Intel® Endpoint Management Assistant (Intel® EMA)

Deployment Guide for Google Cloud Platform

Intel® EMA Version 1.7

February 2022

Legal Disclaimer

This software and the related documents are provided as is, with no express or implied warranties, other than those that are expressly stated in the License.

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses. Check with your system manufacturer or retailer or learn more at http://www.intel.com/technology/vpro.

Contents

1 Introduction

This document describes the procedure to deploy infrastructure to Google Cloud Platform (GCP), a cloud computing platform, needed to support one or more instances of the Intel® Endpoint Management Assistant (Intel® EMA) server. It is intended for IT administrators with intermediate to advanced knowledge of IT infrastructure who may have limited knowledge about cloud computing.

There are several components needed for a complete cloud infrastructure environment, so we recommend that you read this guide carefully to understand how they are configured to work together. A description of each component is provided before the deployment procedure, with a link to the official cloud provider documentation for further information if needed.

1.1 About Cloud Computing

Cloud computing is the on-demand delivery of IT resources over the Internet with pay-as-you-go pricing. Instead of buying, owning, and maintaining physical data centers and servers, you can access technology services, such as computing power, storage, and databases, on an as-needed basis from a cloud provider. You can provision only what you need now and scale capacity to grow and shrink as business needs change.

Large cloud providers have data centers all around the world, allowing you to deploy resources geographically close to where your customers and end users are located.

With fully-managed services like Cloud SQL, you can just focus on your data while the cloud provider manages all of the underlying hardware and software that provide the service. With virtual machines running in the cloud, you manage only the guest operating system and the software installed on it, while the cloud provider manages the underlying hardware and strives to provide you with the best reliability and availability.

1.2 Navigating in the GCP Console

1.2.1 Services Menu

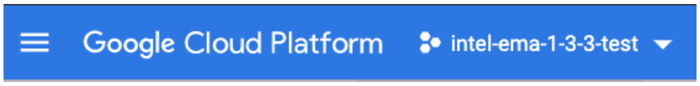	After logging into the GCP console at https://console.cloud.google.com/, you will see a services menu icon in the top left corner. To the right of it you will see a Project menu where you will select the project into which you will deploy your resources, after a Project has been created.

1.2.2 Expanding the Services Menu

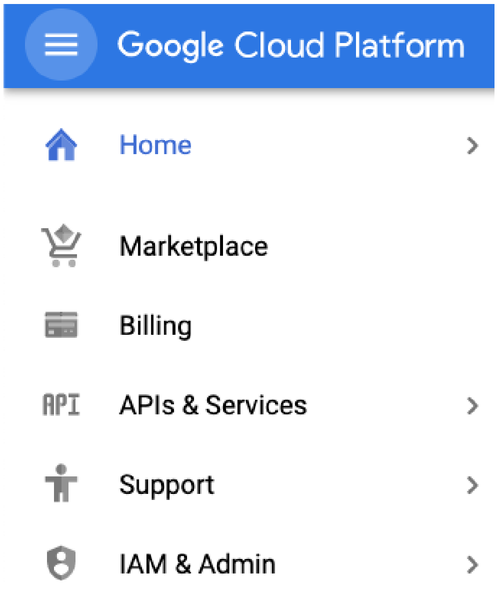

When you click on the services menu icon, you will see a list of services listed below, organized into sections like COMPUTE, STORAGE, and others.

In this guide, we will provide instructions directing you to select a service from this menu when we are deploying the various components that we will need.

1.3 How Resources are Organized in GCP

All resources in GCP are deployed into a Project. If you have an account as an individual, then that is the only structure that you will have. If you have an Organization account, then Projects can be located directly under the Organization node, or they can optionally be grouped together into Folders which are under the Organization node.

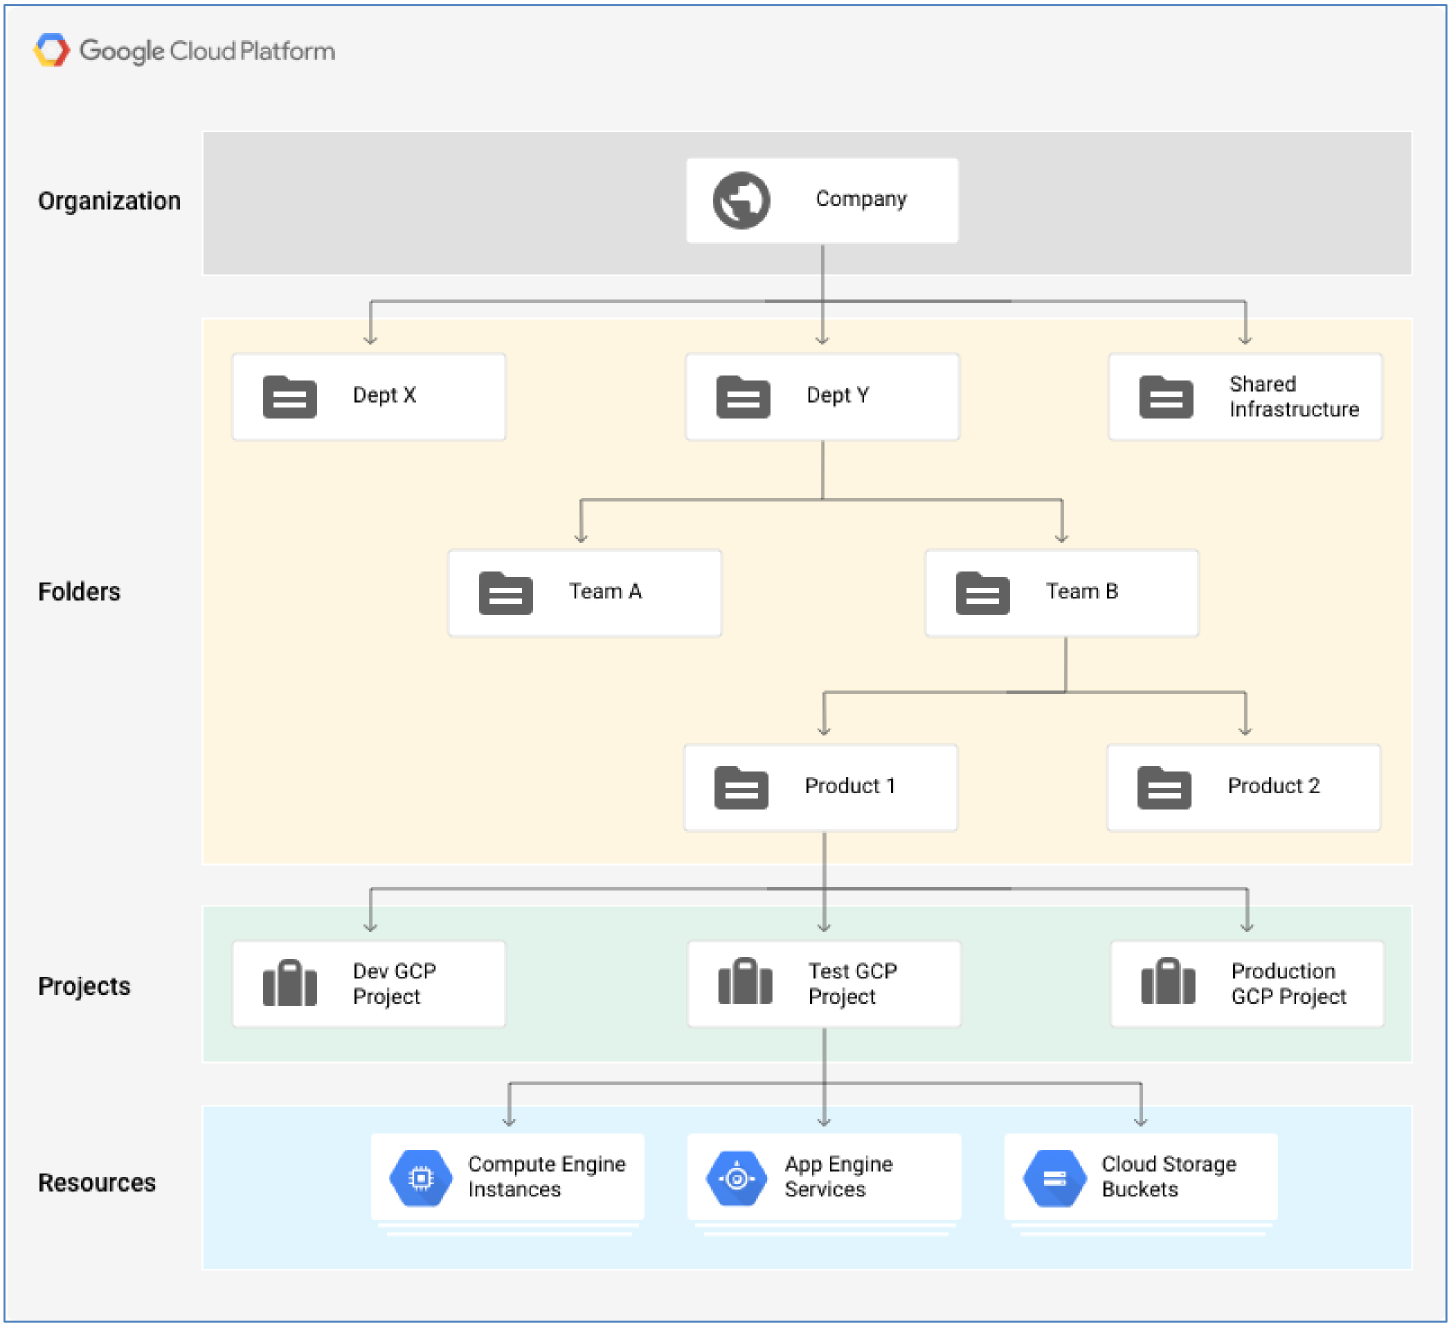

1.4 Before You Begin

If your organization already has a GCP account, then you should ask to have a cloud administrator create a Project for you and give you Project Owner access. If you are the cloud administrator, then you can go to the **IAM & Admin > Manage Resource** menu in GCP to create the project yourself.

If your organization doesn't have a GCP account, or you want to evaluate it as an individual, then you can go to https://console.cloud.google.com/ and sign in with a Google account and then you will be able to start a free trial with promotional credit included.

Check with your network administrator to ask if there is a preferred address space to use. You will want to avoid overlapping with your corporate network to prevent routing issues if you already have a VPN established to the cloud provider, or if you will in the future. You will also want to find out what the source IP address will be for traffic leaving your organization to reach the cloud so that you'll be able to allow only trusted networks to reach the Intel EMA virtual machine from the internet.

2 High-level Architecture Diagrams

2.1 Single Server Deployment

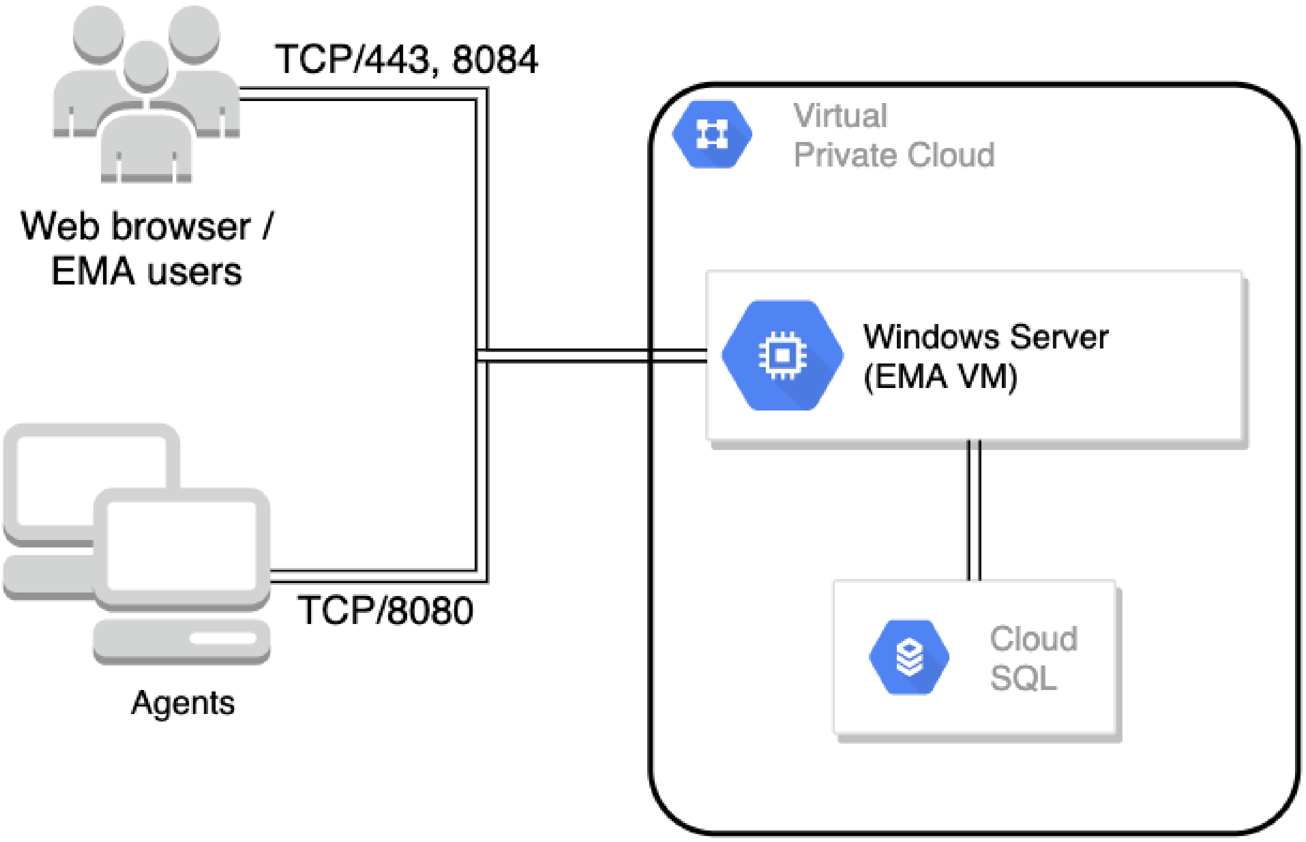

2.2 Distributed Server Deployment

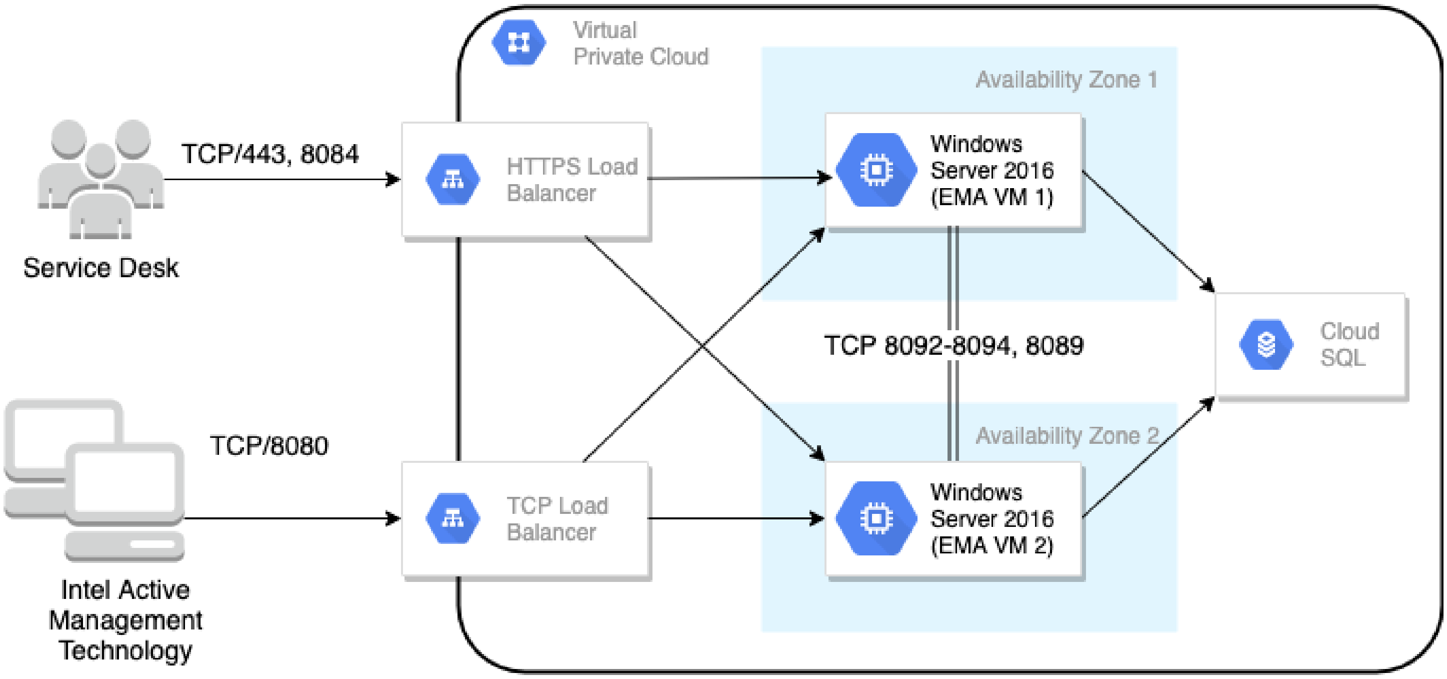

3 Network Deployment

3.1 Overview

In order for virtual machines to communicate with each other, with the cloud provider, or with the internet, we first need to configure a network environment. A Virtual Private Cloud network (VPC network) is the fundamental building block for your private network in GCP, and it closely resembles a traditional network except that it is virtualized within GCP. A VPC network is a global resource that consists of a list of regional virtual subnetworks (subnets) in data centers, all connected by a global wide area network. VPC networks are logically isolated from each other.

When creating a VPC network you will need to provide a custom private IP address space. GCP will assign resources a private IP address from this address space when needed. You should consult with your network engineering team to identify an available IP address block to use to avoid routing conflicts in case your company already has private IP connectivity to the cloud or will in the future.

We will also need to allocate an IP block for private services access to allow virtual machine(s) to access Google services through a private connection rather than through public endpoints.

When we create the VPC network, we'll also need to create at least one subnet. Subnets enable you to segment the virtual network by allocating a portion of the virtual network's address space to each subnet. You can then deploy resources into a specific subnet.

