loadbalancer.org

Enterprise GCP Quick Start Guide

v8.4.3

Rev. 1.0

Table of Contents

1. Introduction....3
2. About Enterprise GCP....3
   - Main Differences to Our Standard (Non-Cloud) Product....3
   - Why use Enterprise GCP?....3
3. Accessing GCP....4
4. GCP Management....4
   - Accessing the GCP Portal....4
   - GCP CLI & GCP API....4
5. Deploying Enterprise GCP from the Marketplace....4
6. Accessing the Appliance....7
   - Accessing the Appliance using the WebUI....8
     - WebUI Menu Options....9
     - Appliance Security....9
   - Checking For Updates....10
   - Appliance Licensing....10
   - Accessing the Appliance using SSH....10
     - Generating SSH Keys....11
     - Update the GCP project metadata....12
     - Accessing the Appliance from Linux....13
     - Accessing the Appliance from Windows using PuTTY....13
     - Accessing the Appliance with the Normal User Credentials....14
7. Deployment examples....14
   - Example 1 – Web Servers: 1 Subnet, 1 Load Balancer Network Interface, Layer 7....14
   - Example 2 – Microsoft Session Host: 1 Subnet, 1 Load Balancer Network Interface, Layer 7....15
8. GCP specific configuration....16
   - Firewall Configuration Notes....16
   - To Create a New Firewall Rule....16
   - Terminating Public IP Addresses Utilizing the GCP Load Balancer....19
     - To Add a GCP Layer 4 Load Balancer....19
     - Apply the Changes to the Loadbalancer.org Appliance....21
9. Testing – General Comments....21
   - Testing Load Balanced Services....21
     - Diagnosing VIP Connection Problems....22
     - Taking Real Servers Offline....23
     - Using Reports & Log Files....23
10. More Information....23
11. Loadbalancer.org Technical Support....23
12. Company Contact Information....24

1. Introduction

Google Cloud Platform (GCP) is a broad suite of cloud-based services running on the same global infrastructure as Googles popular end-user products, such as Search and YouTube. 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 from a variety of management tools, such as WebUI 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 customized versions of Centos 6.x/RHEL 6.x, Linux 4.9.x, LVS, HA-Linux, HAProxy, Pound, STunnel & Ldirectord. 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:
>
> At present, Enterprise GCP can only have a single IP address, so all work-load and management services have to be accessed via the same IP. However, with the utilization of GCP L4 load balancers as front ends to terminate public IP addresses, multiple IP addresses can be attached to the Loadbalancer.org appliance this way. This is discussed in section 8. GCP Specific Configuration.

MAIN DIFFERENCES TO OUR STANDARD (NON-CLOUD) PRODUCT

- Layer 4 DR mode is currently not supported.
- At present, HA (high availability) is not supported due to the nature of Google clouds network design.
- Layer 4 NAT mode where the default gateway on the load balanced real servers is required to be the load balancer is not supported at this time.
- 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.
- Only one interface is allowed per VPC on deployment, due to the nature of Google Clouds network design.

WHY USE ENTERPRISE GCP?

Googles load balancer provides basic load balancing functionality but is limited in several areas. Loadbalancer.org's Enterprise GCP load balancer provides the following additional features & advantages:

1. Supports comprehensive Layer 7 load balancing
2. Load balances both GCP based and non-GCP based servers
3. Supports Round Robin and Least Connection connection distribution algorithms
4. Supports customizable timeouts for custom applications beyond those offered by GCP
5. Supports comprehensive back-end server health-check options
6. Enables fallback servers to be configured and invoked when all load balanced servers/services fail

7. Provides extensive real time and historical statistics reports
8. Supports session distribution based on actual server load (utilizing Loadbalancer.org's feedback agent which is available for both Linux & Windows)
9. Supports SSL Termination
10. Supports Microsoft RDP Cookie based persistence

3. Accessing GCP

To start the Google Cloud Platform, you will 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 Portal
- Gcloud CLI
- GCP API

ACCESSING THE GCP PORTAL

The GCP Portal can be access here.

GCP CLI & GCP API

- Information on how to obtain, install and configure Gcloud CLI is available here.
- Information on how to obtain, install and configure GCP API is available here.

5. Deploying Enterprise GCP from the Marketplace

1. Login into the GCP Portal
2. Select **Marketplace** from the menu and search for Loadbalancer.org, you should see the following three options:
   - **Loadbalancer.org Enterprise GCP MAX** - hourly billing with unlimited VIPs / RIPs
   - **Loadbalancer.org Enterprise GCP R20** - hourly billing with up to 5 VIPs / 4 RIPs
   - **Loadbalancer.org Enterprise GCP BYOL** - for purchasing & applying your own license

   > 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, select the **Launch** button on Compute Engine option.
4. In the deployment section, you can rename the appliance, change zone and disk size/type.

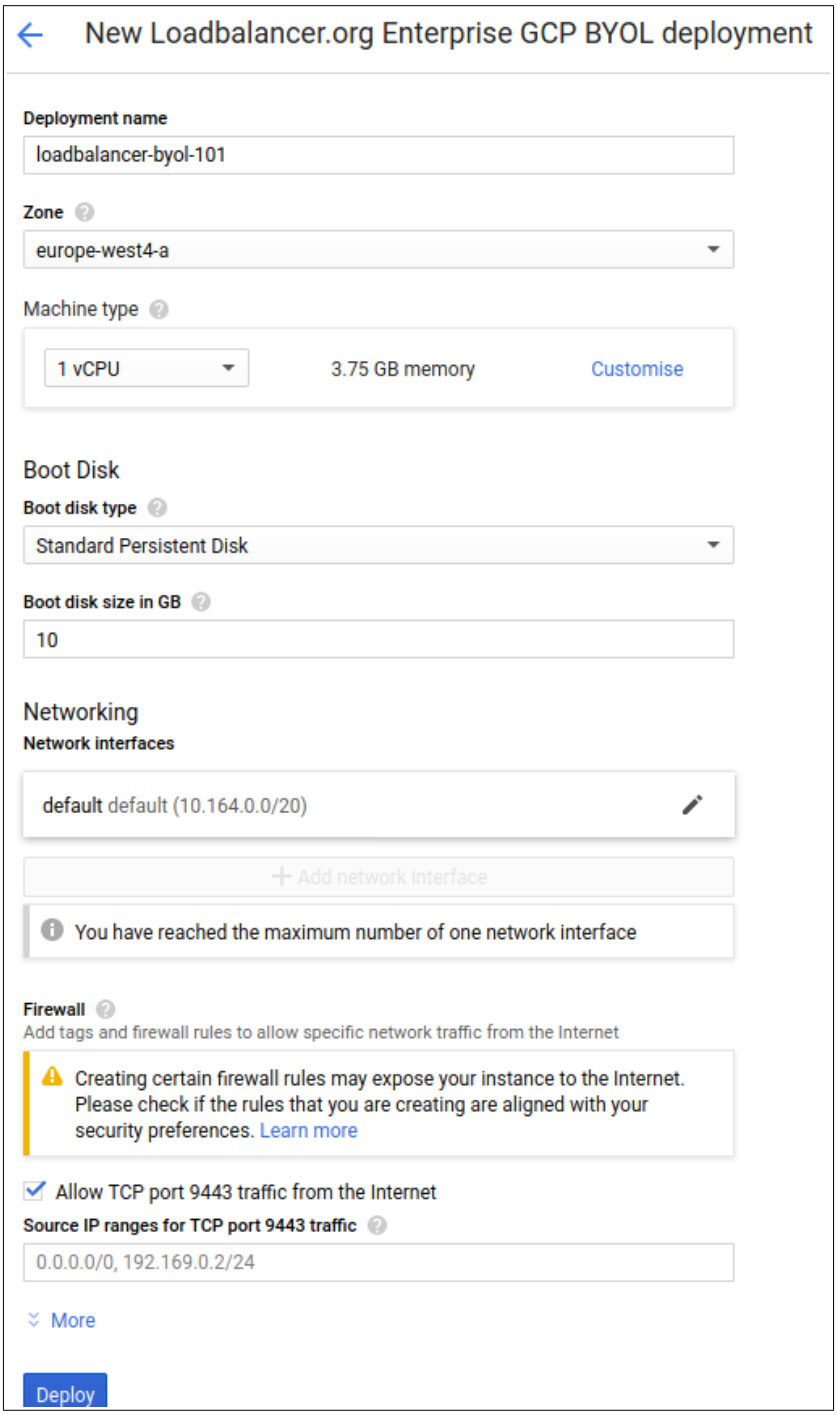

- You can limit access to the load balancer WebUI port 9443 by specifying the sources addresses under Firewall. More ports can be applied to the Firewall settings after the appliance has been deployed. (see section 8. GCP Specific Configuration for more information concerning the firewall rules)
- The standard access ports for the load balancer appliance (9443, 7777, 22) are automatically deployed, these can be locked/secured later.
- You will be able to make the base IP of the load balancer static, rather than ephemeral, after deployment, via the GCP WebUI. This is described further below.
- Please note that you are only allowed one interface on initial deployment, due to the nature of Google clouds VPC Design.

