



Enterprise GCP Configuration Guide **v8.5.8**

Rev. 1.0.0

Table of Contents

1. Introduction.....	3
2. About Enterprise GCP.....	3
Main Differences to Our Standard (Non-Cloud) Product.....	3
Why use Enterprise GCP?.....	4
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
To Reserve a static private IP address for the appliance.....	8
To Reserve A static public IP address for the appliance.....	9
6. Accessing the Appliance.....	10
Accessing the Appliance using the WebUI.....	10
WebUI Menu Options.....	11
Appliance Security.....	11
Checking For Updates.....	12
Appliance Licensing.....	12
Accessing the Appliance using SSH.....	13
Generating SSH Keys.....	13
Update the GCP project metadata.....	15
Accessing the Appliance from Linux.....	15
Accessing the Appliance from Windows using PuTTY.....	15
Accessing the Appliance with the GCP User Credentials.....	16
7. Deployment examples.....	16
Example 1 – Web Servers: Single Subnet, Layer 7, Public facing.....	17
Example 2 – Web Servers: Single Subnet, Layer 7, Internal Facing.....	18
Example 3 – Web Servers: Single Subnet, Layer 4, Public Facing.....	20
Example 4 – using GCP Load balancer To Present Multiple Services on Different Public IPs.....	23
8. Testing – General Comments.....	31
Testing Load Balanced Services.....	31
Diagnosing VIP Connection Problems.....	32
Taking Real Servers Offline.....	33
Using Reports & Log Files.....	34
9. More Information.....	34
10. Loadbalancer.org Technical Support.....	34
11. Company Contact Information.....	35

1. Introduction

Google Cloud Platform (GCP) is a broad suite of cloud-based services running on the same global infrastructure as Google's 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:

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 the section [Example 2 – Web Servers: Single Subnet, Layer 7, Internal Facing](#).

If you want to present additional services on different External IPs, a GCP load balancers (either HTTP(S), TCP or UDP) can be used to terminate additional public IP addresses. For more information and a deployment example, please refer to the section [Example 4 – using GCP Load balancer To Present Multiple Services on Different Public IPs](#).

MAIN DIFFERENCES TO OUR STANDARD (NON-CLOUD) PRODUCT

- Layer 4 DR mode is currently **not** supported.
- HA (high availability) where a clustered pair of appliances is deployed is currently **not** supported.
- 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 the section [Example 3 – Web Servers: Single Subnet, Layer 4, Public Facing](#) for an example of how VPC routing is modified to route traffic back via the load balancer.
- 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 the section [Example 3 – Web Servers: Single Subnet, Layer 4, Public Facing](#) for an example of how VPC routing is modified to route traffic back via the load balancer.
- Enterprise GCP can have only one network interface.

WHY USE ENTERPRISE GCP?

Google's 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 using GCP, 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'll be presented with the following three options:
 - **Loadbalancer.org Enterprise GCP R20** – hourly billing with up to 5 VIPs / 4 RIPs
 - **Loadbalancer.org Enterprise GCP MAX** – hourly billing with unlimited VIPs / 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
4. Click the **Launch** button

New Loadbalancer.org Enterprise GCP BYOL deployment

Deployment name

Zone

Machine type

Machine family

GENERAL-PURPOSE

COMPUTE-OPTIMISED


MEMORY-OPTIMISED

Machine types for common workloads, optimised for cost and flexibility

Series

Powered by Intel Skylake CPU platform or one of its predecessors

Machine type



vCPU	Memory
1	3.75 GB

✓ CPU PLATFORM AND GPU

- Enter an appropriate *Deployment Name*
 - Select the required *Zone*
 - Configure the *Machine Type* settings according to your requirements
5. Scroll down to the *Boot disk & Networking* sections

Boot Disk

Boot disk type *

Standard Persistent Disk

Boot disk size in GB

10

Networking

Network interfaces

Network interface

Network

default

Subnetwork

default

External IP

Ephemeral

DONE

ADD NETWORK INTERFACE

i You have reached the maximum number of one network interface

- Configure the *Boot disk type* according to your requirements, these options can normally be left at their default values
- Configure the *Network Interfaces* 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 described further below.

- Scroll down to the *Firewall* section

Firewall ?

Add tags and firewall rules to allow specific network traffic from the Internet



Creating certain firewall rules may expose your instance to the Internet. Please check if the rules that you are creating are aligned with your security preferences. [Learn more](#)

☒ Allow TCP port 9443 traffic from the Internet

?

IP forwarding

Off ▼ ?