For further information about services deployed in this section, see the following links:

- VPC: https://cloud.google.com/vpc/docs
- Private Google Access: https://cloud.google.com/vpc/docs/configure-private-google-access
- Cloud NAT: https://cloud.google.com/nat/docs/overview
- Cloud Router: https://cloud.google.com/network-connectivity/docs/router

3.2 Virtual Private Cloud Network

Follow this procedure to create a VPC network with a single subnet

3.2.1 Navigate to VPC Networks

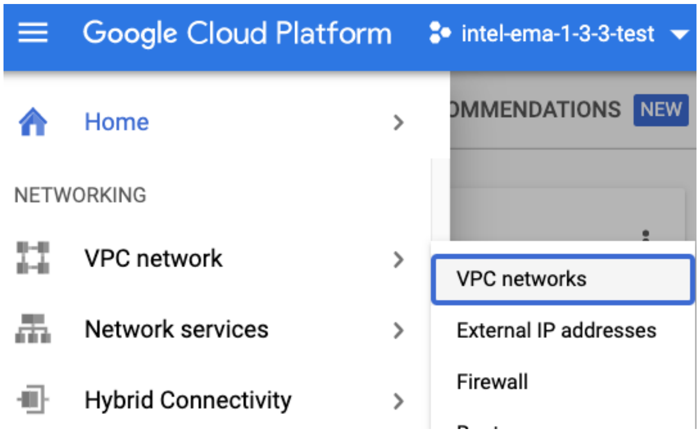 Google Cloud Platform intel-ema-1-3-3-test Home NETWORKING VPC network Network services Hybrid Connectivity OMMENDATIONS NEW VPC networks External IP addresses Firewall	From the service menu, go to **Networking** > **VPC network** > **VPC networks**

3.2.2 Create VPC Network

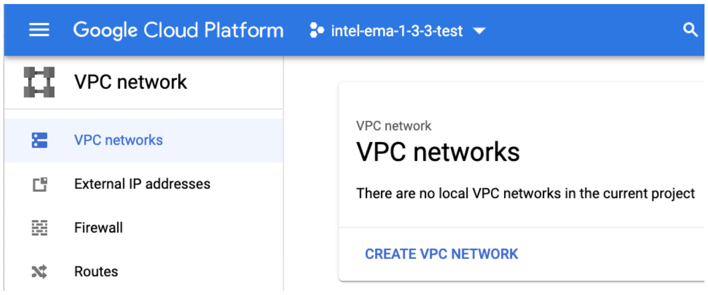 Google Cloud Platform intel-ema-1-3-3-test VPC network VPC networks External IP addresses Firewall Routes VPC network VPC networks There are no local VPC networks in the current project CREATE VPC NETWORK	Click **CREATE VPC NETWORK**

3.2.3 Configure VPC Network

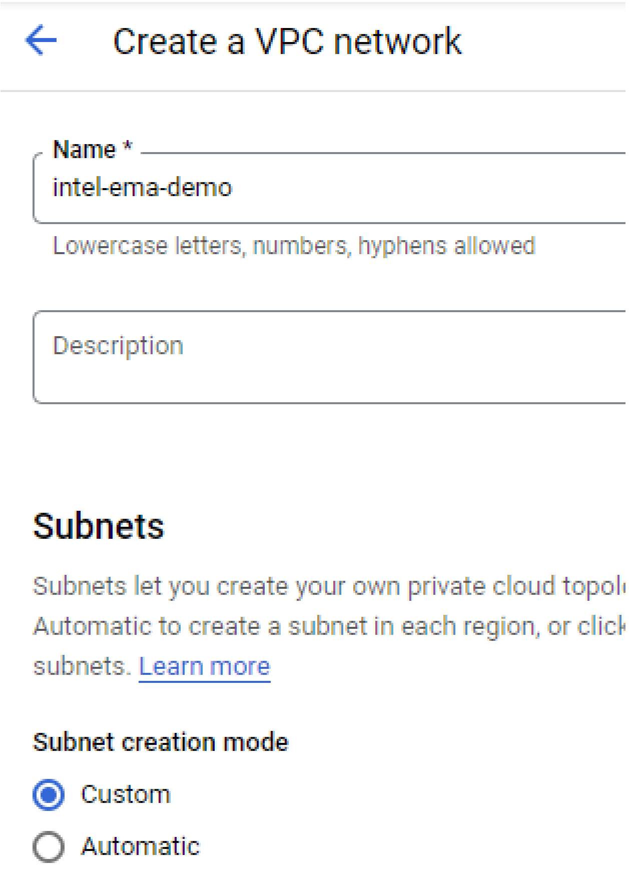 Create a VPC network Name * intel-ema-demo Lowercase letters, numbers, hyphens allowed Description Subnets Subnets let you create your own private cloud topol Automatic to create a subnet in each region, or click subnets. Learn more Subnet creation mode ● Custom ○ Automatic	Configure the VPC as follows: • **Name**: Enter a unique name Example: *intel-ema-demo* • **Subnet creation mode**: *Custom*

3.2.4 Add a Subnet

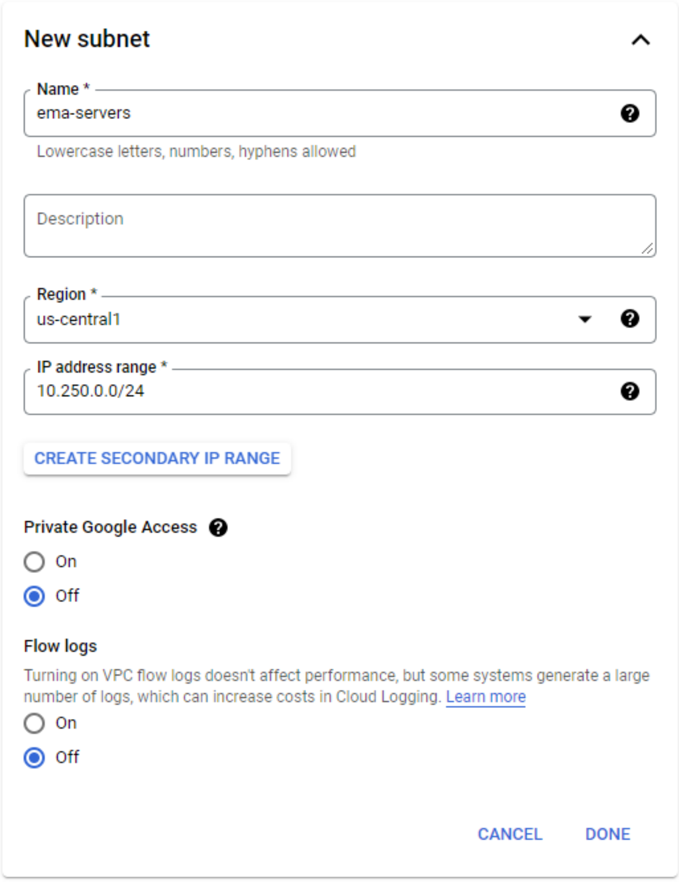

New subnet

Name *: ema-servers

Lowercase letters, numbers, hyphens allowed

Description

Region *: us-central1

IP address range *: 10.250.0.0/24

CREATE SECONDARY IP RANGE

Private Google Access

○ On

● Off

Flow logs

Turning on VPC flow logs doesn't affect performance, but some systems generate a large number of logs, which can increase costs in Cloud Logging. Learn more

○ On

● Off

CANCEL DONE

Configure the **New subnet** section as follows:

- **Name**: Enter a unique subnet name
  Example: *ema-servers*
- **Region**: Choose a region where you want to deploy resources
  Example: *us-central1*
- **IP address range**: Enter an IP address range to use
  Example: *10.250.0.0/24*
- **Private Google access:** *On*

Click the **Done** button.

3.2.5 Finalize VPC

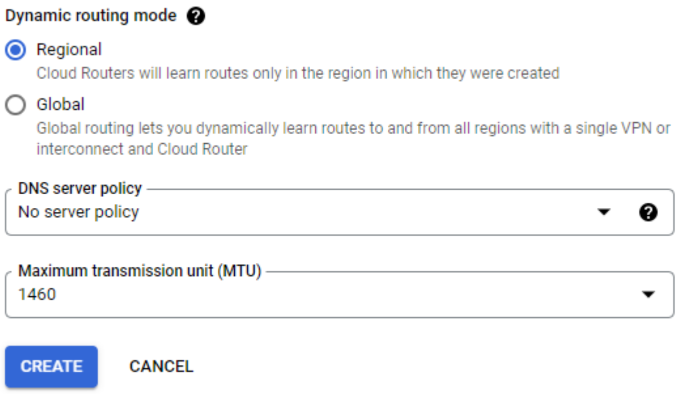

Dynamic routing mode

● Regional
Cloud Routers will learn routes only in the region in which they were created

○ Global
Global routing lets you dynamically learn routes to and from all regions with a single VPN or interconnect and Cloud Router

DNS server policy: No server policy

Maximum transmission unit (MTU): 1460

CREATE CANCEL

You can leave the rest of the settings with default values.

Click the **Create** button to finalize the VPC network.

3.2.6 Go into the VPC Network Details

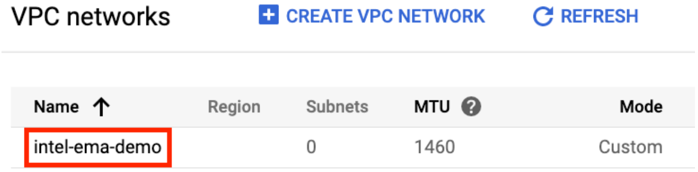

VPC networks

Name ↑	Region	Subnets	MTU	Mode
intel-ema-demo		0	1460	Custom

Click on the name of the newly created VPC to go the details screen.

3.2.7 Allocate Private Service Connection IP Range

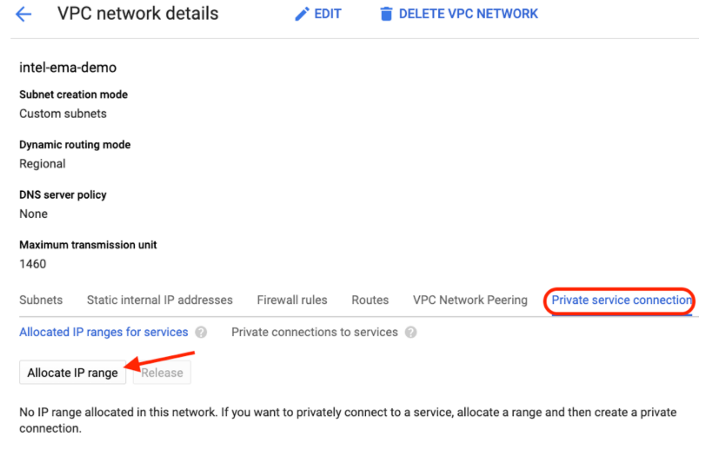
VPC network details
EDIT
DELETE VPC NETWORK

intel-ema-demo

Subnet creation mode
Custom subnets

Dynamic routing mode
Regional

DNS server policy
None

Maximum transmission unit
1460

Subnets
Static internal IP addresses
Firewall rules
Routes
VPC Network Peering
Private service connection
Allocated IP ranges for services
Private connections to services
Allocate IP range
Release

No IP range allocated in this network. If you want to privately connect to a service, allocate a range and then create a private connection.

Click on **Private service connection**.

Click the **Allocate IP range** button.

3.2.8 Enter Private Service IP Range Details

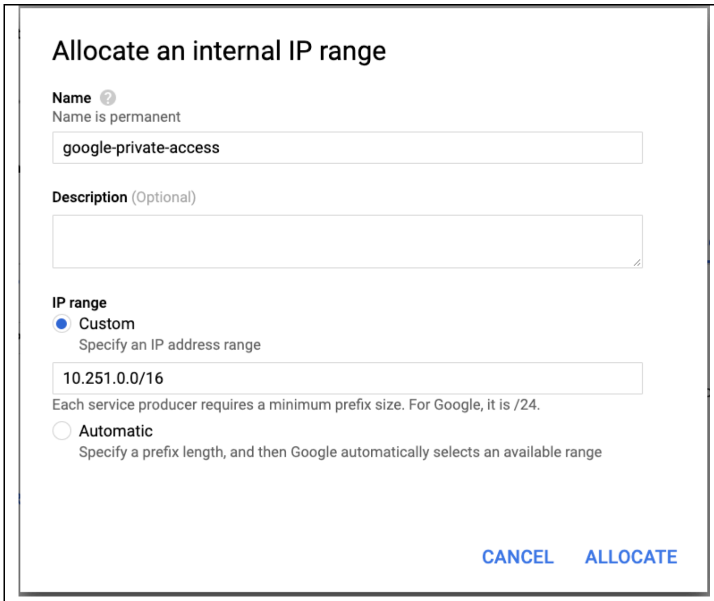
Allocate an internal IP range

Name
Name is permanent
google-private-access

Description (Optional)

IP range
● Custom
Specify an IP address range
10.251.0.0/16
Each service producer requires a minimum prefix size. For Google, it is /24.
○ Automatic
Specify a prefix length, and then Google automatically selects an available range

Configure the IP allocation as follows:

- Name: Enter a unique name for the IP range.
  Example: *google-private-access*
- IP range: *Custom*
  Enter an unused IP address range. Google requires at least a /24 prefix size, but recommends /16.
  Example: *10.251.0.0/16*

Click the **Allocate** button.

3.3 Firewall Rules

Each VPC network implements a distributed virtual firewall that you can configure. Firewall rules allow you to control which packets are allowed to travel to which destinations. Every VPC network has two implied firewall rules that block all incoming connections and allow all outgoing connections.

One of the ways that we will specify a target or destination is using tags, which will later be applied to the virtual machine(s) in order to make the related firewall rules take effect for those VMs.

For more information about using a VPC Firewall, visit the following link: https://cloud.google.com/vpc/docs/firewalls

3.3.1 Navigate to Firewall Rules

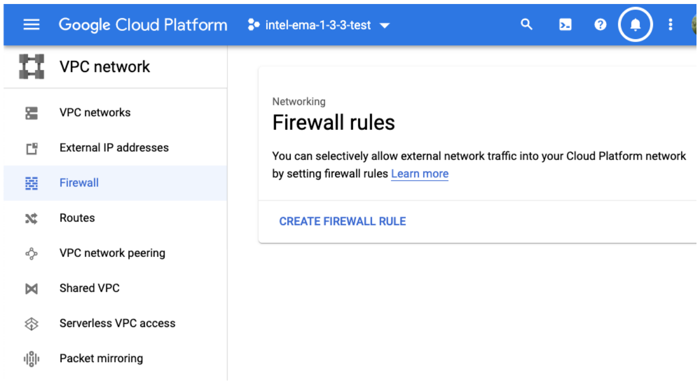 Google Cloud Platform intel-ema-1-3-3-test VPC network VPC networks External IP addresses Firewall Routes VPC network peering Shared VPC Serverless VPC access Packet mirroring Networking Firewall rules You can selectively allow external network traffic into your Cloud Platform network by setting firewall rules Learn more 	From the service menu, go to **Networking** > **VPC network** > **Firewall**. Click **CREATE FIREWALL RULE**.

3.3.2 Create a Firewall Rule for RDP Traffic

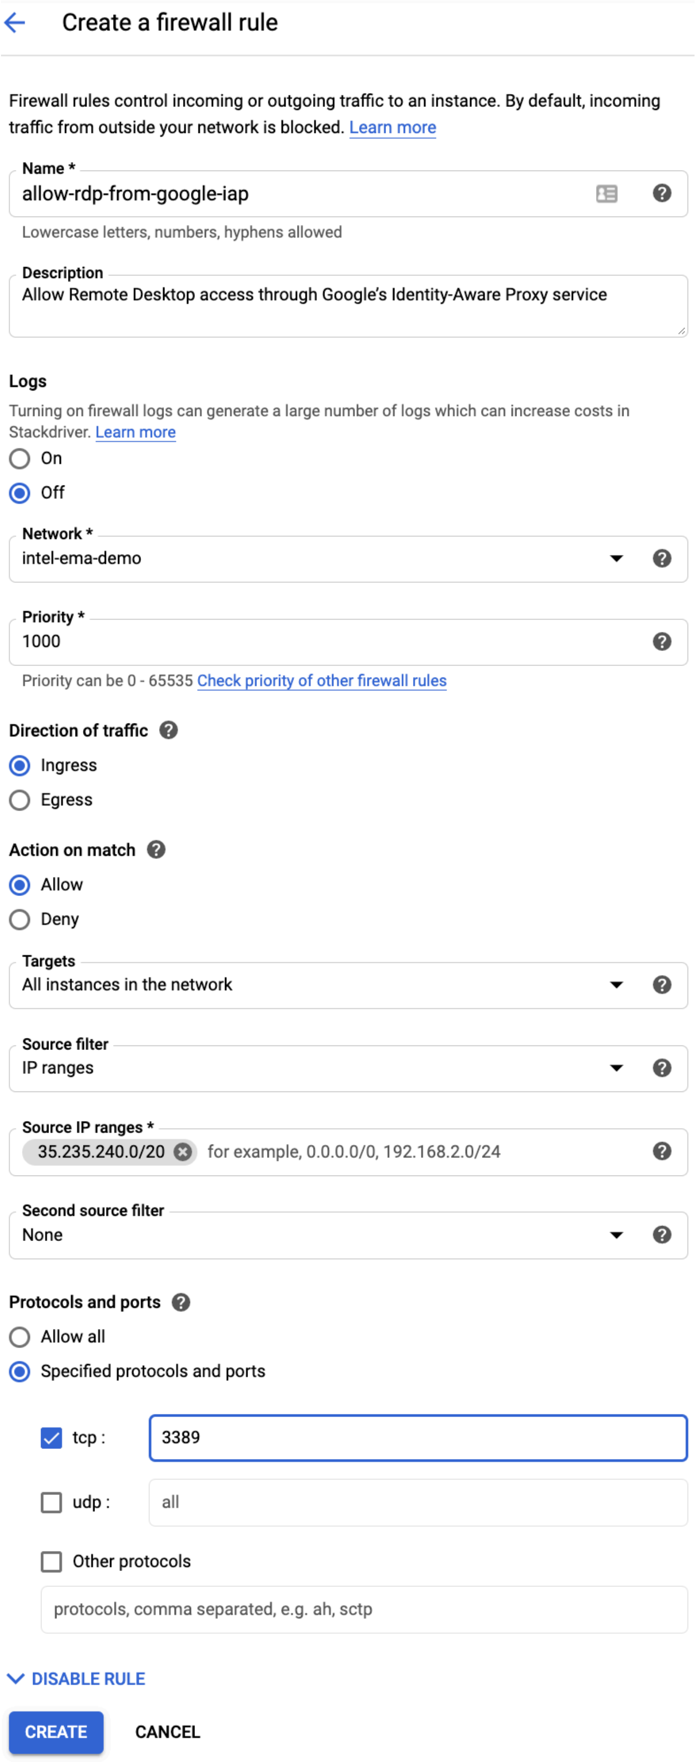

← Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. Learn more

Name *
allow-rdp-from-google-iap
Lowercase letters, numbers, hyphens allowed

Description
Allow Remote Desktop access through Google's Identity-Aware Proxy service

Logs
Turning on firewall logs can generate a large number of logs which can increase costs in Stackdriver. Learn more

○ On
● Off

Network *
intel-ema-demo

Priority *
1000
Priority can be 0 - 65535 Check priority of other firewall rules

Direction of traffic
● Ingress
○ Egress

Action on match
● Allow
○ Deny

Targets
All instances in the network

Source filter
IP ranges

Source IP ranges *
35.235.240.0/20 for example, 0.0.0.0/0, 192.168.2.0/24

Second source filter
None

Protocols and ports
○ Allow all
● Specified protocols and ports

☑ tcp : 3389
□ udp : all
□ Other protocols
protocols, comma separated, e.g. ah, sctp

We need to allow ingress traffic from the Google IP ranged used by the Identity-Aware Proxy (IAP) service, which we will use to log into virtual machines.

