


# Networking in Google Cloud

Module 3: Network Billing and Pricing



Welcome to the Network Billing and Pricing module. This is the third module of the Networking in Google Cloud: Hybrid Connectivity and Network Management course. We will discuss how Google Cloud products and services are billed.

01 Network Pricing

---

02 Network Service Tiers

---


03 Lab: Optimizing Network Spend  
with Network Tiers

---

04 Billing

---

05 Quiz



In this module, we will cover how Google Cloud networking features are charged for, how to leverage Network Service Tiers, and how to administer billing within Google Cloud.

The focus of this module is to help you optimize your network spend in Google Cloud.

You will get to apply this in a lab with Network Service Tiers and you will learn how to use labels to better understand network spend in Google Cloud.

Let's start by talking about networking pricing within Google Cloud.

01 Network Pricing

---

02 Network Service Tiers

---


03 Lab: Optimizing Network Spend  
with Network Tiers

---

04 Billing

---

05 Quiz



01 Network Pricing

---

02 Network Service Tiers

---


**03** Lab: Optimizing Network Spend  
with Network Tiers

---

04 Billing

---

05 Quiz



01 Network Pricing

---

02 Network Service Tiers

---


03 Lab: Optimizing Network Spend  
with Network Tiers

---

**04 Billing**

---

05 Quiz



01 Network Pricing

---

02 Network Service Tiers

---


03 Lab: Optimizing Network Spend  
with Network Tiers

---

04 Billing

---

**05** Quiz



## Each Google Cloud service has its own pricing model



Virtual Private  
Cloud



Cloud Firewall  
Rules



Cloud  
DNS



Cloud  
IAM



Cloud Load  
Balancing



Cloud  
CDN



Cloud External  
IP Addresses



Google Cloud  
Armor



Cloud  
VPN



Cloud  
Router



Dedicated  
Interconnect

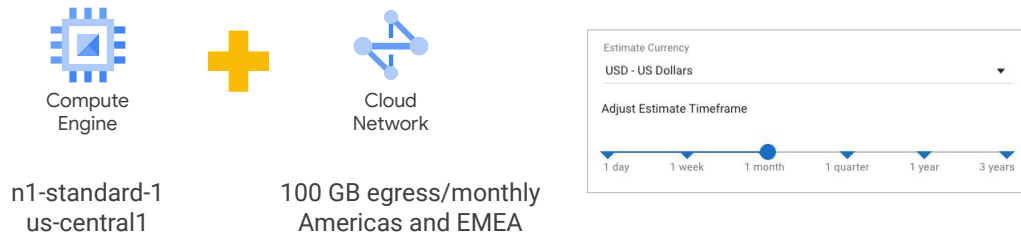


Partner  
Interconnect

Here are some of the many Google Cloud services that we have covered in this course. Each of these services has its own pricing model and dedicated page within the documentation of each service.

Updating a pricing page is much easier than updating the recording of this training. Therefore, for up-to-date pricing, always refer to the documentation.

# Estimate costs with the Google Cloud Pricing Calculator



The diagram illustrates the configuration for the Google Cloud Pricing Calculator. It features two service icons: a blue chip icon for 'Compute Engine' and a blue network icon for 'Cloud Network', separated by a yellow plus sign. Below the Compute Engine icon is the text 'n1-standard-1 us-central1'. Below the Cloud Network icon is the text '100 GB egress/monthly Americas and EMEA'. To the right is a screenshot of the calculator interface, showing a dropdown menu for 'Estimate Currency' set to 'USD - US Dollars' and a timeline for 'Adjust Estimate Timeframe' with '1 month' selected.

<https://cloud.google.com/products/calculator/>

Because each Google Cloud service has its own pricing model, we recommend using the Google Cloud pricing calculator to estimate the cost of a collection of resources.

The pricing calculator is a web-based tool that you use to specify the expected consumption of certain services and resources, and it then provides you with an estimated cost.

For example, you can specify an n1-standard-1 VM instance in us-central1 along with 100 GB of egress traffic to Americas and EMEA. The pricing calculator then returns the total estimated cost.

You can adjust the currency and time frame to meet your needs and when you are done, you can email the estimate or save it to a specific URL for future reference.

To use the pricing calculator today, click the link in the slide.



