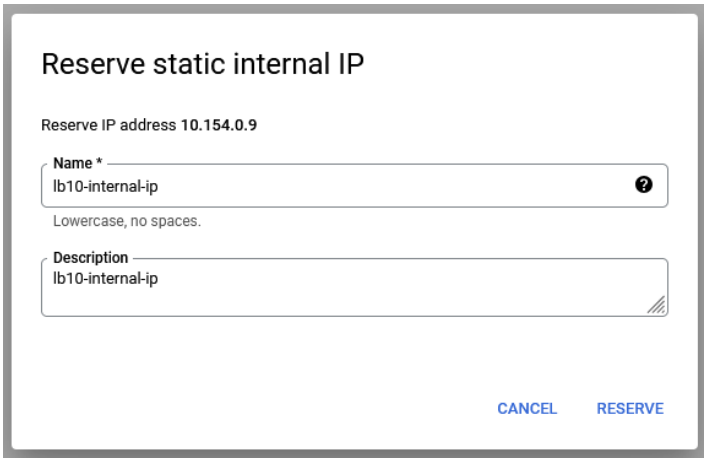

3. In the popup, specify an appropriate ***Name*** & ***Description***.
4. Click **RESERVE**.
5. Scroll to the end of the page and click **Save**.

5.1.2. To Reserve a Static External IP Address for the Appliance

1. Using the GCP console, edit the instance.
2. Expand the ***Network Interfaces*** section and using the dropdown for the external IP address, select the reserve static external address option.

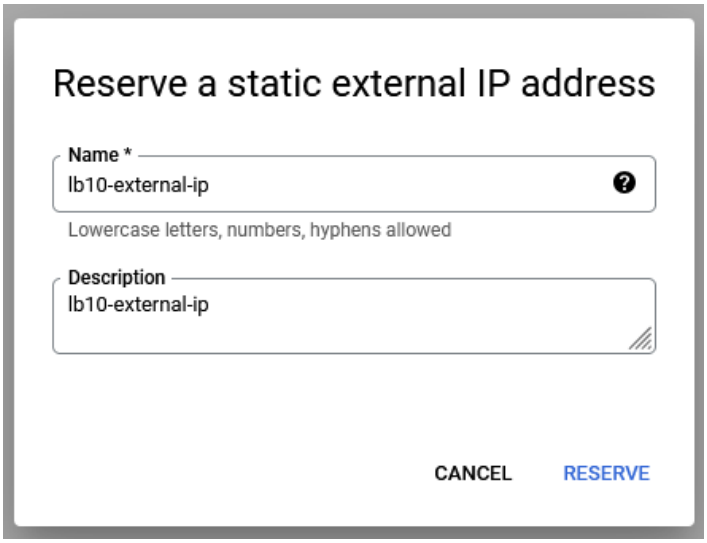

3. In the popup, specify an appropriate *Name* & *Description*.
4. Click **Reserve**.
5. Scroll to the end of the page and click **Save**.

5.2. Configure the GCP Firewall

Ensure that the Firewall policies allow access to the following:

1. The WebUI of the load balancer appliance - by default, this is port TCP/9443.
2. The SSH port of the load balancer appliance - by default this is port TCP/22.
3. The port(s) used to access the load balanced application - in the examples presented in this guide - this is port TCP/80.

6. Accessing the Appliance

6.1. Accessing the Appliance using the WebUI

As mentioned above, you can access the appliance immediately from the deployment screen by clicking the **VISIT THE SITE** button. This will open a new browser window and connect to **https://<instance-public-ip:9443**. Alternatively, open a browser and navigate to the Public IP address on port 9443, i.e.

https://<Public IP Address>:9443
or
https://<FQDN>:9443

> **Note** Google Cloud VPC networks do not automatically support configuring external DNS for a VM.

You'll receive a warning about the certificate as it's a self signed cert not related to an Internet based CA. Confirm you want to continue and a login prompt will be displayed. Use the following default credentials:

Username: loadbalancer

Password: <temporary-password>

The temporary-password is displayed in the deployment summary page after deployment and also in the **custom metadata** section of the VM instance properties in the GCP console:

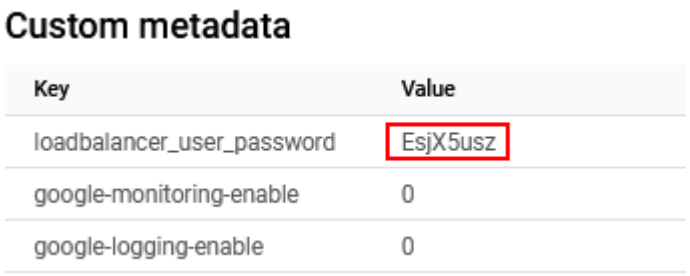

> **Note** To change the password, use the WebUI option: *Maintenance > Passwords*. Changing the password in the appliance will not update the temporary password displayed in the GCP overview.

Once logged in, the WebUI is displayed:

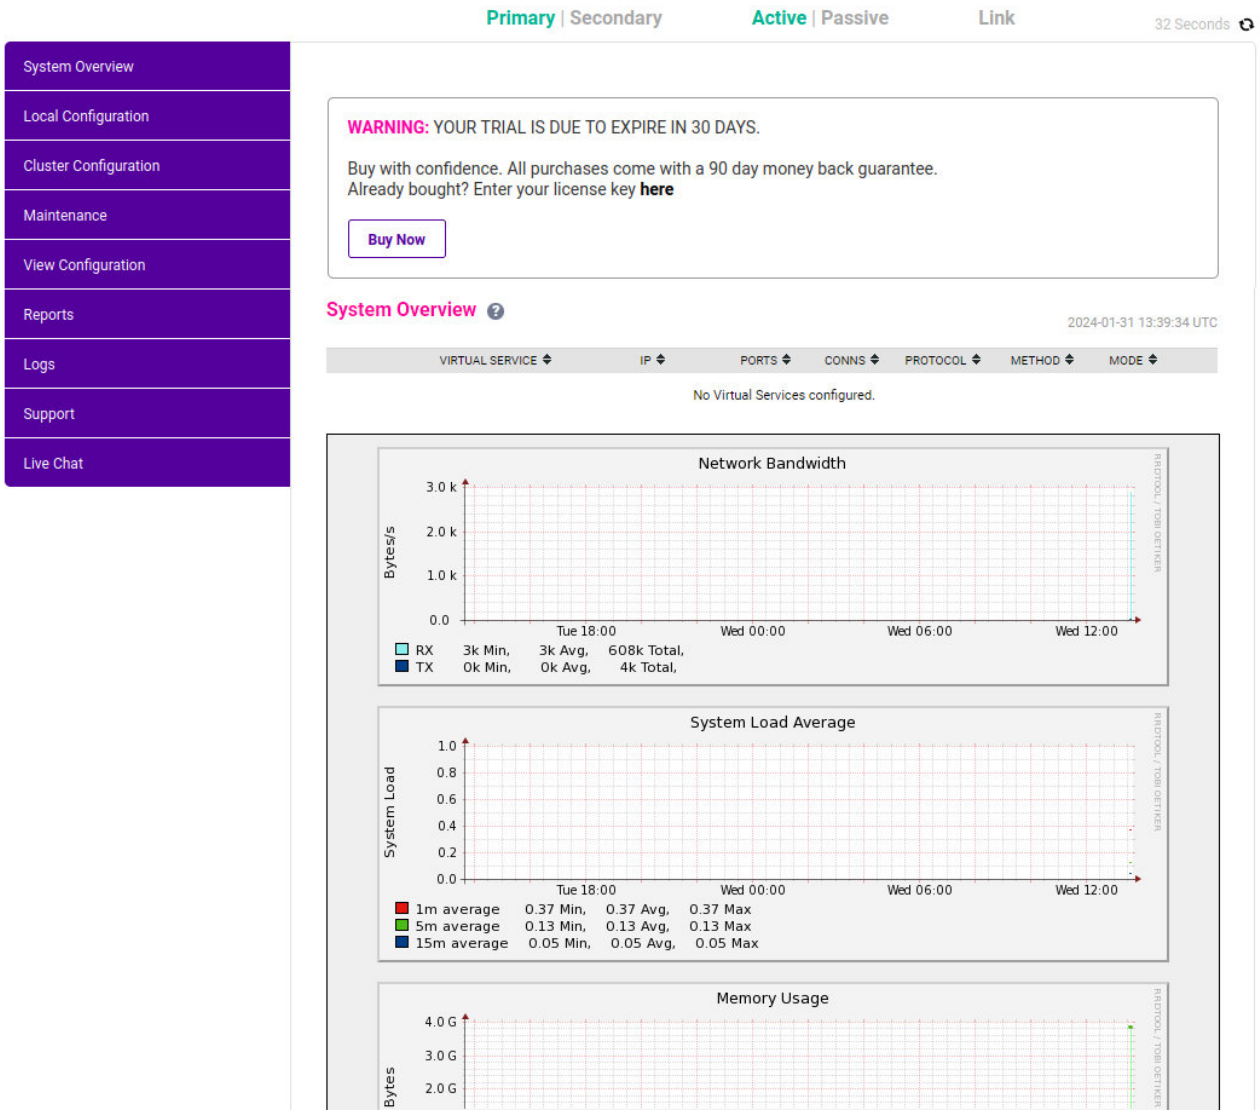

6.1.1. WebUI Menu Options

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

6.2. Appliance Security

> **Note** For full details of each security mode and all other security related features, please refer to Appliance Security Features.