[^ SHOW LESS](#)

- The following firewall rules are automatically configured for each VPC network:

<input type="checkbox"/>	Name	Type	Targets	Filters	Protocols/ports	Action
<input type="checkbox"/>	test-allow-icmp	Ingress	Apply to all	IP ranges: 0.0.0.0/0	icmp	Allow
<input type="checkbox"/>	test-allow-internal	Ingress	Apply to all	IP ranges: 10.128.0.0/9	all	Allow
<input type="checkbox"/>	test-allow-rdp	Ingress	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow
<input type="checkbox"/>	test-allow-ssh	Ingress	Apply to all	IP ranges: 0.0.0.0/0	tcp:22	Allow

- When the Loadbalancer.org appliance is deployed, a new firewall rule is automatically added to enable access to the WebUI on port 9443. This is controlled by the option *Allow TCP port 9443 traffic from the Internet* shown above.
- Filters should be configured where possible to limit access to specific source IP ranges
- If you plan to use Layer 4 NAT or Layer 7 with TProxy ensure that *IP Forwarding* is **On**
- Add additional rules as required by your deployment
- For layer 7 deployments, you'll need to allow access on TCP port 7777 to enable access to the Layer 7 Stats page

Note:

When troubleshooting connection issues, you can disable a firewall rule rather than deleting it. Select the rule, select **EDIT** and navigate to the **Enforcement** section, select **Disabled** and click **SAVE**.

- Click **Deploy** – a deployment summary page will be displayed:

The screenshot shows the Google Cloud Deployment Manager interface. On the left, the 'Deployments' tab is active, showing a tree view of the deployment 'lb1'. The tree includes 'loadbalancer-byol', 'loadbalancer-byol-vm-tmpl', 'lb1-vm', 'generated-password-0', and 'lb1-tcp-9443'. The main panel shows the details for 'loadbalancer-byol', including the 'Site address' (https://34.105.145.149:9443/), 'Admin user' (loadbalancer), 'Admin password (Temporary)' (X7feekYBeuN8), 'Instance' (lb1-vm), 'Instance zone' (europe-west2-c), and 'Instance machine type' (n1-standard-1). Below this, there are links for 'Get started with Loadbalancer.org Enterprise GCP BYOL', 'Suggested next steps' (Request a license, Change the temporary password, Assign a static external IP address), and 'Documentation' (Getting Started Guide, Appliance Administration Manual).

- The auto generated temporary password that is used to login as the WebUI user 'loadbalancer' should be changed once you login
- If the instance has an external IP, you can directly access the WebUI after deployment by clicking the **Visit the Site** button
- If you have deployed a BYOL appliance, you automatically have a 30 day trial. The Request a License link takes you to the loadbalancer.org website to purchase a license
- An ephemeral external IP is assigned to the appliance. For production deployments this should be changed to a static address

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.

TO RESERVE A STATIC PRIVATE IP ADDRESS FOR THE APPLIANCE

1. Using the GCP console, edit the instance
2. In the *Network Interfaces* section change *Internal IP type* to **Static**

Reserve static internal IP address

Reserve IP address 10.154.0.12

Name ?
Name is permanent

lb1-internal-ip

Description (Optional)

LB1 Internal IP

CANCEL RESERVE

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

TO RESERVE A STATIC PUBLIC IP ADDRESS FOR THE APPLIANCE

1. Using the GCP console, edit the instance
2. In the *Network Interfaces* section change *External IP type* to **Create IP address**

Reserve a new static IP address

Name ?
Name is permanent

lb1-external-ip

Description (Optional)

LB1 External IP

Network Service Tier ?
☐ Premium (current project-level tier, [change](#)) ?
☒ Standard ?

Region
europe-west2

i Standard tier uses the same region as your VM instance

CANCEL RESERVE

3. In the popup, specify an appropriate *Name & Description*
4. Set the *Network Service Tier* according to your requirements
5. Click **Reserve**
6. Scroll to the end of the page and click **Save**

6. Accessing the Appliance

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 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>