## General network pricing

Traffic type	Price
Ingress	No charge
Egress to the same zone (internal IP address)	No charge
Egress to Google products (YouTube, Maps, Drive)	No charge
Egress to a different Google Cloud service (within same region, exceptions)	No charge
Egress between zones in the same region (per GB)	\$0.01
Egress to the same zone (external IP address, per GB)	\$0.01
Egress between regions within the US (per GB)	\$0.01
Egress between regions, not including traffic between US regions	At internet egress rates



Because pricing can change over time, I will only focus on general network pricing in this module. Specifically, ingress and egress traffic through internal and external IP addresses.

This table is from the Compute Engine documentation, and it lists the price of each traffic type. These rates do not apply for Cloud CDN, Interconnect, Carrier Peering, Direct Peering, and Cloud Interconnect traffic.

First of all, ingress or traffic coming into Google Cloud's network is not charged. The rest of this table lists egress or traffic leaving a Compute Engine instance. Egress traffic that is not charged for is traffic to the same zone, as long as that egress is through the internal IP address of an instance. Also, egress traffic to Google products, like YouTube, Maps, Drive or traffic to a different Google Cloud service within the same region is not charged for either.

However, there is a charge for egress between zones in the same region, egress within a zone if the traffic is through the external IP address of an instance, and egress between regions. We will cover the internet egress rates shortly because they depend on which Network Service Tier you use.

As for the difference in egress traffic to the same zone, Compute Engine cannot

determine the zone of a virtual machine through the external IP address. Therefore, this traffic is treated like egress between zones in the same region.

Also, there are some exceptions, and pricing can always change, so please refer to the documentation links in the slides.

## External IP address pricing (us-central1) (Subject to change)

Type	Price/Hour (USD)
Static IP address (assigned but unused)	\$0.010
Static and ephemeral IP addresses in use on standard VM instances	\$0.004
Static and ephemeral IP addresses in use on preemptible VM instances	\$0.002
Static and ephemeral IP addresses attached to forwarding rules	No charge



Now, you are charged for static and ephemeral external IP addresses. This table represents the external IP pricing for us-central1 as of this recording.

You can see that if you reserve a static external IP address and do not assign it to a resource, such as a VM instance or a forwarding rule, you are charged at a higher rate than for static and ephemeral external IP addresses that are in use.

Also, external IP addresses on preemptible VMs have a lower charge than for standard VM instances.

Remember, pricing can always change, so please refer to the documentation links in the slides.



# Agenda

01 Networking Pricing

---

02 [Network Service Tiers](#)

---

03 Lab: Optimizing Network Spend  
with Network Tiers

---

04 Billing

---

05 Quiz

---



I mentioned earlier that internet egress rates depend on which Network Service Tier you use. So, what are Network Service Tiers?

## Use Network Service Tiers to optimize your network for performance or cost

Premium Tier	Standard Tier
High performance routing (Google's network)	Lower price and performance than Premium
Unique to Google Cloud	Comparable to other public cloud offerings
Global SLA	No global SLA
Global load balancing, Cloud CDN	Regional load balancing
Performance is main consideration	Cost is main consideration

Network Service Tiers enable you to optimize your cloud network for performance by choosing Premium Tier or for cost with the new Standard Tier. So what is the difference between these two tiers?

The Premium Tier delivers traffic on Google's global network, providing high-performance routing. If you use Google Cloud today, you already use the powerful Premium Tier. The Standard Tier, alternatively, offers an attractively priced network with performance comparable to that of other major public clouds. There are other differences between the two tiers: The Premium Tier offers a global SLA and allows for Global Load Balancing and Cloud CDN, as we explored in a previous module. The Standard Tier does not have a global SLA, and Load Balancing is limited to regional.

So why would you choose the Standard Tier? Well, it all comes back to optimizing your cloud network for performance by choosing Premium Tier or for cost with the Standard Tier. In other words, Network Service Tiers allows you to design your cloud network, your way.

Let's explore each Network Service Tier to better understand the network performance and cost differences.

