

# Cloud cost optimization

## Part 1: Create visibility

There is a growing trend for organizations of all sizes to get out of the business of running their own data centers and move their infrastructure and application workloads to the cloud. This typically includes a transition to software as a service (SaaS) as well as leveraging infrastructure as a service (IaaS). IaaS providers are often referenced as “Hyperscalers” and include Google Cloud Platform (GCP), Amazon AWS, and Microsoft Azure. The advantages of and motivations for adopting cloud applications and infrastructure are often to achieve cost savings through a [shift from a CapEx to an OpEx](#) financial model and the opportunity to only pay for the resources and infrastructure you need or use.

### The problem

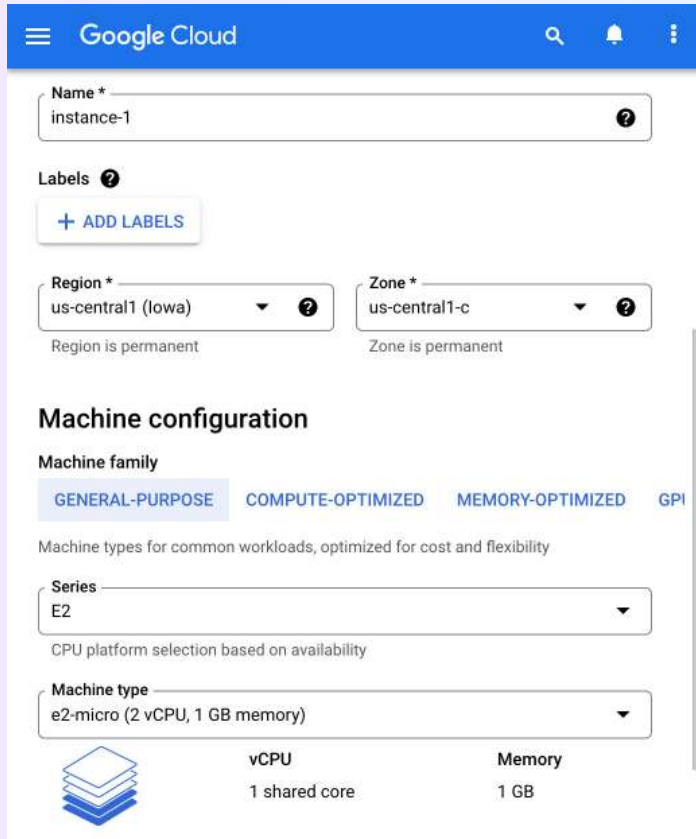
However, most organizations are finding that there is actually a higher cost than expected for running their legacy workloads in a hyperscale environment. In fact, [Gartner estimates that as much as 70% of cloud costs are wasted](#). As a result, there has been a need to adopt best practices for optimizing cloud costs (examples: [What is cloud cost optimization? 15+ best practices for 2022](#), [Cloud Cost optimization: definition and strategies](#), [9 Best practices for cloud cost optimization](#)). While these articles are helpful if you need to drive alignment around an initiative, what they miss are practical steps you can take today to start optimizing your cloud costs.

Whether you are starting your journey to the cloud, or just looking to better optimize your current cloud costs, it helps to break the problem down into two different buckets.

- 1. Infrastructure configuration, optimization, and automation** – Create awareness and visibility into cloud infrastructure usage using infrastructure as code (covered here in part 1). Then review how your infrastructure can be automated and optimized to find cost savings (covered in part 2 – coming soon).
- 2. Application refactoring** – Review application architecture and resource utilization. Often applications are monolithic and not designed to be cloud native. Hyperscalers provide compute, memory, and storage, as well as new standardized services, technologies, and approaches that were not available in the traditional data center environments. This requires insights into the portions of your code that can leverage the newly available technologies and the ability to rapidly update patterns in code across the organization. At the very least, organizations need to identify and optimize their workloads to run more efficiently in the new hyperscale environment (covered in parts 3 and 4 – coming soon).