Configure the firewall rule as follows.

- **Name**: Enter a unique name
  Example: *allow-rdp-from-google-iap*
- **Description**: *Allow Remote Desktop access through Google's Identity-Aware Proxy service*
- **Network**: Select the VPC that you previously created
- **Targets**: *All instances in the network*
- **Source filter**: *IP ranges*
- **Source ip ranges:** *35.235.240.0/20*
- **Specified ports protocols and ports**:
  tcp: *3389*

Click the **Create** button.

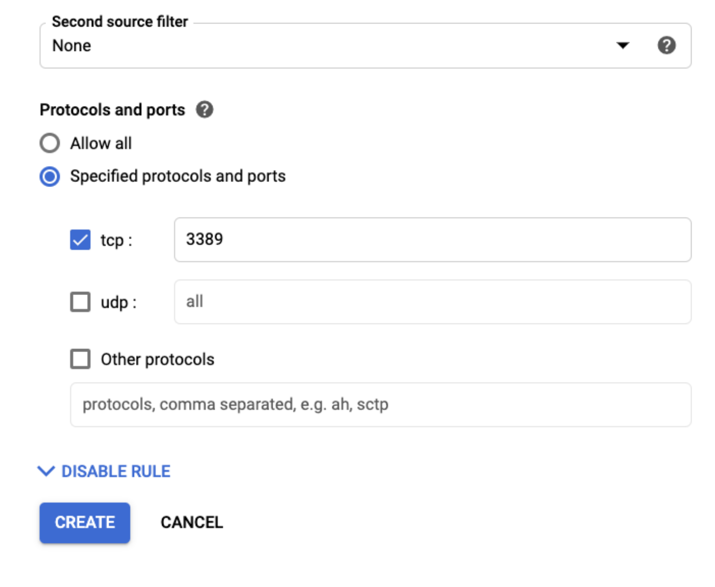

Second source filter
None

Protocols and ports

○ Allow all

● Specified protocols and ports

☑ tcp : 3389

□ udp : all

□ Other protocols

protocols, comma separated, e.g. ah, sctp

3.3.3 Create a Firewall Rule for Web Traffic (Single Server Deployment Only)

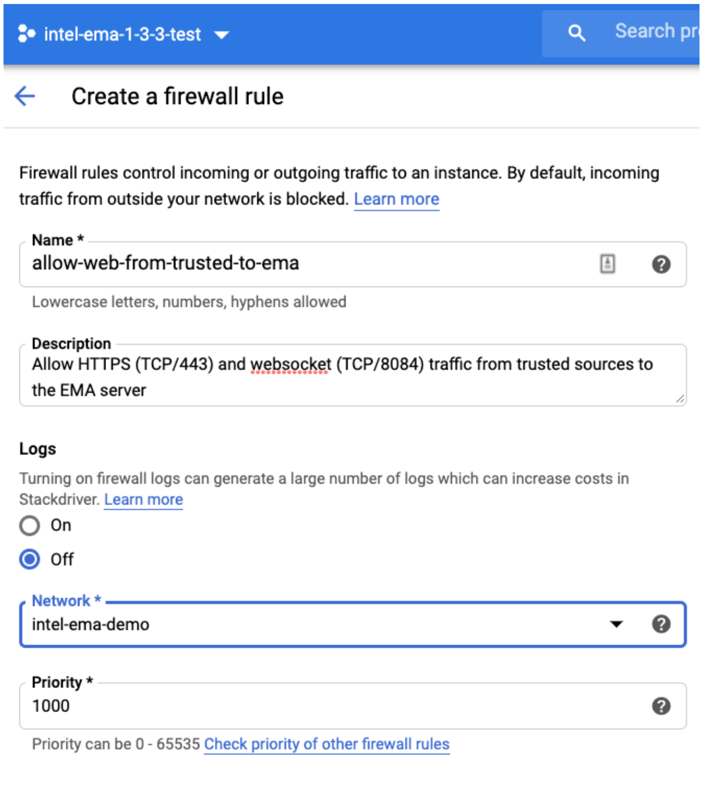

Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. Learn more

Name *
allow-web-from-trusted-to-ema

Lowercase letters, numbers, hyphens allowed

Description
Allow HTTPS (TCP/443) and websocket (TCP/8084) traffic from trusted sources to the EMA server

Logs

Turning on firewall logs can generate a large number of logs which can increase costs in Stackdriver. Learn more

○ On

● Off

Network *
intel-ema-demo

Priority *
1000

Priority can be 0 - 65535 Check priority of other firewall rules

Create a new firewall rule and configure as follows.

- **Name**: Enter a unique name
  Example: *allow-web-from-trusted-to-ema*
- **Description**: *Allow web traffic from trusted sources to the server*
- **Network**: Select the VPC that you previously created
- **Targets**: *Specified target tags*
- **Target tags**: *ema-server*
- **Source filter**: *IP ranges*
- **Source IP ranges:** Enter the trusted network(s) that should have access
- **Specified ports protocols and ports**:
  tcp: *443,8084*

Click the **Create** button.

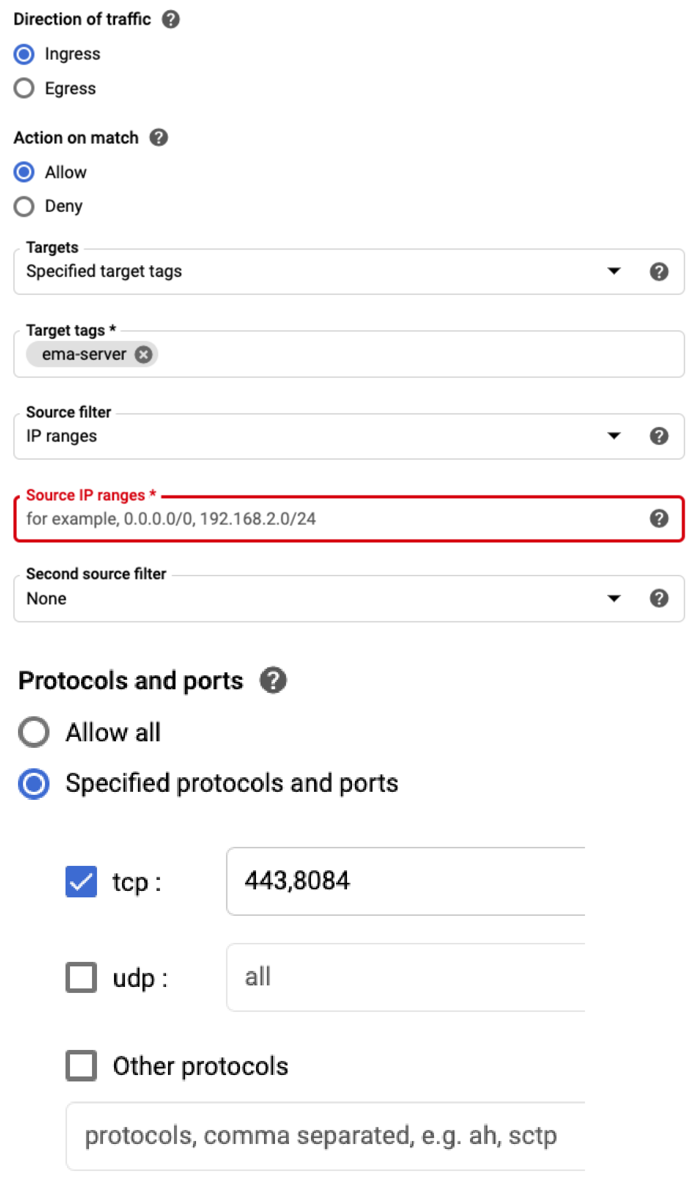 Direction of traffic ● Ingress ○ Egress Action on match ● Allow ○ Deny Targets: Specified target tags Target tags *: ema-server Source filter: IP ranges Source IP ranges *: for example, 0.0.0.0/0, 192.168.2.0/24 Second source filter: None Protocols and ports ○ Allow all ● Specified protocols and ports ☑ tcp : 443,8084 □ udp : all □ Other protocols protocols, comma separated, e.g. ah, sctp	

3.3.4 Create a Firewall Rule for Web Traffic (Distributed Server Deployment Only)

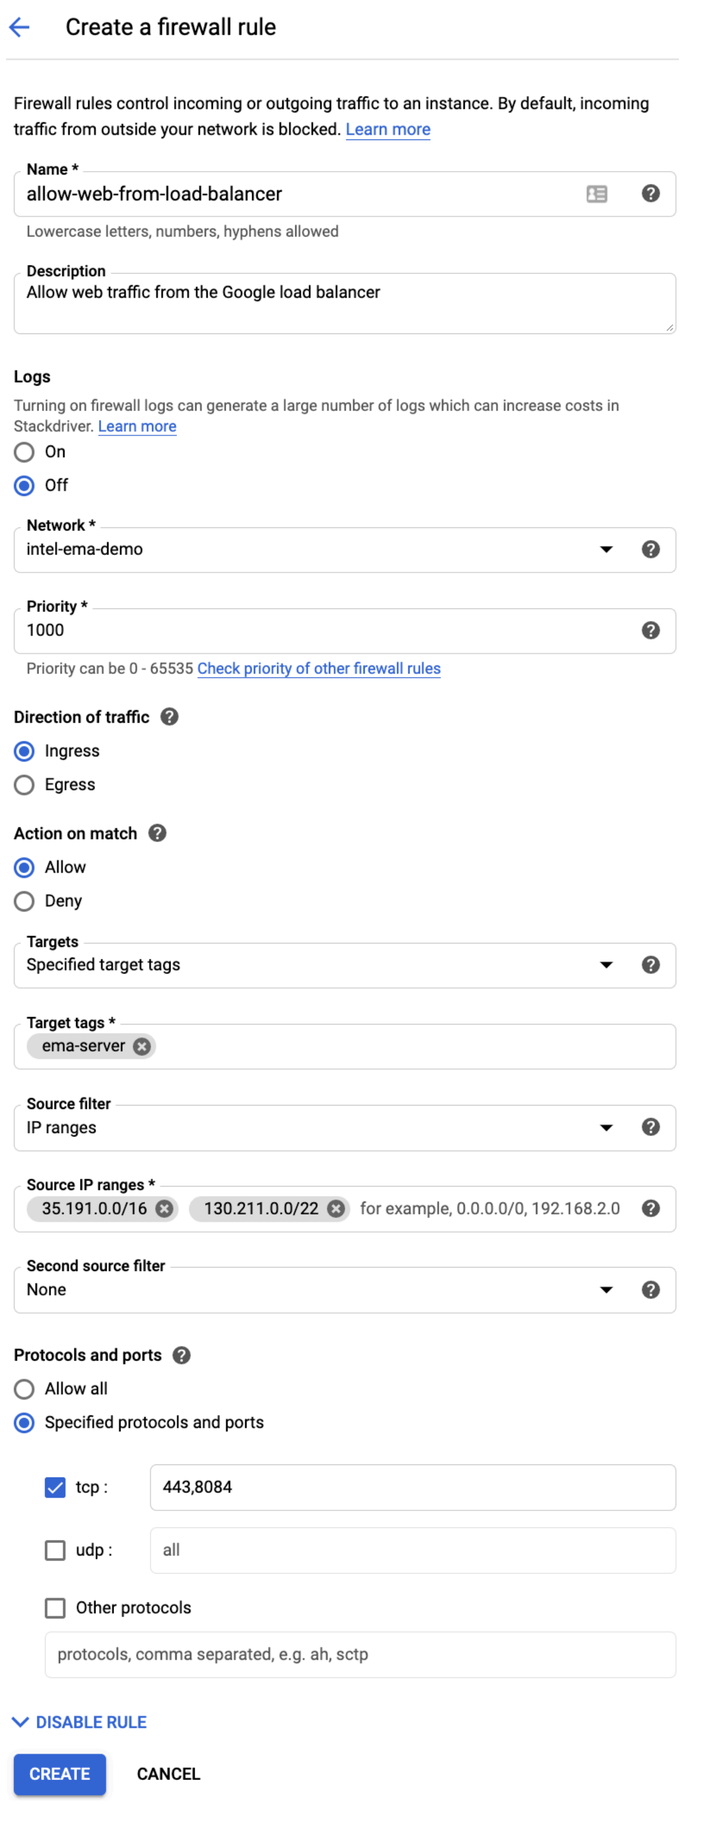

Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. Learn more

Name *
allow-web-from-load-balancer
Lowercase letters, numbers, hyphens allowed

Description
Allow web traffic from the Google load balancer

Logs

Turning on firewall logs can generate a large number of logs which can increase costs in Stackdriver. Learn more

○ On
● Off

Network *
intel-ema-demo

Priority *
1000
Priority can be 0 - 65535 Check priority of other firewall rules

Direction of traffic

● Ingress
○ Egress

Action on match

● Allow
○ Deny

Targets
Specified target tags

Target tags *
ema-server

Source filter
IP ranges

Source IP ranges *
35.191.0.0/16 130.211.0.0/22 for example, 0.0.0.0/0, 192.168.2.0

Second source filter
None

Protocols and ports

○ Allow all
● Specified protocols and ports

☑ tcp : 443,8084

☐ udp : all

☐ Other protocols

protocols, comma separated, e.g. ah, sctp

Create a new firewall rule and configure as follows.

- **Name**: Enter a unique name
  Example: *allow-web-from-load-balancer*
- **Description**: Allow web traffic from the Google load balancer
- **Network**: Select the VPC that you previously created
- **Targets**: *Specified target tags*
- **Target tags**: *ema-server*
- **Source filter**: IP ranges
- **Source IP ranges**:
  - *35.191.0.0/16*
  - *130.211.0.0/22*
- **Specified ports protocols and ports**:
  - tcp: *443,8084*

Click the **Create** button.

3.3.5 Create a Firewall Rule for Swarm Traffic

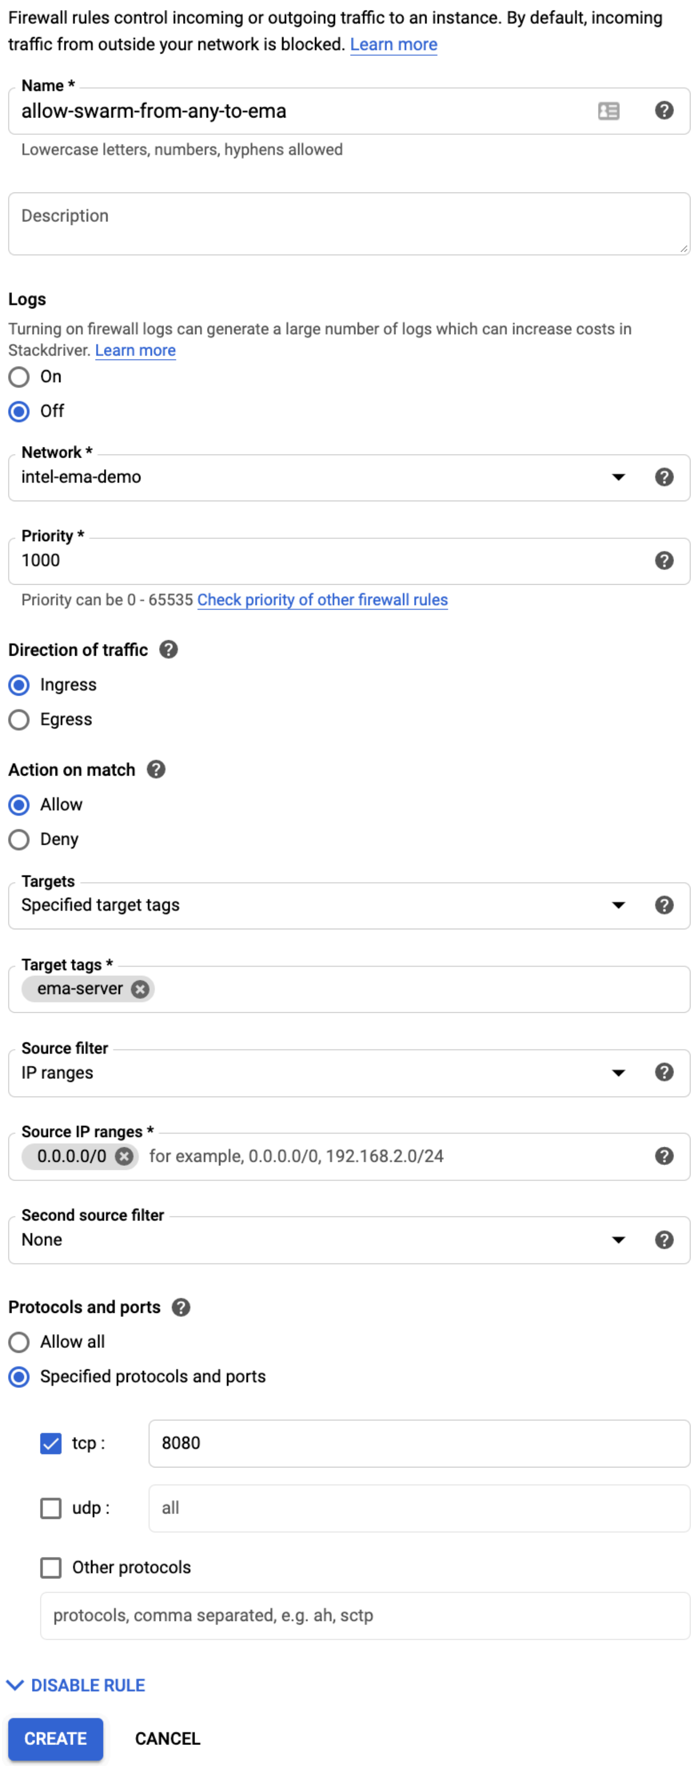

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. Learn more

Name *
allow-swarm-from-any-to-ema
Lowercase letters, numbers, hyphens allowed

Description

Logs
Turning on firewall logs can generate a large number of logs which can increase costs in Stackdriver. Learn more
○ On
● Off

Network *
intel-ema-demo

Priority *
1000
Priority can be 0 - 65535 Check priority of other firewall rules

Direction of traffic
● Ingress
○ Egress

Action on match
● Allow
○ Deny

Targets
Specified target tags

Target tags *
ema-server

Source filter
IP ranges

Source IP ranges *
0.0.0.0/0 for example, 0.0.0.0/0, 192.168.2.0/24

Second source filter
None

Protocols and ports
○ Allow all
● Specified protocols and ports

☑ tcp : 8080
□ udp : all
□ Other protocols
protocols, comma separated, e.g. ah, sctp

Create a new firewall rule and configure as follows.