5. Select **Deploy**.
   You will be presented with a new page from the Deployment Manager specifying the login instructions. Please take note of the suggested next steps from this deployment page:
   - Request a license – If deploying a BYOL appliance, you will have 30 days before the trial expires.
   - Change the temporary password – GCP applies a temporary WebUI password for the Loadbalancer.org appliance, it is recommended that this is changed within the Loadbalancer.org WebUI.
   - Assign a static external IP address – The public IP assigned will be lost and renewed with a different one, each time you restart the instance. It is recommended to set the existing private IP to static and assign a static public IP as well.

> Warning:
>
> Do **NOT** change the private IP of the appliance, within the Loadbalancer.org appliance WebUI, as this will not update the GCP network stack and make the appliance completely unusable. There is no recovery from this as yet, you are strongly advised not to change network settings within the appliance WebUI.

To Reserve the private IP address for the appliance

To reserve a static private IP address, you can edit the VM instance Network Interface in the GCP Console.

1. Select Static from the drop-down menu for **Internal IP type**, specify the **unique name** and description and click reserve.
2. Scroll to the end of the page and select **Save** to make sure the changes are implemented.

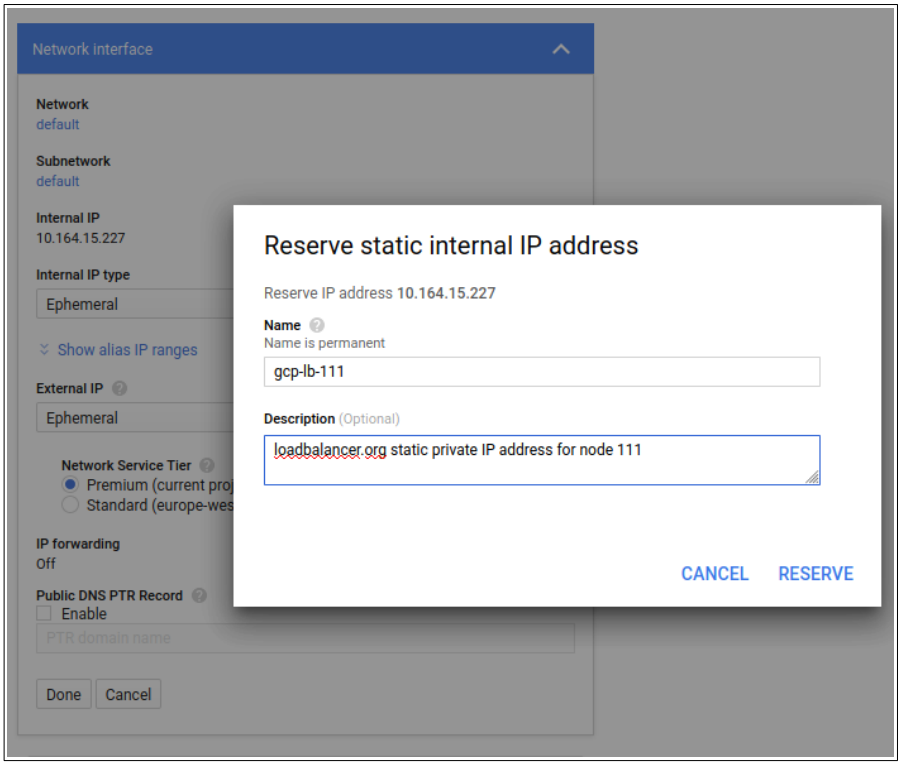

To Reserve A public IP address for the appliance

To reserve a static public IP address, you can edit the VM instance Network Interface in the GCP Console.

1. Select Static from the drop-down menu for **External IP**, you can select a previously created

public IP address (if it is unallocated) or create a new one, specifying the **unique name**, description and Network Service Tier, then select Reserve.

2. Scroll to the end of the page and select **Save** to make sure the changes are implemented.

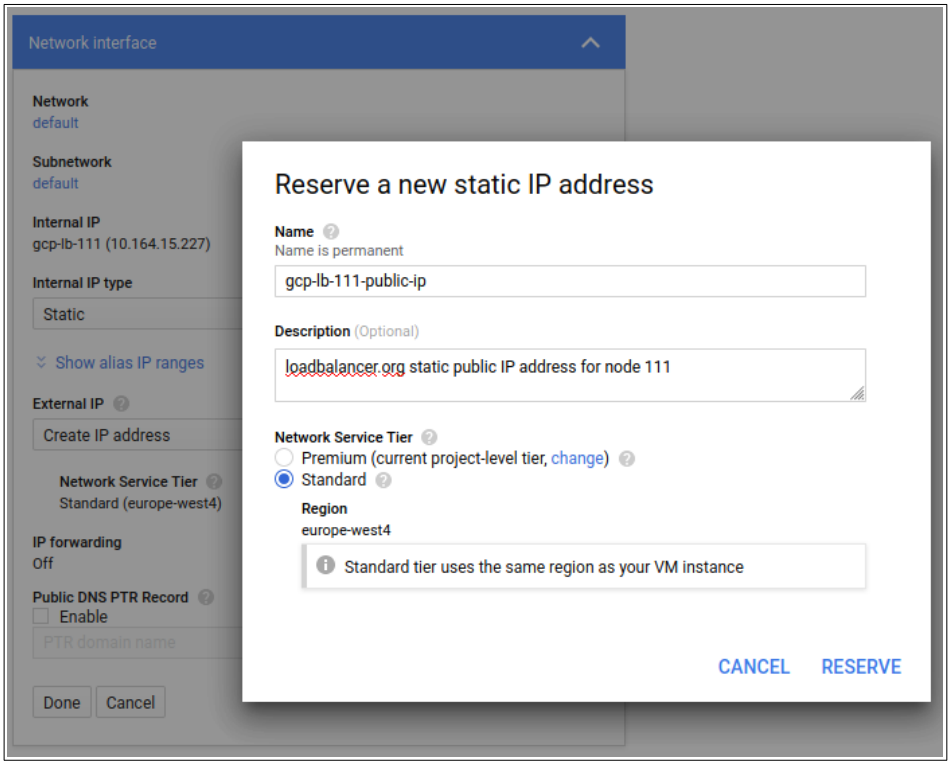

> Warning:
>
> Do **NOT** change the private IP of the appliance, within the Loadbalancer.org appliance WebUI, as this will not update the GCP network stack and make the appliance completely unusable. There is no recovery from this as yet, you are strongly advised not to change network settings within the appliance WebUI.

6. Accessing the Appliance

You can access the appliance immediately from the deployment screen, which gives you the URL to the appliance and the temporary password.

Clicking on the link will open a new browser to the https address, or copy/paste the URl into a new browser window and you can log in with credentials - loadbalancer/<temporary password>

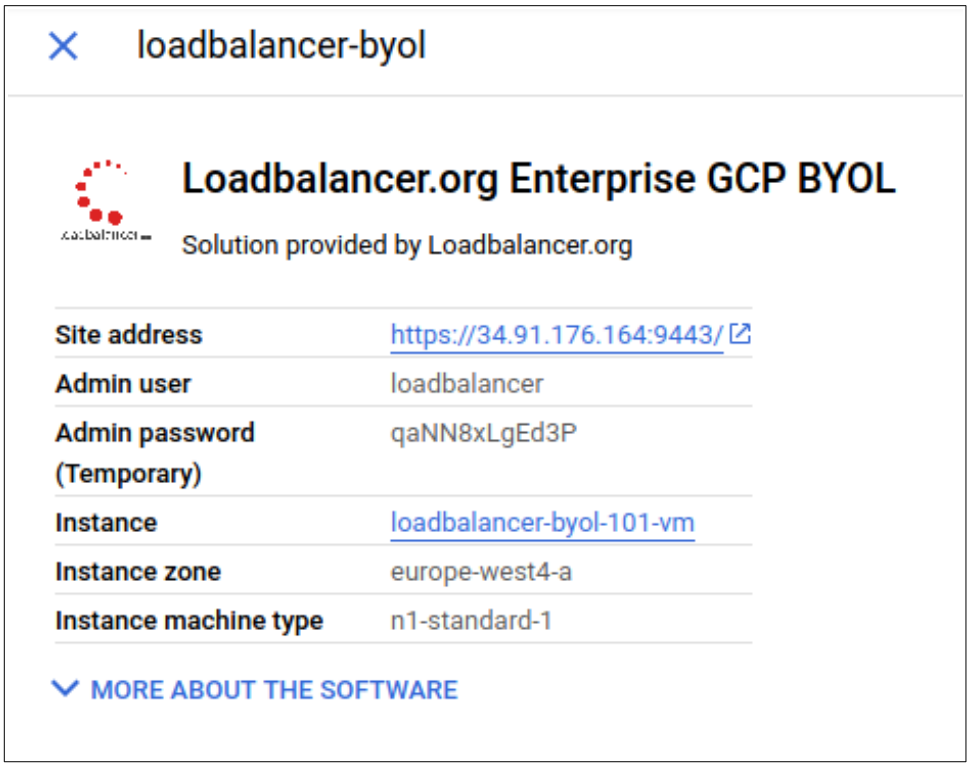

ACCESSING THE APPLIANCE USING THE WEBUI

In a browser, 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 have an internal DNS service and 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>

> Note:
>
> To change the password for the 'loadbalancer' account, use the WebUI option: *Maintenance > Passwords*. Please note that changing the password in the appliance will not update the temporary password in the GCP overview.