6.2.1. Security Mode

To control how the appliance is accessed and which features are enabled, 3 security modes are provided:

- **Custom** - In this mode the security options can be configured to suit your requirements
- **Secure - (Default)** - In this mode:
  - All console access and SSH password access is disabled
  - WebUI connections are forced to use HTTPS
  - Access to the *Local Configuration > Execute shell command* menu option is disabled
  - The Firewall Script & the Firewall Lockdown Wizard Script cannot be edited
- **Secure - Permanent** - This mode is the same as **Secure** but once set it cannot be changed

> **Important** Setting the security mode to **Secure - Permanent** is irreversible.

To configure the Security Mode:

1. Using the WebUI, navigate to: *Local Configuration > Security*.
2. Select the required *Appliance Security Mode*. If **Custom** is selected, configure the required options.
3. Click **Update**.

6.2.2. Passwords

The loadbalancer WebUI account

As mentioned in Accessing the Appliance using the WebUI the password for the **loadbalancer** WebUI user account is set during instance deployment to a random value. This can be changed using the WebUI menu option: *Maintenance > Passwords*.

The root Linux account

The password for the **root** user Linux account is set to "loadbalancer" by default.

As explained in Security Mode above, root user console & SSH password access are disabled by default. If enabled, the root password can be changed at the console, or via an SSH session using the following command:

```
# passwd
```

6.3. Appliance Software Update

For v8.6.0 and later, the appliance periodically contacts the Loadbalancer.org update server (**update.loadbalancer.org**) and checks for updates. If an update is found, a message will be displayed at the top of the screen as shown in the following example:

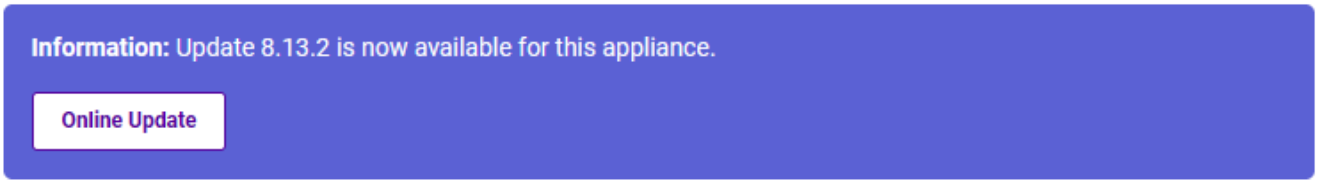

To start the update process click **Online Update**.

The update check can also be initiated manually.

To initiate an online update check:

1. Using the WebUI, navigate to: *Maintenance > Software Update*.
2. Click **Online Update**.

> **Note** Updates are incremental for versions prior to v8.8.0, so repeat the process ignoring calls to restart services or reboot the appliance until you reach v8.8.0. Once at v8.8.0, when the update is next triggered the appliance will be updated to the latest version in a single step.

6.4. Appliance Licensing

If you've deployed the BYOL version of the appliance, by default it runs as a 30 day trial and is completely unrestricted during this time. After 30 days, the appliance continues to work but it's no longer possible to make changes to the configuration. When a license is purchased, you'll be provided with a license key file by our sales team. This must then be installed on your appliance. To install the license, use the WebUI option: *Local Configuration > License Key* to browse to and select the license file provided. Once selected, click **Install License Key** to apply the license. We recommend that you should check for updates ***before*** applying the license key.

6.5. Accessing the Appliance using SSH

When the appliance is deployed, the project's SSH keys are inherited from GCP Compute Engine Metadata, making secure access easier to manage. To SSH into the appliance, you'll need to ensure that the public SSH key is available in the Compute Engines Metadata and that the matching private SSH key is available on the client/application used for SSH access.

More information on managing SSH keys in GCP can be found here.

6.5.1. Generating SSH Keys

The steps below show how to generate SSH key pairs using Linux and Windows.

Using Linux

STEP 1 - Generate a keypair using ssh-keygen

All Distros:

```
# ssh-keygen -q -t rsa -b 2048 -f <output filename>
```

e.g.

```
# ssh-keygen -q -t rsa -b 2048 -f GCPKeys
```

When prompted, enter a pass-phrase, or leave empty for no passphrase:

```
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
```

2 files are created:

- **GCPKeys** – this is the Private Key file and is used on the SSH client machine
- **GCPKeys.pub** – this is the Public Key file, the contents are copied into the *SSH public key* field when the VM is deployed.

Using Windows

STEP 1 – Install PuTTY

1. Download PuTTY from here.
2. Run the installer.

STEP 2 – Use PuTTYgen to generate a Public/Private key pair

1. Browse to the PuTTY program folder and run PuTTYgen.

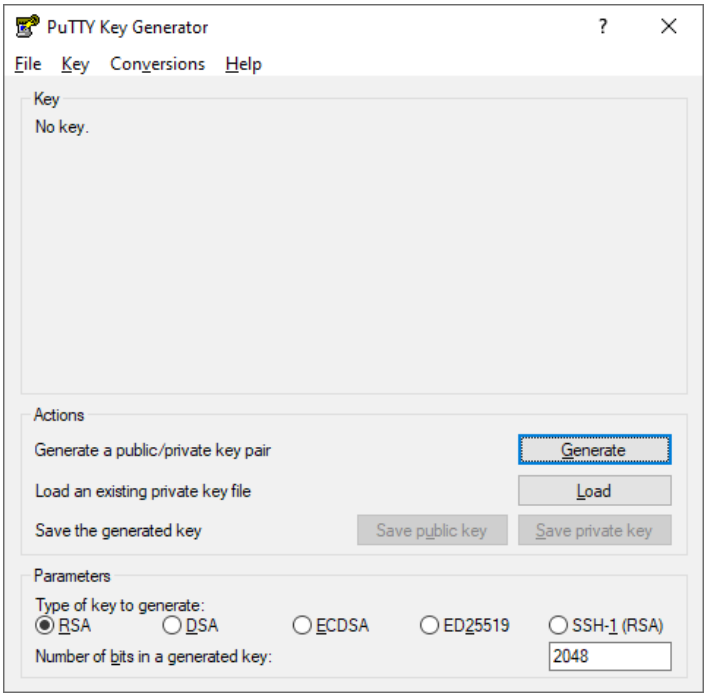

2. Click the **Generate** button.

3. As directed, move the mouse around to create random keys.

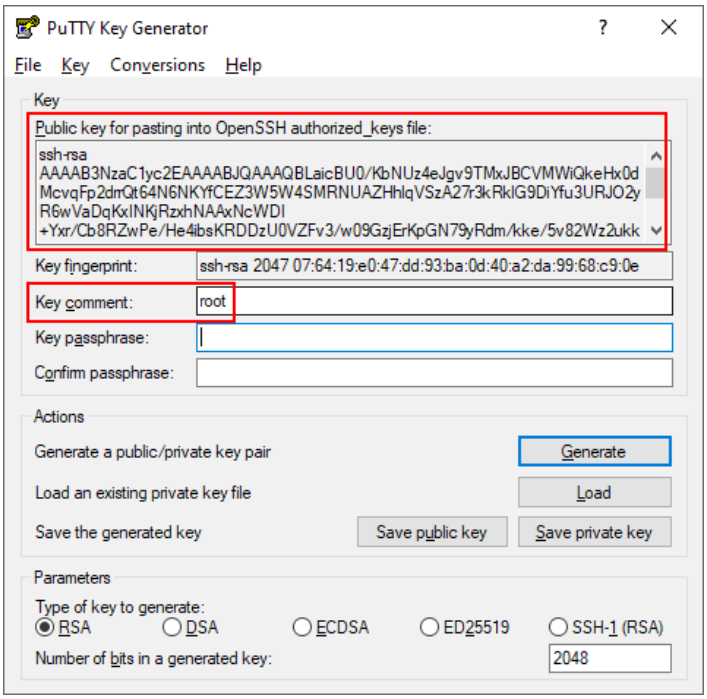

4. Once generated, enter "root" in the ***Key comment*** field as shown above and then copy the public key from the top window and save it to a file.

6.5.2. Add SSH Keys to Project Metadata

Once SSH keys are added, all instances in the project inherit the keys.

