



G Cloud 10 Digital Marketplace

[Google Cloud Pricing Document](#)

Issue date: 14th May 2018

Introduction	2
Google Compute Engine (GCE)	3

Introduction

This document sets out the pricing applicable for Google's services available on the Digital Marketplace under Lot 1 and Lot 2. The pricing will be used as the basis for calculating the fees payable by the Buyer under its contractual arrangements with Google.

This page describes the pricing information for Compute Engine. To see the pricing for other products, read the [Pricing](#) documentation.

Google Compute Engine charges for usage based on the following price sheet. A bill is sent out at the end of each billing cycle, listing previous usage and charges. Prices on this page are listed in US dollars (USD).

Disk size, machine type memory, and network usage are calculated in gigabytes (GB), where 1 GB is 2^{30} bytes. This unit of measurement is also known as a [gibibyte \(GiB\)](#).

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

You can also find pricing information with the following options:

- See the estimated costs of your instances and Compute Engine resources when you create them in the [Google Cloud Platform Console](#).
- Estimate your total project costs with the [Google Cloud Pricing Calculator](#).

Always Free Usage Limits

As part of the [Google Cloud Platform Free Tier](#), Compute Engine offers an amount of usage that is free to use, up to a specific limit. These free usage limits are always available even during and after the free trial period. If you go over these usage limits and are no longer in the free trial period, you will be charged according to the price sheet.

- **1 f1-micro** VM instance per month (US regions, excluding Northern Virginia).
- **30 GB** of [Standard persistent disk](#) storage per month.
- **5 GB** of snapshot storage per month.
- **1 GB** egress from North America to other destinations per month (excluding Australia and China).

For f1-micro instances, you receive free usage equivalent to the number of total hours within the current month, enough to run one instance without interruption for the entire month. For example, March has 744 Hours (31 days x 24 hours). Therefore, you would receive 744 instance hours of free usage for the month.

All usage is aggregated across regions. Always Free Usage Limits are subject to change. Please see our [FAQ](#) for eligibility requirements and other restrictions.

Machine type pricing

Google Compute Engine offers two categories of machine types: **predefined machine types** and **custom machine types**. Predefined machine types have preset virtualized hardware properties and a set price, while custom machine types are priced according to the number of vCPUs and memory that the virtual machine instance uses.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

Machine type billing model

The following billing model applies to all machine types, predefined or custom. The billing model also applies to several [premium images](#) that you run on Compute Engine instances.

1. All machine types are charged a minimum of **1 minute**. For example, if you run your virtual machine for 30 seconds, you will be billed for 1 minute of usage.
2. After 1 minute, instances are charged in **1 second increments**.

★ **Note:** Certain [operating system images](#) can incur additional costs.

Predefined machine types

Google Compute Engine currently offers the following predefined machine types in the US, Europe, and Asia. Compute Engine also provides automatic discounts off these prices for sustained use. You can also use our [Google Cloud Pricing Calculator](#) to better understand price for different configurations.

Standard machine types

London

▼

Monthly

☒

Hourly

☐

Machine type	Virtual CPUs	Memory	Price (USD)	Preemptible price (USD)
n1-standard-1	1	3.75GB	\$0.0612	\$0.01230
n1-standard-2	2	7.5GB	\$0.1224	\$0.02460
n1-standard-4	4	15GB	\$0.2448	\$0.04920
n1-standard-8	8	30GB	\$0.4896	\$0.09840
n1-standard-16	16	60GB	\$0.9792	\$0.19680
n1-standard-32	32	120GB	\$1.9584	\$0.39360

n1-standard-64	64	240GB	\$3.9168	\$0.78720
n1-standard-96 Skylake Platform only	96	360GB	Not available in this region	Not available in this region
Custom machine type	If your ideal machine shape is in between two predefined types, using a custom machine type could save you as much as 40%. Read more about Custom Machine Types .			

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

Shared-core machine types

Shared-core machine types are ideal for applications that don't require a lot of resources. Shared-core machine types are more cost-effective for running small, non-resource intensive applications than standard, high-memory or high-CPU machine types.

f1-micro Bursting

f1-micro machine types offer bursting capabilities that allow instances to use additional physical CPU for short periods of time. Bursting happens automatically when your instance requires more physical CPU than originally allocated. During these spikes, your instance will opportunistically take advantage of available physical CPU in bursts. Note that bursts are not permanent and are only possible periodically.