## Infrastructure as code

Infrastructure as code (IaC) is the process of managing and provisioning cloud infrastructure through machine-readable definition files like Terraform ([open source repositories with examples, example files](#)) and AWS CloudFormation ([open source repositories with examples, example files](#)), rather than physical hardware configuration or interactive configuration tools<sup>1</sup>. These definition files make it possible to have visibility of all your provisioned infrastructure, especially if you have a universal code search solution inside your organization.



The screenshot shows the Google Cloud console interface for creating a new Compute Instance. The 'Name' field is set to 'instance-1'. The 'Region' is set to 'us-central1 (Iowa)' and the 'Zone' is set to 'us-central1-c'. Under 'Machine configuration', the 'Machine family' is set to 'GENERAL-PURPOSE'. The 'Series' is set to 'E2' and the 'Machine type' is set to 'e2-micro (2 vCPU, 1 GB memory)'. A table below the machine type selection shows the specifications: 2 vCPU (1 shared core) and 1 GB Memory.

	vCPU	Memory
	2 shared core	1 GB

```
provider "google" {
  project = "{{YOUR GCP PROJECT}}"
  region  = "us-central1"
  zone    = "us-central1-c"
}

resource "google_compute_instance" "vm_instance" {
  name         = "terraform-instance"
  machine_type = "e2-micro"

  boot_disk {
    initialize_params {
      image = "debian-cloud/debian-9"
    }
  }
}

network_interface {
  # A default network is created for all GCP projects
  network = google_compute_network.vpc_network.self_link
  access_config {
  }
}

resource "google_compute_network" "vpc_network" {
  name                = "terraform-network"
  auto_create_subnetworks = "true"
}
```

According to [Gartner research](#), less than 5% of server provisioning utilized IaC in 2020, and only 40% is expected to do so by 2023. This means that the vast majority of cloud infrastructure is manually provisioned, built on a huge amount of untraceable scripts, or manually configured in the cloud provider interface.

If your organization has not adopted IaC (which is the majority – you're not alone!), then the first step is to create visibility into your cloud infrastructure configuration by exporting it into a version control system so that you can search over the current state and see the history of changes.

