Enterprise GCP Configuration Guide

Version 8.9.1 Revision 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	5
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
Security Mode	12
Passwords	12
Checking For Updates	12
Appliance Licensing	13
Accessing the Appliance using SSH	13
Generating SSH Keys	13
Update the GCP project metadata	15
Accessing the Appliance from Linux	16
Accessing the Appliance from Windows using PuTTY	16
Accessing the Appliance with the GCP User Credentials	17
7. Deployment examples	17
Example 1 – Web Servers: Single Subnet, Layer 7, Public facing	17
Example 2 – Web Servers: Single Subnet, Layer 7, Internal Facing	19
Example 3 – Web Servers: Single Subnet, Layer 4, Public Facing.	21
Example 4 – using GCP Load balancer To Present Multiple Services on Different Public IPs	24
8. Testing & Verification	33
9. More Information	33
10. Loadbalancer.org Technical Support	33
Contacting Support	33

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 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.

	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.
ំ Note	Multiple services on different internal IP addresses can be configured using Alias IPs. For a deployment example please refer to 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 Example 4 – using GCP Load balancer To Present Multiple Services on Different Public IPs.

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 – 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.
- 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 – 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.

5. 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

15

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.

dh.

New Loadbalancer.org Enterprise GCP BYOL deployment

Deployment name Ib1		
Zone		
europe-west2-c		• 0
Machine type		
Machine family		
GENERAL-PURPOSE	COMPUTE-OPTIMISED	MEMORY-OPTIMISED
Aachine types for common Series N1	n workloads, optimised for co	st and flexibility
Machine type n1-standard-1 (1 vCPU,	3.75 GB memory)	▼
	VCPU	Memory
V	1	3.75 GB
CPU PLATFORM AND	GPU	
Enter an appropria	ate Deployment Name	
 Select the require 	d 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	•	0
Boot disk size in GB 10		Ø

Networking

Network interfaces

^
- 0
- 0
- 0
DONE

- 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

និ 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

լեր

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

Source IP ranges for TCP port 9443 traffic		0
IP forwarding		
off	•	0

SHOW LESS



15

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

Name	Туре	Targets	Filters	Protocols/ports	Action
test-allow-icmp	Ingress	Apply to all	IP ranges: 0.0.0.0/0	icmp	Allow
test-allow-internal	Ingress	Apply to all	IP ranges: 10.128.0.0/9	all	Allow
test-allow-rdp	Ingress	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow
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.

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

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



- 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

րել

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.

Network interface	Reserve static internal IP address	
You must stop the VM in Network	Reserve IP address 10.154.0.12 Name @ Name is permanent	
Subnetwork 🔘	Ib1-internal-ip	
default (10.154.0.0/2	Description (Optional)	
Internal IP 10.154.0.12	LB1 Internal IP	14.
Internal IP type		
Ephemeral		

- 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.

8 Jun 2021, 13:58:09	
Network interfaces	Reserve a new static IP address
Network interface	Name 🕖 Name is permanent
	Ib1-external-ip
You must stop the VM in Network	Description (Optional)
default	LB1 External IP
Subnetwork	lin.
default (10.154.0.0/2 Internal IP Ib1-internal-ip (10.154.0.	Network Service Tier Premium (current project-level tier, change) Standard Region europe-west2
Static	Standard tier uses the same region as your VM instance
Show alias IP range External IP Output	CANCEL RESERVE
Create IP address	Reserved Article 2010

- 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

8 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

	To change the password for the 'loadbalancer' account, use the WebUI option: <i>Maintenance ></i>
ß Note	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:

dh.

LOADBALANCER

Enterprise GCP Max



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

Appliance Security

րել

8 Note

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:
 - 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 Local Configuration > Execute shell command menu option is disabled
 - The ability to edit the firewall script & the firewall lockdown wizard script is disabled
 - 'root' user console & SSH password access are disabled
- Secure Permanent This mode is the same as Secure but once set it cannot be changed

(!) 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.
- Select the required Appliance Security Mode if Custom is selected, configure the additional options according to your requirements.
- 3. Click Update.

Passwords

The 'loadbalancer' WebUI account

The password for the '**loadbalance**r' WebUI user account is set during instance deployment. 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

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.