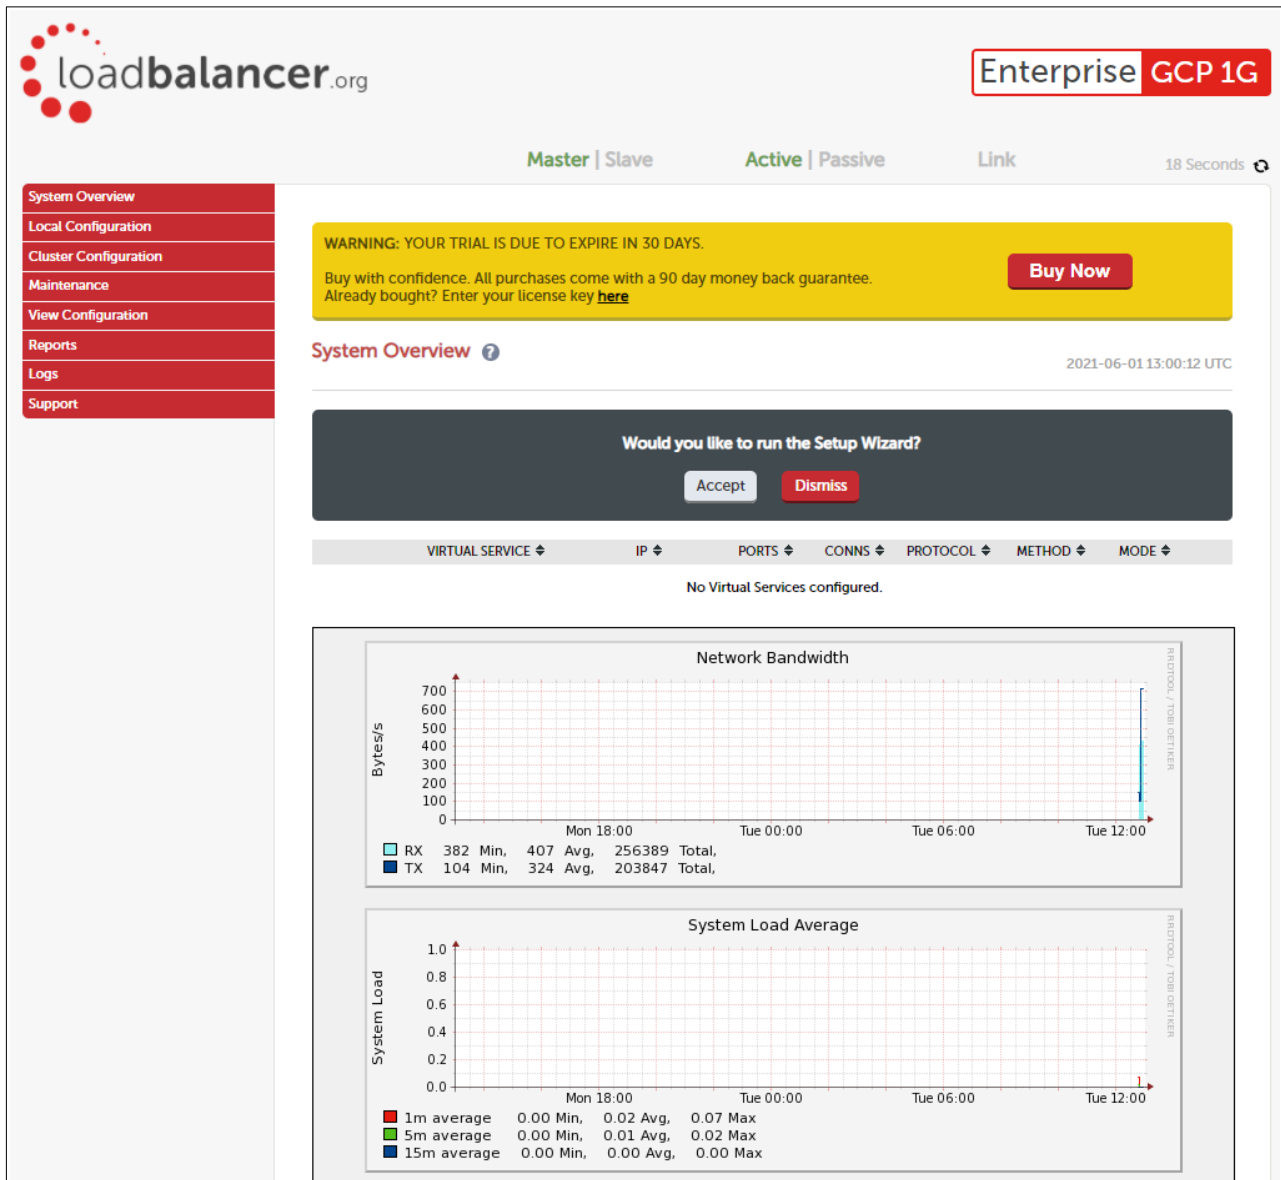
The temporary-password is displayed in the deployment summary page after deployment and also in the **custom metadata** section of the VM instance properties in the GCP console:

Custom metadata	
loadbalancer_user_password	X7feeKYBeuN8
google-monitoring-enable	0
google-logging-enable	0

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 displayed 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 displayed 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.