London ▼		Monthly <input checked="" type="radio"/> Hourly		
Machine type	Virtual CPUs	Memory	Price (USD)	Preemptible price (USD)
f1-micro	1	0.60GB	\$0.0096	\$0.00420
g1-small	1	1.70GB	\$0.0324	\$0.00840

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** f1-micro instances get 0.2 of a vCPU and are allowed to burst up to a full vCPU for short periods. g1-small instances get 0.5 of a vCPU and are allowed to burst up to a full vCPU for short periods.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

High-memory machine types

High-memory machine types have 6.5 GB of system memory per vCPU. High-memory instances are ideal for tasks that require more memory relative to virtual CPUs.

London		Monthly <input checked="" type="radio"/> Hourly <input type="radio"/>		
Machine type	Virtual CPUs	Memory	Price (USD)	Preemptible price (USD)
n1-highmem-2	2	13GB	\$0.1523	\$0.03050
n1-highmem-4	4	26GB	\$0.3046	\$0.06100
n1-highmem-8	8	52GB	\$0.6092	\$0.12200
n1-highmem-16	16	104GB	\$1.2184	\$0.24400
n1-highmem-32	32	208GB	\$2.4368	\$0.48800
n1-highmem-64	64	416GB	\$4.8736	\$0.97600
n1-highmem-96 Skylake Platform only	96	624GB	Not available in this region	Not available in this region
Custom machine type	If your ideal machine shape is in between two predefined types, using a custom machine type could save you as much as 40%. Read more about Custom Machine Types .			

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

High-CPU machine types

High-CPU machine types have one vCPU for every 0.90 GB of system memory. High-CPU machine types are ideal for tasks that require more virtual CPUs relative to memory.

London

Monthly

Hourly

Machine type	Virtual CPUs	Memory	Price (USD)	Preemptible price (USD)
n1-highcpu-2	2	1.80GB	\$0.0912	\$0.0183
n1-highcpu-4	4	3.60GB	\$0.1824	\$0.0366
n1-highcpu-8	8	7.20GB	\$0.3648	\$0.0732
n1-highcpu-16	16	14.40GB	\$0.7296	\$0.1464
n1-highcpu-32	32	28.80GB	\$1.4592	\$0.2928
n1-highcpu-64	64	57.6GB	\$2.9184	\$0.5856
n1-highcpu-96 Skylake Platform only	96	86.4GB	Not available in this region	Not available in this region
Custom machine type	If your ideal machine shape is in between two predefined types, using a custom machine type could save you as much as 40%. Read more about Custom Machine Types.			

If you pay in a currency other than USD, the prices listed in your currency on Cloud Platform SKUs apply.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

Memory-optimized machine types

Memory-optimized machine types are ideal for tasks that require intensive use of memory with higher memory to vCPU ratios than high-memory machine types. Memory-optimized machine types are available in select regions only. [Learn more](#) about memory-optimized machine types.

London

Monthly

Hourly

Machine type	Virtual CPUs	Memory	Price (USD)	Preemptible price (USD)
n1-ultramem-40 ^(Beta)	40	938GB	Not available in this region	Not available in this region
n1-ultramem-80 ^(Beta)	80	1922GB	Not available in this region	Not available in this region
n1-megamem-96 ^(Beta)	96	1433.6GB	Not available in this region	Not available in this region
n1-ultramem-160 ^(Beta)	160	3844GB	Not available in this region	Not available in this region

Committed Use Discounts do not apply to memory-optimized machine types. If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

Custom machine types

Create a [custom machine type](#) with a specific number of vCPUs and amount of memory if predefined machine types are not optimal for your workloads. Custom machine types also save you the cost of running on a larger, more expensive machine type if you don't need to use all the resources of that machine type.

For example, instead of using an n1-standard-8 machine type when you need a machine type with 6 vCPUs, create a custom machine type instance with 6 vCPUs and 22.5 GB of memory. Creating a custom machine type can save you up to 40% compared to selecting a larger machine type.

Custom machine types are billed according to the number of vCPUs and the amount of memory used. The rates for custom machine types are in the table below.