# Google Cloud

Regions, PoPs, and network

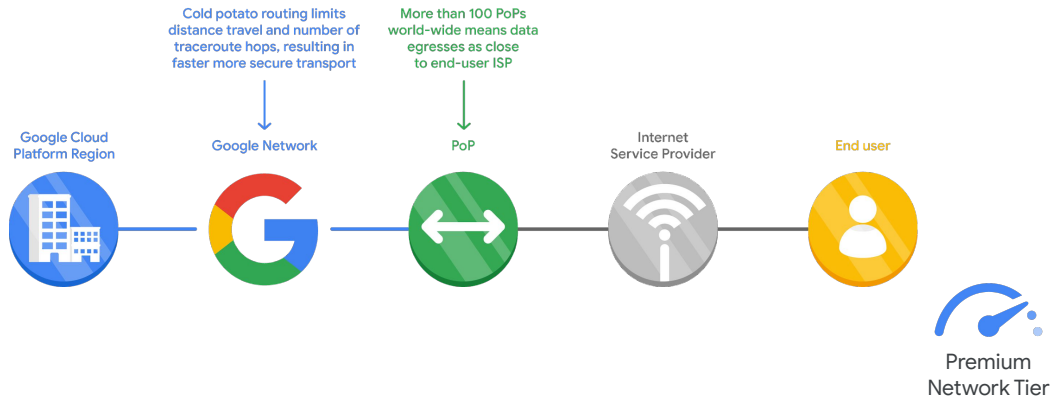


Premium Tier delivers traffic over Google's well-provisioned, low latency, highly reliable global network.

As you can see on this map, this network consists of an extensive global private fiber network with over 100 points of presence across the globe.


# Optimize performance with Premium Tier

## Premium



In Premium Tier, inbound traffic from your end user to your application in Google Cloud enters Google's private, high performance network at the POP closest to your end user, and Google Cloud delivers this traffic to your application over this network. Similarly, Google Cloud delivers outbound traffic from your application to end users on Google's network and exits at the POP closest to them, wherever the end users are across the globe. Which means that most of this traffic will reach its destination with a single hop to the end user's ISP, so it enjoys minimum congestion and maximum performance.

# Optimize cost with Standard Tier



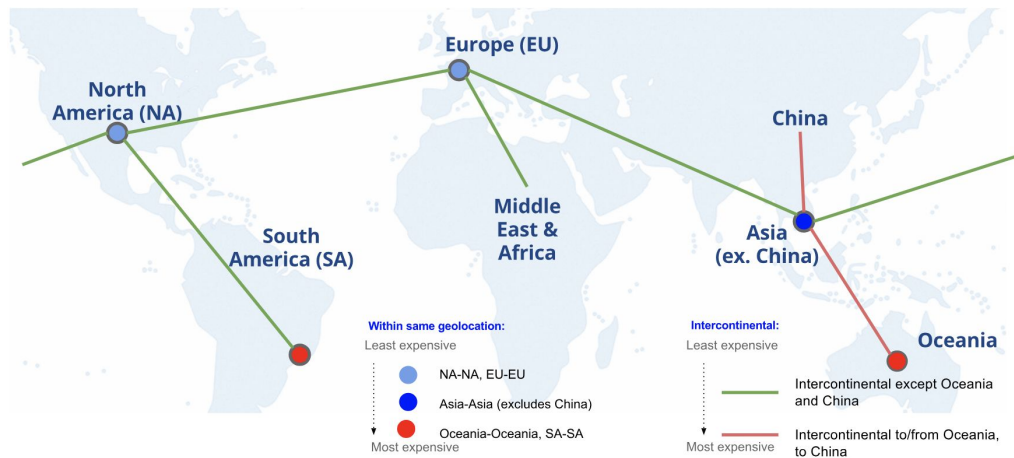
Standard Tier provides network quality that is comparable to other public cloud providers, but lower than Premium Tier. Also, regional network services such as Regional Load Balancing have one VIP per region.

Standard tier is priced lower than Premium because your traffic between Google Cloud and your end user is delivered over ISP networks instead of Google's network.

Now that you understand the differences in performance, let's get into cost.



# Network Service Tiers pricing



Premium Tier pricing is based on both source and destination of traffic. This is because the cost of network traffic varies with the distance your traffic travels over Google's network. In contrast, Standard Tier traffic is source-based because it does not travel much over Google's network.