To add keys to the project:

1. Access the GCP Console and navigate to ***Compute Engine > Metadata*** and select the ***SSH Keys*** tab.
2. Click **Edit**.
3. Click **Add Item** and paste the contents of the public key file (e.g GCPKeys.pub or the data displayed at the top in PuTTYgen) into the window, as shown in the following example:

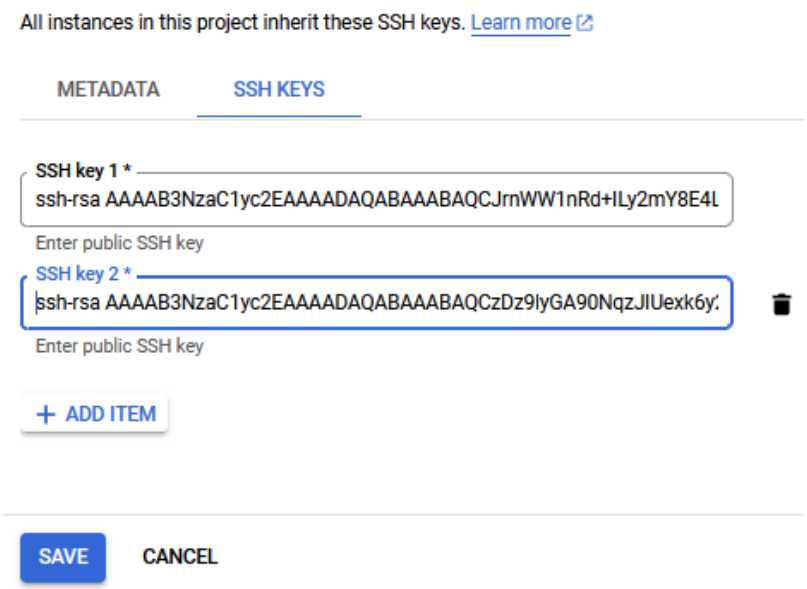

4. Click **Save** – the updated public key will be added to the metadata for the project.

6.5.3. Accessing the Appliance from Linux

Start SSH specifying the private key file and login as the "root" user, e.g.:

Using the IP address:

```
# ssh -i GCPKeys root@1.2.3.4
```

6.5.4. Accessing the Appliance from Windows using PuTTY

1. Run PuTTY.
2. Expand the *SSH* and *Auth* sections and select *Credentials* as shown below:

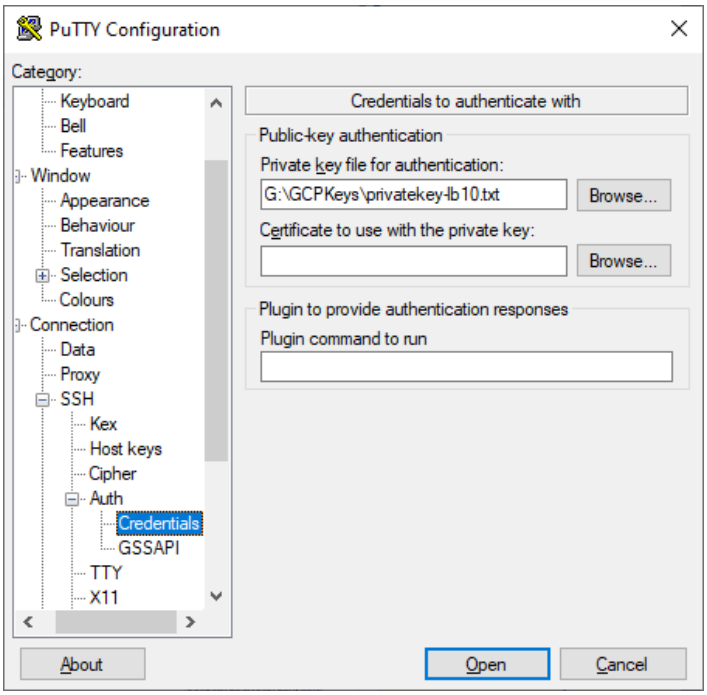

3. Click **Browse** and select the private key created earlier.
4. Click **Open** to start the SSH session.
5. Login as "root", no password should be required.

7. Deployment examples

The following section provides four examples to help illustrate how the load balancer can be deployed. It is important to consider that when configured at layer 7, the load balancer acts as a proxy and is not transparent which means that the source IP address of packets reaching the real servers will be the load balancer's own IP address.

7.1. Example 1 – Web Servers, Layer 7, Public facing

This simple example shows how to configure a layer 7 VIP on the appliance's internal IP address that is accessible via the appliance's external public IP address.

Step 1 - Deploy the GCP instances

1. Deploy the load balancer instance as described in Deploying Enterprise GCP from the Marketplace and configure a static internal and external IP address as described in Configuring Static IP Addresses.
2. Deploy the web server instances to the same VPC as the load balancer and configure a static internal IP address for each as described in Configuring Static IP Addresses.
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80.

Step 2 - Configure the Virtual Service

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

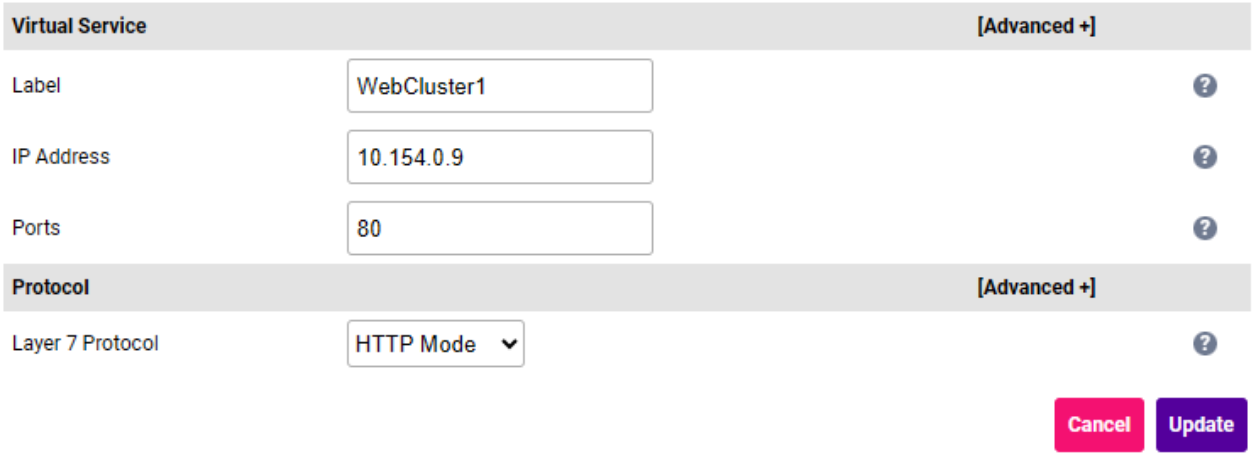

3. Enter an appropriate label for the VIP, e.g. **WebCluster1**.
4. Set the Virtual Service IP address field to the **Base IP address**, e.g. **10.154.0.9**.

5. Set the Virtual Service Ports field to the required port, e.g. **80**.
6. Leave Layer 7 Protocol set to **HTTP Mode**.
7. Click **Update**.

Step 3 - Configure the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

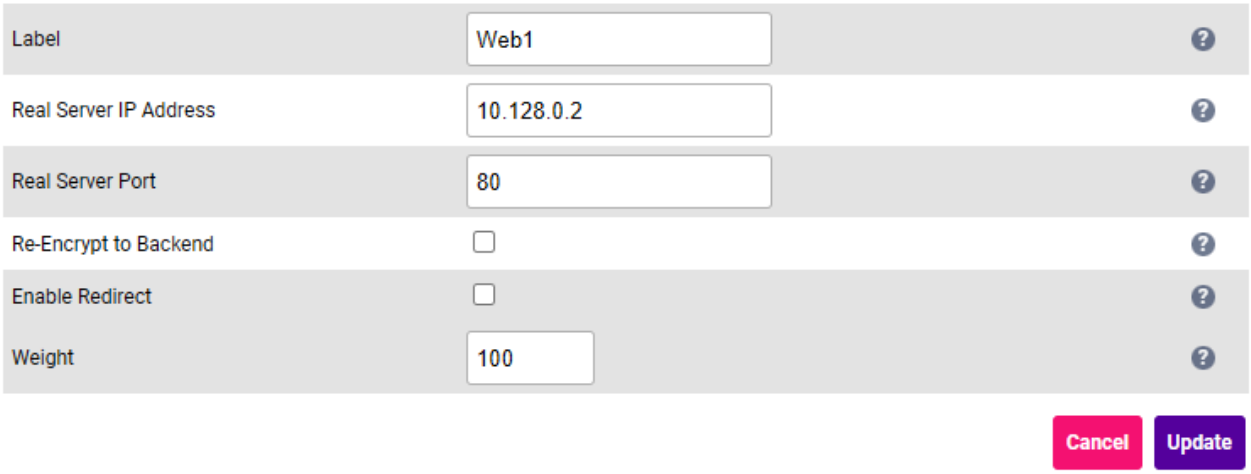

3. Enter an appropriate label for the RIP, e.g. **Web1**.
4. Set the Real Server IP Address field to the required IP address, e.g. **10.128.0.2**.
5. Set the Real Server Port field to the required port, e.g. **80**.
6. Click **Update**.
7. Repeat the above steps to add additional web server(s).

Step 4 - Apply the new Layer 7 Settings

1. Once the configuration is complete, use the **Reload HAProxy** button at the top of the screen to apply the changes.

Step 5 - Testing

1. Connect to the external IP address of the load balancer instance on port 80 to verify that the web page is displayed.

7.2. Example 2 – Web Servers, Layer 7, Internal Facing (Using an Alias IP)

This example shows how a VIP can be configured on an **Alias IP**. This allows multiple internal/private VIPs to be configured on different IP addresses.

Step 1 - Deploy the GCP Instances

1. Deploy the load balancer instance as described in Deploying Enterprise GCP from the Marketplace and configure a static external IP address as described in Configuring Static IP Addresses.
2. Deploy the web server instances to the same VPC as the load balancer and configure a static internal IP address for each as described in Configuring Static IP Addresses.
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80.

Step 2 - Configuring the Alias IP

1. Using the GCP console, edit the load balancer instance.
2. Scroll down to the *Networking Interfaces* section and expand the network interface section.
3. Scroll down to the *Alias IP ranges* section and click **ADD IP RANGE**.
4. Click **Show alias IP ranges**.