For an accurate estimate of your billing with custom machine types, use the [Google Cloud Pricing Calculator](#).

★ **Note:** Certain specific [CPU selections](#) and [operating system images](#) may incur additional costs.

<div> <div>London ▼</div> <div> Monthly <input checked="" type="checkbox"/> Hourly <input type="checkbox"/> </div> </div>		
Item	Price (USD)	Preemptible price (USD)
vCPU	\$0.040692 / vCPU hour	\$0.00815 / vCPU hour

Memory	\$0.005453 / GB hour	\$0.00109 / GB hour
--------	----------------------	---------------------

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

★ **Note:** Google offers discounted predefined machine types so if you create an instance with a custom machine type that is equivalent to a predefined machine type in the GCP Console, the instance automatically uses the predefined machine type. However, the `gcloud` command-line tool and the API do not support this functionality yet.

Extended memory

When using a custom machine type, any memory up to and including 6.5 GB of memory per vCPU is charged at the standard [custom machine type pricing](#). Any memory above the 6.5 GB per vCPU will be charged according to the extended memory prices that are described in detail below. Learn more about [Extended Memory](#).

London		Monthly <input checked="" type="radio"/> Hourly
Item	Price (USD)	Preemptible price (USD)
Extended Memory	\$0.01171 / GB hour	\$0.00235 / GB hour

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Sustained use discounts

If you run an instance for a significant portion of the billing month, your VM instances will automatically qualify for a **sustained use discount**. When you use an instance for more than 25% of a month, Compute Engine automatically gives you a discount for every incremental second you use for that instance. The discount increases with usage and you can get up to a 30% net discount for instances that run the entire month.

Sustained use discounts are applied automatically and will be calculated and added to your bill as your project earns them. There is no action needed on your part to enable sustained use discounts.

To learn more about sustained use discounts, see the [Sustained Use Discounts](#) documentation.

Committed use discounts

Compute Engine offers the ability to purchase a [committed use contract](#) in return for deeply discounted prices for VM usage. These discounts are known as **committed use discounts**. You can purchase a committed use contract by creating a [commitment](#). Commitments are appropriate for predictable and steady state usage where you will use a specific amount of vCPUs and system memory for future workloads. Commitments allow you to purchase a specific number of vCPUs and amount of memory at up to a 57% discount over full prices. You commit to the entire usage term and are billed for each month regardless of whether usage has occurred.

You can only use committed use discounts for predefined machine types and custom machine types. Small machine types, such as f1-micro and g1-small, are not eligible for committed use discounts. Memory-optimized machine types are not eligible for committed use discounts.

London

Monthly

Hourly

Item	1 year commitment price (USD)	3 year commitment price (USD)
vCPU	\$0.025636 / vCPU hour	\$0.0183114 / vCPU hour
Memory	\$0.003435 / GB hour	\$0.0024539 / GB hour

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** Listed monthly pricing assumes the instance runs for a 730 hour month.

With committed use discounts, VM prices can be up to 57% less expensive than regular VM prices. For example, the table below compares the 1 and 3 year prices to regular prices in the `us-west1`, `us-central1`, and `us-east1` regions.

Machine Type	Item	Full Price	1 year commitment	3 year commitment
Custom machine types	vCPU	\$0.033174 / vCPU hour	\$0.019915 / vCPU hour	\$0.014225 / vCPU hour
	Memory	\$0.004446 / GB hour	\$0.002669 / GB hour	\$0.001907 / GB hour
Predefined machine types*	n1-standard-16	\$0.7600 / hour	\$0.47878 / hour	\$0.34202 / hour
	n1-highmem-16	\$0.9472 / hour	\$0.59622 / hour	\$0.42593 / hour

n1-highcpu-16	\$0.5672 / hour	\$0.35707 / hour	\$0.25506 / hour
----------------------	-----------------	------------------	------------------

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Committed use discounts can apply to all predefined n1- machine types and custom machine types. They do not apply to small machine types, f1-micro and g1-small. For comparison, the tables use prices for 16 vCPU machine-types.