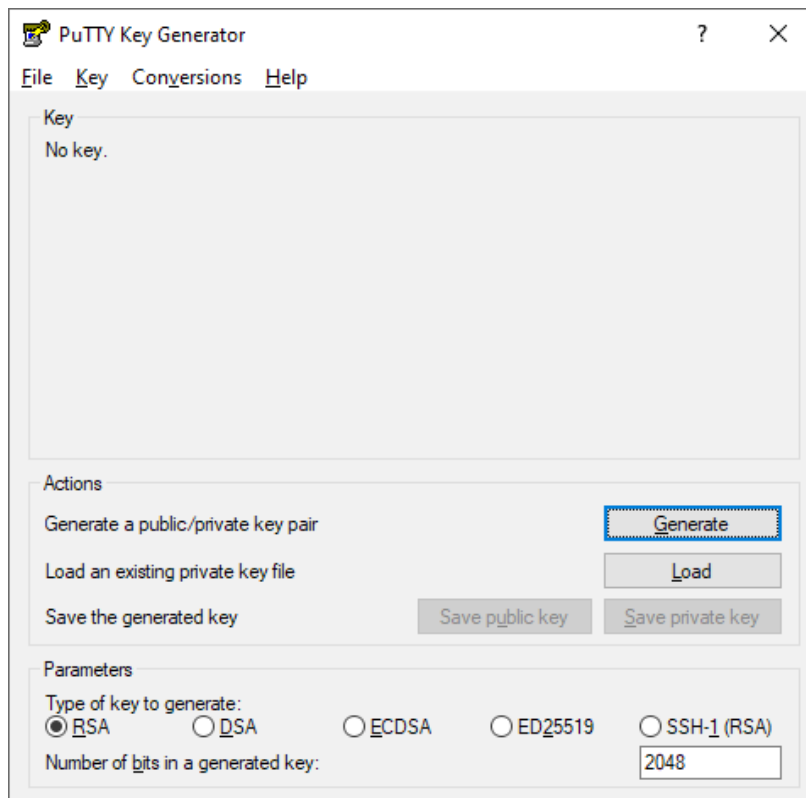
Using Windows

STEP 1 – Install PuTTY

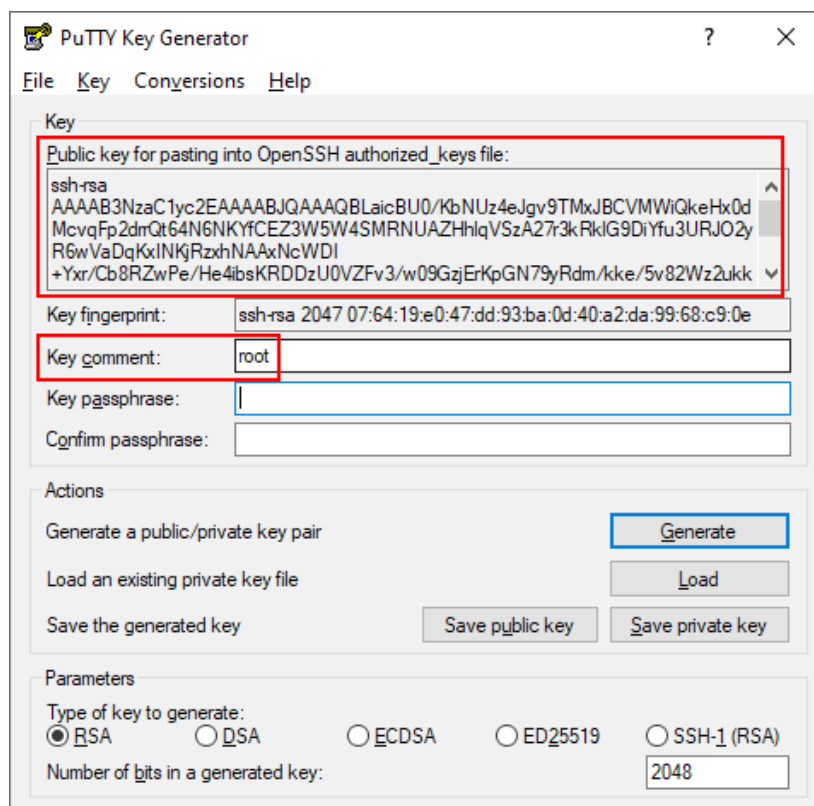
1. Download PuTTY from [here](#)
2. Run the installer

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

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 or the data displayed at the top in PuTTYgen) into the window, as shown in the following example:

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

ACCESSING THE APPLIANCE FROM LINUX

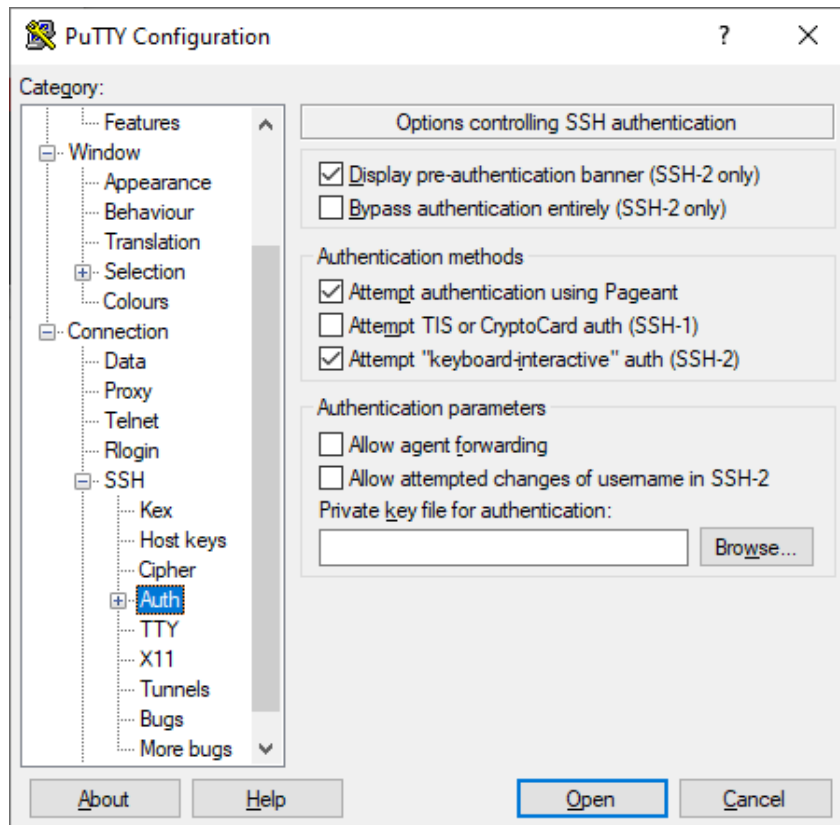
Start SSH specifying the private key file and login as the user “root”, e.g.:

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 GCP 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 GCP 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 provides 3 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.

