

# 8 Best Practices for Reducing Spend In Google Cloud Platform



POWERED BY **CloudHealth**


# INTRODUCTION

Google Cloud Platform (GCP) adoption is growing at a staggering rate. Customers are moving to GCP for their multi cloud strategy and the rich features around Machine Learning, Artificial Intelligence, Big Data and Containers along with IaaS, PaaS and SaaS offerings. Competitive pricing and increased credibility as an enterprise cloud provider are also helping drive adoption.

Customers leverage the public cloud environment with certain set of objectives in mind. Typical drivers include leveraging agility, flexibility, security and governance from cloud providers at a lower cost. No one comes to the public cloud with the intention of spending more money on wasted resources, although it is a common occurrence.

Reducing spend in GCP doesn't need to be a process of trial and error — there are proven ways to save money in your GCP environment without negatively impacting desired outcomes.

In this eBook we will talk about how to drive cost efficiencies with more accountability and financial transparency. It is easy to keep your GCP environment optimized with just a few observations and process changes.



Google Cloud Platform

# 1


## DELETE UNATTACHED PERSISTENT DISK

It's common to see thousands of dollars in unattached persistent disk being spent within GCP accounts. Generally, these are disks that are costing money but aren't being used for anything. When a Compute Engine VM is launched, a disk is usually attached to act as the local block storage for the application.

When the Compute Engine VM is terminated, it's possible that the unattached disk will be left running. GCP continues to charge for the full list price of the disk, despite the fact that the disk is no longer in use.

### *Are you spending without your knowledge?*

By continuously checking for unattached disks in your infrastructure, you can cut thousands of dollars from your monthly bill.



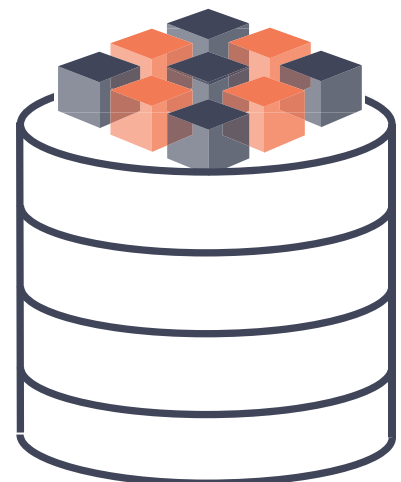
By continuously checking for unattached disks in your infrastructure, you can cut thousands of dollars from your monthly bill.

# 2

## DELETE SNAPSHOTS WHEN NO LONGER NEEDED

Persistent Disk snapshots are used to create point in time backup of disks in case of data loss. However, these costs can quickly get out of control if not closely monitored. Individual snapshots are not costly, but when hundreds are provisioned the cost adds up.

Organizations can control snapshots by monitoring snapshot cost and usage per Compute Engine VM to make sure they do not spike out of control. Set a standard in your organization for how many snapshots should be retained per Compute Engine VM. Remember that the majority of the time, a recovery will occur from the most recent snapshot.



# 3

## DELETE DISASSOCIATED NETWORK IP ADDRESSES

An external IP address is usually associated with a Compute Engine VM and allows the Compute Engine VM to be reachable via the internet. Customers can also reserve IP addresses to have them dedicated for a project, which is more costly. As a rule, ingress for networking is free and won't impact your bill. Egress, however, under certain conditions is charged which should not be overlooked.

When you reserve a static external IP address, there is no charge when you are using it; but, as soon as it is no longer being used, Google charges for it. 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, enabling you to delete those that are show a "reserved" status rather than an "in use" status.

From a best practice standpoint, monthly Network IP Address charges should be as close to zero as possible. If disassociated Network IP Addresses are within the Google Cloud accounts, they should either be reassociated to a Compute Engine VM or outright deleted in order to avoid wasted cost. Size flexibility is offered on Linux/UNIX Regional RIs and is available at no additional cost.

# zero

Monthly Network  
IP Address charges  
should be as close  
to zero as possible.




# 4

## TERMINATE ZOMBIE ASSETS

Zombie assets are infrastructure components that are running in your cloud environment but not being used for any purpose. Zombie assets come in many forms. For example, they could be Compute Engine VMs that were once used for a particular purpose, but are no longer in use and have not been turned off. Or they could have been protected through 'deletionProtection' flag for certain SQL servers, but are no longer in active use.

Zombie Compute Engine VMs also can occur when Compute Engine VMs fail during the launch process or because of errors in script that fail to deprovision Compute Engine VMs. Zombie assets can also come in the form of idle Load Balancers that aren't being used effectively, or an idle SQL Database.