Once logged in, the WebUI is displayed:

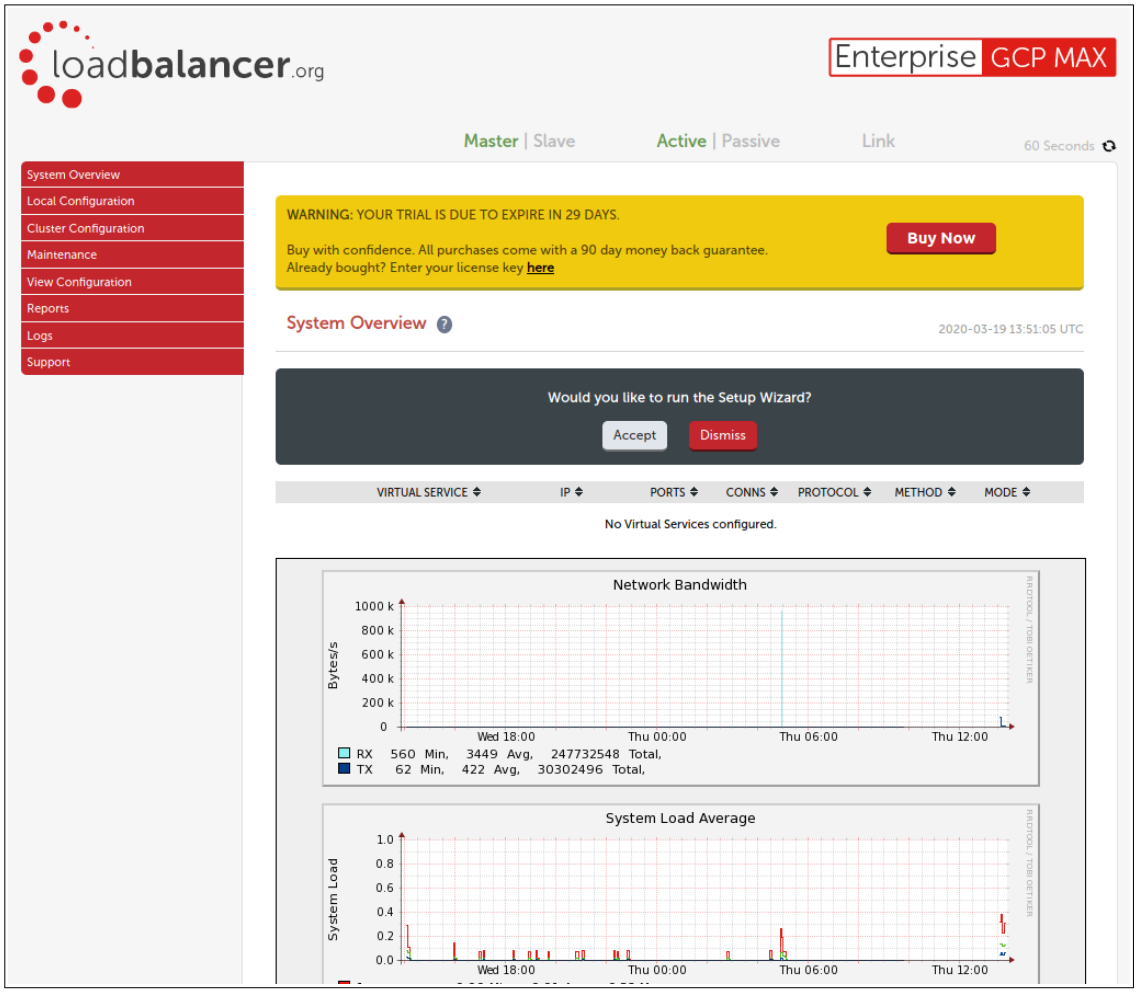

WEBUI MENU OPTIONS

The main menu options are as follows:

System Overview – Displays a graphical summary of all VIPs, RIPS and key appliance statistics
Local Configuration – Configure local host settings such as DNS, Date & 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

APPLIANCE SECURITY

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

- **Secure –** this is the default mode. In this mode:
  - the WebUI is accessible on HTTPS port **9443.** If you attempt to access the WebUI on HTTP port **9080** you will be redirected to HTTPS port **9443**
  - access to the "Execute Shell Command" menu option is disabled

- the ability to edit the firewall script & the lockdown wizard is disabled
- 'root' user console & SSH password access are disabled

- **Custom** – In this mode, the security options can be configured to suit your requirements
- **Secure – Permanent** - this mode is the same as Secure, but the change is ***irreversible***

> IMPORTANT:
>
> Only set the security mode to **Secure - Permanent** if you are 100% sure this is what you want!

To configure the Security Mode:

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

> Note:
>
> For full details of all options, please refer to the Administration Manual and search for "Appliance Security Options".

Default Password

We strongly recommend that the default 'loadbalancer' WebUI account password is changed as soon as the appliance is deployed. This can be changed using the WebUI menu option: *Maintenance > Passwords*.

Please note that changing the password in the appliance will not update the temporary password in the GCP overview.

CHECKING FOR UPDATES

Once you have access to the WebUI, we recommend that you use the online update feature to ensure that you're running the very latest version of the appliance. To check for updates, use the WebUI option: *Maintenance > Software Update* and click the **Online Update** button. If updates are available, you'll be presented with a list of changes that are included in the update. To start the update, click the second **Online Update** button at the bottom of the screen. Updates are incremental, so repeat the process until you're informed that no more updates are available.

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.

ACCESSING THE APPLIANCE USING SSH

When the appliance is deployed, the projects users and SSH keys are inherited from GCP Compute Engine Metadata, making secure access easier to manage. To SSH into the appliance, you will need to ensure that the public SSH key file in the Compute Engines Metadata is correct and that the matching private SSH key

file is on the device you are using.

More information on managing SSH keys in GCP can be found here - https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys

To access the appliance via SSH, either via Windows or Linux, it is recommended to use a client application rather than from the GCP browser.

GENERATING SSH KEYS

The steps below show how to generate SSH key pairs using Linux and Windows, before copying the public key to GCP Project Metadata. If you already have keys setup in GCP you can skip this step.

Using Linux

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

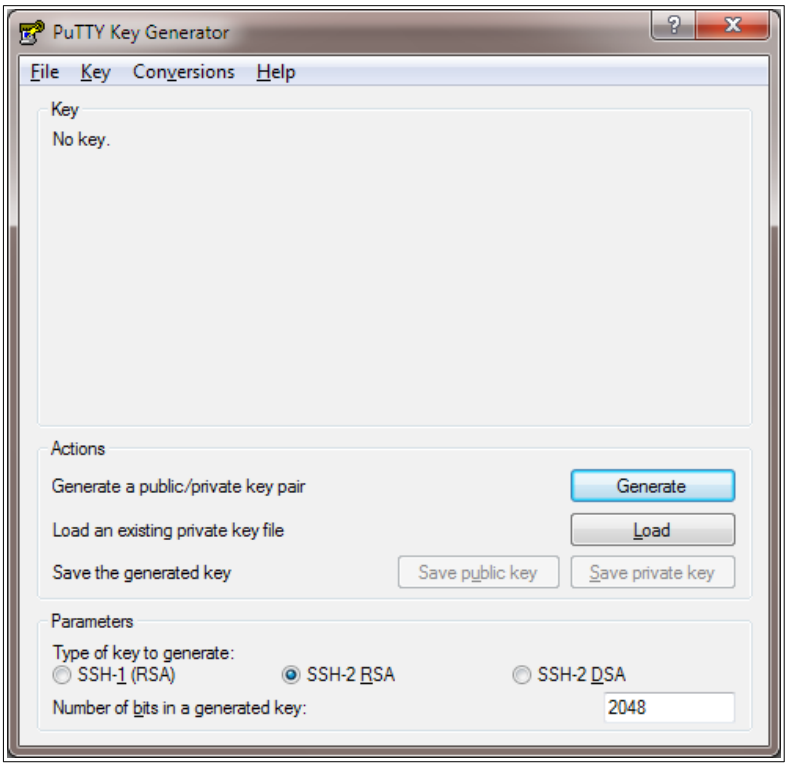

2. Click the **Generate** button
3. As directed, move the mouse around to create random keys
4. Once generated, click the **Save public key** and **Save private key** buttons to save the key

UPDATE THE GCP PROJECT METADATA

If you are an existing user in the relevant GCP Project then you should be able to SSH into the appliance. However, if you have generated new SSH keys or are adding a new client device, you will need to update the GCP Project Metadata keys.