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

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.

-								
PuTT	Y Key Generator						?	×
le <u>K</u> ey	Con <u>v</u> ersions	<u>H</u> elp						
Key								
No key.								
• •								
Actions								
Actions Generat	e a public/private	key pair				Ge	nerate	
Actions Generat Load an	e a public/private n existing private ka	key pair ev file				Ge	enerate Load	
Actions Generat Load an	e a public/private a existing private ka	key pair ey file		0		Ge	nerate Load	
Actions Generat Load an Save th	e a public/private e existing private ke e generated key	key pair ey file		Save	p <u>u</u> blic key	<u>G</u> e	enerate Load private k	еу
Actions Generat Load an Save the Paramet	e a public/private n existing private ke e generated key ers	key pair ey file		Save	p <u>u</u> blic key	<u>G</u> e <u>I</u> Save p	enerate Load private k	еу
Actions Generat Load an Save th Paramet Type of	e a public/private e existing private ke e generated key ers key to generate:	key pair ey file	[Save	p <u>u</u> blic key	<u>G</u> e <u>I</u> <u>S</u> ave p	enerate Load	ey
Actions Generat Load an Save th Paramet Type of (• <u>R</u> SA	e a public/private a existing private ka e generated key ers key to generate: <u>D</u> SA	key pair ey file	⊖ <u>e</u> cds/	Save	public key () ED <u>2</u> 5519	<u>G</u> e <u>I</u> <u>S</u> ave p	enerate Load private k SSH- <u>1</u> (R	ey (SA)

2. Click the **Generate** button.

լեր

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

Ľ	PuTTY Key Generate	or				?	\times
F	ile <u>K</u> ey Con <u>v</u> ersion	s <u>F</u>	lelp				
	Key Public key for pasting in ssh-rsa	ito Op	enSSH authorize	d_keys	file:		^
	AddaB3NzaC19C2EA McvqFp2drrQt64N6Nk R6wVaDqKxINKjRzxhl +Yxr/Cb8RZwPe/He4	WfCE NAAx bsKR	Z3W5W4SMRN NeWDI DDzU0VZFv3/w	UAZHH 09GzjE	rKpGN79yRdm/k	ke/5v82Wz2ukk	~
	Key fingerprint:	ssh-n	sa 2047 07:64:19):e0:47:	dd:93:ba:0d:40:a	2:da:99:68:c9:0e	
	Key <u>c</u> omment:	root					
	Key p <u>a</u> ssphrase:	l					
	Confirm passphrase:						
	Actions						
	Generate a public/priva	te ke	y pair			<u>G</u> enerate	
	Load an existing private	keyf	ile		[<u>L</u> oad	
	Save the generated ke	y		Sav	e p <u>u</u> blic key	<u>S</u> ave private key	
	Parameters						
	Type of key to generate <u>R</u> SA <u>D</u>	e: SA		SA	◯ ED <u>2</u> 5519	○ SSH- <u>1</u> (RS/	4)
	Number of <u>b</u> its in a gen	erated	l key:			2048	

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

15

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:

Metadata SSH Keys	
robert	ssh-rsa AAAAB3NzaClyc2EAAAADAQABAAABAQDJ8C68SFzaKcS1n7uF2W5BGfggdEncQWf7/iofW99JFQ jES&178TwDfT/ifCEUCK6rNYUu4ipJgxqw8IeNIT1sfY/HsGMI1uWVJY/6S5BTbQaQYzYcE54mkLIH+JzZ lxc7VH03WoLwbs+f4bgtDZanAnhiDq2cOPdxB0Y2qcXVNh0LM00Dc3pbV/M845DSY2Qu5BDNJQHgmmxVAK oE1p0kET30zmFbRYmJuzcbazSTnPdotnFUHqvBtvSZAor7ViewYevXgVcXYFJLCgEeWP/+tvbGUsJ2McO0 oDqklQjT/fZNeH3EHiqtSGx8W/jPQp/s0YqBEXQi1aDL0gP5+nJv robert@cs-407518999553-defaul t-boost-xf2wp
oot	ssh-rsa AAAAB3NzaClyc2EAAAABJQAAAQEAwToiuyqYV1JM/lJm7/wId8Jk1hAZjCkE2bL1Tbpr6sn//u JJwUhPFGVZDTYKcXT1b0Wz7+/8LocRLHT9pq1lp+qvJmoU5jyxdgA6Gpg+wV2y1SvzMPw+Un2M1yIG38y ma5a24R/tPpt8wL72gX87SSvTBSbjaSOEUnGeBpTjISKVokyAgurdKhGqz8Rz5Dsjf9IXzAq0rEeRHwO4h WR/6KIC4+ObTtN28C2T6hMYBYnAu8HjgyMbi4fJ0DZZxZC12b5x1pHZ8wIjKsRwKUw4nukRbv+bGjdypY1 eu7dzbn8zY9pTDtkRcHJgvATcdnqQ6zUs3xEC9S4zrSiTeo5oQ== root
	+ Add item