5. Enter the required Alias IP – this will be used for the new VIP.
6. Click **Done**, click **Save**.

Step 3 - Configure the Virtual Service

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

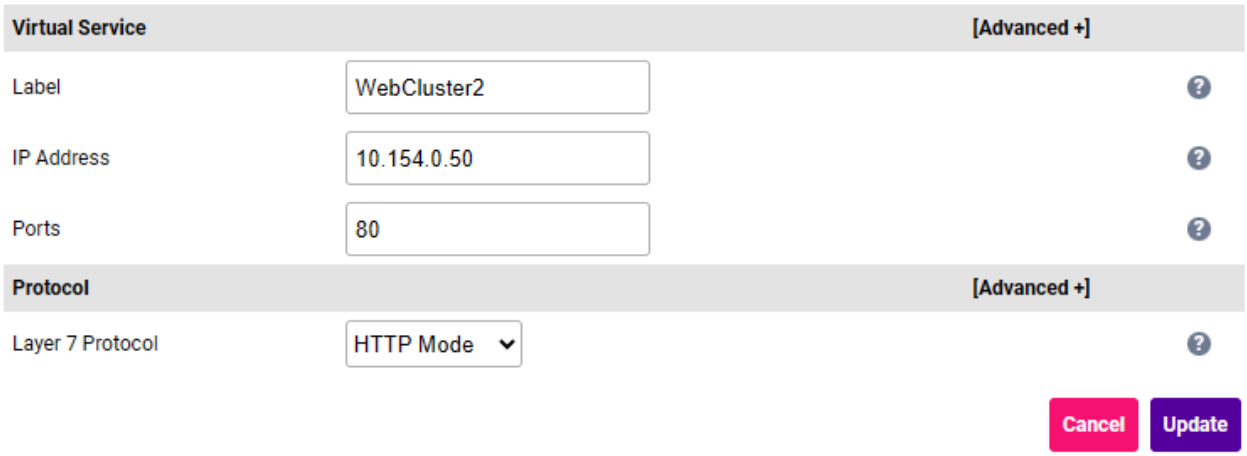

3. Enter an appropriate label for the VIP, e.g. **WebCluster2**.

4. Set the Virtual Service IP address field to the **Alias IP address**, e.g. **10.154.0.50**.
5. Set the Virtual Service Ports field to the required port, e.g. **80**.
6. Leave Layer 7 Protocol set to **HTTP Mode**.
7. Click **Update**.

Step 4 - Configure the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

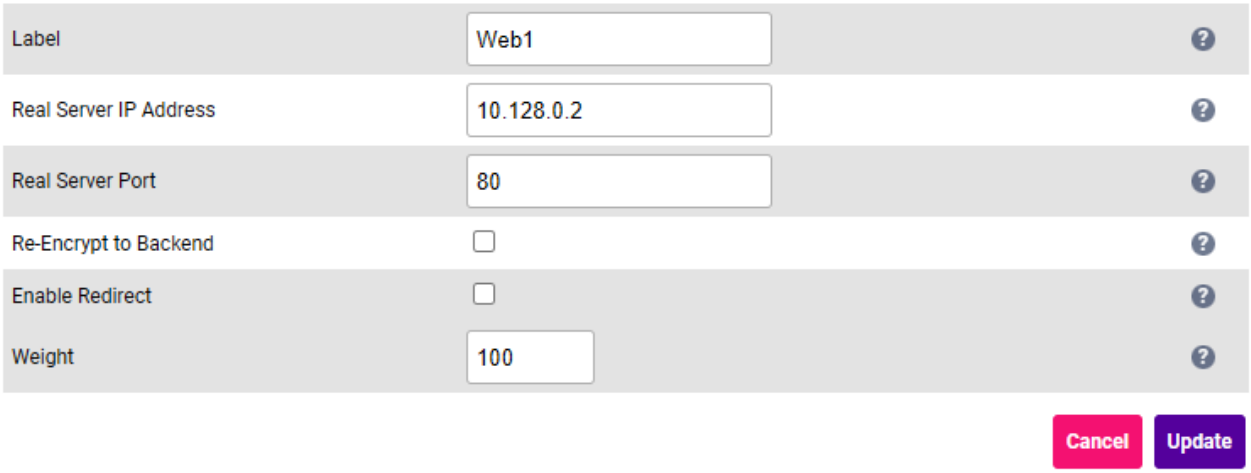

3. Enter an appropriate label for the RIP, e.g. **Web1**.
4. Set the Real Server IP Address field to the required IP address, e.g. **10.128.0.2**.
5. Set the Real Server Port field to the required port, e.g. **80**.
6. Click **Update**.
7. Repeat the above steps to add additional web server(s).

Step 5 - Apply the new Layer 7 Settings

1. Once the configuration is complete, use the **Reload HAProxy** button at the top of the screen to apply the changes.

Step 6 - Testing

1. Connect to the IP Alias address of the load balancer instance on port 80 to verify that the web page is displayed.

7.3. Example 3 – Web Servers, Layer 4, Public Facing

In this example, VPC routing is modified to route return traffic from the web servers via the load balancer - this is a requirement for layer 4 NAT mode. This is achieved by adding a route for all destination IP ranges (0.0.0.0/0), setting the next hop as the load balancer and using tags to ensure this route only applies to the load balanced

web servers.

Step 1 - Deploying the GCP Instances

1. Deploy the load balancer instance as described in Deploying Enterprise GCP from the Marketplace and configure a static internal and external IP address as described in Configuring Static IP Addresses.
2. Deploy the web server instances to the same VPC as the load balancer and configure a static internal IP address for each as described in Configuring Static IP Addresses.
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80.

Step 2 - Enable IP Forwarding for the Load balancer VM

Use the following GCP CLI command to retrieve the current VM settings:

```
gcloud compute instances export INSTANCE_NAME --project PROJECT_ID --zone ZONE
--destination=FILE_PATH
```

e.g.

```
gcloud compute instances export lb10-vm --project lbtesting --zone europe-west2-a
--destination=C:\GCP\lb10-vm.txt
```

Next, edit the output file, in this case *lb10-vm.txt* and change **canIpForward: false** to **canIpForward: true**

Now use the following command to push the updated settings back to the VM:

```
gcloud compute instances update-from-file INSTANCE_NAME --project PROJECT_ID --zone ZONE
--source=FILE_PATH --most-disruptive-allowed-action REFRESH
```

e.g.

```
gcloud compute instances update-from-file lb10-vm --project lbtesting --zone europe-west2-a
--source=C:\GCP\lb10-vm.txt --most-disruptive-allowed-action REFRESH
```

Step 3 - Configure GCP Routing

1. Using the GCP console, select **Routes** in the VPC Network menu.
2. Select the *ROUTE MANAGEMENT* tab.
3. Click **CREATE ROUTE**.