1. Access the GCP portal and navigate to Compute Engine → Metadata → SSH Keys and click the Edit button.
2. Click Add Item and paste the contents of the public key file (e.g GCPKeys.pub) into the window, as shown in the following example:

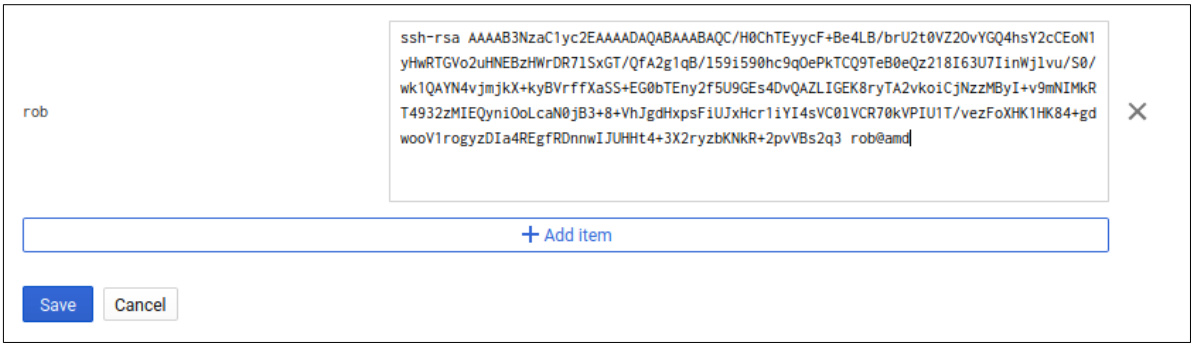

3. Click Save and the updated public key will be added to the metadata for that project.

ACCESSING THE APPLIANCE FROM LINUX

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

Using the IP address:

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

ACCESSING THE APPLIANCE FROM WINDOWS USING PUTTY

1. Run PuTTY
2. Expand the SSH section and select *Auth* as shown below

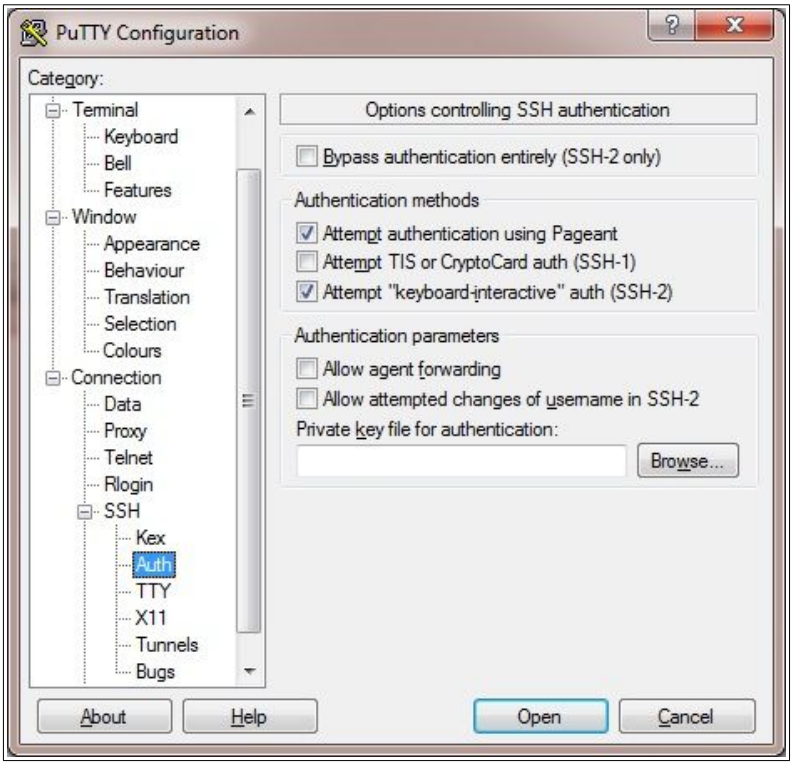

3. Click **Browse** and select the private key created earlier
4. Click **Open** to start the SSH session
5. Login using root and the password you have specified (default is loadbalancer)

> Note:
>
> It is highly recommend to change the default root password and regenerate the ssh keys after deployment of the appliance.
>
> To change the root password at the command line:
>
> **passwd root**
>
> To regenerate the ssh keys at the command line:
>
> **lbsecure ssh**

ACCESSING THE APPLIANCE WITH THE NORMAL USER CREDENTIALS

It is possible to log into the appliance via SSH with the GCP user credentials, if the users public key is correctly recognized in the Compute Engine Metadata, as described previously.

If using the normal user credentials via SSH, some commands may need root credentials to run, and can be executed using "sudo" to achieve elevated privileges.

7. Deployment examples

The following section provide a number of 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 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.

EXAMPLE 1 – WEB SERVERS: 1 SUBNET, 1 LOAD BALANCER NETWORK INTERFACE, LAYER 7

This is a simple layer 7 example using one subnet for both the load balancer and the web servers. The load balancer has a single network interface.

a) Setting up GCP

1. Deploy the load balancer instance as described in section 5. Deploying Enterprise GCP.
2. Deploy your required web server instances into the same VPC & subnet as the load balancer, configure firewall rules accordingly.

b) Setting up the Virtual Service

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

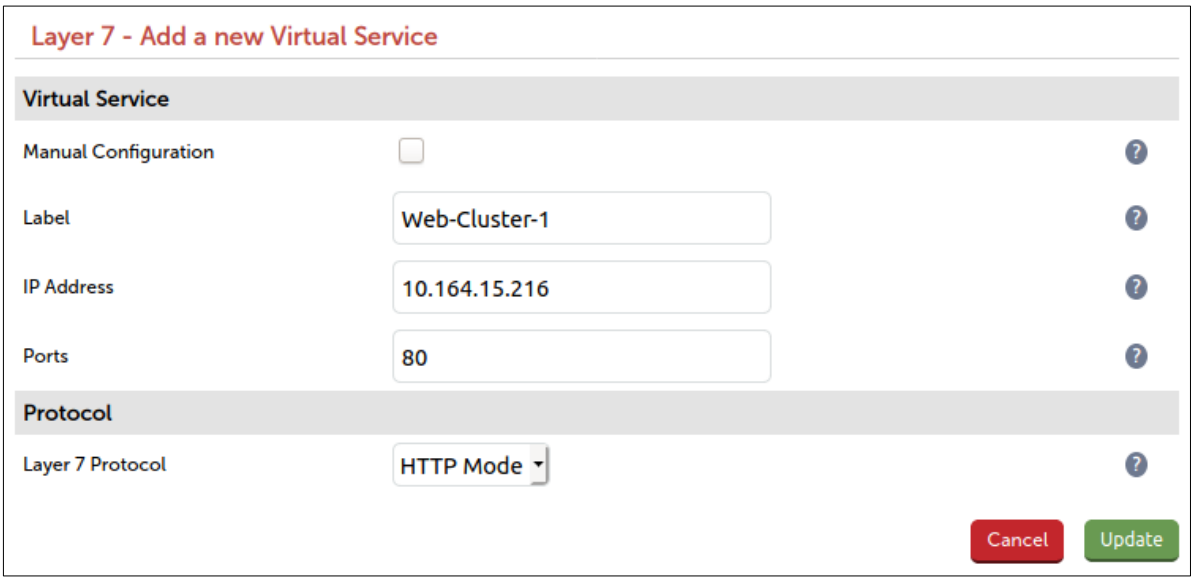

3. Enter an appropriate label for the VIP, e.g. **Web-Cluster1**
4. Set the Virtual Service IP address field to the **Base IP address**, e.g. **10.164.15.216**
5. Set the Virtual Service Ports field to the required port, e.g. **80**

6. Leave Layer 7 Protocol set to **HTTP Mode**
7. Click **Update**

c) Setting up 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:

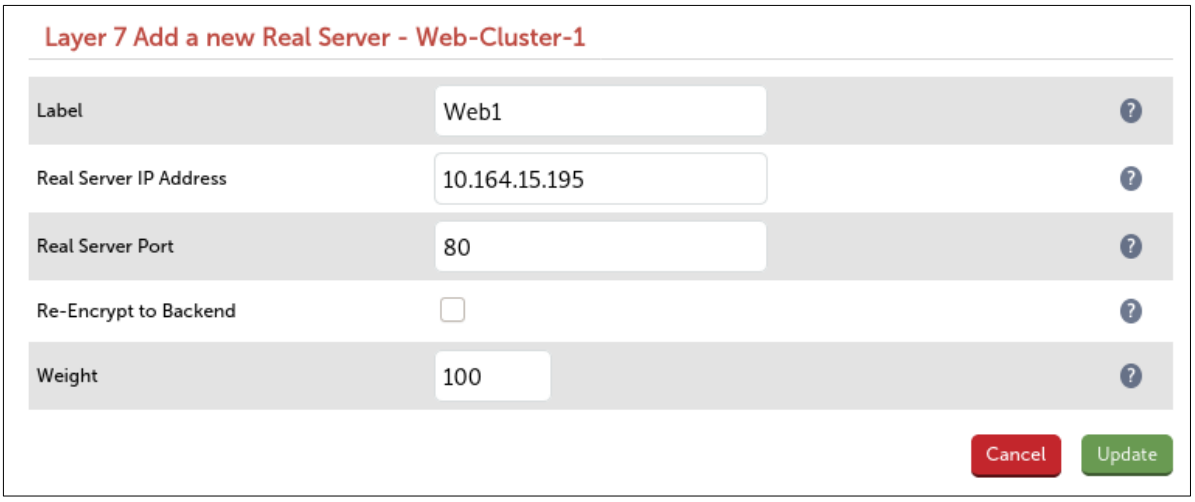

3. Enter an appropriate label for the RIP, e.g. **Web1**
4. Set the Real Server IP Address field to the required IP address, e.g. **10.164.15.95**
5. Set the Real Server Port field to the required port, e.g. **80**
6. Click Update
7. Repeat the above steps to add your other web server(s)

d) Applying 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.
2. Test the services accordingly.

EXAMPLE 2 – MICROSOFT SESSION HOST: 1 SUBNET, 1 LOAD BALANCER NETWORK INTERFACE, LAYER 7

This is a simple Layer 4 example using one subnet for both the load balancer and the RDS servers. The load balancer has a single network interface.

a) Setting up GCP

1. Deploy the load balancer instance as described in section 5. Deploying Enterprise GCP.
2. Deploy your required web server instances into the same VPC & subnet as the load balancer

b) Setting up the Virtual Service

1. Using the WebUI, navigate to: Cluster Configuration > Layer 7 – Virtual Service and click **Add a**

New Virtual Service

2. Enter the following details:

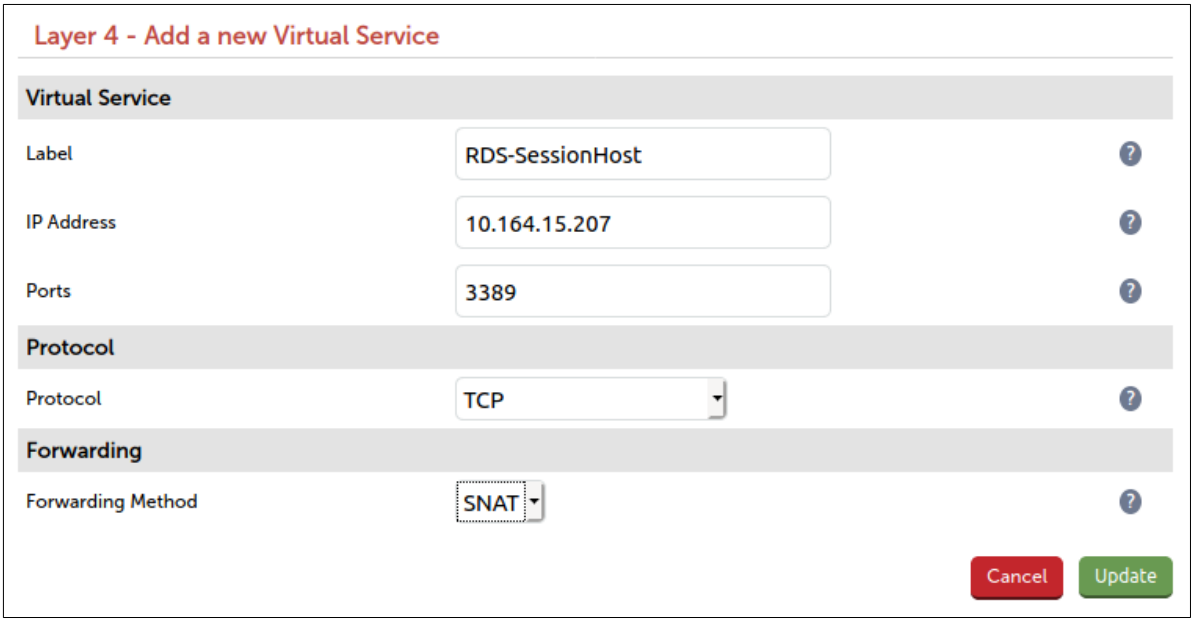

3. Enter an appropriate label for the VIP, e.g. **RDS-SessionHost**
4. Set the Virtual Service IP address field to the Base IP address, e.g. **10.164.15.207**
5. Set the Virtual Service Ports field to the required port, e.g. **3389**
6. Leave Layer 4 Protocol set to **TCP** and set Forwarding Method to **SNAT**
7. Click **Update**
8. Repeat the above steps to add your other server(s), with Layer 4 no services need reloading.

8. GCP specific configuration

In this section we'll cover some of the Google cloud specific configurations that impact the deployment:

- Firewall Configuration notes
- Using the GCP load balancers to be front end for multiple Public IP addresses.

FIREWALL CONFIGURATION NOTES

When the appliance is deployed, a firewall rule is automatically created to provide access to the WebUI and maintenance ports., and referenced under **Network Tags** in the appliance settings.

To allow other services through the load balancer, a new rule is required with the necessary protocols and ports, then added to the Network Tags in the appliance settings.

TO CREATE A NEW FIREWALL RULE

In this example we will add access to HTTPS/HTTPS from a specific IP network.

1. Using the GCP Console, navigate to VPC network → Firewall rules and select CREATE FIREWALL RULE from the top of the page.
2. Apply a Name and Description, select the appropriate Network that the load balancer resides.

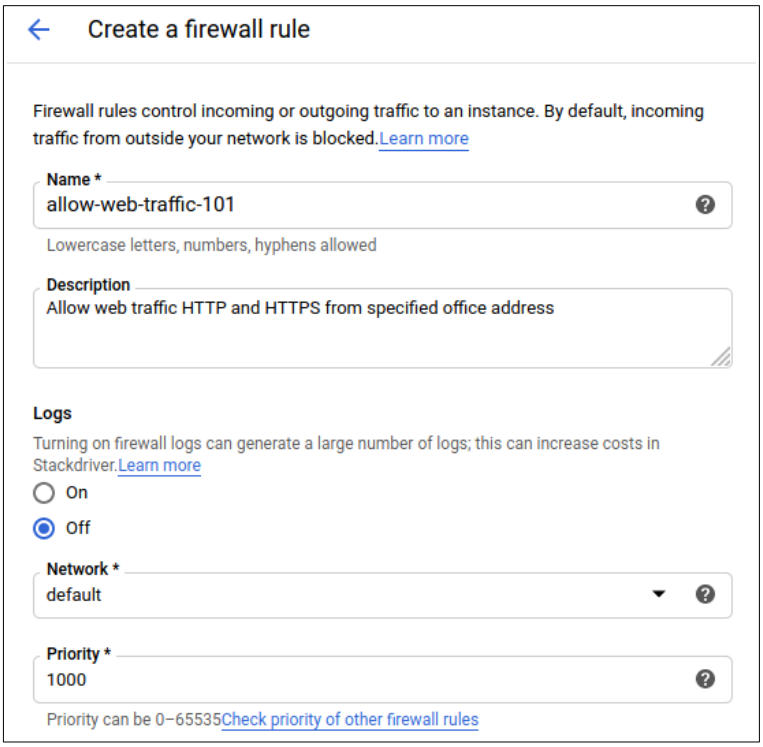

3. Select **Specified target tags** under Targets. Fill in a **Target Tag** with the relevant identifier you wish and specify a **Source filter** as **IP ranges,** then enter the source IP range you require to be allowed access.
4. Under Protocols and ports select Specified protocol and ports and enter the relevant data, which in this example is TCP 80/443.