EXAMPLE 1 – WEB SERVERS: SINGLE SUBNET, LAYER 7, PUBLIC FACING

This is a simple layer 7 example using one subnet for both the load balancer and the web servers.

a) Deploying the GCP instances

1. Deploy the load balancer instance as described in the section [Deploying Enterprise GCP from the Marketplace](#), configure a static External IP address – this must be done after deployment
2. Deploy your web server instances into the same VPC & subnet as the load balancer, use a static Internal IP address – this can either be reserved in advance and assigned during deployment or promoted after deployment
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80

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:

Virtual Service	
Manual Configuration	<input type="checkbox"/>
Label	WebCluster1
IP Address	10.154.0.17
Ports	80
Protocol	
Layer 7 Protocol	HTTP Mode
Cancel Update	

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. **10.154.0.17**
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:

Label	<input type="text" value="Web1"/>	?
Real Server IP Address	<input type="text" value="10.154.0.18"/>	?
Real Server Port	<input type="text" value="80"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

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. **10.154.0.18**
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.

e) Testing

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

EXAMPLE 2 – WEB SERVERS: SINGLE SUBNET, LAYER 7, INTERNAL FACING

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

a) Deploying the GCP Instances

1. Deploy the load balancer instance as described in the section [Deploying Enterprise GCP from the Marketplace](#), configure a static Internal IP address – this must be done after deployment
2. Deploy your web server instances into the same VPC & subnet as the load balancer, use a static Internal IP address – this can either be reserved in advance and assigned during deployment or promoted after deployment
3. Ensure that firewall rules allow internal access to the load balancer on HTTP port 80

b) Configuring the Alias IP

1. Using the GCP console, edit the load balancer instance
2. Scroll down to the *Network Interfaces* section and click the edit button
3. Click **Show alias IP ranges**.

Alias IP ranges

Subnet range: Primary (10.154.0.0/20)

Alias IP range: 10.154.0.50

+ Add IP range

[Hide alias IP ranges](#)

4. Enter the required Alias IP – this will be used for the new VIP (the /32 will be added automatically)
5. Click **Done**, click **Save**

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:

Virtual Service

Manual Configuration ☐

Label: WebCluster2

IP Address: 10.154.0.50

Ports: 80

Protocol

Layer 7 Protocol: HTTP Mode

Cancel Update

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**

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:

Label	<input type="text" value="Web1"/>	?
Real Server IP Address	<input type="text" value="10.154.0.18"/>	?
Real Server Port	<input type="text" value="80"/>	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	<input type="text" value="100"/>	?

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. **10.154.0.18**
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.

e) Testing

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

EXAMPLE 3 – WEB SERVERS: SINGLE SUBNET, LAYER 4, PUBLIC FACING

This is a layer 4 (NAT mode) example using one subnet for both the load balancer and the web servers. In this example, VPC routing must be configured to route return traffic from the web servers via the load balancer. 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.

a) Deploying the GCP Instances

1. Deploy the load balancer instance as described in the section [Deploying Enterprise GCP from the Marketplace](#), configure a static External IP address – this must be done after deployment
2. Deploy your web server instances into the same VPC & subnet as the load balancer, use a static Internal IP address – this can either be reserved in advance and assigned during deployment or promoted after deployment
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80

b) Configuring GCP Routing

1. Using the GCP console, select **Routes** in the VPC menu
2. Click **CREATE ROUTE**

Name *
route-via-lb

Lowercase letters, numbers, hyphens allowed

Description
routes return traffic via the load balancer

Network *
default

Destination IP range *
0.0.0.0/0

E.g. 10.0.0.0/16

Priority *
900

Priority should be a positive integer (lower values take precedence)

Instance tags
route-via-lb

Next hop
Specify an instance

Next hop instance *
lb100-vm

CREATE **CANCEL**

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

c) Tagging the Web Server instances

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

c) Setting up the Virtual Service

1. Using the WebUI, navigate to: Cluster Configuration > Layer 4 – Virtual Service and click **Add a**

New Virtual Service

2. Enter the following details:

Virtual Service		
Label	<input type="text" value="WebCluster1"/>	?
IP Address	<input type="text" value="10.154.0.17"/>	?
Ports	<input type="text" value="80"/>	?
Protocol		
Protocol	<input type="text" value="TCP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="NAT"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

3. Enter an appropriate *Label* for the VIP, e.g. **WebCluster1**
4. Set the *IP address* to the **Base IP address**, e.g. **10.154.0.17**
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**

d) Setting up 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:

Label	<input type="text" value="Web1"/>	?
Real Server IP Address	<input type="text" value="10.154.0.18"/>	?
Real Server Port	<input type="text" value="80"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?
		<input type="button" value="Cancel"/> <input type="button" value="Update"/>

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. **10.154.0.18**
5. Set the *Real Server Port* to the required port, e.g. **80**
6. Click **Update**

7. Repeat the above steps to add your other web server(s)

e) Testing

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

EXAMPLE 4 – USING GCP LOAD BALANCER TO PRESENT MULTIPLE SERVICES ON DIFFERENT PUBLIC IPS

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. In this case there is no need to configure a health-check on the GCP load balancer, traffic will always be forwarded to the Loadbalancer.org appliance.

The following 3 services will be configured:

- **Service 1** – Web Cluster 1 presented on public IP 1 (the External IP of the instance)
- **Service 2** – Web Cluster 2 presented on public IP 2 (additional External IP reserved for service 2)
- **Service 3** – Custom application 1 presented on public IP 3, TCP/UDP port 2050 (additional External IP reserved for service 3)

a) Deploying the GCP Instances

1. Deploy the load balancer instance as described in the section [Deploying Enterprise GCP from the Marketplace](#), configure a static External IP address – this must be done after deployment
2. Deploy your web server instances and custom application instances to the same VPC & subnet as the load balancer, use a static Internal IP address – this can either be reserved in advance and assigned during deployment or promoted after deployment
3. Ensure that firewall rules allow external access to the load balancer on HTTP port 80

b) Configure Service 1

1. Configure service 1 using [Example 1 – Web Servers: Single Subnet, Layer 7, Public facing](#) as a guide, Service 1 will then be presented on the Loadbalancer.org appliance's public IP address

c) Add a GCP Load Balancer & Configure Service 2

1. Navigate in the GCP Console to *Network Services > Load balancing* and click **Create Load balancer**
2. Under *TCP Load Balancing* click **Start Configuration**

Internet facing or internal only

Do you want to load balance traffic from the Internet to your VMs or only between VMs in your network?

☒ From Internet to my VMs
☐ Only between my VMs

Multiple regions or single region

Do you want to place the backends for your load balancer in a single region or across multiple regions?

☐ Multiple regions (or not sure yet)
☒ Single region only

Backend type

Do you want to use target pool or Regional Backend Service as a backend?

☒ Target pool or target instance
☐ Backend service ?

[Continue](#)

3. Leave the defaults and click **Continue**
4. Type a *Name* for the load balancer, e.g. **gcplb1**
5. Click **Backend Configuration**

Name ?

gcplb1

Region ?

europa-west2

Backends ?

[Select existing instance groups](#)
[Select existing instances](#)

lb100-vm (europa-west2-c) ✕

Add an instance

Backup pool ? (Optional)

None

Failover ratio ?

10 %

Health check ?

No health check

Session affinity ?

None

6. Set the *Region* according to your requirements
7. Under *Backends* click **Select existing instances**
8. In the *Add an instance* drop-down, select the Loadbalancer.org load balancer instance
9. Leave the remaining settings at their default values
10. Click **Frontend Configuration** and then click **Add Frontend IP and port**

11. Enter a suitable *Name*, e.g. **service2**
12. Under *IP*, select **Create IP address**
13. Enter an appropriate *Name*, e.g. **service2**
14. Click **Reserve**
15. Specify the required *Port*, e.g. **80**
16. Click **Done**
17. Click **Review and Finalize** – there will be a warning concerning no health checks but that can be ignored as we want all traffic to be sent to the Loadbalancer.org appliance
18. Click **Create**

d) Configure Service 3 (TCP)

1. Under the load balancing menu, click **advanced menu**
2. Click **CREATE FORWARDING RULE**

Name ?
Name is permanent

service3-tcp

Description (Optional)

Region ?

europe-west2

External IP ?

service3 (34.89.121.70)

Network Service Tier ?

Premium

Protocol ?

TCP

Port/range ?

80

Target pool ?

gcplb1

Create Cancel

3. Enter a suitable *Name*, e.g. **service3-tcp**
4. Select the appropriate *Region*
5. Under *External IP*, select **Create IP address**
6. Enter an appropriate *Name*, e.g. **service3**
7. Click **Reserve**
8. Set the *Protocol* to **TCP**
9. Specify the required *Port*, e.g. **2050**
10. Under *Target pool* select the GCP load balancer just created, e.g **gcplb1**
11. Click **Create**

d) Configure Service 3 (UDP)

1. Click **CREATE FORWARDING RULE**
2. Enter a suitable *Name*, e.g. **service3-udp**
3. Select the appropriate *Region*
4. Under *External IP*, select **service3**
5. Set the *Protocol* to **UDP**
6. Specify the required *Port*, e.g. **2050**
7. Under *Target pool* select the GCP load balancer just created, e.g **gcplb1**

8. Click **Create**

e) Verify forwarding Rules

The following forwarding rules will be displayed:

<input type="checkbox"/> Name ^	Description	Type	Region	Address	Protocol	Target
<input type="checkbox"/> service2		Regional	europe-west2	34.89.33.219:80-80	tcp	gcplb1
<input type="checkbox"/> service3-tcp		Regional	europe-west2	34.89.121.70:2050-2050	tcp	gcplb1
<input type="checkbox"/> service3-udp		Regional	europe-west2	34.89.121.70:2050-2050	udp	gcplb1

f) Configure Routing to ensure return traffic for service3 routes back form the application servers via the loadbalancer.org instance as required by Layer 4 NAT mode

1. Configure routing for the application servers using [Example 3 – Web Servers: Single Subnet, Layer 4, Public Facing](#) as a guide

g) Configure VIPs on the Loadbalancer.org appliance

1. Connect to the Loadbalancer.org appliance WebUI

Configure Service 1

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

Virtual Service

Manual Configuration ☐ ?

Label ?

IP Address ?

Ports ?

Protocol

Layer 7 Protocol ?

Cancel

Update

2. Enter an appropriate *label*, e.g. **Service1**
3. Enter the appliance's Internal IP address, e.g. **10.154.0.17**
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

Label	Web1	?
Real Server IP Address	10.154.0.18	?
Real Server Port	80	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	100	?

Cancel
Update

8. Enter an appropriate *label*. e.g. **Web1**
9. Set the *Real Server IP Address* field to the required IP address, e.g. **10.154.0.18**
10. Set the *Real Server port* field to **80**
11. Click **Update**
12. Repeat the above steps to add your other real server(s)

Configure Service 2

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

Virtual Service		
Manual Configuration	<input type="checkbox"/>	?
Label	service2	?
IP Address	34.89.33.219	?
Ports	80	?
Protocol		
Layer 7 Protocol	HTTP Mode ▼	?

Cancel
Update

2. Enter an appropriate *label*, e.g. **Service2**
3. Enter the External IP that was reserved for Service2 in the GCP console, e.g. **34.89.33.219**
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

Label	Web1	?
Real Server IP Address	10.154.0.19	?
Real Server Port	80	?
Re-Encrypt to Backend	<input type="checkbox"/>	?
Weight	100	?

Cancel Update

- Enter an appropriate *label*. e.g. **Web1**
- Set the *Real Server IP Address* field to the required IP address, e.g. **10.154.0.19**
- Set the *Real Server port* field to **80**
- Click **Update**
- Repeat the above steps to add your other real server(s)

Configure Service 3 – TCP

- Navigate to *Cluster Configuration > Layer 4 Virtual Services* and click **Add a New Virtual Service**

Virtual Service		
Label	Service3-tcp	?
IP Address	34.80.121.70	?
Ports	2050	?
Protocol		
Protocol	TCP	?
Forwarding		
Forwarding Method	NAT	?

Cancel Update

- Enter an appropriate *label*, e.g. **Service3-tcp**
- Enter the External IP that was reserved for Service2 in the GCP console, e.g. **34.89.121.70**
- Set the *Ports* field to **2050**
- Set the *Protocol* set to **TCP**
- Set the *Forwarding Method* to **NAT**
- Click **Update**
- Navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created VIP

Label	<input type="text" value="app1"/>	?
Real Server IP Address	<input type="text" value="19.154.0.20"/>	?
Real Server Port	<input type="text" value="2050"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

Cancel Update

9. Enter an appropriate *label*. e.g. **app1**
10. Set the *Real Server IP Address* field to the required IP address, e.g. **10.154.0.20**
11. Set the *Real Server port* field to **80**
12. Click **Update**
13. Repeat the above steps to add your other real server(s)

Configure Service 3 – UDP

1. Navigate to *Cluster Configuration > Layer 4 Virtual Services* and click **Add a New Virtual Service**

Virtual Service		
Label	<input type="text" value="Service3-udp"/>	?
IP Address	<input type="text" value="34.80.121.70"/>	?
Ports	<input type="text" value="2050"/>	?
Protocol		
Protocol	<input type="text" value="UDP"/>	?
Forwarding		
Forwarding Method	<input type="text" value="NAT"/>	?

Cancel Update

2. Enter an appropriate *label*, e.g. **Service3**
3. Enter the External IP that was reserved for Service2 in the GCP console, e.g. **34.89.121.70**
4. Set the *Ports* field to **2050**
5. Set the *Protocol* set to **UDP**
6. Set the *Forwarding Method* to **NAT**
7. Click **Update**
8. Navigate to: *Cluster Configuration > Layer 4 – Real Servers* and click **Add a new Real Server** next to the newly created VIP

Label	<input type="text" value="app1"/>	?
Real Server IP Address	<input type="text" value="19.154.0.20"/>	?
Real Server Port	<input type="text" value="2050"/>	?
Weight	<input type="text" value="100"/>	?
Minimum Connections	<input type="text" value="0"/>	?
Maximum Connections	<input type="text" value="0"/>	?