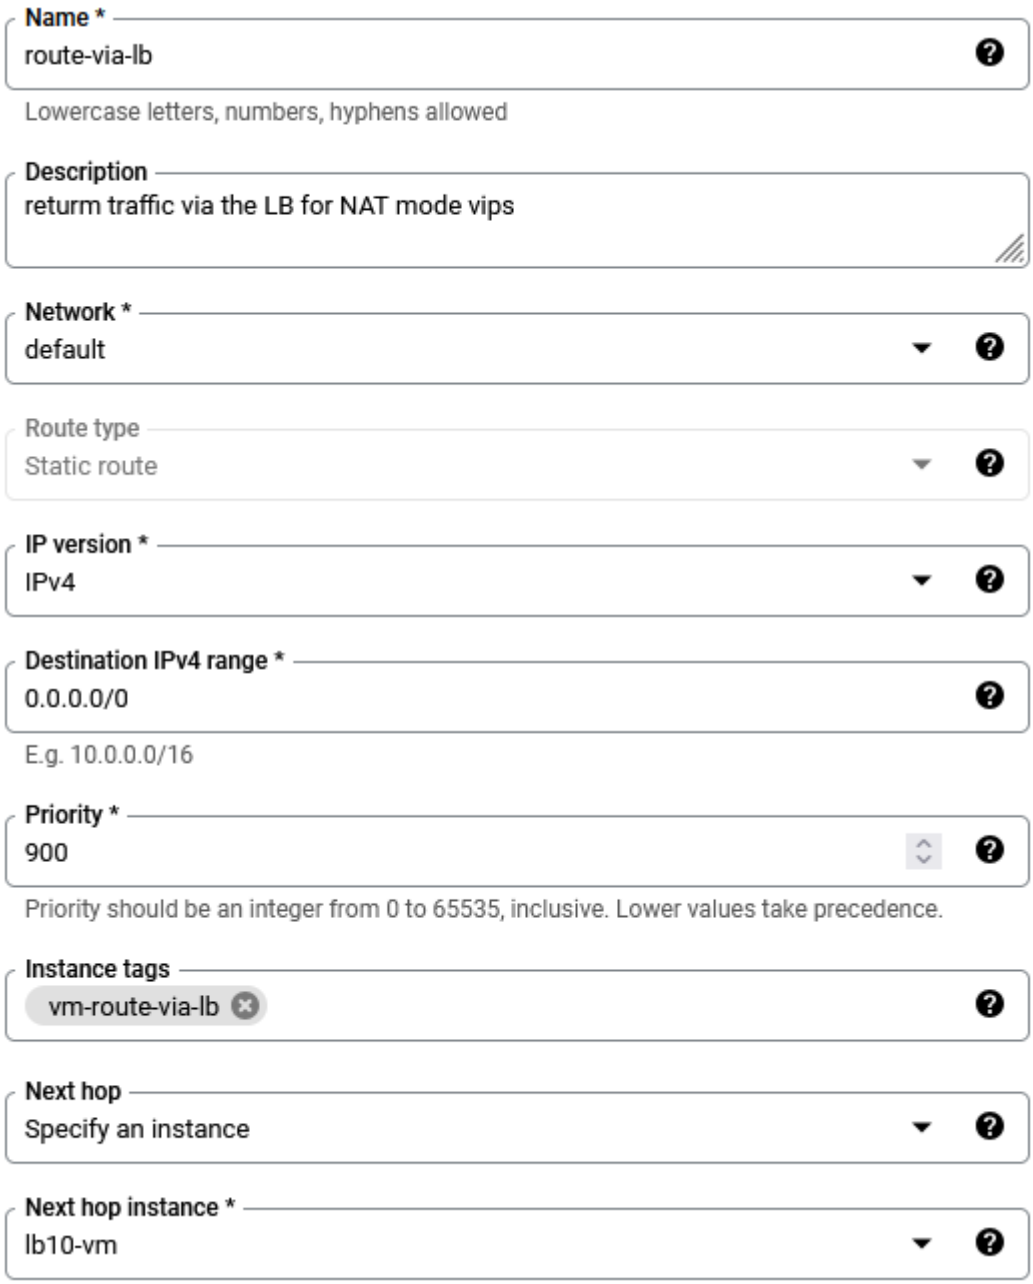

- Enter an appropriate *Name*, e.g. **route-via-lb**.
- Select the relevant *Network*, e.g. **default**.
- Set the *Destination Ip Range* to **0.0.0.0/0**.
- Set the priority higher (lower number) than the default 0.0.0.0/0 route, e.g. **900**.
- Enter an *Instance tag* to filter which instances the rule applies to, e.g. **vm-route-via-lb**.
- Set *Next Hop* to **Specify an instance**.
- Set the *Next hop instance* to the load balancer.
- Click **CREATE**.

Step 4 - Tag the Web Server instances

1. Using the GCP console, Edit each web server instance.
2. In the *Network tags* section add the same tag used for the route above, e.g. **vm-route-via-lb**.
3. Click **Save**.

Step 5 - Configure the Virtual Service

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Virtual Services* and click **Add a New Virtual Service**.
2. Enter the following details:

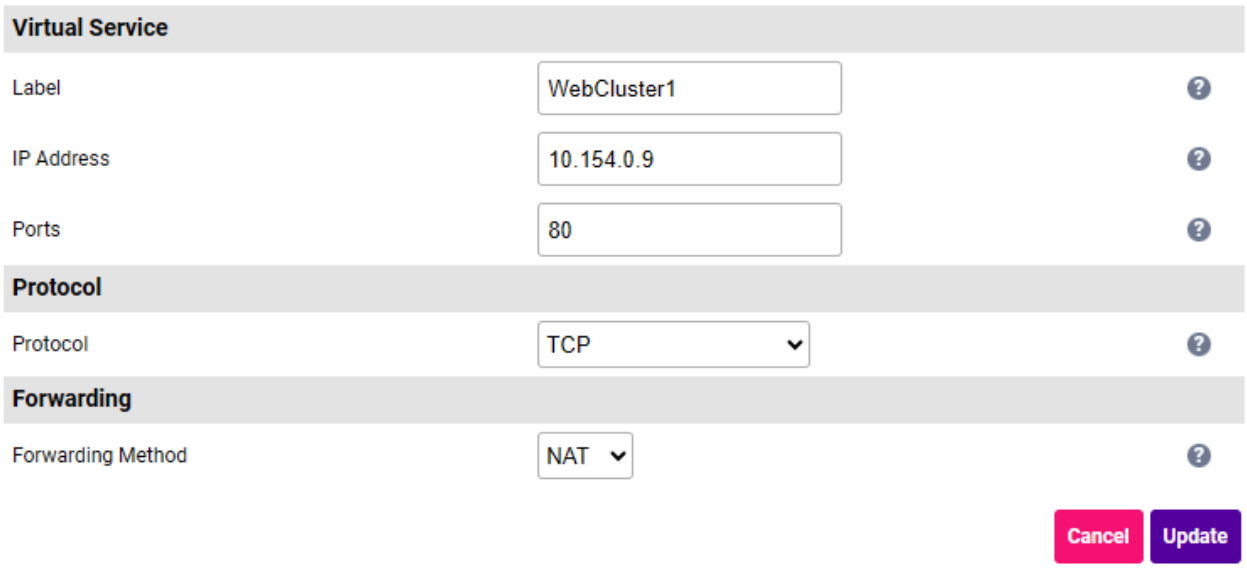

3. Enter an appropriate *Label* for the VIP, e.g. **WebCluster1**.
4. Set the *IP address* to the internal IP address of the load balancer, e.g. **10.154.0.9**.
5. Set the *Ports* to the required port, e.g. **80**.
6. Set Protocol to **TCP**.
7. Set Forwarding Method to **NAT**.
8. Click **Update**.

Step 6 - Configure the Real Servers

1. Using the WebUI, navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.
2. Enter the following details:

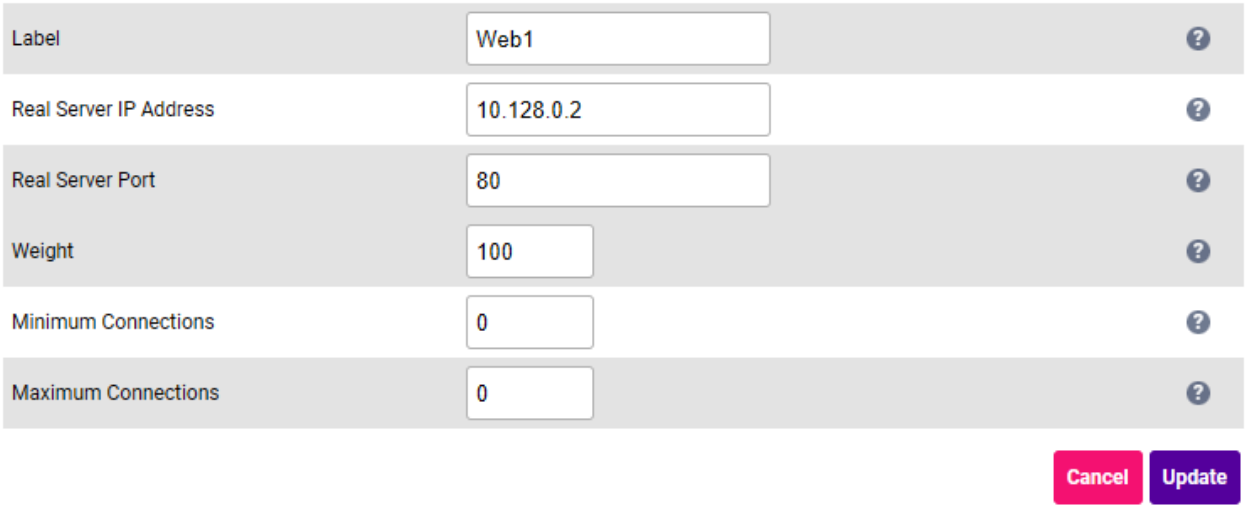

3. Enter an appropriate *label* for the RIP, e.g. **Web1**.
4. Set the *Real Server IP Address* to the required IP value, e.g. **10.128.0.2**.
5. Set the *Real Server Port* to the required port, e.g. **80**.
6. Click **Update**.
7. Repeat the above steps to add additional web server(s).

Step 7 - Test the VIP

1. Connect to the external IP address of the load balancer instance on port 80 to verify that the web page is displayed.

7.4. Example 4 - Web Severs, Layer 7, on Multiple Public IPs (Using GCP Load Balancer)

This example shows how a GCP load balancer can be used to present multiple services on different public IP addresses. The first service can be presented in the load balancer's public IP address, other services are presented via the GCP load balancer.

The following services will be configured:

- **Service-1 -** Web Cluster 1 presented on public IP 1 (the External IP of the load balancer instance)
- **Service-2 -** Web Cluster 2 presented on public IP 2 (additional External IP reserved for service-2)

Step 1 - Deploy the GCP instances

1. Deploy the load balancer instance as described in Deploying Enterprise GCP from the Marketplace and configure a static internal and external IP address as described in Configuring Static IP Addresses.
2. Deploy the web server instances to the same VPC as the load balancer and configure a static internal IP address for each as described in Configuring Static IP Addresses.
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80.

Step 2 - Configure a GCP Load Balancer

a) Create an Instance Group

An instance group containing the load balancer VM must be created. This is used in the GCP load balancer definition when the backend is created.