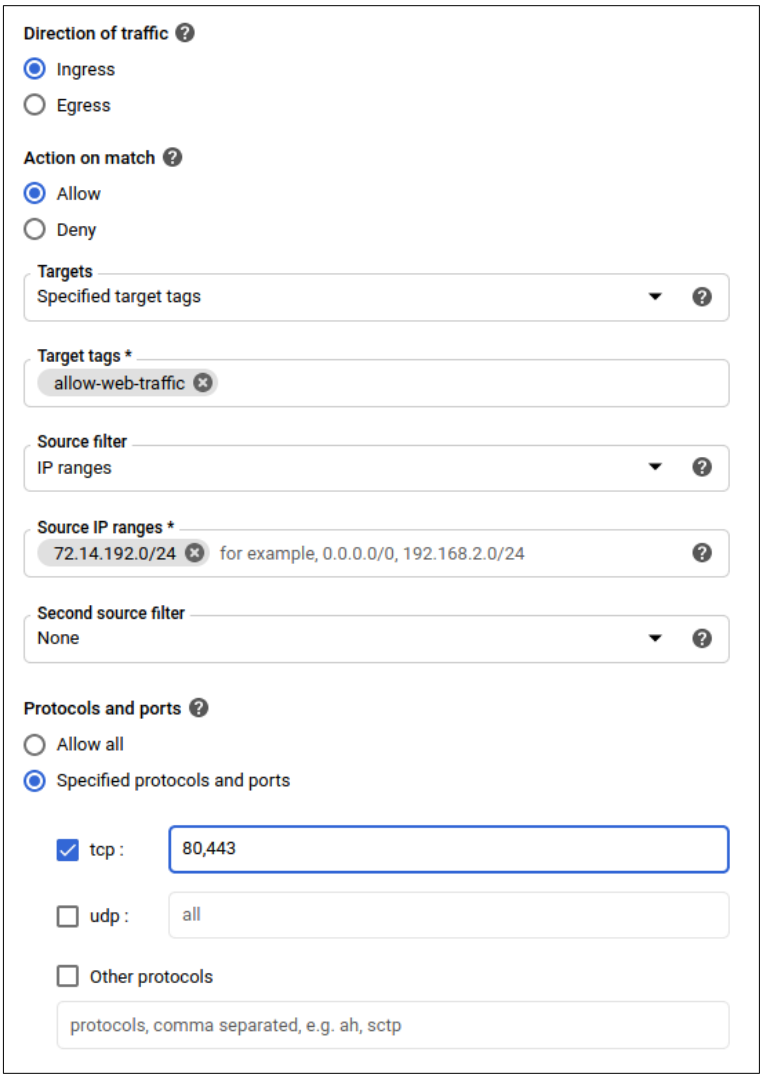

5. When complete and there are no validation issues, click Create.

> Note:
>
> When trouble shooting a connection issue, you can edit the firewall rule to disable it, rather than delete it. Select the rule, select **EDIT** and navigate to the **Enforcement** section, select **Disabled** as your requirement and select **SAVE**. The rule will now be disabled and will be displayed differently in the Firewall rules table.

Then apply the new firewall rule to an alliance:

1. Navigate to **Compute Engine** → **VM instances** and select the load balancer instance, select **EDIT.**
2. Under **Network tags** enter the name of your **firewall target tag**, in our example allow-web-traffic, as shown below:

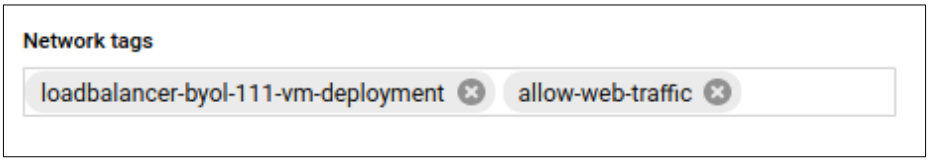

3. Select **Save** to complete the configuration.

TERMINATING PUBLIC IP ADDRESSES UTILIZING THE GCP LOAD BALANCER

Terminating Public IP addresses directly on the Loadbalancer.org appliance is not possible due the the nature of the Google Cloud design. However, t is possible to utilize the GCP Layer 4 load balancer as a front-end termination points, for the Public IP addresses and forward them to the Loadbalancer.org appliance.

There is no requirement for health-checks on the GCP load balancer as it always forwards traffic to the Loadbalancer.org appliance.

The same Public IP can be used on the GCP Layer 4 TCP and UDP load balancers, for a service that requires both (e,g DNS TCP/UDP port 53)

TO ADD A GCP LAYER 4 LOAD BALANCER

1. Navigate in the GCP Console to **Network Services** → **Load balancing** and click **Create Load balancer**.
2. Select the type of protocol you require, TCP or UDP load balancer by clicking **Start configuration**. Please note that you can have both for certain services (such as DNS) by terminating the same Public IP on the TCP and UDP load balancers at the same time.
3. Select **From Internet to my VMs** and **Single region only** (or as appropriate to your deployment) and then **Continue**.
4. Enter a **Name** for the load balancer and then select **Backend configuration**, enter the appropriate **Region**, then under **Backends** click **Select existing instances** and select the Loadbalancer.org appliance from the drop-down menu.

   Please note that we do not recommend configuring a health-check between the GCP load balancer and the Loadbalancer.org appliance as it provides no further functionality in failover.

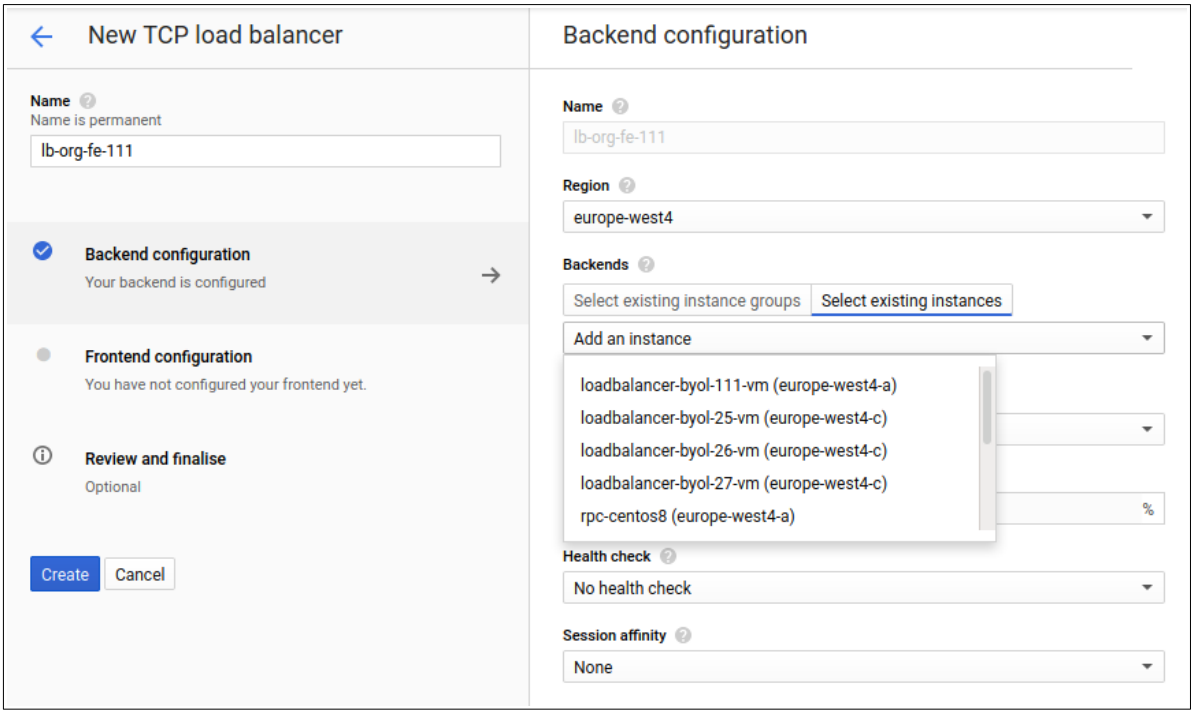

5. Now select **Frontend configuration** and give a **Name** to the New Frontend IP and port, select the **Network Service Tier** you require, the click on the IP drop-down menu to select the IP address you require, or click **Reserve a new static IP address** to create a new one.