# Create visibility

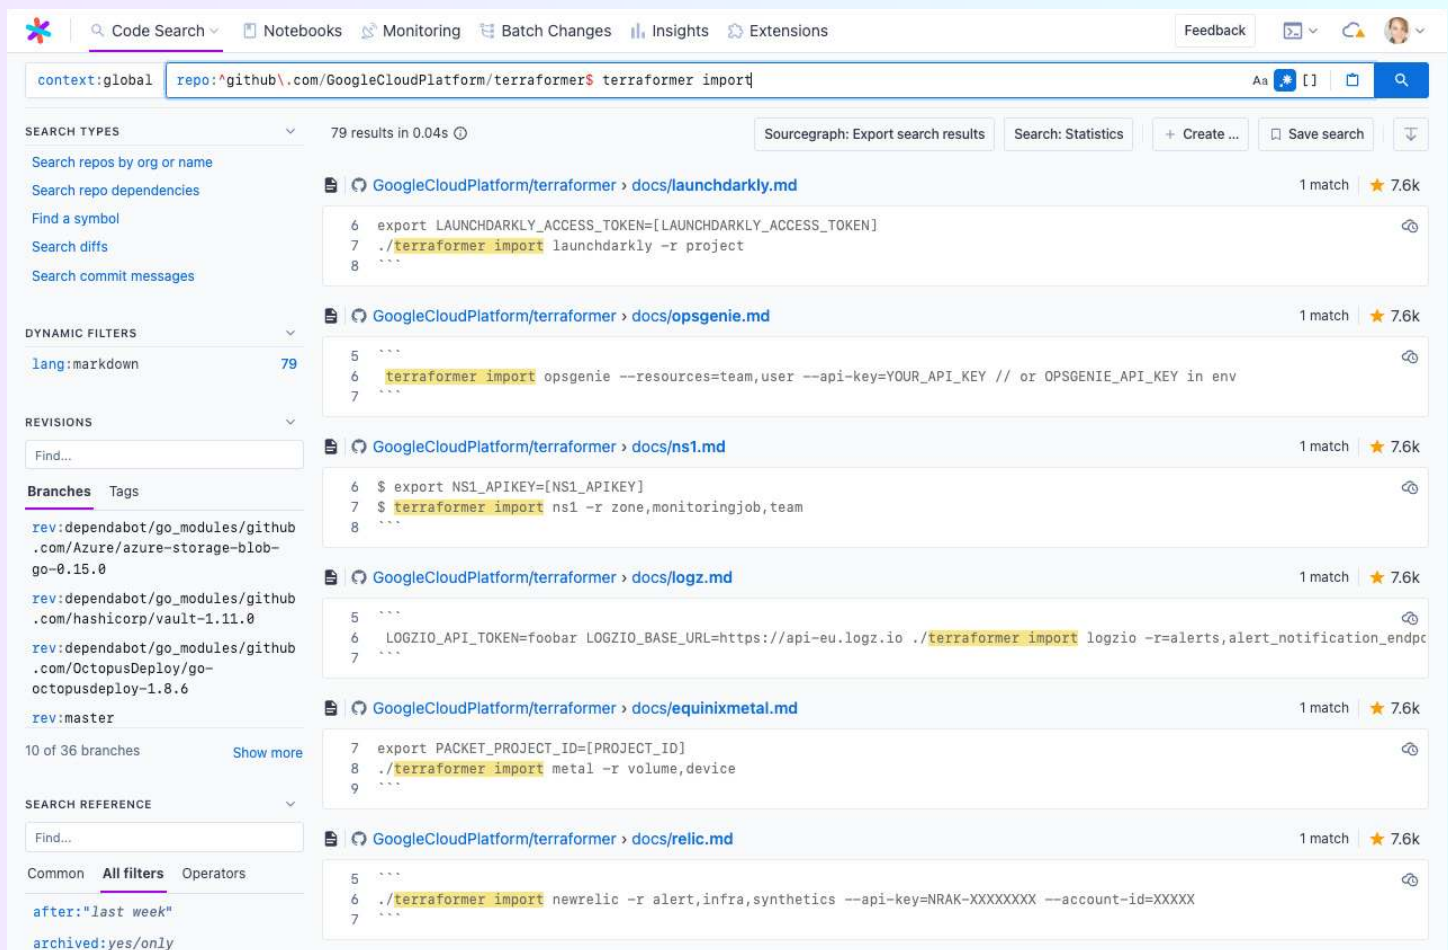
Doing a one-time audit of your cloud infrastructure is beneficial for cutting costs, but if you want to continually review for cost reduction, you'll want to have automation in place to regularly export your config into a Git repository. We've pulled together some helpful resources that should make this setup a bit easier.

1. Export all cluster config into a Git repository – use a tool like CloudFormation or Terraform (provisioning tools).

a. First, choose the provisioning tool you want to export to. We recommend Terraform. Here is a helpful blog series on choosing Terraform over other provisioning tools.

b. Terraformer (<https://github.com/GoogleCloudPlatform/terraformer>) is a CLI to export config to Terraform on any cloud provider.

## Example export commands for all cloud providers from Terraformer docs



The screenshot shows a GitHub search interface with the following details:

- Search Query:** `repo:^github\.com/GoogleCloudPlatform/terraformer$ terraformer import`
- Results:** 79 results in 0.04s
- Filters:** `lang:markdown` (79 results)
- Search Reference:** `after:"last week"`, `archived:yes/only`

The search results list several documents from the `GoogleCloudPlatform/terraformer` repository, each with a snippet of a `terraformer import` command:

- docs/launchdarkly.md** (1 match, 7.6k stars):

```
6 export LAUNCHDARKLY_ACCESS_TOKEN=[LAUNCHDARKLY_ACCESS_TOKEN]
7 ./terraformer import launchdarkly -r project
8 ...
```
- docs/opsgenie.md** (1 match, 7.6k stars):

```
5 ...
6 terraformer import opsgenie --resources=team,user --api-key=YOUR_API_KEY // or OPSGENIE_API_KEY in env
7 ...
```
- docs/ns1.md** (1 match, 7.6k stars):

```
6 $ export NS1_APIKEY=[NS1_APIKEY]
7 $ terraformer import ns1 -r zone,monitoringjob,team
8 ...
```
- docs/logz.md** (1 match, 7.6k stars):

```
5 ...
6 LOGZIO_API_TOKEN=foobar LOGZIO_BASE_URL=https://api-eu.logz.io ./terraformer import logzio -r=alerts,alert_notification_endpc
7 ...
```
- docs/equinixmetal.md** (1 match, 7.6k stars):

```
7 export PACKET_PROJECT_ID=[PROJECT_ID]
8 ./terraformer import metal -r volume,device
9 ...
```
- docs/relic.md** (1 match, 7.6k stars):

```
5 ...
6 ./terraformer import newrelic -r alert,infra,synthetics --api-key=NRAK-XXXXXXXX --account-id=XXXXX
7 ...
```