Instance uptime

Instance uptime is measured as the number of seconds between when you start an instance and when you stop an instance, the latter being when the instance state is `TERMINATED` . In some cases, your instance can suffer from a failure and be marked as `TERMINATED` by the system; in these cases, you will not be charged for usage after the instance reaches the `TERMINATED` state. If an instance is idle, but still has a state of `RUNNING` , it will be charged for instance uptime. The easiest way to determine the status of an instance is to use `gcloud compute instances list` command or to visit the [Google Cloud Platform Console](#).

Note that Google Compute Engine bills for a minimum of 1 minute of usage, so if you run an instance for 30 seconds of uptime, you are billed for 1 minute. After 1 minute, your instance is billed on a per-second basis. For more information, see the [billing model](#).

GPU pricing

Attach one or more GPUs to your instances to accelerate specific workloads or offload work from your vCPUs. Each GPU adds to the cost of your instance in addition to the cost of the [machine type](#).

GPU devices are available only in specific regions and zones. Read [GPUs on Compute Engine](#) to see a complete list of regions and zones where GPU devices are available.

GPU devices attached to non-preemptible instances receive [sustained use discounts](#) similar to vCPUs. For GPUs attached to [preemptible instances](#), you will be charged at the preemptible prices for GPUs but will not receive sustained use discounts.

★ **Note:** [Preemptible GPUs](#) are currently in Beta.

London

Monthly ☒ Hourly

Model	GPUs	GPU memory	GPU price	Preemptible GPU (Beta) price
NVIDIA® Tesla® V100	1 GPU	16 GB HBM2	Not available in this region	Not available in this region

	8 GPUs	128 GB HBM2		
NVIDIA® Tesla® P100	1 GPU	16 GB HBM2	Not available in this region	Not available in this region
	2 GPUs	32 GB HBM2		
	4 GPUs	64 GB HBM2		
NVIDIA® Tesla® K80	1 GPU	12 GB GDDR5	Not available in this region	Not available in this region
	2 GPUs	24 GB GDDR5		
	4 GPUs	48 GB GDDR5		
	8 GPUs	96 GB GDDR5		

★ **Note:** NVIDIA® K80® boards contain two GPUs each. The pricing for K80 GPUs is by GPU, not by the board.

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

★ **Note:** Listed monthly pricing includes applicable, automatic [sustained use discounts](#), assuming the instance runs for a 730 hour month.

Use the [Google Cloud Platform Pricing Calculator](#) to help determine the total cost of your instances including both the cost of GPUs and machine type configurations.

To learn more about how you can use GPUs to accelerate your applications, see [GPUs on Compute Engine](#).

Premium images

The price for a premium image is different depending on which machine type you use. For example, a standard SUSE image costs \$0.02 per hour to run on an `f1-micro` instance, but the same image costs \$0.11 per hour to run on an `n1-standard-8` instance. The prices for premium images are the same worldwide and do not differ based on zones or regions.

All prices for premium images are in addition to charges for using a machine type. For example, the total price to use an `n1-standard-8` instance with a SUSE image would be the sum of the machine type cost and the image cost:

`n1-standard-8` cost + SUSE image cost = \$0.3800 + \$0.11 = \$0.49 per hour

★ **Note:** [Preemptible instances](#) do not reduce the costs of premium images and do not change the way that you are billed for the use of those images. If Compute Engine terminates a preemptible instance that runs a premium image, you are billed for that image as if you terminated the instance yourself. The charges for minimum usage still apply, and bills for premium images are still calculated by rounding up to the nearest usage increment.

Red Hat Enterprise Linux (RHEL) and RHEL for SAP images

- **\$0.06 USD/hour** for instances with 4 or fewer vCPUs
- **\$0.13 USD/hour** for instances with more than 4 vCPUs

All RHEL and RHEL for SAP images are charged a **1 minute** minimum. After 1 minute, RHEL images are charged in **1 second increments**.

Google reports your billing entity name and total hours of Red Hat premium OS usage on Compute Engine to Red Hat, which complies with the Red Hat licensing requirements.

SUSE images and SLES for SAP images

SLES images:

- **\$0.02 USD/hour** for `f1-micro` and `g1-small` machine types
- **\$0.11 USD/hour** for all other machine types

SLES for SAP images:

- **\$0.17 USD/hour** for instances with 1 - 2 vCPUs
- **\$0.34 USD/hour** for instances with 3 - 4 vCPUs
- **\$0.41 USD/hour** for instances with 5 or more vCPUs

All SUSE images are charged a **1 minute** minimum. After 1 minute, SUSE images are charged in **1 second increments**.