1. In the GCP Console navigate to *Compute Engine > Instance Groups*.
2. Click **CREATE INSTANCE GROUP**.

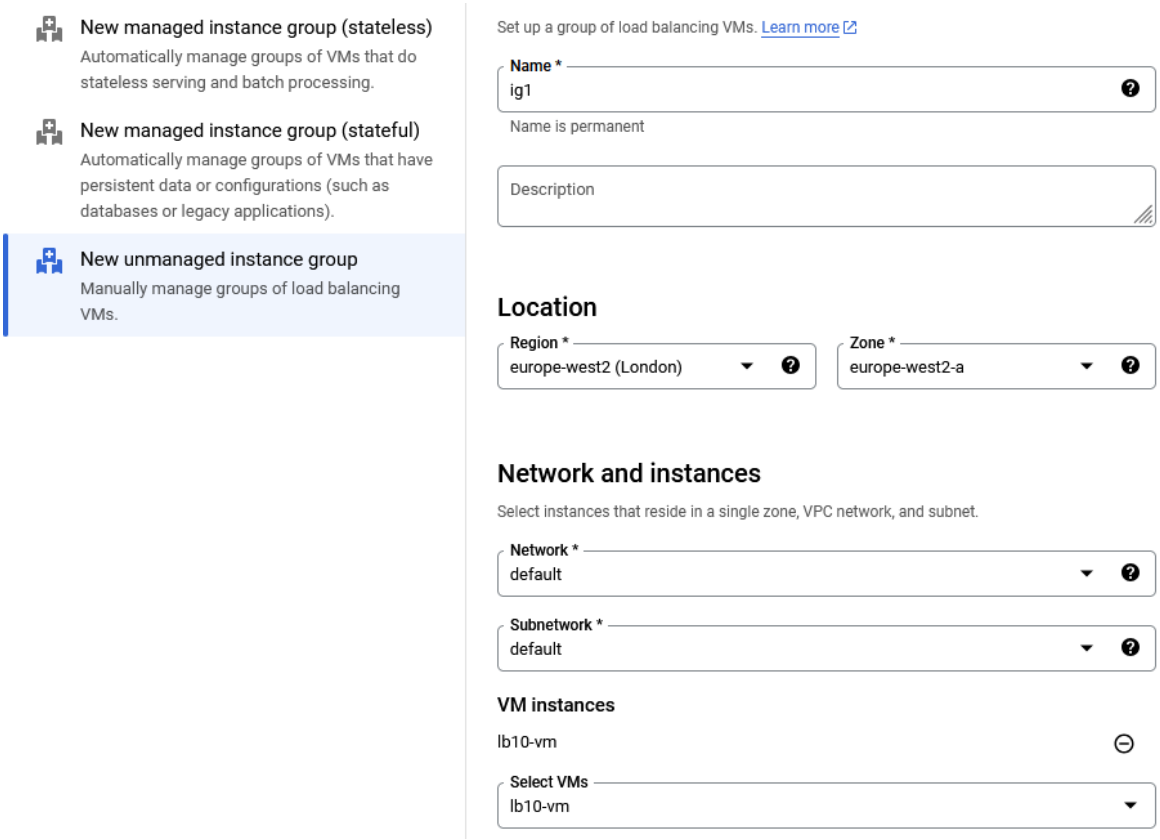

3. Select *New unmanaged instance group*.
4. Enter a suitable *Name*.
5. Select the required *Region* and *Zone*.
6. Select the required *Network* and *Subnetwork*.
7. Using the *Select VMs* dropdown in the *VM Instances* section, select the load balancer VM.
8. Click **Create**.

b) Create a Health Check

The health check is used by the GCP load balancer.

> **Important** Make sure that you modify the ingress firewall rules to allow the health checks to have the required access to the load balancer VM. For details, see Probe IP ranges and firewall rules.

1. In the GCP Console navigate to *Compute Engine > Health checks*.
2. Click **CREATE A HEALTH CHECK**.

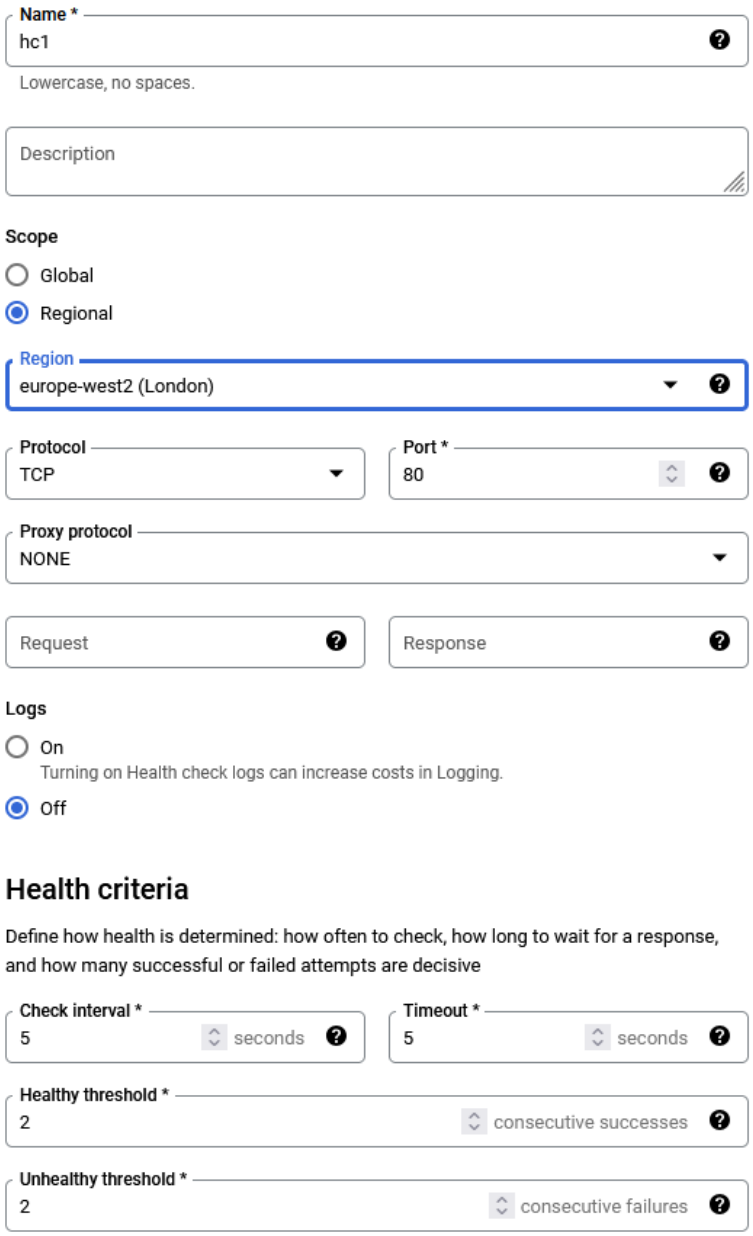

3. Enter a suitable *Name*.
4. Select the required *Region*.
5. Set the *Protocol* to **TCP** and the *Port* to **80**.
6. Leave all other settings at their default value.
7. Click **CREATE**.

c) Add the GCP load balancer

1. In the GCP Console navigate to *Network Services > Load balancing*.
2. Select **Network Load Balancer** and click **NEXT**.

3. Select **Passthrough load balancer** and click **NEXT**.
4. Select **Public Facing (external)** and click **NEXT**.
5. Click **CONFIGURE** to create the load balancer.