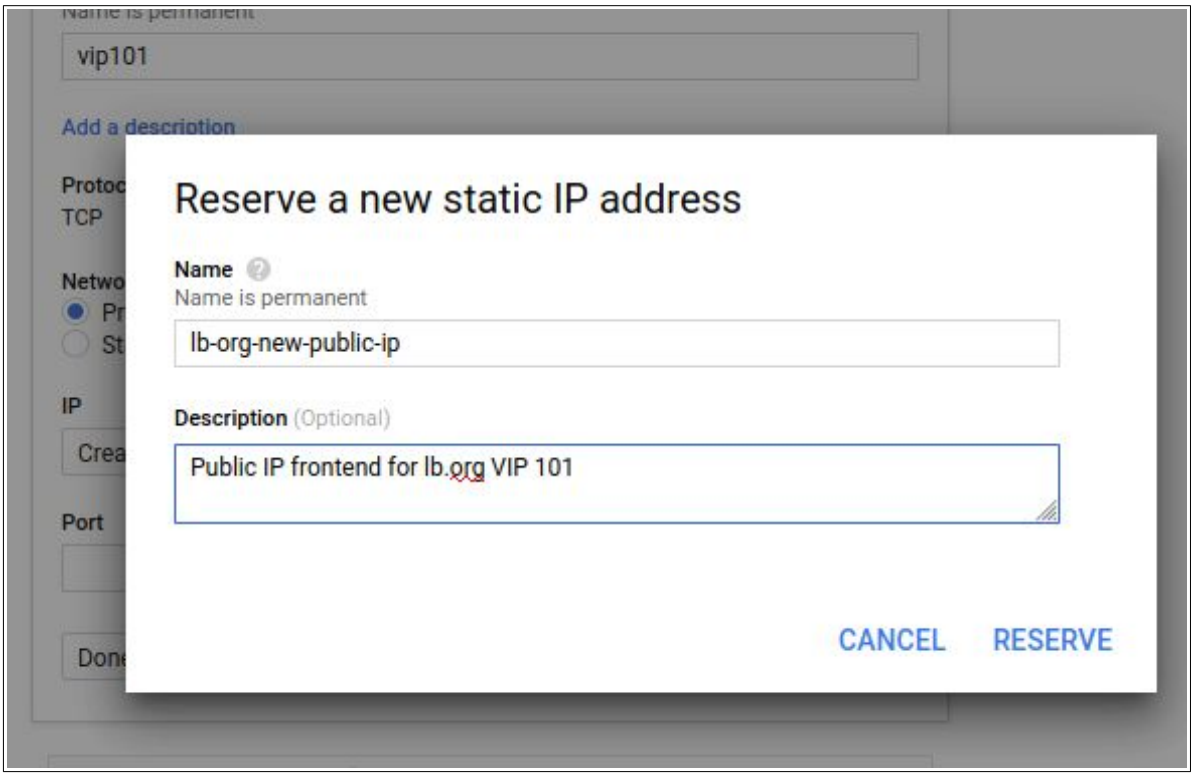

6. The **Port** can be specified as a single port or two port numbers separated by a dash. A port range is not possible due to the constraints of the GCP load balancer.
7. When all of the information has been entered, click **Done.** You can Add another Frontend IP from the same menu or continue to finalise the configuration.

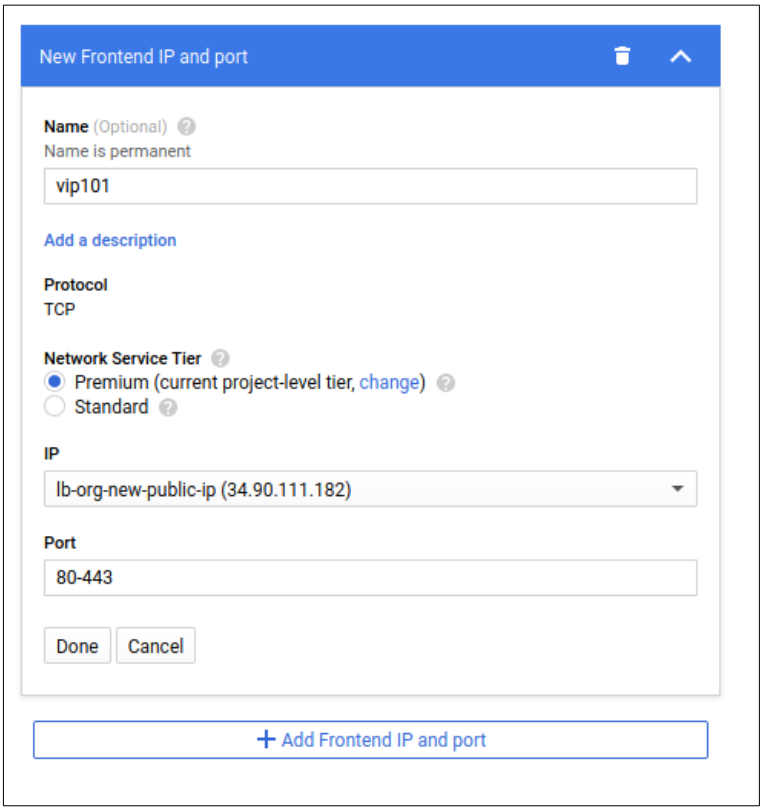

8. Once you have completed that configuration, select **Review and finalize.** There will be an warning

concerning no health checks but that can be ignored as we want all traffic to be sent to the Loadbalancer.org appliance.

9. Select **Create** to finalize the configuration of the GCP load balancer as the front end.

APPLY THE CHANGES TO THE LOADBALANCER.ORG APPLIANCE

1. Connect to the Loadbalancer.org appliance WebUI and select **Cluster Configuration**.
2. Choose the service you wish to configure, in this example **Layer 7 – Virtual Services**.
3. Select **Add a new Virtual Service,** enter the Public IP address that was specified in the Frontend configuration of the GCP load balancer and then the port(s).

   In this example, ports 80 and 443 are passed through to the back-end servers.

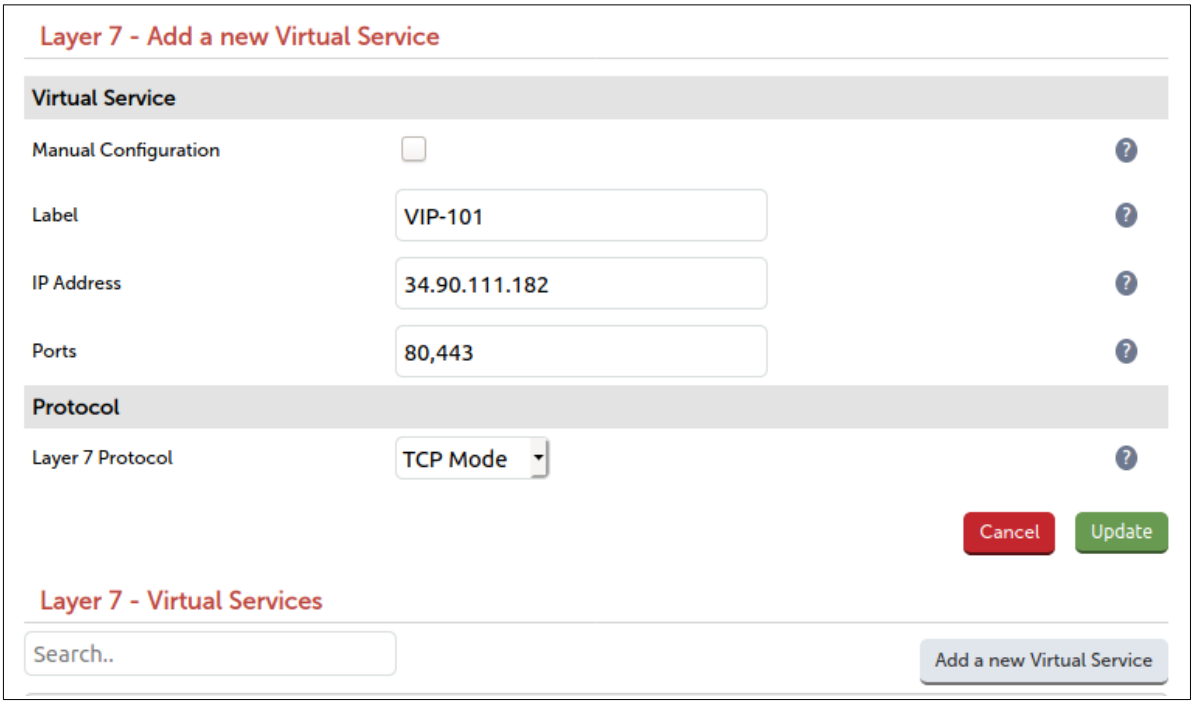

4. Select Update and then navigate to Layer 7 – Real Servers to add your backend servers.
5. Select **Reload HAProxy** at the top of the page when prompted.

9. Testing – General Comments

TESTING LOAD BALANCED SERVICES

For example, to test a web server based configuration, add a page to each web servers root directory e.g. ***test.html*** and put the server name on this page for easy identification during the tests.

Use two or more clients to do the testing. Open up a web browser on each test clients and enter the URL for the VIP e.g. **http://104.40.133.119**

Provided that persistence is disabled, each client should see a different server name because of the load balancing algorithm in use , i.e. they are being load balanced across the cluster.

Why test using two clients? If you use a single client it will most likely keep on hitting the same server for multiple requests. This is to do with the way that the load balancing algorithms are optimized.