This map illustrates that Network Service Tiers categorizes all countries and continents into the listed geolocations. Depending from where to where traffic travels, costs will vary.

# Premium Tier pricing



Iowa (us-central1) ▾
 

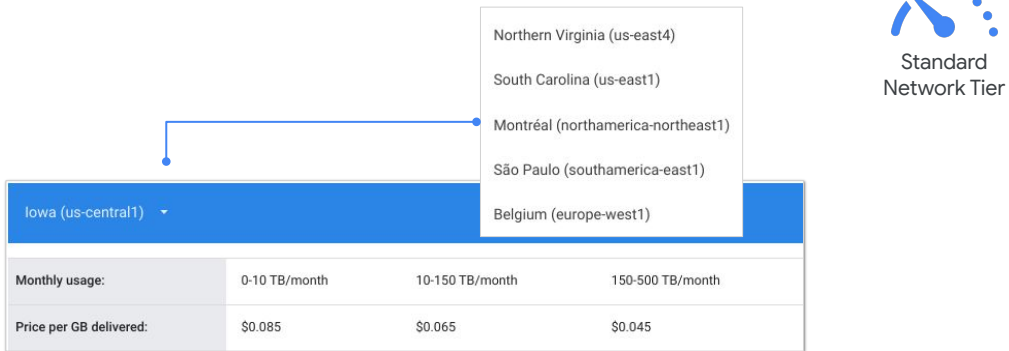
- Northern Virginia (us-east4)
- South Carolina (us-east1)
- Montréal (northamerica-northeast1)
- São Paulo (southamerica-east1)
- Belgium (europe-west1)

Monthly Usage	Network (Egress) Worldwide Destinations (excluding China & Australia, but including Hong Kong) (per GB in USD)	Network (Egress) China Destinations (excluding Hong Kong) (per GB in USD)	Network (Egress) Australia Destinations (per GB in USD)	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

This Premium Tier pricing table differentiates traffic into 3 monthly usage levels: 0-1 TB, 1-10 TB, and greater than 10 TB. Ingress pricing is free. These prices are applied both during and after the Google Cloud Free Tier period. During the Free Tier period, the prices are charged against the Free Tier credit amount.

For the latest details, and to access the interactive pricing matrix, refer to Google Cloud's pricing documentation.

# Standard Tier pricing

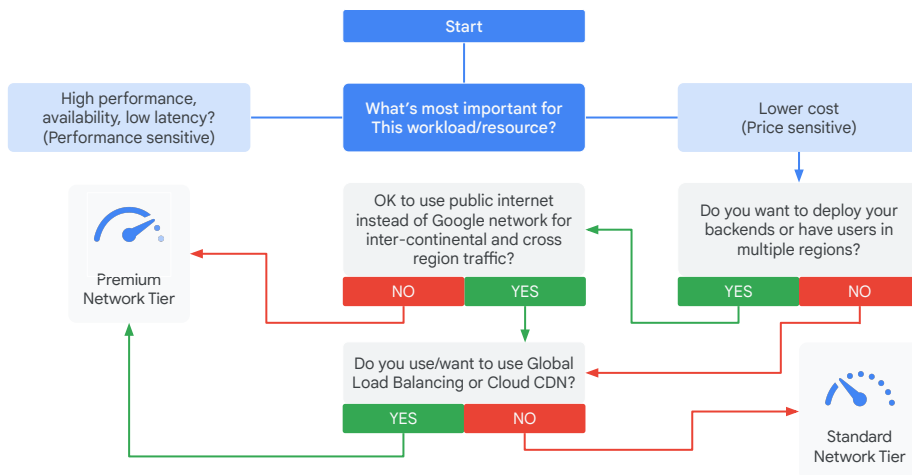


Monthly usage:	0-10 TB/month	10-150 TB/month	150-500 TB/month
Price per GB delivered:	\$0.085	\$0.065	\$0.045

This Premium Tier pricing table differentiates traffic into 3 monthly usage levels: 0-1 TB, 1-10 TB, and greater than 10 TB. Ingress pricing is free. These prices are applied both during and after the Google Cloud Free Tier period. During the Free Tier period, the prices are charged against the Free Tier credit amount.

For the latest details, and to access the interactive pricing matrix, refer to Google Cloud's pricing documentation.