No matter the cause, you will be charged for them as long as these assets are running. They must be isolated, evaluated, and immediately terminated if deemed nonessential. Take a backup of the asset before terminating or stopping it to ensure you can recover it if the asset is needed again.




# 5

## RIGHTSIZE COMPUTE ENGINE VMS

Rightsizing Compute Engine VMs is the cost reduction initiative with the potential for the biggest impact. It's common for developers to spin up new Compute Engine VMs that are substantially larger than necessary.

This may be intentional, to give themselves extra headroom, or accidental since they don't know the performance requirements of the new workload yet. If you are coming from an on-premises environment, the experience of delay in infrastructure delivery can be another reason for over-provisioning.

Over-provisioning a Compute Engine VM can lead to exponentially higher costs and inefficiencies in the environment. Without performance monitoring or cloud management tools, it's hard to tell when assets are over-or under-provisioned..



# 6

## MANAGING COMPUTE ENGINE VMs ON A SCHEDULE

Google will bill for a Compute Engine VM as long as its running. Inversely, if a Compute Engine VM is in a stopped state, there is no charge associated with it. For Compute Engine VMs that are running 24/7, Google Cloud will bill for 672 to 744 hours per Compute Engine VM per month, depending on the usage. If a Compute Engine VM is turned off between 5pm and 9am on weekdays and stopped weekends and holidays, then total billable hours per month would range from 152 to 184 hours per Compute Engine VM, saving you 488 to 592 Compute Engine VM hours per month.

This is an extreme example, as having flexible workweeks and global teams means that you can't just power down Compute Engine VMs outside normal working hours. However, outside of production, you'll likely find many Compute Engine VMs that do not need to truly run 24/7/365.

The most cost-efficient environments dynamically stop and start Compute Engine VMs based on a set schedule. Each cluster of Compute Engine VMs can be treated a different way.





# 7

## TAKE ADVANTAGE OF COMMITTED USE DISCOUNTS & SUSTAINED USE DISCOUNTS

For stable and predictable workloads, Google Cloud Platform gives you the option to purchase a specific amount of Compute and Memory for discounts. At no upfront payment and a commitment of either 1 year or 3 years customers can save up to 57% of the normal price and thus are a no-brainer for any company with sustained virtual machine usage.

These discounts can be used for standard, highmem, highcpu, custom machine types and sole-tenant node groups. When they expire Compute Engine VMs get charged at the normal price. Another thing to note is that once commitment discounts are purchased, customers cannot cancel them.

Even if you don't make the commitment, you can still benefit from discounts for prolonged usage, in a program Google calls "Sustained Use Discounts." Sustained use discounts are given by Google when you consume certain resources for better part of the billing month. They are applicable to resources like custom machines, sole-tenant nodes, GPU devices, etc. These discounts are given automatically and customers don't have to do anything additional to take advantage of them.

At no upfront payment  
and a commitment of  
1 -3 years, customers  
can save up to

**57%**

# 7

## MIGRATE OBJECT STORAGE TO LOWER COST OFFERINGS

For stable and predictable workloads, Google Cloud Platform gives you the option to purchase a specific amount of Compute and Memory for discounts. At no upfront payment and a commitment of either 1 year or 3 years customers can save up to 57% of the normal price and thus are a no-brainer for any company with sustained virtual machine usage.

These discounts can be used for standard, highmem, highcpu, custom machine types and sole-tenant node groups. When they expire Compute Engine VMs get charged at the normal price. Another thing to note is that once commitment discounts are purchased, customers cannot cancel them.

Even if you don't make the commitment, you can still benefit from discounts for prolonged usage, in a program Google calls "Sustained Use Discounts." Sustained use discounts are given by Google when you consume certain resources for better part of the billing month. They are applicable to resources like custom machines, sole-tenant nodes, GPU devices, etc. These discounts are given automatically and customers don't have to do anything additional to take advantage of them.

It is very important to keep track of this since the price differential from the top tier to the bottom is quite high

**35%**



# CONCLUSION

It's important to remember that these best practices are not meant to be one-time activities, but ongoing processes. Because of the dynamic and ever changing nature of the cloud, cost optimization activities should ideally take place continuously.

Learn more about how you can help you automate the continuous optimization by visiting [www.vastITservices.com](http://www.vastITservices.com)

POWERED BY **CloudHealth**

**VAST View™** is a trademark of VAST IT Services.