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 "root" user, 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.

րել։

5. Login using root and the password you have specified (default is loadbalancer).



 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.

Step 1 - 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.
- 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.

Step 2 - Setting up the Virtual Service

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

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]	
Label	WebCluster1		?
IP Address	10.154.0.17		7
Ports	80		?
Protocol			
Layer 7 Protocol	HTTP Mode 🗸		0
		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.

Step 3 - Setting up the Real Servers

- 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:

Layer 7 Add a new Real Server

Label	Web1		?
Real Server IP Address	10.154.0.18		9
Real Server Port	80		?
Re-Encrypt to Backend			9
Enable Redirect			?
Weight	100		0
		Cancel	Update

- 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.

րել։

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

Step 4 - 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.

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.

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.

Step 1 - 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.
- 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.

Step 2 - 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		Alias IP range 📀	
Primary (10.154.0.0/20)	*	10.154.0.50	×
· · ·	+ Add	IP range	

- A 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**.

լեր

Step 3 - Setting up the Virtual Service

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

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]	
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.

լեր

Step 4 - Setting up the Real Servers

- 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:

Layer 7 Add a new Real Server

Label	Web1		•
Real Server IP Address	10.154.0.18		0
Real Server Port	80		?
Re-Encrypt to Backend			9
Enable Redirect			?
Weight	100		0
		Cancel	Update

- 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).

Step 5 - 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.

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.

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.

Step 1 - 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.
- 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.

Step 2 - Configuring GCP Routing

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

Name *		
route-via-lb		e
Lowercase letters, numbers, hyphens allowed		
Description		
routes return traffic via the load balancer		
Network *		
default	•	(
Destination IP range *		
0.0.0/0		(
E.g. 10.0.0/16		
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.

Step 3 - 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.

րել։

Step 4 - Setting up the Virtual Service

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

Layer 4 - Add a new Virtual Service

Virtual Service				
Label	WebCluster1			?
IP Address	10.154.0.17			?
Ports	80			?
Protocol				
Protocol	ТСР	~		?
Forwarding				
Forwarding Method	NAT 🗸			?
			Cancel	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.

Step 5 - Setting up the Real Servers

- 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	Web1		?
Real Server IP Address	10.154.0.18		?
Real Server Port	80		?
Weight	100		?
Minimum Connections	0		?
Maximum Connections	0		?
		Cancel	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).

Step 6 - 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)

Step 1 - 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.
- 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.

Step 2 - Configure Service 1

15

 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.

Step 3 - 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 ②
	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 🕘		
europe-west2		•
Backends 📀		
Select existing instance groups	Select existing instances]
lb100-vm (europe-west2-c)		×
Add an instance		-
Backup pool ② (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.

New Frontend IP and port	Ē	^
Name (Optional) 📀 Name is permanent		
service2		
Add a description		
Protocol TCP		
Network Service Tier Premium (current project-level tier, change) Standard		
IP		
service2 (34.89.33.219)		-
Port		
80		
Done Cancel		

- 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.

15

Step 4 - Configure Service 3 (TCP)

1. Under the load balancing menu, click advanced menu.

2. Click CREATE FORWARDING RULE.

Name 🛞	
service3-tcp	
Description (Optional)	
	//
Region Ø	
europe-west2	•
External IP 📀	
service3 (34.89.121.70)	-
Premium Protocol ②	
TCP	•
Port/range 🕢	
80	
Target pool 💿	
gcplb1	*
Create Cancel	
nter a suitable <i>Name</i> , 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.