DIAGNOSING VIP CONNECTION PROBLEMS

1. ***Make sure that the device is active*** – this can be checked in the WebUI. For a single appliance, the status bar should report **Master** & **Active** as shown below:

2. ***Check that the Real Servers are up*** – Using *System Overview* make sure that none of your VIPs are colored red. If they are, the entire cluster is down (i.e. all Real Servers). Green indicates a healthy cluster, yellow indicates that your cluster may need attention (one or more of the Real Servers may be down), and blue indicates all Real Server have been deliberately taken offline (by using either Halt or Drain).

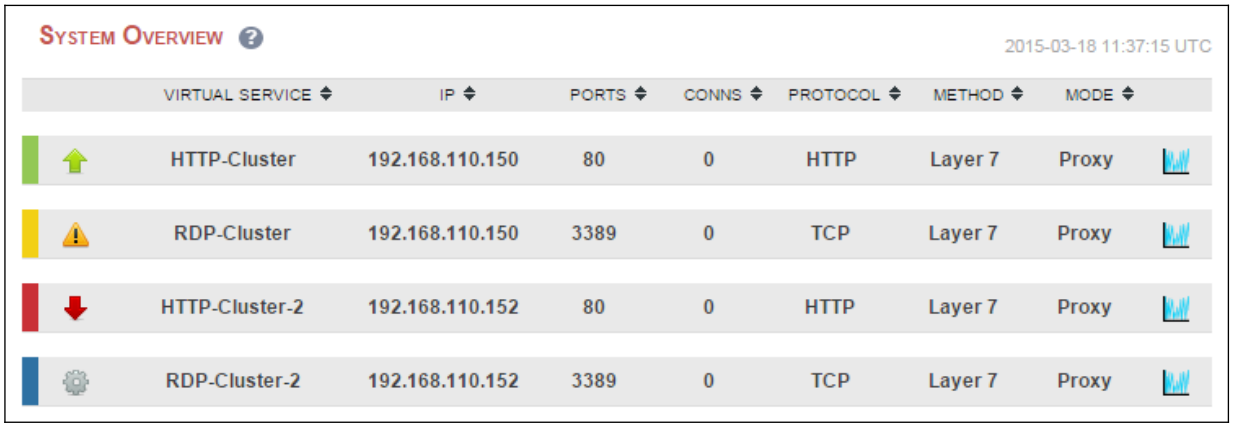

3. ***Check the connection state***

For Layer 7 VIPs, check *Reports > Layer 7 Status.* The default credentials required are:

username: loadbalancer
password: loadbalancer (or the **password specified in Layer 7 Advanced**)

This will open a second tab in the browser and display a statistics/status report as shown in the example below (this is accessed on port TCP/7777 so make sure that the inbound rules allow connections on this port):

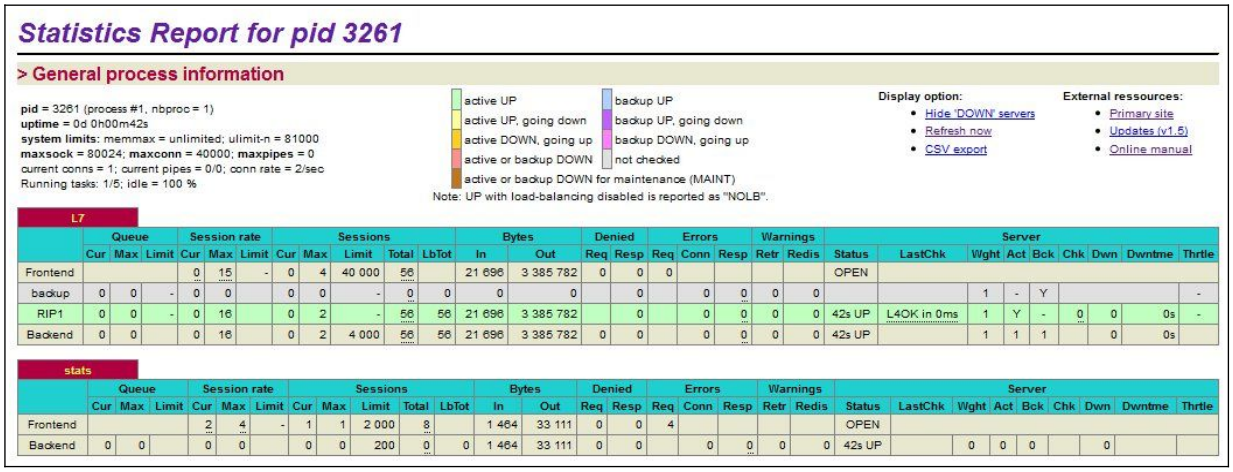

TAKING REAL SERVERS OFFLINE

1) Using the *System Overview* check that when you Halt one of the Real Servers the connections are redirected to the other server in the cluster.

2) Stop the web service/process on one of the servers, wait a few seconds (for the load balancer to detect the change) and then refresh the browsers on both clients. They should now both switch to the same server (since one has been removed from the load balancing list). Also check that the server is shown red (down) in the system overview.

3) Start the web service/process on the server, wait a few seconds and then refresh the browsers again. After a few refreshes they should again show different web servers. Also check that the server is shown green (up) in the system overview.

The *System Overview* shows the status as these tests are performed:

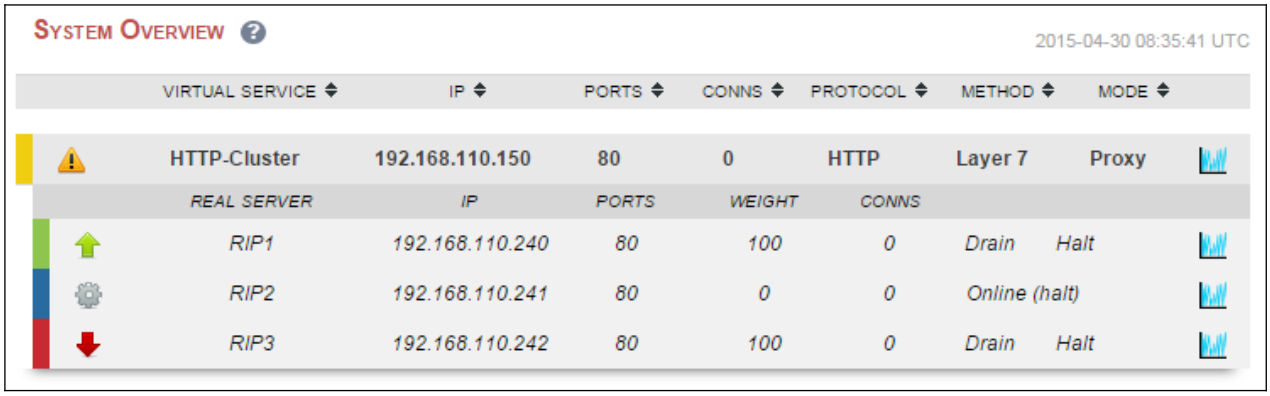

In this example:

- **RIP1** is green, this indicates that it's operating normally
- **RIP2** is blue, this indicates that it has been either Halted or Drained. in this example Halt has been used as indicated by *Online (Halt)* being displayed. If it had been drained it would show as *Online (Drain)*
- **RIP3** is red, this indicates that it has failed a health check

USING REPORTS & LOG FILES

The appliance includes several logs and reports that are very useful when diagnosing issues. Both are available as main menu options in the WebUI. Details of both can be found in the administration manual.

10. More Information

Please refer to our website for the latest administration manual, deployment guides and all other documentation: https://www.Loadbalancer.org/uk/resources/manuals

11. Loadbalancer.org Technical Support

If you have any questions regarding the appliance or how to load balance your application, please don't hesitate to contact our support team using the following email address: support@Loadbalancer.org

12. Company Contact Information

Website	URL:	www.loadbalancer.org
North America (US)		Loadbalancer.org, Inc. 4550 Linden Hill Road, Suite 201 Wilmington, DE 19808 USA
	Tel: Email (sales): Email (support):	+1 833.274.2566 sales@loadbalancer.org support@loadbalancer.org
North America (Canada)		Loadbalancer.org Appliances Ltd. 300-422 Richards Street Vancouver, BC V6B 2Z4 Canada
	Tel: Email (sales): Email (support):	+1 866.998.0508 sales@loadbalancer.org support@loadbalancer.org
Europe (UK)		Loadbalancer.org Ltd. Compass House North Harbour Business Park Portsmouth, PO6 4PS UK
	Tel: Email (sales): Email (support):	+44 (0)330 380 1064 sales@loadbalancer.org support@loadbalancer.org
Europe (Germany)		Loadbalancer.org GmbH Tengstraße 27 80798 München Germany
	Tel: Email (sales): Email (support):	+49 (0)89 2000 2179 vertrieb@loadbalancer.org support@loadbalancer.org