- **Name**: Enter a unique name
  Example: *allow-swarm-from-any-to-ema*
- **Description**: *Allow EMA agent traffic from anywhere to the server*
- **Network**: Select the VPC that you previously created
- **Targets**: *Specified target tags*
- **Target tags**: *ema-server*
- **Source filter**: *IP ranges*
- **Source IP ranges:** *0.0.0.0/0*
- **Specified ports protocols and ports**:
  tcp: *8080*

Click the **Create** button.

3.3.6 Create a Firewall rule for Server-to-server Traffic (Distributed Server Deployment Only)

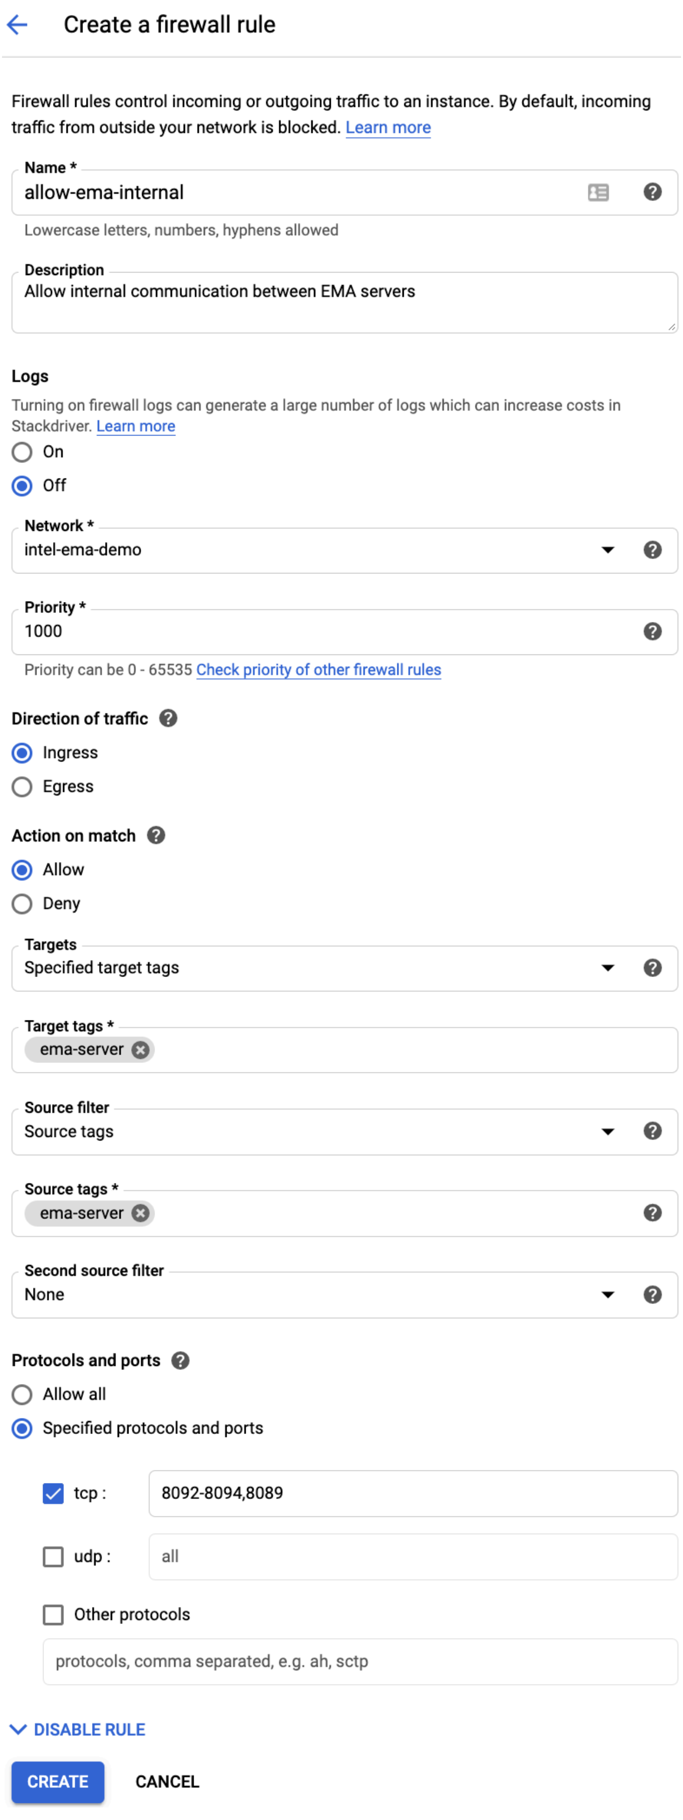

← Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. Learn more

Name *
allow-ema-internal
Lowercase letters, numbers, hyphens allowed

Description
Allow internal communication between EMA servers

Logs
Turning on firewall logs can generate a large number of logs which can increase costs in Stackdriver. Learn more
○ On
● Off

Network *
intel-ema-demo

Priority *
1000
Priority can be 0 - 65535 Check priority of other firewall rules

Direction of traffic
● Ingress
○ Egress

Action on match
● Allow
○ Deny

Targets
Specified target tags

Target tags *
ema-server

Source filter
Source tags

Source tags *
ema-server

Second source filter
None

Protocols and ports
○ Allow all
● Specified protocols and ports

☑ tcp : 8092-8094,8089
□ udp : all
□ Other protocols
protocols, comma separated, e.g. ah, sctp

Create a new firewall rule and configure as follows.

- **Name**: Enter a unique name
  Example: *allow-ema-internal*
- **Description**: *Allow internal communication between EMA servers*
- **Network**: Select the VPC that you previously created
- **Targets**: *Specified target tags*
- **Target tags**: *ema-server*
- **Source filter**: *Source tags*
- **Source tags:** *ema-server*
- **Specified ports protocols and ports**:
  tcp: *8092-8094,8089*

Click the **Create** button.

3.4 Deploy Cloud NAT and Cloud Router

3.4.1 Navigate to Cloud NAT

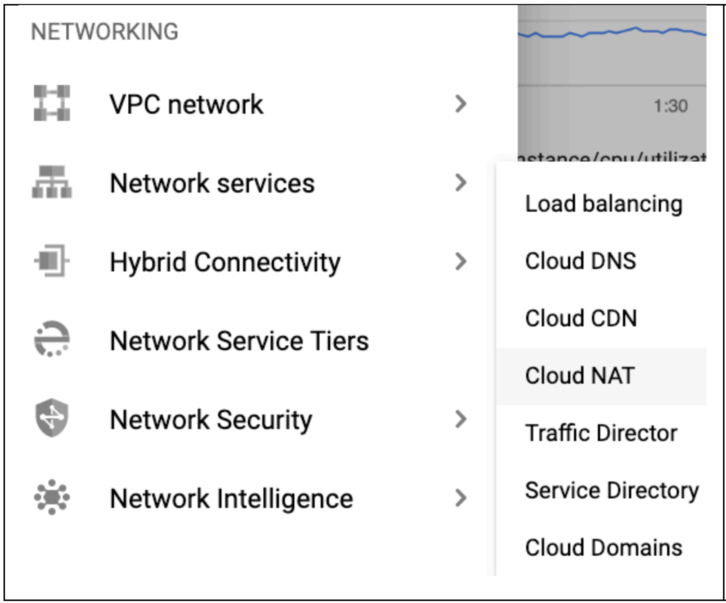

From the service menu, go to **Networking** > **Network services**, > **Cloud NAT**

Click “**Get started**”

3.4.2 Configure Cloud NAT Details and Create Cloud Router

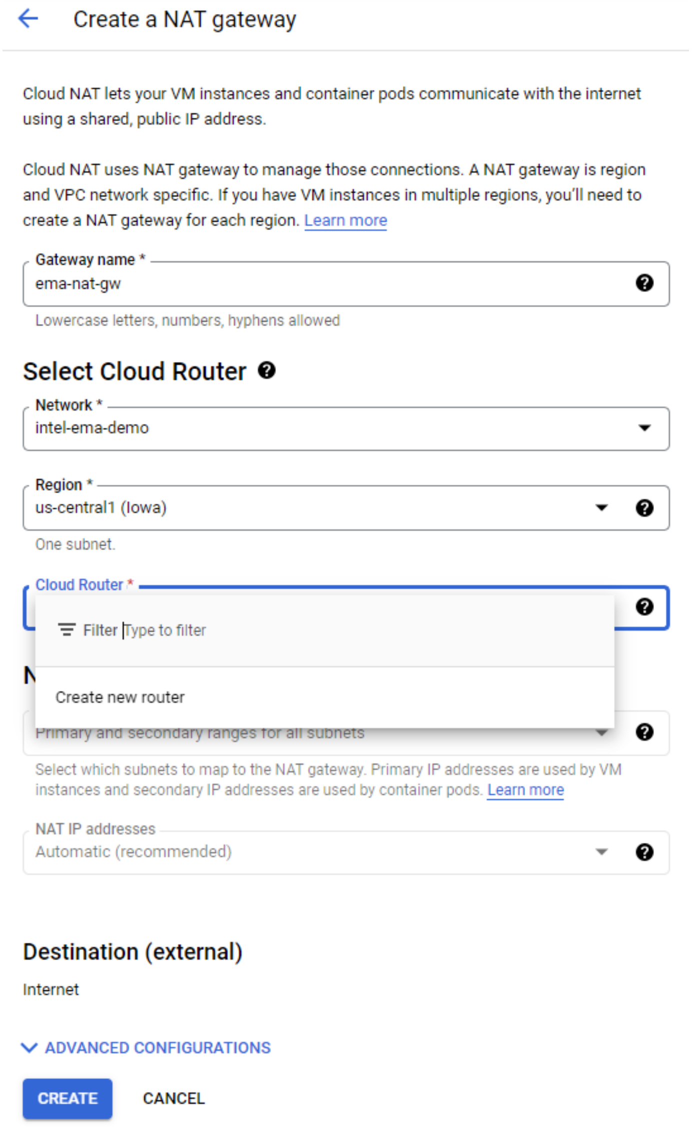

Create a NAT gateway

Cloud NAT lets your VM instances and container pods communicate with the internet using a shared, public IP address.

Cloud NAT uses NAT gateway to manage those connections. A NAT gateway is region and VPC network specific. If you have VM instances in multiple regions, you'll need to create a NAT gateway for each region. Learn more

Gateway name *: ema-nat-gw

Lowercase letters, numbers, hyphens allowed

Select Cloud Router

Network *: intel-ema-demo

Region *: us-central1 (Iowa)

One subnet.

Cloud Router *

Filter Type to filter

Create new router

Primary and secondary ranges for all subnets

Select which subnets to map to the NAT gateway. Primary IP addresses are used by VM instances and secondary IP addresses are used by container pods. Learn more

NAT IP addresses: Automatic (recommended)

Destination (external)

Internet

Configure the NAT gateway as follows:

- **Gateway name**: Enter a unique name
  Example: *ema-usc1-nat-gw*
- **VPC network:** Select the previously created VPC
- **Region:** Choose the region where you're deploying your virtual machines.
- **Cloud Router:** Select *Create new router* from the dropdown menu.
  - Enter a unique name for the Cloud Router
    Example: *ema-usc1-router*
  - Click the **Create** button to finalize the Cloud Router.

Click the **Create** button to finalize the Cloud NAT gateway.

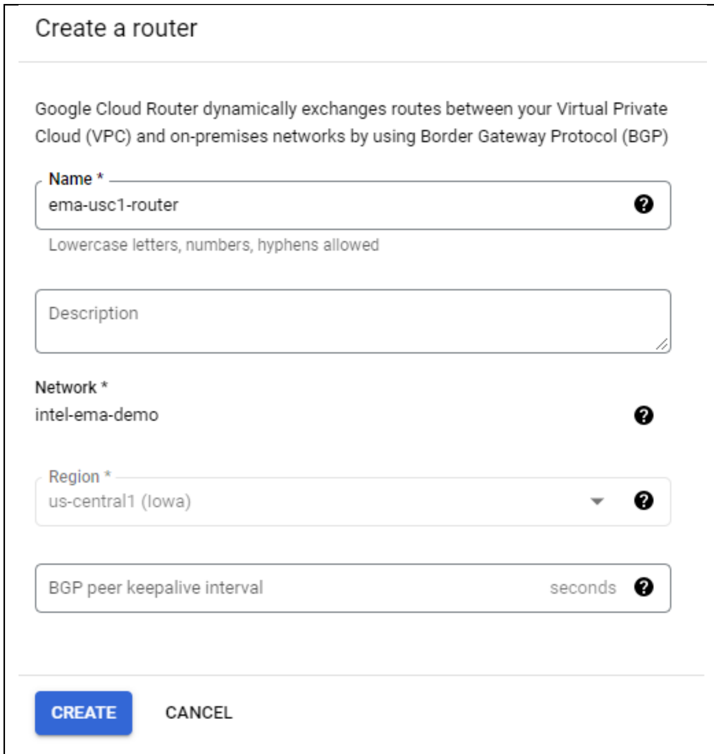 Create a router Google Cloud Router dynamically exchanges routes between your Virtual Private Cloud (VPC) and on-premises networks by using Border Gateway Protocol (BGP) Name * ema-usc1-router Lowercase letters, numbers, hyphens allowed Description Network * intel-ema-demo Region * us-central1 (Iowa) BGP peer keepalive interval seconds CREATE CANCEL	

4 Cloud SQL Deployment

Google Cloud SQL for SQL Server is a fully managed platform-as-a-service (PaaS) database engine, with features including:

- Custom machine types with up to 624 GB of RAM and 96 CPUs.
- Up to 30 TB of storage available, with the ability to automatically increase storage size as needed.
- Create and manage instances in the Google Cloud Console.
- Instances available in US, EU, Asia or Australia.
- Customer data encrypted on Google's internal networks and in database tables, temporary files, and backups.
- Support for secure external connections with the Cloud SQL Proxy or with the SSL/TLS protocol.
- Import databases using BAK and SQL files.
- Export databases using BAK files.
- Automated and on-demand backups.
- Integration with Stackdriver logging and monitoring.
- SQL Server Agent enabled to facilitate replication and other jobs.

Note: Cloud SQL Does NOT support AD Authentication.

For more information about Cloud SQL, including a full list of features that are not supported, visit the following link: https://cloud.google.com/sql/docs/sqlserver

4.1 Create the Cloud SQL Server

Follow this procedure to create a SQL Server database and grant access to your virtual machine(s).

4.1.1 Navigate to the SQL Service

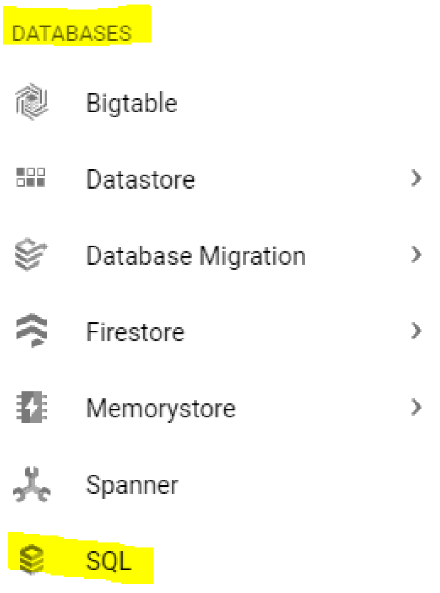

From the service menu, go to **DATABASES** > **SQL**

4.1.2 Create the SQL Server Instance

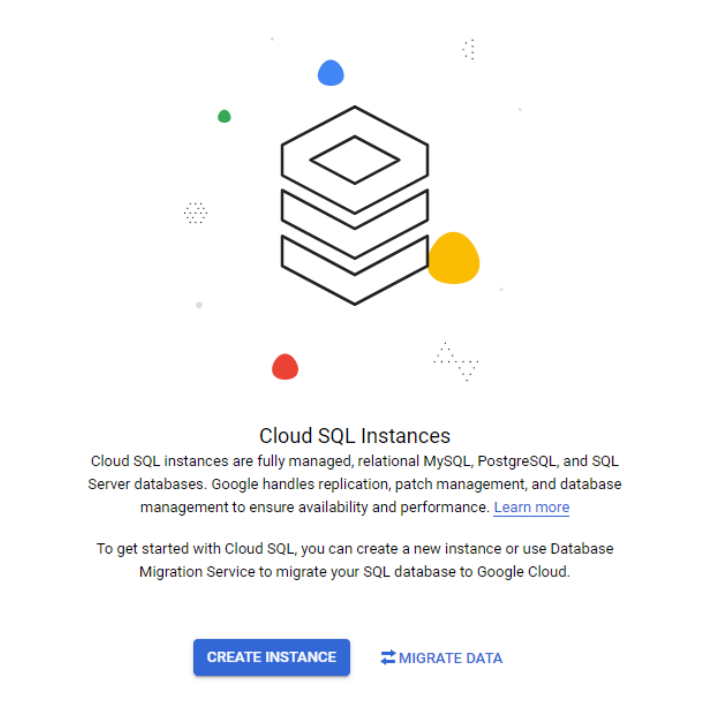 Cloud SQL Instances Cloud SQL instances are fully managed, relational MySQL, PostgreSQL, and SQL Server databases. Google handles replication, patch management, and database management to ensure availability and performance. Learn more To get started with Cloud SQL, you can create a new instance or use Database Migration Service to migrate your SQL database to Google Cloud. CREATE INSTANCE ⇄ MIGRATE DATA	Click **CREATE INSTANCE**

4.1.3 4.1.3 Select the Database Engine

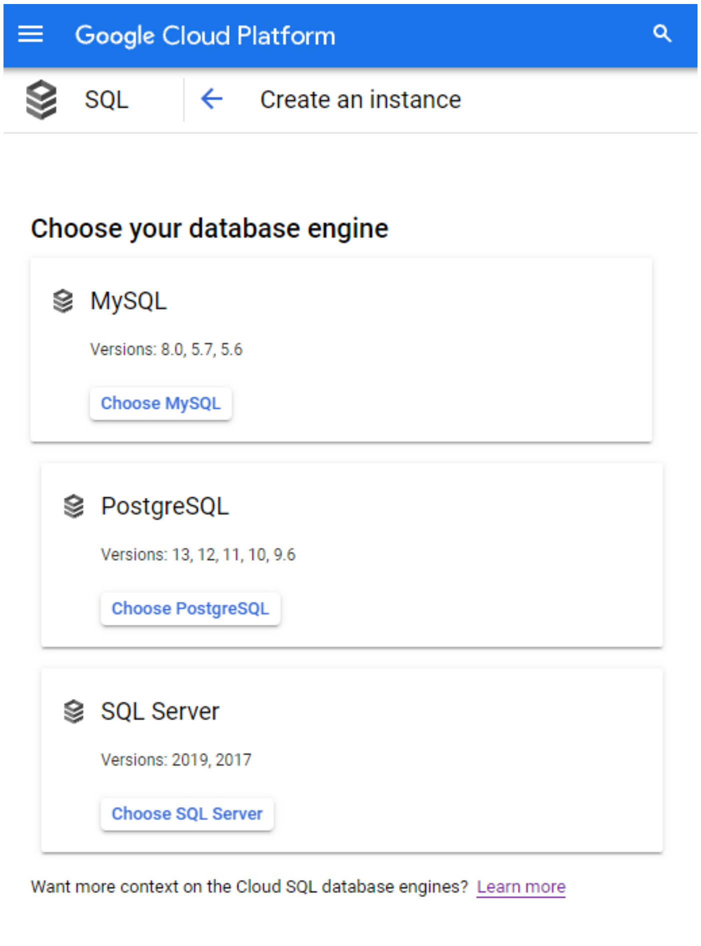 Google Cloud Platform SQL ← Create an instance Choose your database engine MySQL Versions: 8.0, 5.7, 5.6 Choose MySQL PostgreSQL Versions: 13, 12, 11, 10, 9.6 Choose PostgreSQL SQL Server Versions: 2019, 2017 Choose SQL Server Want more context on the Cloud SQL database engines? Learn more	Select *SQL Server* as the database engine.