З.

- 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.

Step 5 - 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.

Step 6 - Verify forwarding Rules

The following forwarding rules will be displayed:

Name ^	Description	Туре	Region	Address	Protocol	Target
service2		Regional	europe-west2	34.89.33.219:80-80	tcp	gcplb1
service3-tcp		Regional	europe-west2	34.89.121.70:2050-2050	tcp	gcplb1
service3-udp		Regional	europe-west2	34.89.121.70:2050-2050	udp	gcplb1

Step 7 - 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.

Step 8 - Configure VIPs on the Loadbalancer.org appliance

1. Connect to the Loadbalancer.org appliance WebUI.

Configure Service 1

լեղ,

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

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]	
Label	Service1		?
IP Address	10.154.0.17		?
Ports	80		?
Protocol			
Layer 7 Protocol	HTTP Mode 🗸		?
		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.

Layer 7 Add a new Real Server

Label	Web1		?
Real Server IP Address	10.154.0.18		0
Real Server Port	80		?
Re-Encrypt to Backend			?
Enable Redirect			?
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**.

Layer 7 - Add a new Virtual Service

Virtual Service		[Advanced +]	
Label	service2		2
IP Address	34.89.33.219		?
Ports	80		0
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.

Layer 7 Add a new Real Server

Label	Web1		?
Real Server IP Address	10.154.0.19		9
Real Server Port	80		?
Re-Encrypt to Backend			0
Enable Redirect			?
Weight	100		0
		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. 10.154.0.19.
- 10. Set the *Real Server port* field to 80.
- 11. Click Update.

15

12. Repeat the above steps to add your other real server(s).

Configure Service 3 – TCP

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

Layer 4 - Add a new Virtual Service

Virtual Service				
Label	Service3-tcp			0
IP Address	34.80.121.70			U
Ports	2050			?
Protocol				
Protocol	ТСР	~		?
Forwarding				
Forwarding Method	NAT 🗸			?
			Cancel	Update

- 2. Enter an appropriate *label*, e.g. Service3-tcp.
- 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* to **TCP**.
- 6. Set the *Forwarding Method* to **NAT**.
- 7. Click Update.
- Navigate to: Cluster Configuration > Layer 4 Real Servers and click Add a new *Real Server next to the newly created VIP.

Label	app1		?
Real Server IP Address	19.154.0.20		?
Real Server Port	2050		9
Weight	100		0
Minimum Connections	0		9
Maximum Connections	0		9
		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.

լեղ,

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.

							· ·
Lay	/er 4	4 -	Add	а	new	virtual	Service

Virtual Service				
Label	Service3-udp			9
IP Address	34.80.121.70			?
Ports	2050			•
Protocol				
Protocol	UDP	~		?
Forwarding				
Forwarding Method	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* to **UDP**.
- 6. Set the *Forwarding Method* to **NAT**.
- 7. Click Update.
- Navigate to: Cluster Configuration > Layer 4 Real Servers and click Add a new Real Server next to the newly created VIP.

Label	app1		2
Real Server IP Address	19.154.0.20		9
Real Server Port	2050		0
Weight	100		0
Minimum Connections	0		9
Maximum Connections	0		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.

լեղ,

13. Repeat the above steps to add your other real server(s).

Step 9 - Verify VIPs

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

System Overview 👔 2021-06-15 15:19:01 UTC									:01 UTC
		VIRTUAL SERVICE 🗢	IP 🗢	PORTS 🗢	CONNS 🗢	PROTOCOL 🗢	METHOD 🗢	MODE 🗢	
	1	Service1	10.154.0.17	80	0	HTTP	Layer 7	Proxy	<u>8.41</u>
	1	Service2	34.89.33.219	80	0	HTTP	Layer 7	Proxy	<u>8.41</u>
	1	Service3-tcp	34.89.121.70	2050	0	ТСР	Layer 4	NAT	<u>8.41</u>
	1	Service3-udp	34.89.121.70	2050	0	UDP	Layer 4	NAT	<u></u>

Step 10 - 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 & 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.

Contacting Support

dh.

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.

IL 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.