## Examples of scripts that reference `terraform import`

The screenshot shows a code search interface with the following components:

- Search Bar:** context:global "terraform import" -lang:Markdown
- Search Types:** 14 results in 1.65s. Some results excluded. Search: Statistics. Sourcegraph: Export search results. + Create ... Save search.
- DYNAMIC FILTERS:** fork:yes 797900, archived:yes 152116, lang:go 5, lang:shell 5, lang:ruby 2, lang:python 1.
- REPOSITORIES:** Find... r: Checkmarx/Kics 4, r: Houssemdellai/terraform-course 3, r: ned1313/terraform-tuesdays 2, r: CiscoDevNet/terraform-provider-aci 1, r: Homebrew/homebrew-core 1.
- SEARCH REFERENCE:** Find... Common All filters Operators. after:"last week", archived:yes/only, author:name, before:"last thursday", case:yes, content:"pattern", count:N/all, file:regex-pattern.

The search results list several repositories with code snippets:

- Homebrew/homebrew-core > Formula/terraform.rb** (1 match, 11.3k stars):

```
31     assert_match "aaa",  
32     shell_output("#{bin}/terraform import google --resources=gcs --projects=aaa 2>&1", 1)  
33 end
```
- Checkmarx/kics > pkg/terraform/azure/azure\_cloud\_provider.go** (1 match, 1.1k stars):

```
48     case <-ctxT.Done():  
49         return errors.New("terraform import execution timeout")  
50     }
```
- Checkmarx/kics > pkg/terraform/aws/aws\_cloud\_provider.go** (1 match, 1.1k stars):

```
50     case <-ctxT.Done():  
51         return errors.New("terraform import execution timeout")  
52     }
```
- Checkmarx/kics > pkg/terraform/gcp/gcp\_cloud\_provider.go** (1 match, 1.1k stars):

```
52     case <-ctxT.Done():  
53         return errors.New("terraform import execution timeout")  
54     }
```
- Checkmarx/kics > pkg/terraform/terraform\_alt.go** (1 match, 1.1k stars):

```
6 /*  
7     Since terraform import is very big, this creates problems for the dev env  
8     regarding CPU and MEM usage.
```
- ned1313/terraform-tuesdays > 2021-03-01-Terraformer/commands.sh** (2 matches, 209 stars):

```
50  
51 terraform import plan generated/azurerem/terraform/plan.json  
52  
58 terraform plan azure --output hcl --resource-group taconet --resources="*" --compact --path-pattern "{output}/{provider}/"  
59 terraform import plan generated/azurerem/plan.json  
60
```
- Houssemdellai/terraform-course > 92\_import\_terraform/commands.sh** (3 matches, 138 stars):

```
9  
10 terraform import azure -r resource_group  
11  
12 terraform import azure -R my_resource_group -r virtual_network,resource_group  
13  
14 terraform import azure -r resource_group --filter=resource_group=/subscriptions/<Subscription id>/resourceGroups/<RGNAME>  
15
```

2. Export all services config into a Git repository – use a tool like Chef, Puppet, Ansible, or Kubernetes (configuration management tools) to export the services configuration.

a. We deploy Sourcegraph.com with Kubernetes, and use kube-backup (<https://github.com/pieterlange/kube-backup>) to export and backup this configuration.

3. Set up a CronJob to regularly snapshot the configuration files from the first two steps into the Git repos. By regularly updating, you will be able to have better traceability of changes over time, and it will allow you to better understand the entire system. Additionally, you have now moved one small step closer to infrastructure as code.