Windows Server images

Public images for several versions of Windows Server are available in either the Server Core configuration or the Server with Desktop Experience configuration. Both configurations are available at the following prices:

- **\$0.02 USD/hour** for `f1-micro` and `g1-small` machine types
- **\$0.04 USD per core/hour** for all other machine types

Standard machine types, high-CPU machine types, and high-memory machine types are charged based on the number of CPUs. For example, `n1-standard-4`, `n1-highcpu-4`, and `n1-highmem-4` are machine-types with 4 vCPUs, and are charged at \$0.16 USD/hour (4 x \$0.04 USD/hour).

Windows Server images are charged a **1 minute minimum**. After 1 minute, Windows images are charged in **1 second**

increments. [SQL Server images](#) are charged a **10 minute minimum**. After 10 minutes, SQL Server images are charged in **1 minute increments**.

SQL Server images

SQL Server images incur costs in addition to the base cost for normal [Windows Server images](#).

- **\$0.399 USD per core/hour** for SQL Server Enterprise
- **\$0.1645 USD per core/hour** for SQL Server Standard
- **\$0.011 USD per core/hour** for SQL Server Web
- No additional charge for SQL Server Express

Microsoft SQL Server licensing requires a core license to be applied to each virtual CPU on your virtual machine instance with a four core minimum for each instance. For example, on instances with fewer than 4 vCPUs, Compute Engine charges for SQL Server Standard at 4 x \$0.1645 USD/hour (\$0.658 USD/hour).

For instances with 4 or more vCPUs, Compute Engine charges you for Microsoft SQL Server licenses in increments of 2. Instances with [custom machine types](#) receive a number of SQL Server licenses that is equal to the number of vCPUs.

Google recommends that you do not use SQL Server images on [f1-micro](#) or [g1-small](#) machine types based on Microsoft's [minimum hardware and software recommendations](#) [\[2\]](#).

Unlike other premium images, SQL Server images are charged a **10 minute minimum**. After 10 minutes, SQL Server images are charged in **1 minute increments**.

Network pricing

General network pricing

Traffic type	Price
Ingress	No charge
Egress* to the same zone	No charge
Egress to Google products (such as YouTube, Maps, Drive), whether from a VM in GCP with a public (external) IP address or a private (internal) IP address	No charge
Egress to a different Google Cloud Platform service within the same region, except for Cloud Memorystore for Redis and for Cloud SQL	No charge

Egress* between zones in the same region (per GB)	\$0.01
Egress to Cloud Memorystore for Redis is charged at the rate of "Egress between zones in the same region"	
Egress to Cloud SQL is charged at the rates described in Traffic through external IP addresses	
Egress between regions within the US (per GB)	\$0.01
Egress between regions, not including traffic between US regions	At Internet egress rates

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

*These charges are for egress through [internal IP addresses](#). There are different [charges](#) for egress through [external IP addresses](#), regardless of whether the instances are in the same zone. The prices are used regardless of network or subnet. The price for traffic within a zone using internal IP addresses is the same even if the traffic is to a different subnet or network. The price on traffic between zones in the same region is the same if the two instances are in the same subnet, different subnets, or different networks. Pricing is the same whether the instances are in a VPC network or a legacy network.

*Promotional pricing

Internet egress rates*

London				
Monthly Usage	Network (Egress) Worldwide Destinations (excluding China & Australia, but including Hong Kong) (per GB)	Network (Egress) China Destinations (excluding Hong Kong) (per GB)	Network (Egress) Australia Destinations (per GB)	Network (Ingress)
0-1 TB	\$0.12	\$0.23	\$0.19	Free
1-10 TB	\$0.11	\$0.22	\$0.18	Free
10+ TB	\$0.08	\$0.20	\$0.15	Free

★ **Note:** The monthly usage is counted by each SKU, while the "Worldwide Destinations" may contain multiple SKUs. For more details please refer [Cloud Platform SKUs](#)

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

*These rates do not apply for [Cloud CDN](#), [CDN Interconnect](#), [Carrier Peering](#), [Direct Peering](#), and [Cloud Interconnect](#) traffic. Instead, see pricing for [Cloud CDN](#), [CDN Interconnect](#), [Carrier Peering](#), [Direct Peering](#), or [Cloud Interconnect](#).