4.1.4 Configure Basic Instance Information

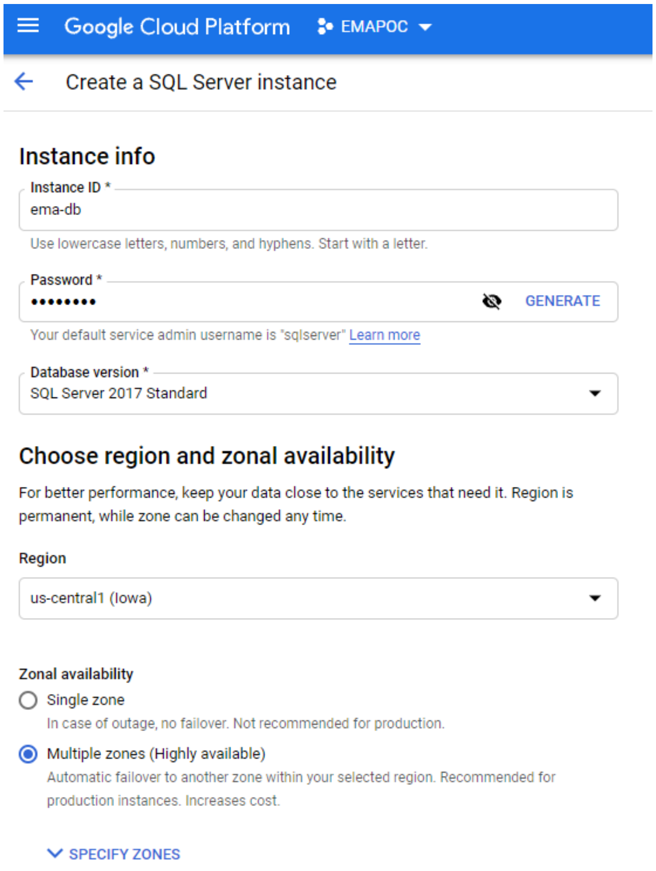

Create a SQL Server instance

Instance info

Instance ID *: ema-db

Use lowercase letters, numbers, and hyphens. Start with a letter.

Password *: ••••••••

Your default service admin username is "sqlserver" Learn more

Database version *: SQL Server 2017 Standard

Choose region and zonal availability

For better performance, keep your data close to the services that need it. Region is permanent, while zone can be changed any time.

Region: us-central1 (Iowa)

Zonal availability

○ Single zone
In case of outage, no failover. Not recommended for production.

● Multiple zones (Highly available)
Automatic failover to another zone within your selected region. Recommended for production instances. Increases cost.

Configure basic details as follows.

- **Instance ID**: Enter a unique name
  Example: *ema-db*
- **Password**: Create a password. Note that the default service admin username is '*sqlserver*'.
- **Database version**: SQL Server 2017 Standard
- **Region**: Use the same region as your subnet.
- **Zonal availability**: *Multiple zones*

4.1.5 Configure Machine Type and Storage

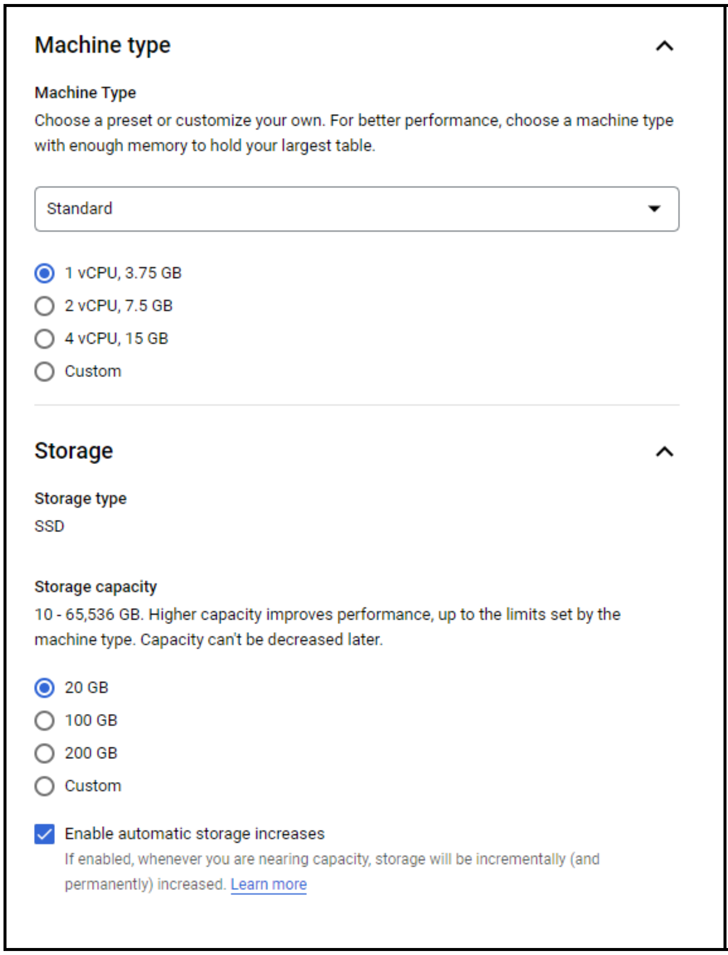

Machine type

Machine Type

Choose a preset or customize your own. For better performance, choose a machine type with enough memory to hold your largest table.

Standard ▾

- ● 1 vCPU, 3.75 GB
- ○ 2 vCPU, 7.5 GB
- ○ 4 vCPU, 15 GB
- ○ Custom

Storage

Storage type

SSD

Storage capacity

10 - 65,536 GB. Higher capacity improves performance, up to the limits set by the machine type. Capacity can't be decreased later.

- ● 20 GB
- ○ 100 GB
- ○ 200 GB
- ○ Custom

☑ Enable automatic storage increases

If enabled, whenever you are nearing capacity, storage will be incrementally (and permanently) increased. Learn more

Under **Configuration options**, configure **Machine type and storage** as follows.

- **Machine type:** *Standard*

Storage can be left at default values unless you have specific needs.

See the Intel® Endpoint Management Assistant Server Installation Guide for advice about system requirements.

4.1.6 Configure Connectivity

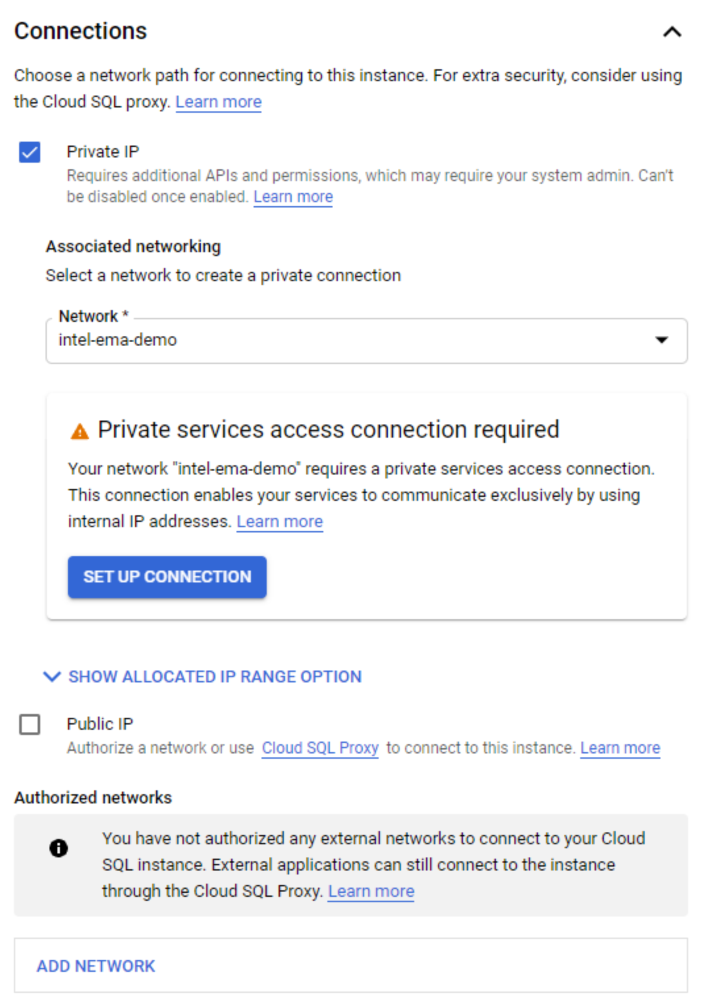

Connections

Choose a network path for connecting to this instance. For extra security, consider using the Cloud SQL proxy. Learn more

☑ Private IP
Requires additional APIs and permissions, which may require your system admin. Can't be disabled once enabled. Learn more

Associated networking
Select a network to create a private connection

Network *: intel-ema-demo

Private services access connection required

Your network "intel-ema-demo" requires a private services access connection. This connection enables your services to communicate exclusively by using internal IP addresses. Learn more

□ Public IP
Authorize a network or use Cloud SQL Proxy to connect to this instance. Learn more

Authorized networks

You have not authorized any external networks to connect to your Cloud SQL instance. External applications can still connect to the instance through the Cloud SQL Proxy. Learn more

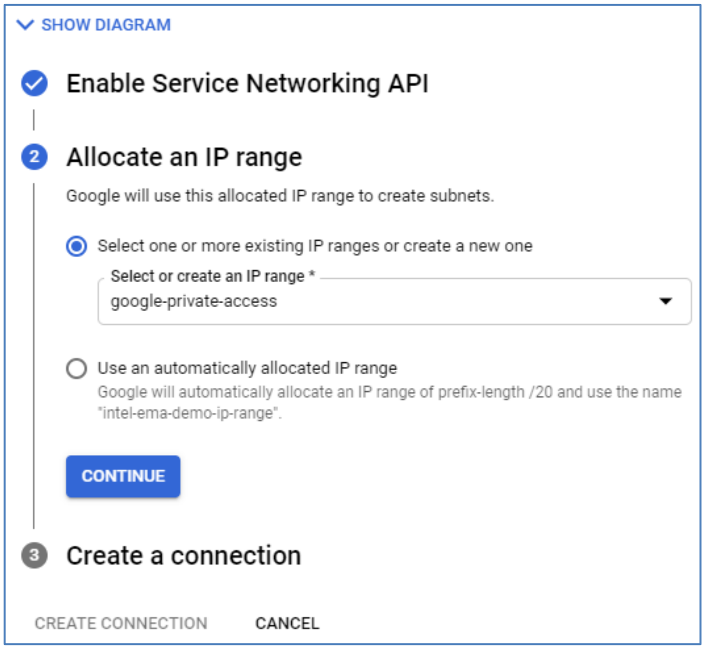

✓ Enable Service Networking API

2 Allocate an IP range

Google will use this allocated IP range to create subnets.

● Select one or more existing IP ranges or create a new one
Select or create an IP range *: google-private-access

○ Use an automatically allocated IP range
Google will automatically allocate an IP range of prefix-length /20 and use the name "intel-ema-demo-ip-range".

3 Create a connection

Under **Configuration options**, configure C**onnections** as follows.

- Select **Private IP**
- **Associated networking**: Select the VPC that you previously created
- Click **SET UP CONNECTION**
  - See second (lower) image at left
  - Select **Select one or more...** and select *google-private-access* for **IP range**
  - Click **CONTINUE**
  - Click **Create Connection**

4.1.7 Configure Backups

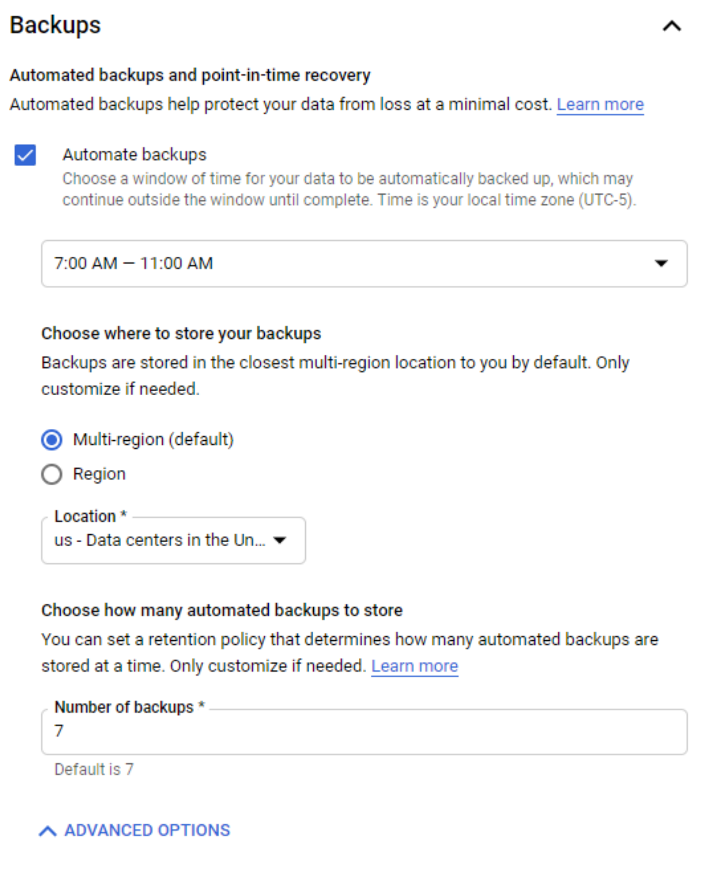

Backups

Automated backups and point-in-time recovery

Automated backups help protect your data from loss at a minimal cost. Learn more

☑ Automate backups

Choose a window of time for your data to be automatically backed up, which may continue outside the window until complete. Time is your local time zone (UTC-5).

7:00 AM – 11:00 AM

Choose where to store your backups

Backups are stored in the closest multi-region location to you by default. Only customize if needed.

● Multi-region (default)

○ Region

Location *: us - Data centers in the Un...

Choose how many automated backups to store

You can set a retention policy that determines how many automated backups are stored at a time. Only customize if needed. Learn more

Number of backups *: 7

Default is 7

Under **Configuration options**, configure **Backups** as follows.

The backup settings are at your discretion.

It is recommended to enable **Choose where to store your backups** for production deployments.

Click the **Create Instance** button to finalize the database creation.

4.1.8 Get the Database IP Address

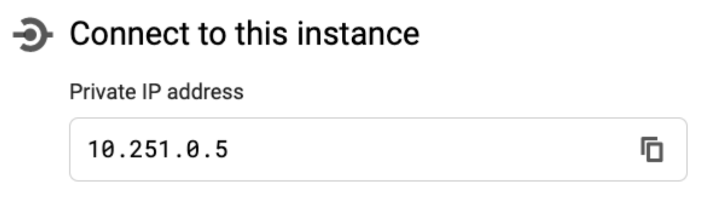

Connect to this instance

Private IP address

10.251.0.5

Once the database is created, the Overview page will show its private IP address in the “Connect to this instance” section.

5 Virtual Machine Deployment

5.1 Overview

Google Compute Engine (GCE) gives you the flexibility of compute virtualization without having to buy and maintain the physical hardware that runs it. However, you are still responsible for maintaining the guest operating system and the software that runs on it.

When you create an instance in a project, you specify the zone, operating system, and machine type of that instance. When you delete an instance, it is removed from the project. The machine type is what determines the CPU and Memory to allocated to the GCE virtual machines (VMs) at the time of creation, with Storage being a separate option, but you also change the machine type of a stopped instance or increase the amount of storage at a later time.

Each Compute Engine instance belongs to one VPC network. Instances in the same network communicate with each other through a local area network protocol. An instance uses the internet to communicate with any machine, virtual or physical, outside of its own network.

For more information about Google Compute Engine, visit the following links:
https://cloud.google.com/compute
https://cloud.google.com/compute/docs/concepts

5.2 Create a GCE VM Instance

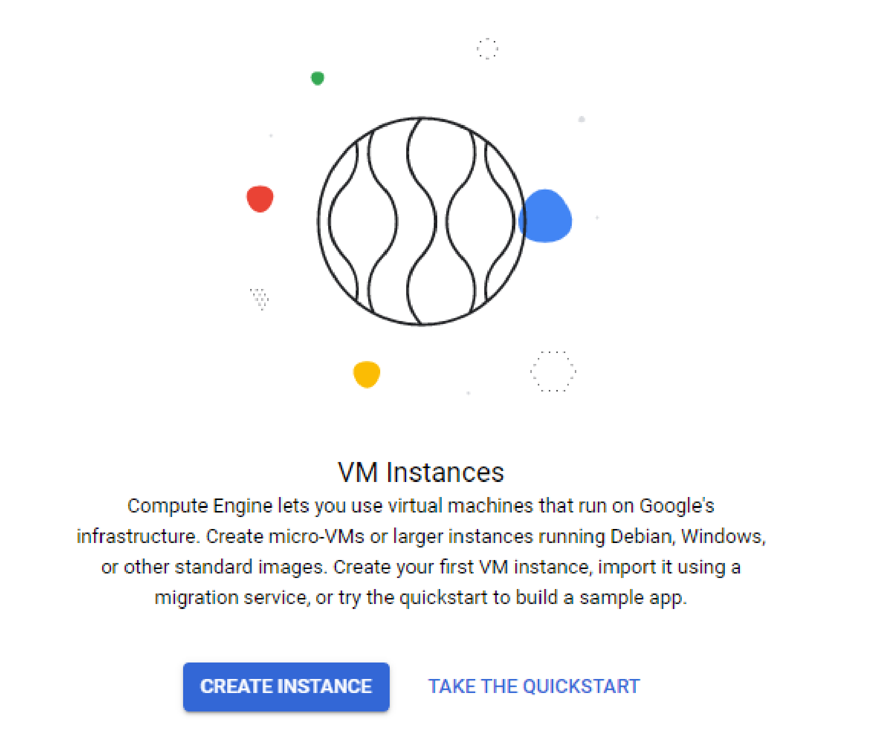 VM Instances Compute Engine lets you use virtual machines that run on Google's infrastructure. Create micro-VMs or larger instances running Debian, Windows, or other standard images. Create your first VM instance, import it using a migration service, or try the quickstart to build a sample app. 	From the service menu, go to **Compute** > **Compute Engine** > **VM instances.** Click the **Create Instance** button.