## Network Service Tiers decision tree




We have gone over both the performance and cost differences between the Network Service Tiers. If you need more guidance on choosing the tier that best meets your needs, use this decision tree.

Ask yourself whether high performance or lower cost is most important for your workload or resource. The Premium Tier is the clear choice for performance. If cost is your main consideration, remember that the Standard Tier has other restrictions in addition to network performance. If you want to deploy your backends or have users in multiple regions, but don't want to use the public internet over Google's network for inter-continental and cross region traffic, you want to choose the Premium Tier. Also, if you want Global Load Balancing or Cloud CDN, you need to use the Premium Tier.

Otherwise, the Standard Tier is a great choice if you don't need any of those services and are okay using the public internet instead of Google's network.

## Lab intro

Optimizing Network Spend  
with Network Tiers



Let's apply what we just covered to optimize network spend with Network Service Tiers.

In this lab, you create one VM using the Premium tier, which is the default, and one VM using the Standard tier. Then, you compare the latency and network paths for each VM instance.



# Today's agenda

- 01 Networking Pricing

---

- 02 Network Service Tiers

---

- 03 Lab: Optimizing Network Spend with Network Tiers

---

- 04 **Billing**

---

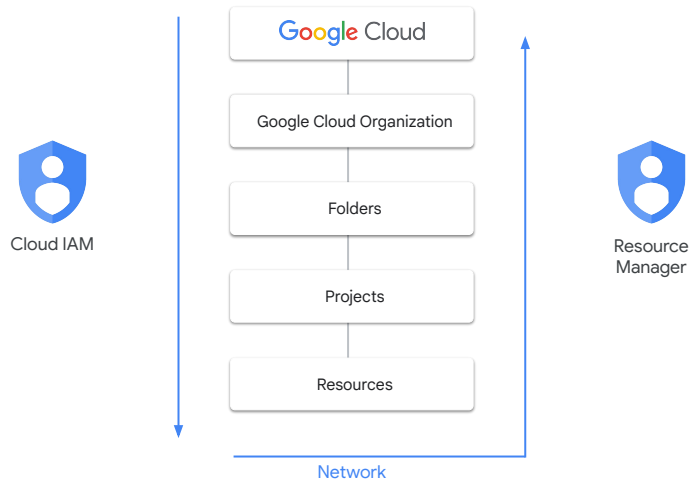
- 05 Quiz

---




Now, that we have covered the pricing of different network services, let's get into how those services are billed.

## Billing is accumulated from the bottom up

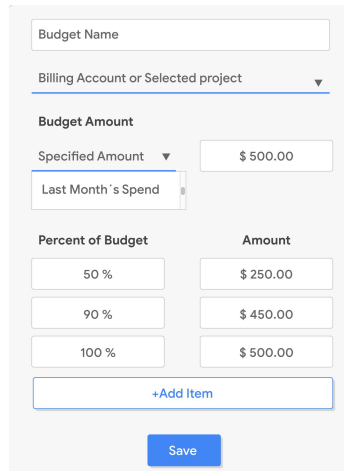


In a previous module, we mentioned that Cloud IAM policies are inherited top-to-bottom, as we can see on the left-hand side. Billing, on the other hand, is accumulated from the bottom up, as we can see on the right-hand side. Because a resource belongs to only one project, a project accumulates the consumption of all its resources.

This is even true for Shared VPC. Billing for resources that participate in a Shared VPC network is attributed to the service project where the resource is located, even though the resource uses the Shared VPC network within the host project.

Now, each project is associated with one billing account, meaning that an organization contains all billing accounts.

## Configure budgets and alerts to control costs



The screenshot shows a form for creating a budget. It includes a text input for 'Budget Name', a dropdown for 'Billing Account or Selected project', and a 'Budget Amount' section. Under 'Budget Amount', there is a 'Specified Amount' dropdown set to '\$ 500.00' and a 'Last Month's Spend' dropdown. Below this is a table for 'Percent of Budget' alerts:

Percent of Budget	Amount
50 %	\$ 250.00
90 %	\$ 450.00
100 %	\$ 500.00

At the bottom of the form, there is a '+Add Item' button and a 'Save' button.