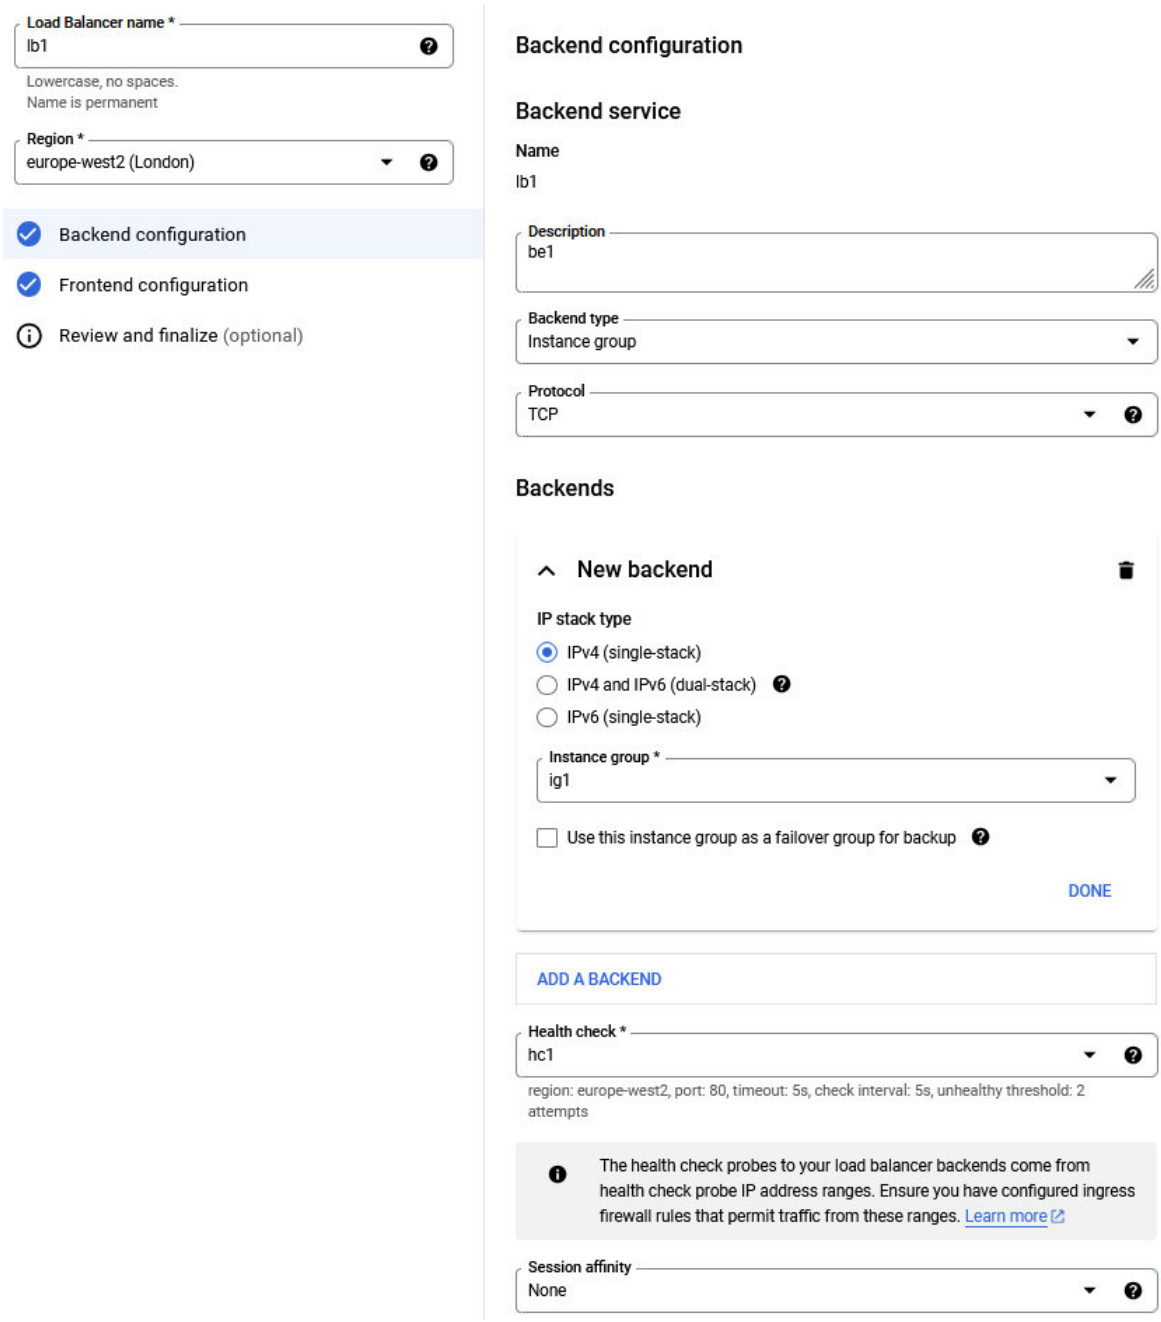

6. Enter a suitable *Load Balancer name*.
7. Select the required *Region*.

Configure the Backend

1. Click *Backend Configuration*.
2. Enter a suitable *Description*.
3. Set the *Backend type* to **Instance Group**.

4. Set the *Protocol* to **TCP**.
5. Set the required *IP Stack Type*.
6. Select the *instance group* created previously.
7. Select the *health check* created previously.
8. Leave all other settings at their default value.

Configure the Frontend

1. Click *Frontend Configuration*.

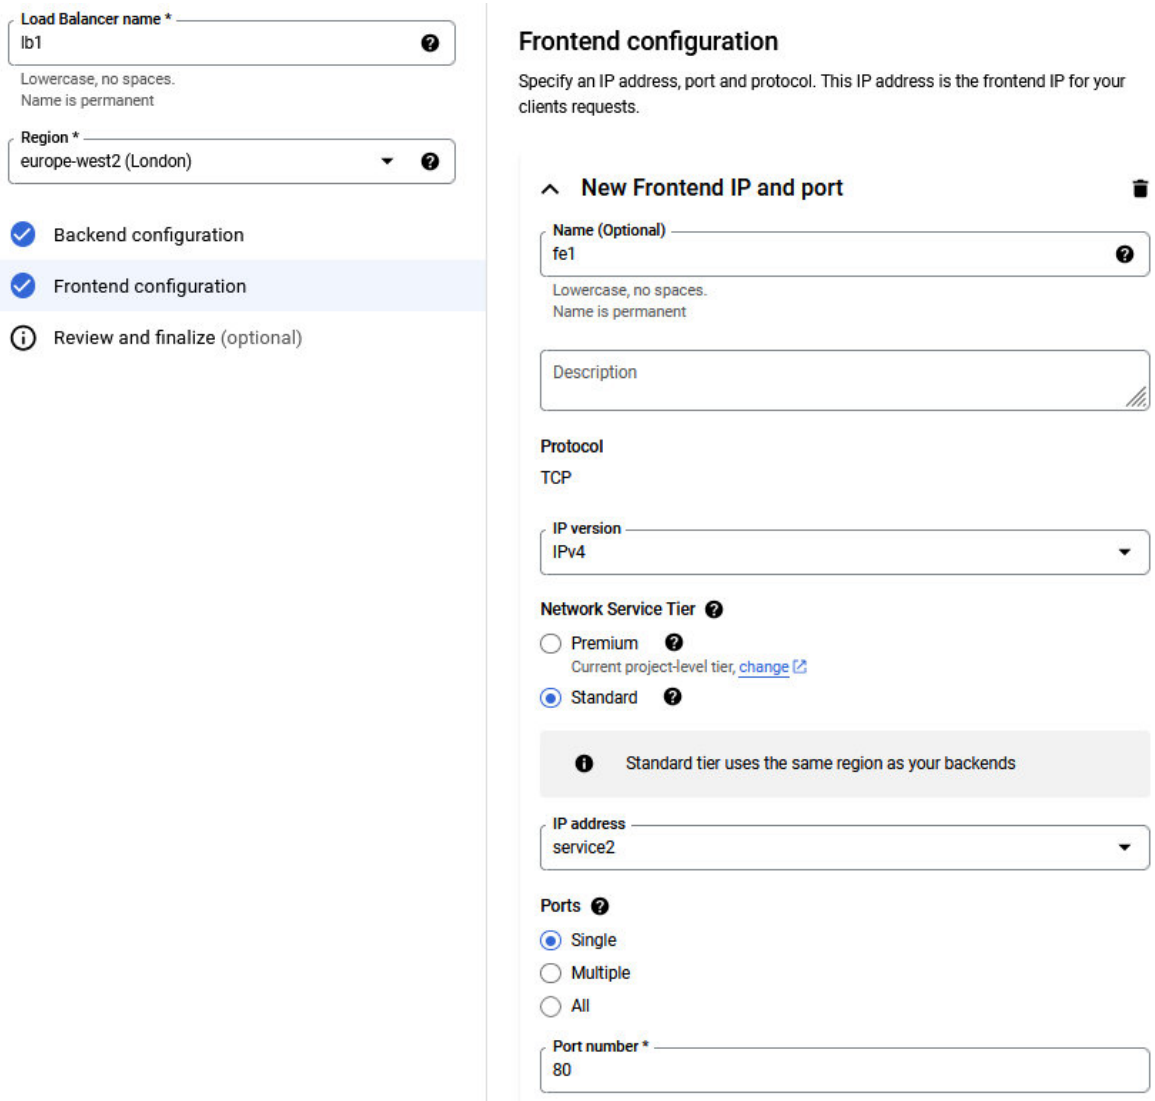

2. Enter a suitable *Name*.
3. Set the required *IP version*.
4. Select the required *Network Service Tier*.
5. Using the *IP address* dropdown, select *CREATE IP ADDRESS*.
6. Enter a suitable name and click **RESERVE**.

7. In the *Ports* section, select single and set the *Port number* to **80**.
8. Click **CREATE**.

Step 3 - Configure the Virtual Services & Real Servers

1. Connect to the Loadbalancer.org appliance WebUI.

Service-1

1. Navigate to *Cluster Configuration > Layer 7 Virtual Services* and click **Add a New Virtual Service**.

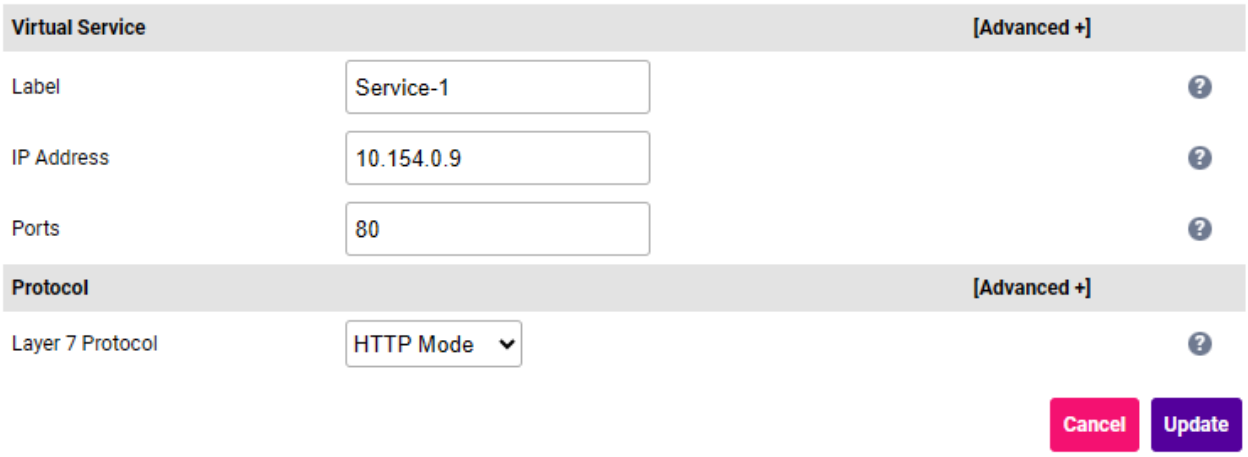

2. Enter an appropriate *label*, e.g. **Service-1**.
3. Enter the appliance's Internal IP address, e.g. **10.154.0.9**.
4. Set the *Ports* field to **80**.
5. Leave *Layer 7 Protocol* set to **HTTP Mode**.
6. Click **Update**.
7. Navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.