5.2.1 Configure the VM Basic Details

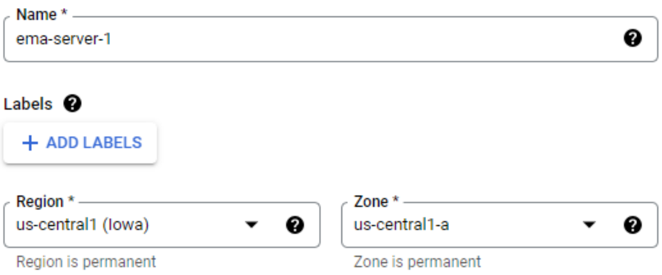

Name *
ema-server-1

Labels

Region *
us-central1 (Iowa)
Region is permanent

Zone *
us-central1-a
Zone is permanent

Configure VM basic details as follows.

- **Name:** Enter a unique name
  Example: *ema-server-1*
- **Region**: Select the same region that we've used previously.
- **Zone**: Choose a zone that will be different from the other EMA VM

5.2.2 Configure the VM Machine Type

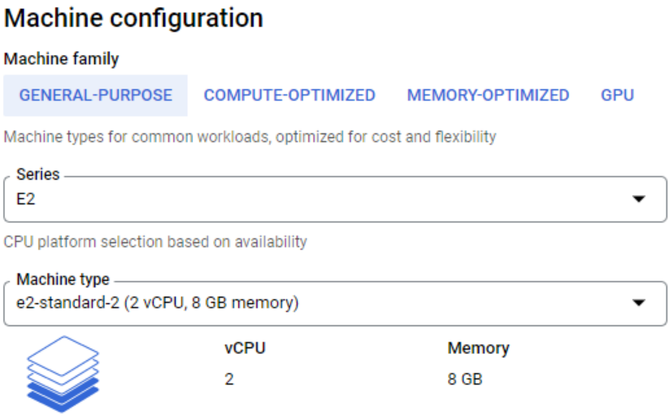

Machine configuration

Machine family

Machine types for common workloads, optimized for cost and flexibility

Series
E2

CPU platform selection based on availability

Machine type
e2-standard-2 (2 vCPU, 8 GB memory)

vCPU	Memory
2	8 GB

Choose the appropriate machine type. See the Intel® Endpoint Management Assistant Server Installation Guide for system requirements.

You can change this at a later time when the VM is powered down.

5.2.3 Configure the VM Boot Image

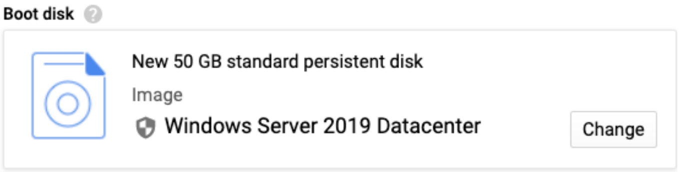

Boot disk

New 50 GB standard persistent disk
Image
Windows Server 2019 Datacenter

Set **Boot disk** to the latest version of Windows Server Datacenter supported by Intel EMA.

See the Intel® Endpoint Management Assistant Server Installation Guide for supported operating systems.

5.2.4 Configure VM Access and Firewall

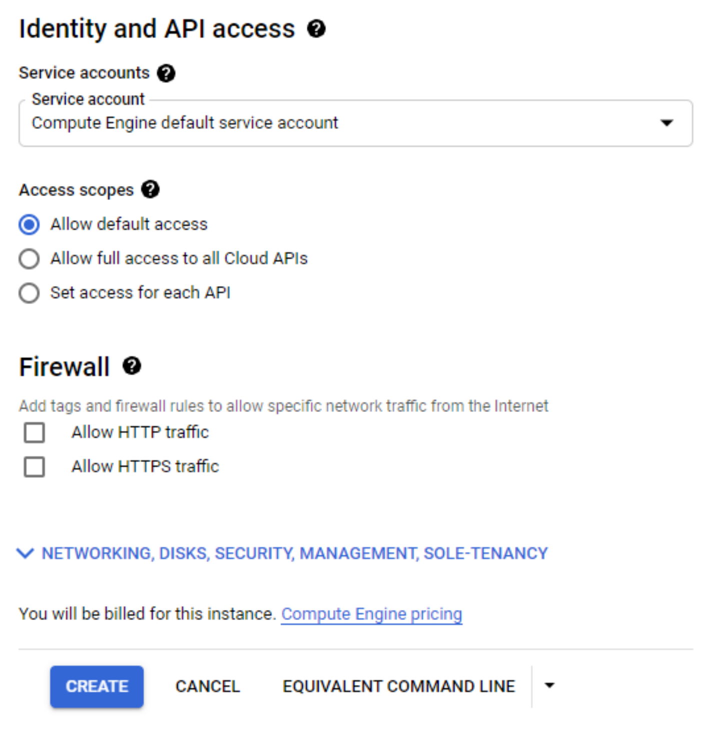

Identity and API access

Service accounts

Service account: Compute Engine default service account

Access scopes

● Allow default access

○ Allow full access to all Cloud APIs

○ Set access for each API

Firewall

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

□ Allow HTTP traffic

□ Allow HTTPS traffic

NETWORKING, DISKS, SECURITY, MANAGEMENT, SOLE-TENANCY
You will be billed for this instance. Compute Engine pricing
CREATE
CANCEL
EQUIVALENT COMMAND LINE

In the **Identity and API access** section, you can leave the defaults which give the VM permission to write logs to Google Cloud Logging among a few other things.

In the Firewall section, both checkboxes should be clear because we are going to permit network access using Network tags, which will be configured in the next step.

Click the **NETWORKING, DISKS, SECURITY, MANAGEMENT, SOLE-TENANCY** link to expand that section before continuing to the next step.

5.2.5 Configure VM Networking

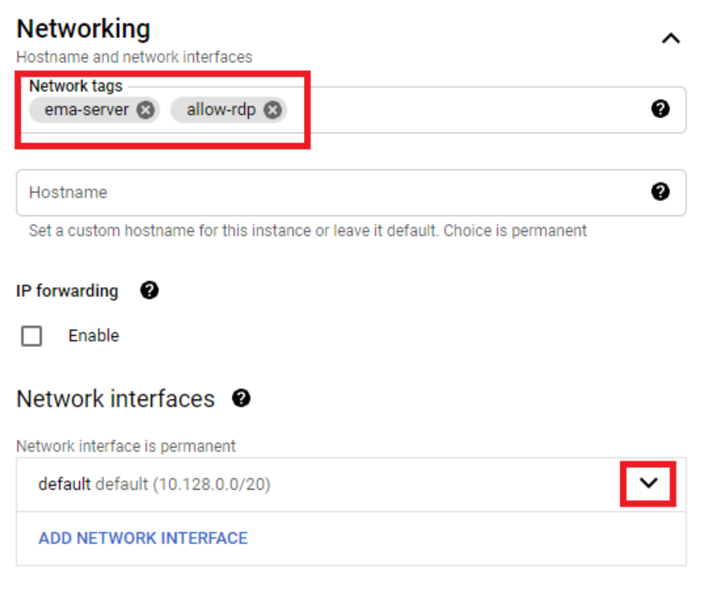

Networking

Hostname and network interfaces

Network tags: ema-server, allow-rdp

Hostname

Set a custom hostname for this instance or leave it default. Choice is permanent

IP forwarding

□ Enable

Network interfaces

Network interface is permanent

default default (10.128.0.0/20)

Select the **Networking** tab.

Set the following **Network tags**:

- ema-server

Click the pencil icon on the network interface before continuing to the next step.

5.2.6 Configure the VM Network Interface (Single Server Deployment)

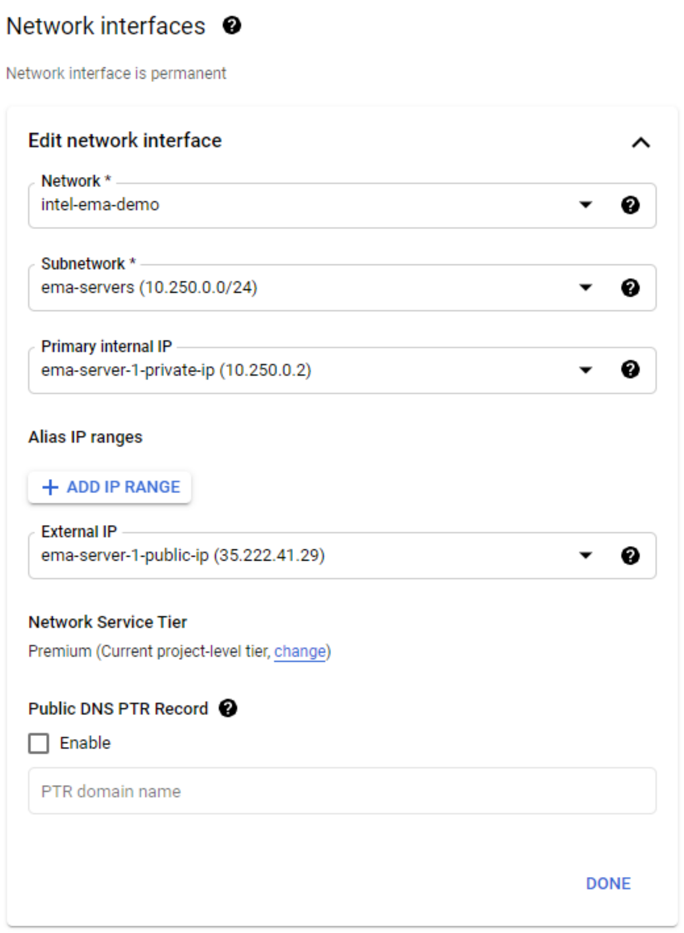

Network interfaces

Network interface is permanent

Edit network interface

Network *: intel-ema-demo

Subnetwork *: ema-servers (10.250.0.0/24)

Primary internal IP: ema-server-1-private-ip (10.250.0.2)

Alias IP ranges

External IP: ema-server-1-public-ip (35.222.41.29)

Network Service Tier

Premium (Current project-level tier, change)

Public DNS PTR Record

□ Enable

Set **Primary internal IP** to *Reserve static internal IP address*

Enter a unique name for the IP reservation
Example: *ema-server-1-private-ip*

Click the **Reserve** button.

Set **External IP** to Create IP address.

Enter a unique name for the IP reservation
Example: *ema-server-1-public-ip*

Click the **Reserve** button.

Click the **Done** button.

5.2.7 Configure the VM Network Interface (Distributed Server Deployment)

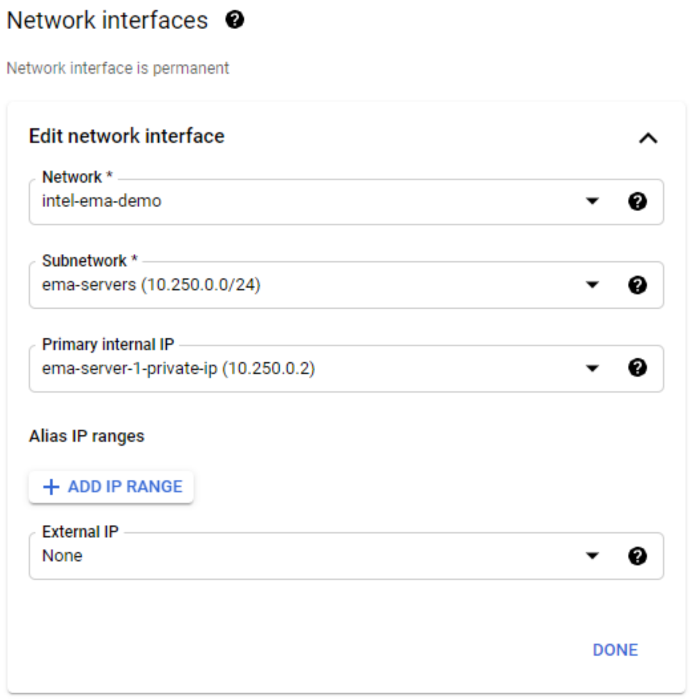

Network interfaces

Network interface is permanent

Edit network interface

Network *: intel-ema-demo

Subnetwork *: ema-servers (10.250.0.0/24)

Primary internal IP: ema-server-1-private-ip (10.250.0.2)

Alias IP ranges

External IP: None

Set **Primary internal IP** to *Reserve static internal IP address*

Enter a unique name for the IP reservation
Example: *ema-server-1-private-ip*

Click the **Reserve** button.

Set **External IP** to None.

Click the **Done** button.

5.2.8 Finalize VM Creation

Click the **Create** button at the bottom of the screen to finalize the VM creation.

5.2.9 Set Windows Password

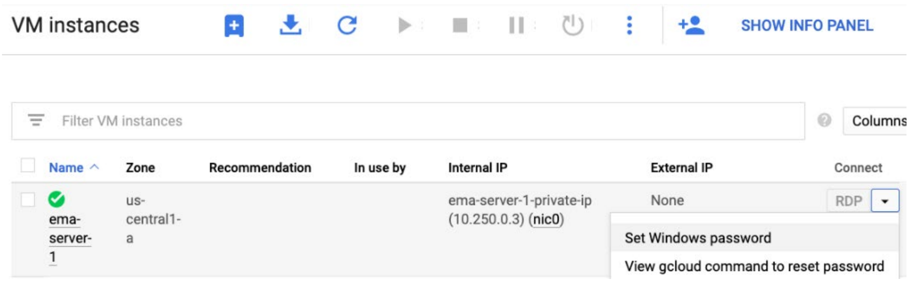

Name	Zone	Recommendation	In use by	Internal IP	External IP	Connect
ema-server-1	us-central1-a			ema-server-1-private-ip (10.250.0.3) (nic0)	None	RDP

After the VM is created then from the VM instance list you can click the Connect arrow button to set a Windows password for it.

5.3 Create a Second GCE VM Instance (Distributed Server Deployment Only)

For a distributed server deployment, repeat the previous steps to create another VM. It is recommended that you deploy to a different zone to mitigate the impact of a zone outage.

5.4 Logging into virtual machines with RDP

For virtual machines that do not have a public IP address, this section describes a method of tunneling a RDP connection to your VMs using Google's Identity-Aware Proxy (IAP).

This section requires that you have the Cloud SDK installed so that you have access to the gcloud command-line utility. For installation instructions, see https://cloud.google.com/sdk/docs/install

Once you have the gcloud utility installed and configured, then you will be able to start an IAP tunnel to your virtual machine in order to forward a local port of your choosing to the RDP port of the VM. Example command:

```
gcloud compute start-iap-tunnel ema-server-1 3389 --local-host-port=localhost:33389 --zone=us-central1-a
```

You will need to adjust the command to have the correct server name and zone in order to work.

For more information about using IAP for TCP forwarding, visit the following link: https://cloud.google.com/iap/docs/using-tcp-forwarding

6 Load Balancer Deployment (Distributed Server Deployment Only)

A load balancer distributes user traffic across multiple instances of your applications. By spreading the load, load balancing reduces the risk that your applications become overburdened, slow, or nonfunctional.

We will use a HTTPS Load Balancer for web traffic, and a TCP Proxy Load Balancer for swarm traffic. You will need to have a SSL/TLS certificate during HTTPS LB creation.

The backend of the load balancer an instance group. Our virtual machines need manual configuration and so not support auto-scaling, so we will be using Unmanaged Instance Groups. These are zonal resources, so we will need to create separate instance groups for each zone that you have deployed Intel EMA VMs into.

One other important note is that the TCP load balancer we are using only accepts traffic on certain well-known ports on the front-end, so this will require you to update some settings after you have installed Intel EMA on the server. There are instructions on how to do this in the Intel EMA Server Installation Guide.

For more information about Google load balancing, visit the following link:
https://cloud.google.com/load-balancing/docs

6.1 Create Unmanaged Instance Group(s)

6.1.1 Navigate to Instance Groups

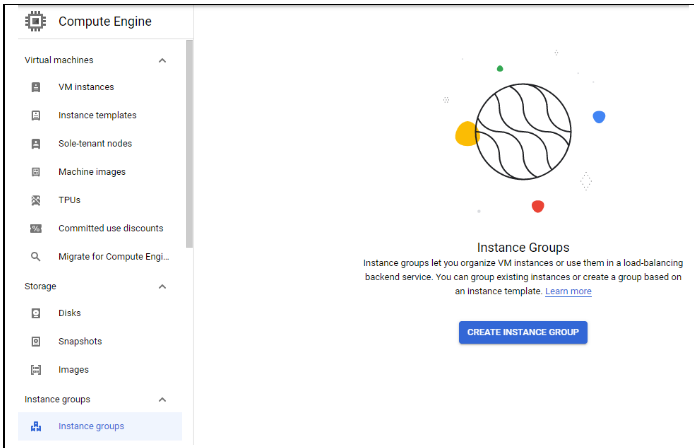 Compute Engine Virtual machines VM instances Instance templates Sole-tenant nodes Machine images TPUs Committed use discounts Migrate for Compute Engi... Storage Disks Snapshots Images Instance groups Instance groups Instance Groups CREATE INSTANCE GROUP Instance groups let you organize VM instances or use them in a load-balancing backend service. You can group existing instances or create a group based on an instance template. Learn more	From the service menu, go to **Compute** > **Compute Engine** > **VM instances.**

6.1.2 Create an Unmanaged Instance Group

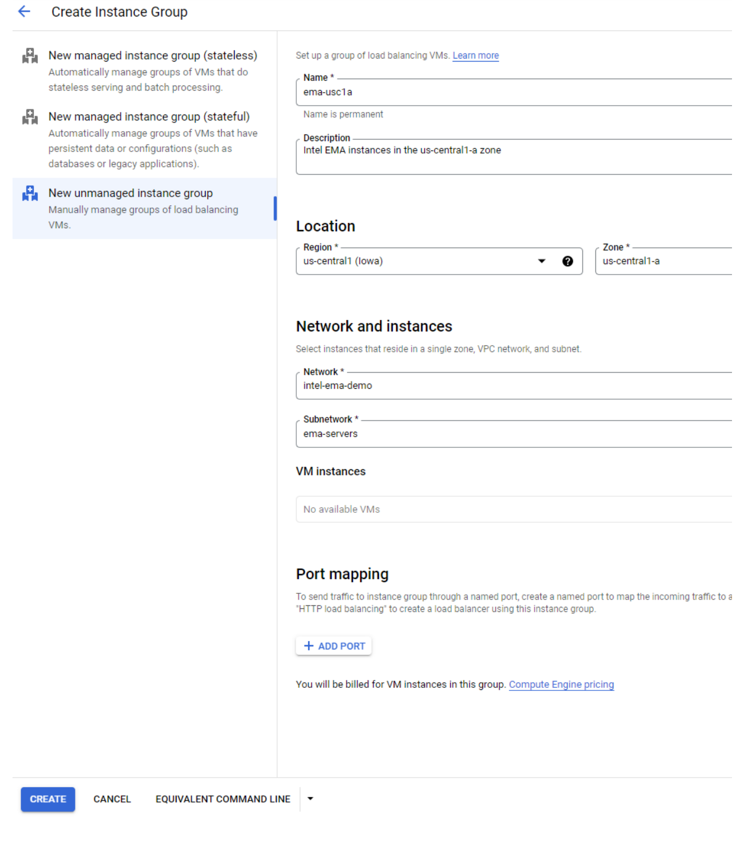

Create Instance Group

New managed instance group (stateless)
Automatically manage groups of VMs that do stateless serving and batch processing.