To help you with project planning and controlling costs, you can set a budget in Google Cloud. Setting a budget lets you track how your spend is growing toward that amount.

This screenshot shows the budget creation interface. It allows you to set a budget on either a billing account or a project. You can set the budget at a specific amount or match it to the previous month's spend.

After you determine your budget amount, you can set budget alerts. These alerts send emails to billing admins after spend exceeds a percent of the budget or a specified amount. In our case, it would send an email when spending reaches 50%, 90%, and 100% of the budget amount.

Please note that these alerts are based on estimated expenses, so actual expenses may be greater.



## Labels can help you optimize network spend

```
1 SELECT
2   TO_JSON_STRING(labels) as labels,
3   sum(cost) as cost
4 FROM `project.dataset.table`
5 GROUP BY labels;
```



BigQuery

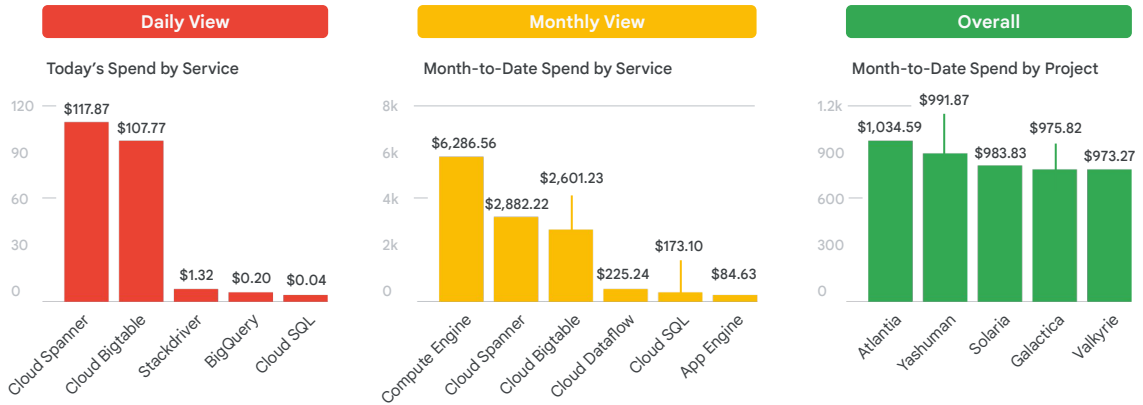
Another way to help optimize your Google Cloud network spend is to use labels. A label is a key-value pair that helps you organize your Google Cloud resources, such as instances. You can attach a label to each resource, which is forwarded to the billing system, so you can break down your billing charges by label.

For example, you could label your backends to better understand the cost of your load balancer. Maybe your backends are sending most of their traffic to a different continent, which could incur higher costs. In that case, you might consider relocating some your backends or using Cloud CDN to cache content closer to your users, which reduces your networking spend.

I recommend labeling all your resources and exporting your billing data to BigQuery to analyze your spend. BigQuery is Google's scalable, fully managed Enterprise Data Warehouse with SQL and fast response times. Creating a query is as simple as shown in this screenshot.

I will illustrate how to apply labels to resources and analyze network spend with BigQuery in the next demo.

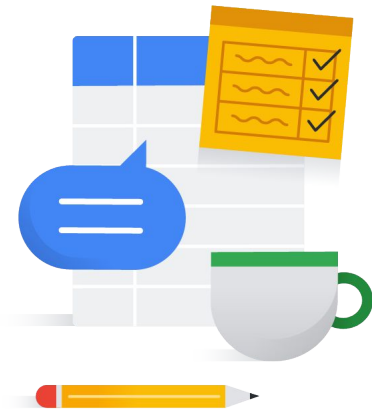
# Visualize network spend with Data Studio



You can even visualize spend over time with Data Studio. Data Studio turns your data into informative dashboards and reports that are easy to read, easy to share, and fully customizable.

For example, you can slice and dice your billing reports using your labels.

# Debrief



In this module, we covered how Google Cloud networking features are charged for, how to choose between the premium and standard tier, and how billing data can be analyzed within Google Cloud.

The intent of this module was to help you optimize your network spend.

You explored this in the lab by creating VMs with different Network Service Tiers and by getting an overview on analyzing labels in BigQuery and Data Studio.