Cancel Update

9. Enter an appropriate *label*. e.g. **app1**
10. Set the *Real Server IP Address* field to the required IP address, e.g. **10.154.0.20**
11. Set the *Real Server port* field to **80**
12. Click **Update**
13. Repeat the above steps to add your other real server(s)

g) Verify VIPs

The VIPs will be displayed in the *System Overview* of the WebUI

System Overview ?								2021-06-15 15:19:01 UTC	
	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE		
↑	Service1	10.154.0.17	80	0	HTTP	Layer 7	Proxy		
↑	Service2	34.89.33.219	80	0	HTTP	Layer 7	Proxy		
↑	Service3-tcp	34.89.121.70	2050	0	TCP	Layer 4	NAT		
↑	Service3-udp	34.89.121.70	2050	0	UDP	Layer 4	NAT		

h) Testing

1. Verify that **service1:80** is available on the external IP of the Loadbalancer.org appliance
2. Verify that **service2:80** is available on the IP address reserved for service 2
3. Verify that **service3:2050 (TCP & UDP)** is available on the IP address reserved for service 3

8. 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).

SYSTEM OVERVIEW ?							2015-03-18 11:37:15 UTC
	VIRTUAL SERVICE	IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy
	RDP-Cluster	192.168.110.150	3389	0	TCP	Layer 7	Proxy
	HTTP-Cluster-2	192.168.110.152	80	0	HTTP	Layer 7	Proxy
	RDP-Cluster-2	192.168.110.152	3389	0	TCP	Layer 7	Proxy

3. **Check the connection state**

For layer 4 NAT mode VIPs, check *Reports > Layer 4 Current Connections* to view the current traffic in detail. Any packets with state **SYN_RECV** often implies a return traffic routing issue:

- for single subnet Layer 4 mode make sure that the default gateway on all real servers is set to be the load balancer
- for dual subnet Layer 4 mode make sure that routing on the second subnet has been configured correctly

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

username: loadbalancer
password: <password>

The password is set to a random arbitrary value when first booted. To view the password,

using the WebUI navigate to: View Configuration > Layer 7, scroll to the “listen stats” section at the bottom of the screen and the password can be found after the colon on the “stats auth” line, e.g.

```
stats auth loadbalancer:EBAIA2h2i2NJ
```

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):

Statistics Report for pid 3261

> General process information

pid = 3261 (process #1, nbproc = 1)

uptime = 0d 0h00m42s

system limits: memmax = unlimited; ulimit-n = 81000

maxsock = 80024; maxconn = 40000; maxpipes = 0

current conns = 1; current pipes = 0/0; conn rate = 2/sec

Running tasks: 1/5; idle = 100 %

active UP

active UP, going down

active DOWN, going up

active or backup DOWN

active or backup DOWN for maintenance (MAINT)

backup UP

backup UP, going down

backup DOWN, going up

not checked

Display option:

- Hide 'DOWN' servers
- Refresh now
- CSV export

External resources:

- Primary site
- Updates (v1.5)
- Online manual

Note: UP with load-balancing disabled is reported as "NOLB".

L7

	Queue			Session rate			Sessions			Bytes		Denied		Errors		Warnings		Server												
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntime	Thrtle	
Frontend				0	15	-	0	4		40 000	56	21 696	3 385 782	0	0	0	0	0	0	0	OPEN									
backup	0	0	-	0	0	0	0	4		0	0	0	0	0	0	0	0	0	0	0			1	-	Y					-
RIP1	0	0	-	0	16	-	0	2		56	56	21 696	3 385 782	0	0	0	0	0	0	0	42s UP	L4OK in 0ms	1	Y	-	0	0	0s	-	
Backend	0	0		0	16	-	0	2		4 000	56	21 696	3 385 782	0	0	0	0	0	0	0	42s UP		1	1	1		0	0s		

stats

	Queue			Session rate			Sessions			Bytes		Denied		Errors		Warnings		Server												
	Cur	Max	Limit	Cur	Max	Limit	Cur	Max	Limit	Total	LbTot	In	Out	Req	Resp	Req	Conn	Resp	Retr	Redis	Status	LastChk	Wght	Act	Bck	Chk	Dwn	Dwntime	Thrtle	
Frontend				2	4	-	1	1		2 000	8	1 464	33 111	0	0	4	0	0	0	0	OPEN									
Backend	0	0		0	0	-	0	0		200	0	1 464	33 111	0	0	0	0	0	0	0	42s UP		0	0	0		0	0		

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:

SYSTEM OVERVIEW ? 2015-04-30 08:35:41 UTC

VIRTUAL SERVICE		IP	PORTS	CONNS	PROTOCOL	METHOD	MODE
!	HTTP-Cluster	192.168.110.150	80	0	HTTP	Layer 7	Proxy
REAL SERVER		IP	PORTS	WEIGHT	CONNS		
↑	RIP1	192.168.110.240	80	100	0	Drain	Halt
⚙	RIP2	192.168.110.241	80	0	0	Online (halt)	
↓	RIP3	192.168.110.242	80	100	0	Drain	Halt

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.

9. 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>

10. 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

11. Company Contact Information

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