Load balancing and forwarding rules

The following applies to all types of load balancing and forwarding rules (protocol forwarding).

London		
Item	Price per Unit (USD)	Pricing Unit
First 5 forwarding rules	\$0.030	Per Hour
Per additional forwarding rule	\$0.012	Per Hour
Ingress data processed by load balancer	\$0.010	Per GB

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Load balancing ingress and egress charges

- The Load balancing and forwarding table above contains the charge for ingress data processed by load balancers.
- Normal egress rates are charged for traffic outbound from a load balancer.
- There is no additional load balancer egress cost beyond normal egress rates.

Forwarding rules charges

Compute Engine charges for forwarding rules that are created for [load balancing](#) or other uses of [forwarding rules](#).

The following examples use US pricing:

Up to 5 forwarding rules you create are charged at \$0.025/hour. For example, if you create one forwarding rule, you will be charged \$0.025/hour. If you have 3 forwarding rules, you will still be charged \$0.025/hour. However, if you have 10 rules, you will be charged:

- 5 forwarding rules = \$0.025/hour
- Each additional forwarding rule = \$0.01/hour

$\$0.025/\text{hour for 5 rules} + (5 \text{ additional rules} * \$0.01/\text{hour}) = \$0.075/\text{hour}$

Traffic through external IP addresses

When you send traffic between virtual machines through the external IP addresses, it will be charged as follows:

- Traffic sent between the external IP address of two virtual machines in the same region will be charged as [Egress between zones in the same region](#), even if the virtual machines are in the same zone.
- Traffic sent between the external IP address of two virtual machines in different regions within the US will be charged as [Egress between regions within the US](#).
- Traffic sent between the external IP address of two virtual machines in different regions, not including traffic between US regions, will be charged at [Internet Egress pricing](#).

It is not possible to charge traffic through the external IP address of two virtual machine instances in the same zone as **Egress to the same zone** because Compute Engine cannot determine the zone of a virtual machine through the external IP address.

VPN

Cloud VPN pricing is based on the location of the VPN endpoint and the number of tunnels per hour.

London	
Component Billed	Price (USD)
Per tunnel (per hour)	\$0.060
IPsec traffic	Charged the same as if it wasn't in a VPN tunnel. See General network pricing .
Public IP for VPN Gateway	Charged according to IP address pricing .

Google does not charge for forwarding rules that send traffic to the VPN gateway.

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Disk pricing

All disk-related charges in this section are prorated based on a granularity of seconds. This includes both types of persistent disk types, snapshot storage, and local SSD pricing.

For example, based off US pricing, a 200 GB standard persistent disk volume would cost \$8.00 for the whole month. If you only provisioned a 200 GB volume for half a month, it would cost \$4.00. Likewise, a 200 GB SSD persistent disk volume would cost \$34.00 for the whole month. If you only provisioned a 200 GB volume for half a month, it would cost \$17.00.

Persistent disk pricing

Persistent disks are charged for the amount of provisioned space per disk. Persistent disk I/O operations are included in the charges for provisioned space and persistent disk performance grows linearly to the size of the persistent disk volume, so you might want to create a larger or smaller persistent disk to account for your I/O needs. For more information, see the [persistent disk specifications](#).

After you successfully delete a persistent disk, you will no longer be charged for that disk.

London	
Type	Price (per GB / month)
Standard provisioned space	\$0.048
SSD provisioned space	\$0.204
Snapshot storage	\$0.031
IO operations	No additional charge

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Persistent disk snapshots

Persistent disk snapshots are only charged for the total size of the snapshot. For example, if you only used 2TB of disk space on a 5TB persistent disk, your snapshot size will be around 2TB, rather than the full 5TB of provisioned disk space. Google Compute Engine also provides [incremental snapshots](#), which means that after the initial snapshot, subsequent snapshots only contain data that has changed since the previous snapshot, providing for a generally lower cost for snapshot storage.

Note that snapshot storage is also calculated based on a prorated rate.

Network charges for snapshot restoration

Standard [network pricing](#) applies to snapshot restore. For example, a snapshot restore across regions or continents incurs the same fees as Google Compute Engine egress.