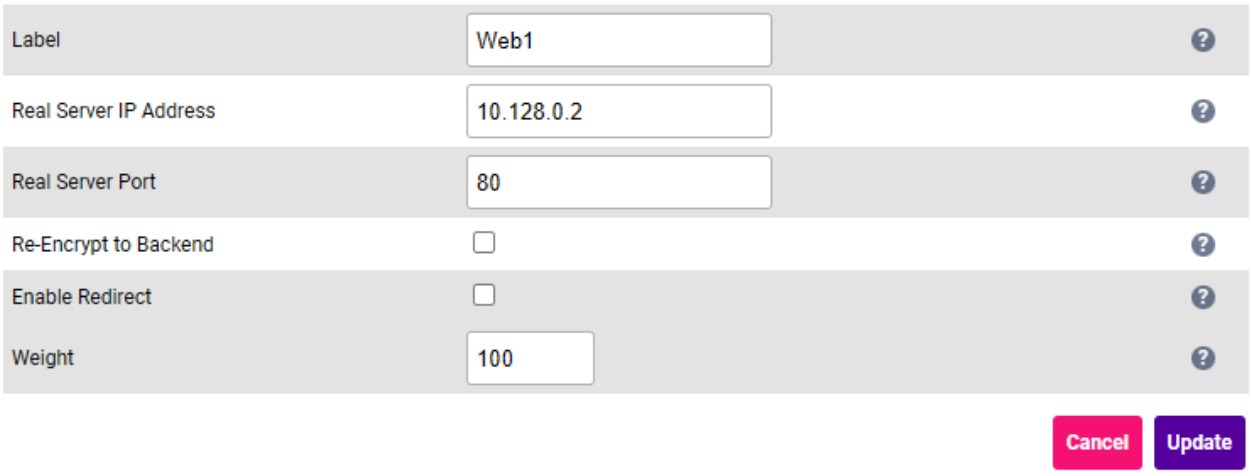

8. Enter an appropriate *label*. e.g. **Web1**.
9. Set the *Real Server IP Address* field to the required IP address, e.g. **10.128.0.2**.
10. Set the *Real Server port* field to **80**.
11. Click **Update**.
12. Repeat the above steps to add additional real server(s).

Service-2

1. Navigate to *Cluster Configuration > Layer 7 Virtual Services* and click **Add a New Virtual Service**.

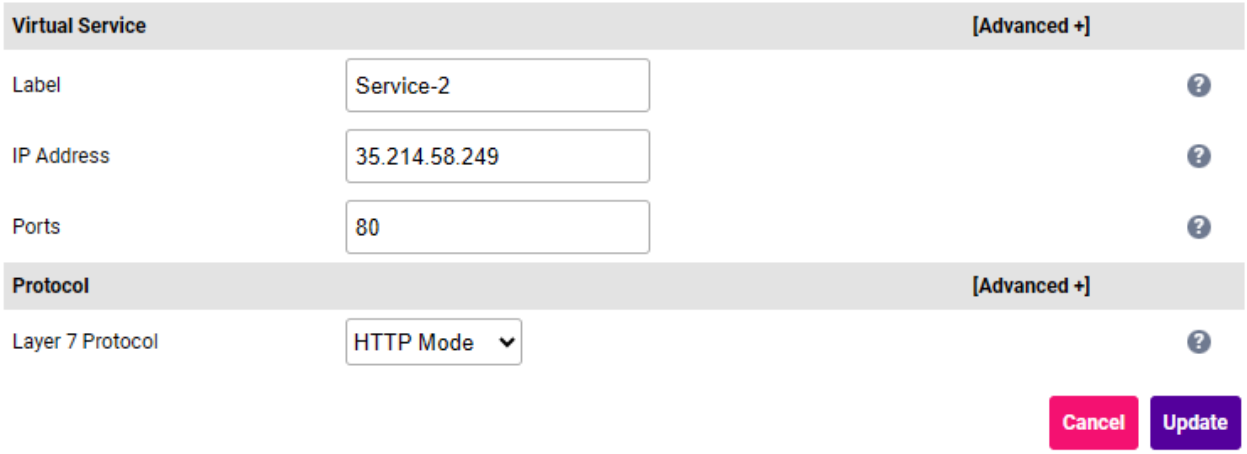

2. Enter an appropriate *label*, e.g. **Service-2**.
3. Enter the External IP that was reserved for Service-2 when setting up the frontend of the GCP load balancer, e.g. **35.214.58.249**.
4. Set the *Ports* field to **80**.
5. Leave *Layer 7 Protocol* set to **HTTP Mode**.
6. Click **Update**.
7. Navigate to: *Cluster Configuration > Layer 7 – Real Servers* and click **Add a new Real Server** next to the newly created VIP.

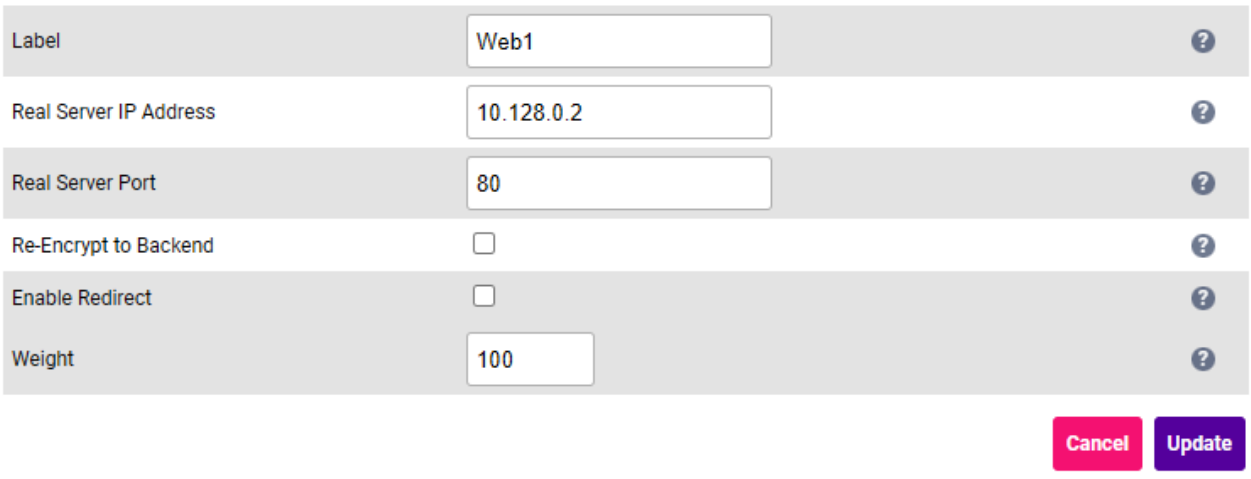

8. Enter an appropriate *label*. e.g. **Web1**.
9. Set the *Real Server IP Address* field to the required IP address, e.g. **10.128.0.2**.
10. Set the *Real Server port* field to **80**.
11. Click **Update**.
12. Repeat the above steps to add additional real server(s).

Step 4 - Apply the new Layer 7 Settings

1. Once the configuration is complete, use the **Reload HAProxy** button at the top of the screen to apply the changes.

Step 5 - Verify the VIPs

The VIPs will be displayed in the *System Overview* of the WebUI. All should be green indicating that all real servers are passing the health checks:

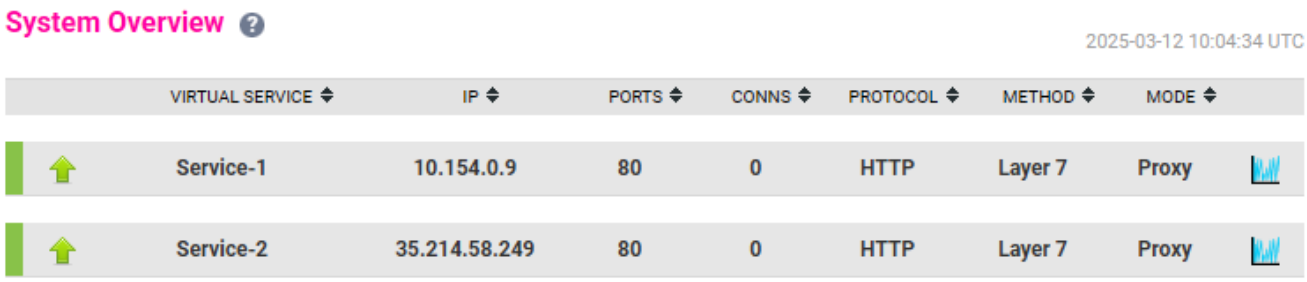

Step 6 - Test the VIPs

1. Verify that **service-1:80** is available on the external IP of the Loadbalancer.org appliance.
2. Verify that **service-2:80** is available on the IP address reserved for Service-2.

8. Testing & Verification

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

9. More Information

Please refer to our website for all the latest Manuals and Deployment Guides.

10. Loadbalancer.org Technical Support

Our highly experienced Support Engineers are on hand to help 24 hours a day, 365 days a year.

10.1. Contacting Support

If you have any questions regarding the appliance or need assistance with load balancing your application, please

don’t hesitate to contact support@loadbalancer.org.

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.