New managed instance group (stateful)
Automatically manage groups of VMs that have persistent data or configurations (such as databases or legacy applications).

New unmanaged instance group
Manually manage groups of load balancing VMs.

Set up a group of load balancing VMs. Learn more

Name *: ema-usc1a
Name is permanent

Description: Intel EMA instances in the us-central1-a zone

Location

Region *: us-central1 (Iowa)

Zone *: us-central1-a

Network and instances

Select instances that reside in a single zone, VPC network, and subnet.

Network *: intel-ema-demo

Subnetwork *: ema-servers

VM instances

No available VMs

Port mapping

To send traffic to instance group through a named port, create a named port to map the incoming traffic to a "HTTP load balancing" to create a load balancer using this instance group.

You will be billed for VM instances in this group. Compute Engine pricing

Click the **Create Instance Group** button.

Configure the instance group as follows:

- **Name**: Enter a unique name for the instance group Example: *ema-usc1a*
- **Description (optional)**: *Intel EMA instances in the us-central1-a zone*
- **Location:** Choose your preferred region and zone Example: *us-central1-a*
- **Port name mapping:** Add the following items
    - *web : 443*
    - *redirection : 8084*
    - *swarm : 8080*
- **Network**: Select your VPC network
- **Subnetwork**: Select your subnet
- **VM instances**: Select all of the VMs in this zone. There should be at least one.

Click the **Create** button

6.1.3 Create additional Instance Groups

Follow the previous steps to create an unmanaged instance group for each other zone in which you have deployed an Intel EMA VM.

6.2 Create Health Checks

We need to create health checks so that the load balancers will be able to determine which instances a healthy and can receive traffic.

6.2.1 Create a Health Check for the Web Backend

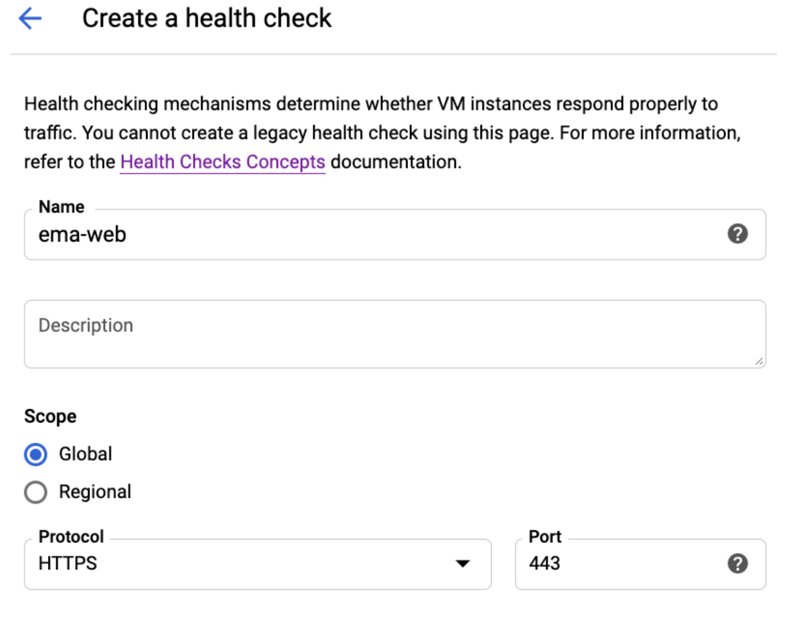

Create a health check

Health checking mechanisms determine whether VM instances respond properly to traffic. You cannot create a legacy health check using this page. For more information, refer to the Health Checks Concepts documentation.

Name: ema-web

Description

Scope

● Global

○ Regional

Protocol: HTTPS

Port: 443

From the Compute Engine sidebar, navigate to **Instance Group > Health checks**.

Click **Create Health Check**.

Configure the health check as follows:

- Name: Enter a unique name for the health check
  Example: *ema-web*
- Scope: Global
- Protocol: HTTPS
- Port: 443

You can accept the rest of the default values.

Click Create to finalize the health check.

6.2.2 Create a Health Check for the Swarm Backend

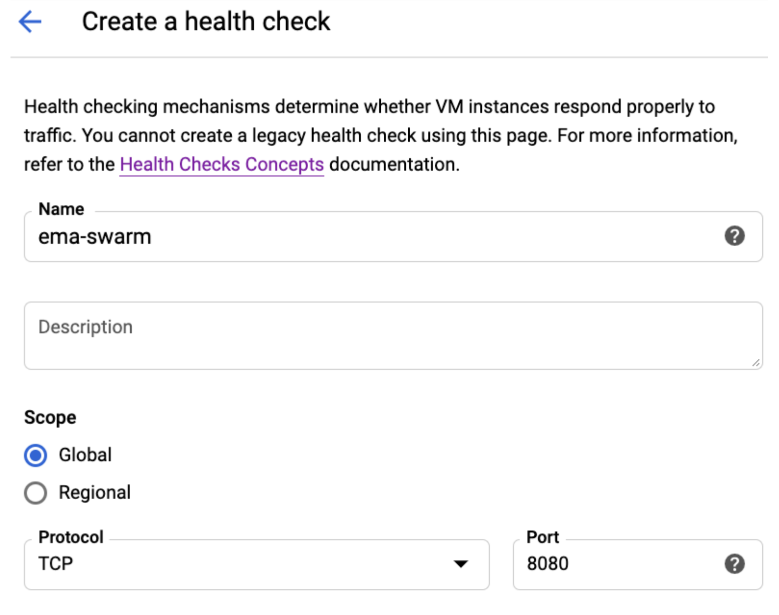

Create a health check

Health checking mechanisms determine whether VM instances respond properly to traffic. You cannot create a legacy health check using this page. For more information, refer to the Health Checks Concepts documentation.

Name: ema-swarm

Description

Scope

● Global

○ Regional

Protocol: TCP

Port: 8080

From the Compute Engine sidebar, select **Health checks**.

Click **Create Health Check**.

Configure the health check as follows:

- Name: Enter a unique name for the health check
  Example: *ema-swarm*
- Scope: Global
- Protocol: TCP
- Port: 8080

You can accept the rest of the default values.

Click Create to finalize the health check.

6.3 Navigate to Load Balancing

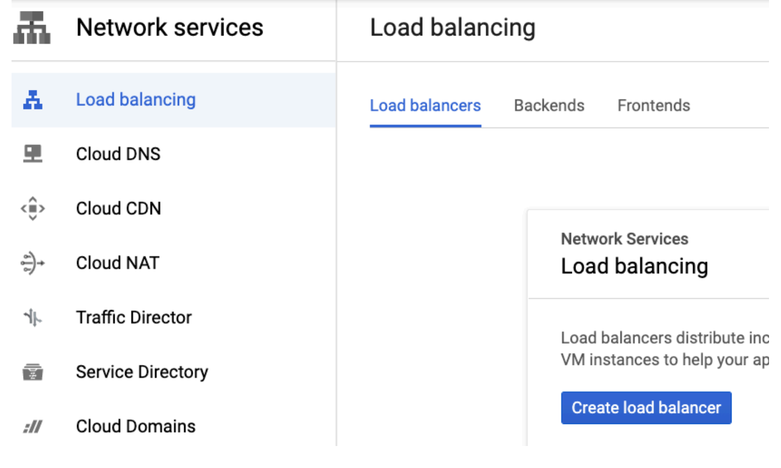

Load balancers distribute inc
VM instances to help your ap

From the service menu, go to **Networking** > **Network services > Load Balancing**

6.4 Create the HTTPS Load Balancer

6.4.1 Choose HTTP(S) Load Balancing

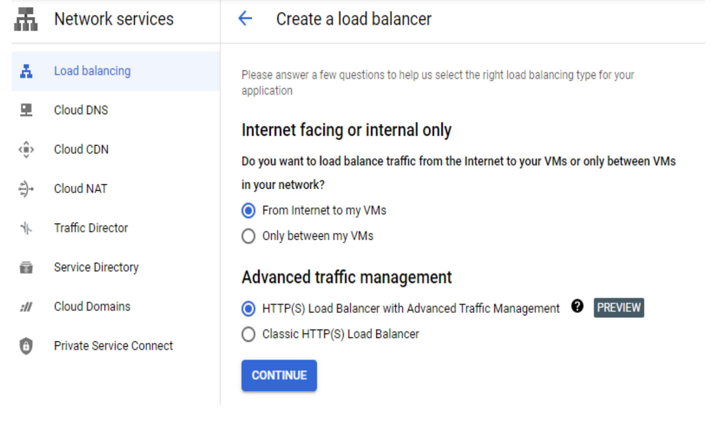

Please answer a few questions to help us select the right load balancing type for your application

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

Advanced traffic management

● HTTP(S) Load Balancer with Advanced Traffic Management

○ Classic HTTP(S) Load Balancer

Click **Create load balancer**.

Under **HTTP(S) Load Balancing**, click the **Start configuration** button.

Select **From Internet to my VMs**.

Select **HTTP(S) Load Balancer with Advanced Traffic Management**.

Click the **Continue** button.

6.4.2 Set a Name for the Load Balancer

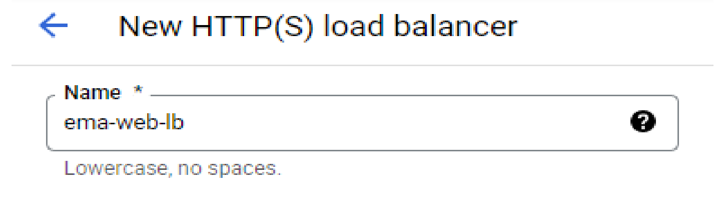

Name *
ema-web-lb

Lowercase, no spaces.

Enter a unique name for the load balancer
Example: **ema-web-lb**

6.4.3 Backend Service Configuration

6.4.3.1 Create a backend service

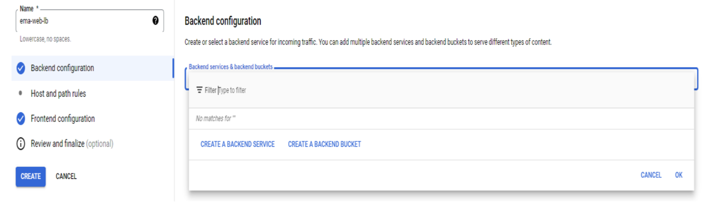

Name *
ema-web-lb

Lowercase, no spaces.

Backend configuration

Create or select a backend service for incoming traffic. You can add multiple backend services and backend buckets to serve different types of content.

Backend services & backend buckets

No matches for ""

Click on **Backend configuration**

From the drop-down menu, navigate to **Backend services** > **Create a backend service**.

6.4.3.2 Configure backend service, basic details

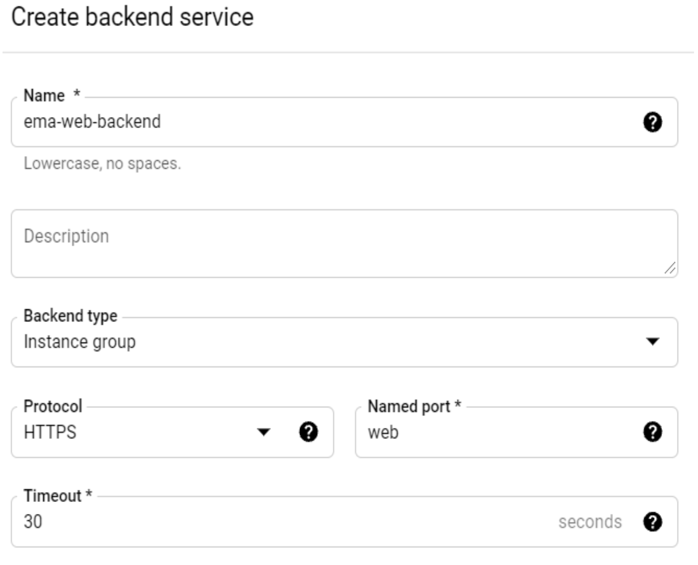

Create backend service

Name *
ema-web-backend
Lowercase, no spaces.

Description

Backend type
Instance group

Protocol
HTTPS

Named port *
web

Timeout *
30 seconds

Configure the backend service as follows:

- **Name**: Enter a unique name for the backend service Example: *ema-web-backend*
- **Backend type**: *Instance group*
- **Protocol**: *HTTPS*
- **Named port**: *web*

6.4.3.3 Add Backends

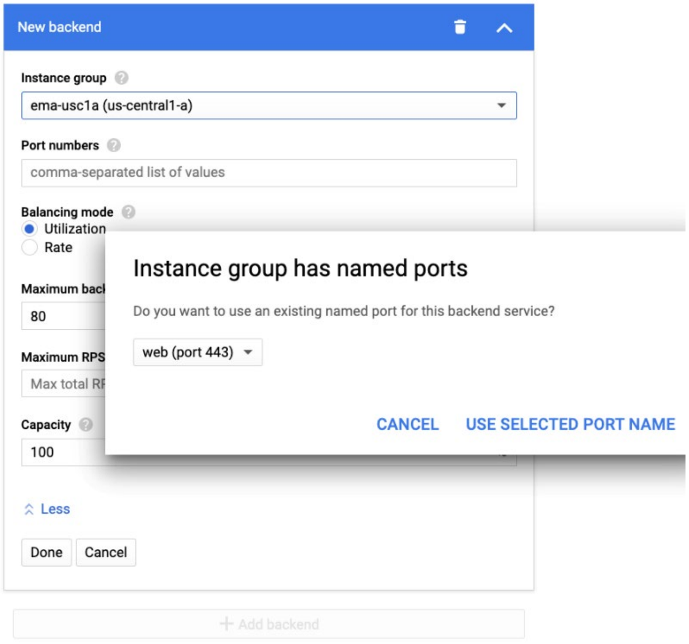

New backend

Instance group
ema-usc1a (us-central1-a)

Port numbers
comma-separated list of values

Balancing mode
● Utilization
○ Rate

Maximum bac
80

Maximum RPS
Max total R

Capacity
100

Less

Done Cancel

+ Add backend

Instance group has named ports

Do you want to use an existing named port for this backend service?

web (port 443)

CANCEL USE SELECTED PORT NAME

In the **New backend** section, select your first instance group that you previously created.

You will get a pop-up window asking if you want to use an existing named port. Select **web (port 443)** and click **Use Selected Port Name**.

Click Done.

For each additional unmanaged instance group that you previously created, click the Add Backend button and then repeat these instructions.

6.4.3.4 Set Health Check

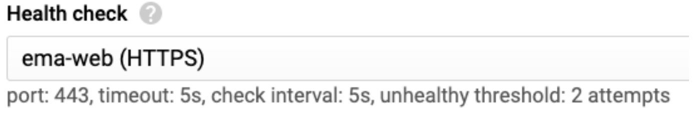

Health check

ema-web (HTTPS)

port: 443, timeout: 5s, check interval: 5s, unhealthy threshold: 2 attempts

From the **Health check** dropdown menu, select the **ema-web (HTTPS)** health check that you previously created.

6.4.3.5 Enable Session Affinity

Advanced configurations (Session affinity, connection draining timeout)

Click on **Advanced configurations** to show additional options.

Set **Sessional affinity** to *Generated cookie*.

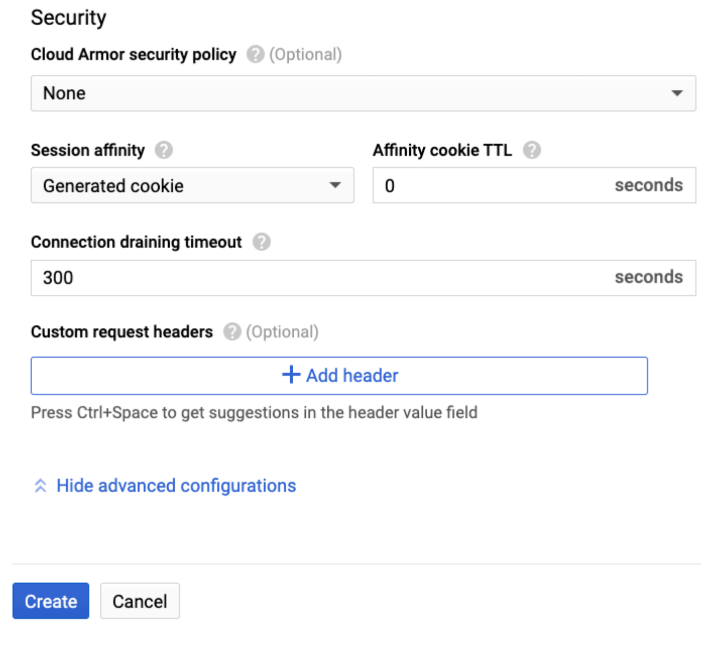 Security Cloud Armor security policy (Optional) None Session affinity Generated cookie Affinity cookie TTL 0 seconds Connection draining timeout 300 seconds Custom request headers (Optional) + Add header Press Ctrl+Space to get suggestions in the header value field Hide advanced configurations Create Cancel	Click the **Create** button.