Local SSD pricing

Local SSD devices are charged for the amount of provisioned space per device for the lifetime of the instance it is attached to. The prices for local SSDs differ depending on the region. For example, in the Iowa, Oregon, Taiwan, and Belgium regions, local SSDs cost \$0.080 per GB per month. As mentioned earlier, local SSD charges are prorated to a granularity of seconds.

Since local SSDs can only be purchased in 375 GB increments, the cost per month for a single device is the monthly rate multiplied by 375 GB. For example, at a monthly rate of \$0.080, the cost would be \$30.00 per device per month. Actual data storage and usage are included in that price and there is no additional charge for local traffic between the virtual machine and the local SSD device.

London		Monthly <input checked="" type="radio"/> Hourly	
Type	Price (per GB / month)	Preemptible price (per GB / month)	Estimated cost per device*
Local SSD provisioned space	\$0.096	\$0.058	\$0.049

*Based on a 730 hour month on standard (non-preemptible) rate; actual monthly costs may vary.
**Local SSD devices are only offered in one size, at 375 GB per device.
If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Image storage

London	
Type	Price

Image Storage (per GB / month)

\$0.102

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

Unused IP address pricing

If you reserve a [static external IP address](#) but do not use it, you will be charged for the IP address according to the table below. If you reserve a static external IP address and use it with a Compute Engine resource, such as VM instance or a forwarding rule, the address is considered in use and you will not be charged for it.

London

Type	Price/Hour
Static IP address (assigned but unused)	\$0.012
Static IP address (assigned and in use)	No charge
Ephemeral IP address (attached to instance or forwarding rule)	No charge

If you pay in a currency other than USD, the prices listed in your currency on [Cloud Platform SKUs](#) apply.

You can check whether a static external IP address is in use by making a `gcloud compute addresses list` request. This command returns a list of static external IP addresses and their statuses:

```
gcloud compute addresses list
```

NAME	REGION	ADDRESS	STATUS
address-1		130.211.8.68	IN_USE
address-2		35.186.217.84	RESERVED

In this example, IPv4 `address-1` is in use and would not be charged while IPv4 `address-2` is reserved but not being used. `address-2` would be charged according to the unused IPv4 address pricing.

Simulated maintenance event pricing

★ Beta

This is a Beta release of Simulated Maintenance Events. This feature is not covered by any SLA or deprecation policy and may be subject to backward-incompatible changes.

Running [simulated maintenance events](#) incurs the following charges:

- Simulated maintenance on instances configured for [live migration](#) incur costs for each of the following instance resources:
 - Price per vCPU on the instance, where f1-micro and g1-small are each equivalent to 1 vCPU: \$0.040
 - Price per GB of system memory: \$0.010
 - Price per GB of local SSD space: \$0.001
- Simulated maintenance on [preemptible VM instances](#): Free
- Simulated maintenance on instances configured to [terminate and restart](#): Free

★ **Note:** Normal 1 minute minimum usage charges for machine types and premium images still apply to instances that you terminate or preempt during a simulated maintenance event. See the [machine type billing model](#) and [\[premium image pricing\]\(premium images\)](#) for details.

Viewing usage

The Google Cloud Platform Console provides a transaction history for each of your projects. This history describes your current balance and estimated resource usage for that particular project.

To view a project's transaction history, go to the [estimated billing invoice page](#).

Pricing calculator

To help you understand how your bill is calculated, use the [Google Cloud Pricing Calculator](#).

What's next?

- Refer to the [Pricing Overview](#) documentation.
- Try the [Pricing Calculator](#).
- Learn more about Compute Engine [storage options](#).
- Learn more about Compute Engine [machine types](#).
- Learn more about Compute Engine [operating Systems](#).
- Learn more about Compute Engine [IP addresses](#).
- Learn more about Compute Engine [load balancing](#).

Was this page helpful? Let us know how we did:



[SEND FEEDBACK](#)

Except as otherwise noted, the content of this page is licensed under the [Creative Commons Attribution 3.0 License](#), and code samples are licensed under the [Apache 2.0 License](#). For details, see our [Site Policies](#). Java is a registered trademark of Oracle and/or its affiliates.

Last updated May 11, 2018.