6.4.4 Frontend Configuration

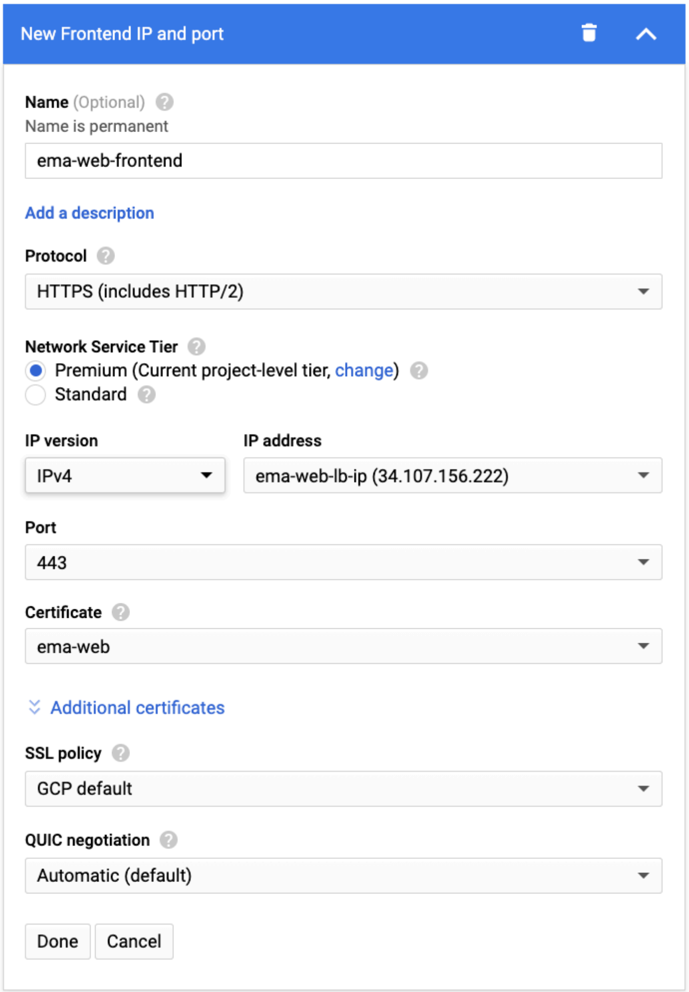 New Frontend IP and port Name (Optional) Name is permanent ema-web-frontend Add a description Protocol HTTPS (includes HTTP/2) Network Service Tier ● Premium (Current project-level tier, change) ○ Standard IP version IPv4 IP address ema-web-lb-ip (34.107.156.222) Port 443 Certificate ema-web Additional certificates SSL policy GCP default QUIC negotiation Automatic (default) Done Cancel	Click **Frontend configuration**. Configure the Frontend as follows: - **Name**: Enter a unique name for the frontend Example: *ema-web-frontend* - **Protocol**: *HTTPS* - **IP address**: Select Create IP address from the menu - Enter a unique name for the IP address Example: *ema-web-lb-ip* - Click **Reserve** - **Port**: *443* - **Certificate**: Select create a new certificate and input your SSL certificate information Click the **Done** button.

6.4.5 Review and Finalize

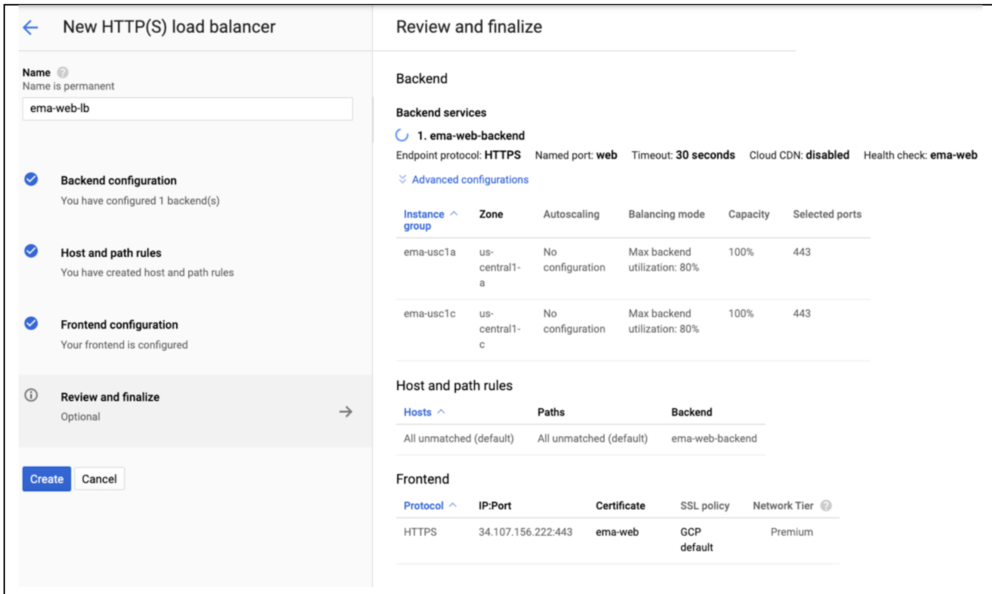

Name
Name is permanent
ema-web-lb

✓ Backend configuration
You have configured 1 backend(s)

✓ Host and path rules
You have created host and path rules

✓ Frontend configuration
Your frontend is configured

Review and finalize
Optional

Review and finalize

Backend

Backend services

1. ema-web-backend

Endpoint protocol: **HTTPS** Named port: **web** Timeout: **30 seconds** Cloud CDN: **disabled** Health check: **ema-web**

Advanced configurations

Instance group	Zone	Autoscaling	Balancing mode	Capacity	Selected ports
ema-usc1a	us-central1-a	No configuration	Max backend utilization: 80%	100%	443
ema-usc1c	us-central1-c	No configuration	Max backend utilization: 80%	100%	443

Host and path rules

Hosts	Paths	Backend
All unmatched (default)	All unmatched (default)	ema-web-backend

Frontend

Protocol	IP:Port	Certificate	SSL policy	Network Tier
HTTPS	34.107.156.222:443	ema-web	GCP default	Premium

Click **Review and finalize**.

Review the information on the screen and then click the **Create** button.

6.5 Create the TCP Load Balancer

6.5.1 Choose TCP Load Balancing

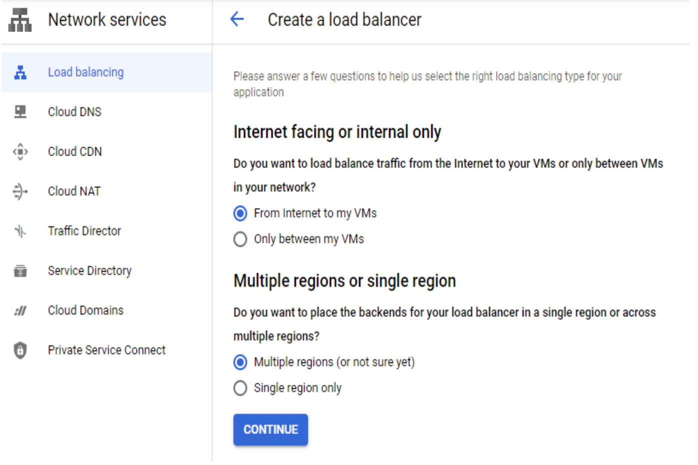

Please answer a few questions to help us select the right load balancing type for your application

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

Click **Create load balancer**.

Under **TCP Load Balancing**, click the **Start configuration** button.

Select **From Internet to my VMs**.

Select **Multiple regions.**

Click the **Continue** button.

6.5.2 Set a Name for the Load Balancer

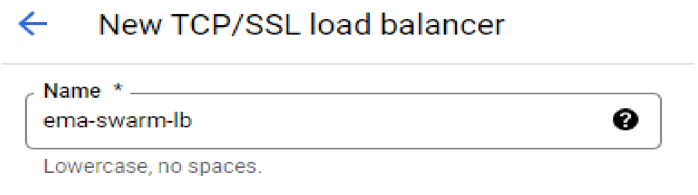

Name *

ema-swarm-lb

Lowercase, no spaces.

Enter a unique name for this load balancer.
Example: *ema-swarm-lb*

6.5.3 Backend Service Configuration

6.5.3.1 Configure backend service, basic details

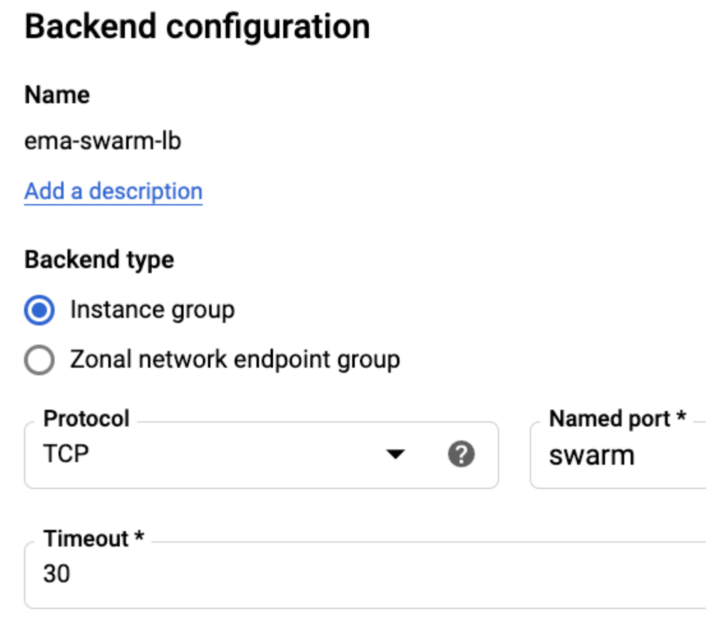

Backend configuration

Name
ema-swarm-lb

Add a description

Backend type
● Instance group
○ Zonal network endpoint group

Protocol: TCP

Named port *: swarm

Timeout *: 30

Configure the backend service as follows:

- **Backend type**: *Instance group*
- **Protocol**: *TCP*
- **Named port**: *swarm*

6.5.3.2 Add Backends

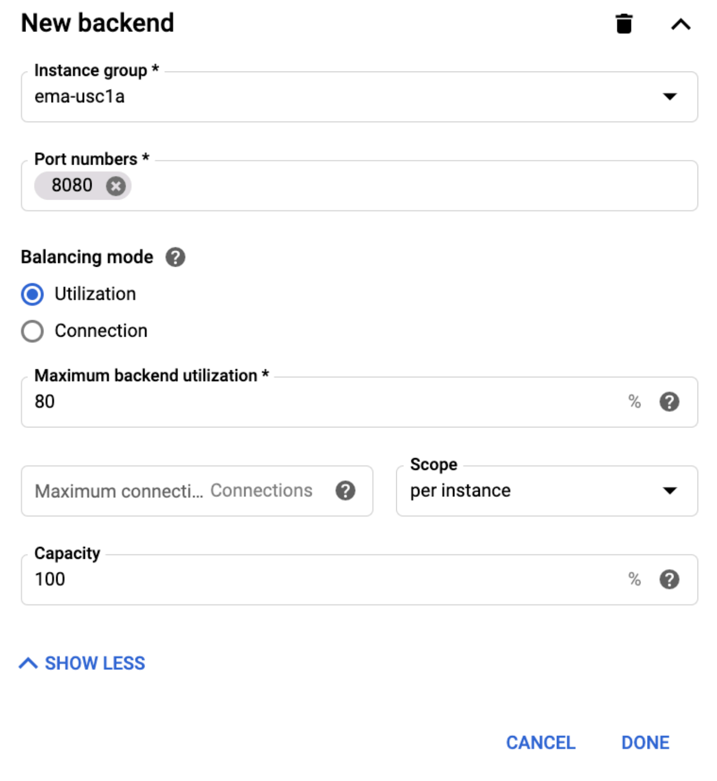

New backend

Instance group *: ema-usc1a

Port numbers *: 8080

Balancing mode
● Utilization
○ Connection

Maximum backend utilization *: 80 %

Maximum connecti... Connections

Scope: per instance

Capacity: 100 %

In the **New backend** section, select your first instance group that you previously created.

You will get a pop-up window asking if you want to use an existing named port. Select **swarm (port 8080)** and click **Use Selected Port Name**.

The rest of the settings can be left at default values.

Click **Done**.

For each additional unmanaged instance group that you previously created, click the Add Backend button and then repeat these instructions.

6.5.3.3 Set Health Check

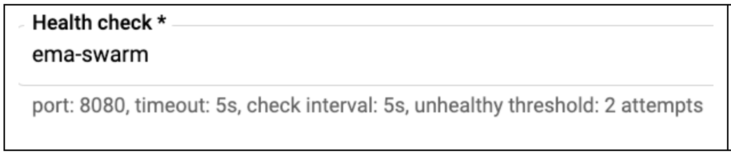

Health check *: ema-swarm

port: 8080, timeout: 5s, check interval: 5s, unhealthy threshold: 2 attempts

From the **Health check** dropdown menu, select the **ema-swarm** health check that you previously created.

6.5.4 Frontend Configuration

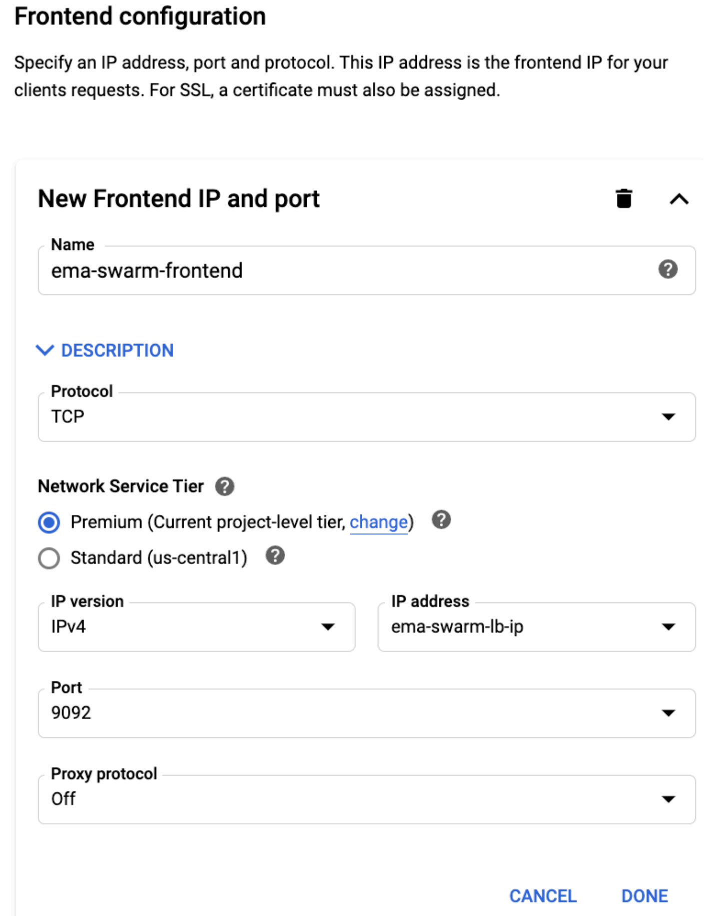

Frontend configuration

Specify an IP address, port and protocol. This IP address is the frontend IP for your clients requests. For SSL, a certificate must also be assigned.

New Frontend IP and port

Name: ema-swarm-frontend

DESCRIPTION

Protocol: TCP

Network Service Tier

- ● Premium (Current project-level tier, change)
- ○ Standard (us-central1)

IP version: IPv4

IP address: ema-swarm-lb-ip

Port: 9092

Proxy protocol: Off

Click **Frontend configuration**.

Configure the Frontend as follows:

- **Name**: Enter a unique name for the frontend
  Example: *ema-swarm-frontend*
- **Protocol**: *TCP*
- **IP address**: Select Create IP address from the menu
  - Enter a unique name for the IP address
    Example: *ema-swarm-lb-ip*
  - Click **Reserve**
- **Port**: *9092*
  You can choose an alternate port from the list. The important thing is that you follow the instructions in the Intel EMA Server Install Guide later to tell the server to advertise the matching port.

Click the **Done** button.

6.5.5 Review and finalize

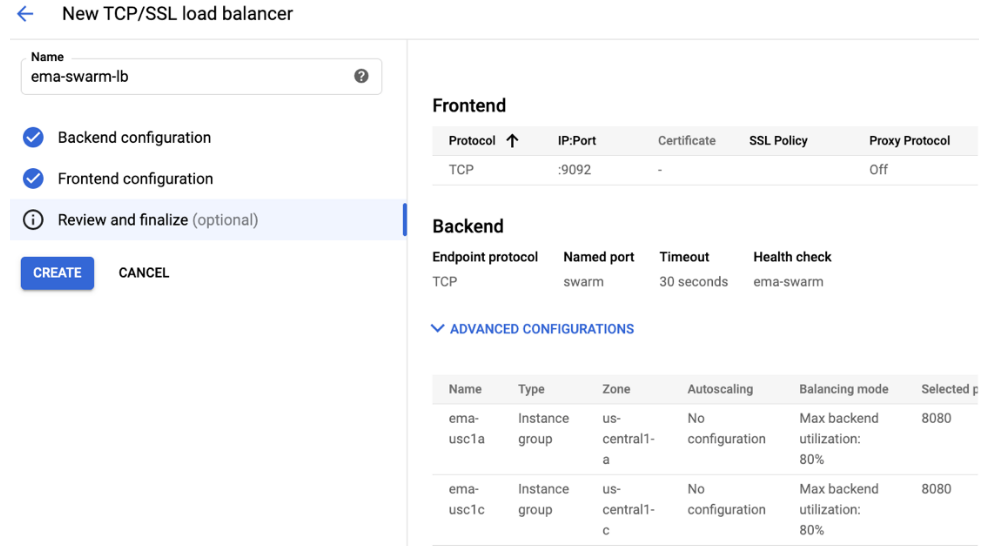

New TCP/SSL load balancer

Name: ema-swarm-lb

Frontend

Protocol ↑	IP:Port	Certificate	SSL Policy	Proxy Protocol
TCP	:9092	-		Off

Backend

Endpoint protocol	Named port	Timeout	Health check
TCP	swarm	30 seconds	ema-swarm

ADVANCED CONFIGURATIONS

Name	Type	Zone	Autoscaling	Balancing mode	Selected p
ema-usc1a	Instance group	us-central1-a	No configuration	Max backend utilization: 80%	8080
ema-usc1c	Instance group	us-central1-c	No configuration	Max backend utilization: 80%	8080

Click **Review and finalize**.

Review the information on the screen and then click the **Create** button.

6.6 DNS for Your Intel EMA Server

For a single-server deployment, if you have your own domain then you will want to create a DNS record pointing to the public IP address that was reserved for the Intel EMA virtual machine.

For a distributed server deployment, then you'll want to create DNS records pointing to the public IP addresses of the load balancers.

Consult with your DNS administrator on this task.

7 Appendix B - Notes on Active Directory Integration

As of February 2020, Managed Service for Microsoft AD has become generally available. We have not tested an Intel EMA deployment using this new service, but some links for further reading are included here.

https://cloud.google.com/blog/products/identity-security/managed-service-for-microsoft-active-directory-is-ga
https://cloud.google.com/managed-microsoft-ad/?hl=en_US