Engineer

Google Cloud

Google's cloud platform

What it is

Google Cloud is the part of Google that rents out computing from its own data centres: virtual machines on Compute Engine, object storage in Cloud Storage, managed databases in Cloud SQL, Kubernetes clusters in GKE, the BigQuery data warehouse and many more services, all paid for by use. Its older name, Google Cloud Platform or GCP, is still in wide use, even in parts of Google's own documentation.

Everything Google Cloud runs for a customer lives inside a project. A project is the container for resources, for the APIs switched on for them, for the quotas they draw on and for the link to the billing account that pays. It has an editable display name and a project ID that is unique across all of Google Cloud and permanent once the project exists. Projects can be grouped into folders, folders can hold other folders, and the whole tree hangs from an organisation resource that belongs to a company's Google Workspace or Cloud Identity account. Roles and organisation policies set at the organisation or on a folder are inherited by every project beneath it, which makes a well-drawn tree easy to govern and a badly drawn one hard to untangle.

The network is drawn differently from the other large clouds. A VPC network in Google Cloud, together with its routes and firewall rules, is a global resource that belongs to no region; only its subnets are regional. A virtual machine in Belgium and another in Iowa can therefore share one private network and reach each other on internal addresses, with no VPN or peering between the two regions. On AWS, by contrast, each VPC lives in a single region.

Identity comes in two kinds. People sign in with Google accounts and receive IAM roles on a project, a folder or the whole organisation. Programs use service accounts: identities that belong to one project and that a virtual machine, a

container or a build job can act as. Attach a service account to a Compute Engine instance, and the code running there can call Google APIs with that account's roles without a password or key file on the disk.

Use the concept in a decision

Put "What is Google Cloud?" to work in a real decision. We can help relate the concept to the system you are building or buying. Servers sized for what you actually use, moved into the cloud or back out of it, with the monthly bill worked out first.

Explore our services Get in touch →

Test the assumptions behind the term

Use the explanation and linked work to frame questions for your project. We can examine the assumptions with you before a tool is chosen. Ran 20 Docker apps on Google Cloud and AWS — 20 containerised applications deployed and watched

Ran 20 Docker apps on Google Cloud and AWS Get in touch →

Why it matters

For many teams the reason to pick Google Cloud is data. BigQuery is a warehouse with no servers to size: tables sit in datasets inside a project, queries are written in standard SQL, and on the on-demand model each query is billed by the bytes it reads, while capacity pricing pays for a reserved amount of processing, measured in slots, instead. Storage is columnar, so a query pays for the columns it touches rather than for whole rows. Cloud Storage sits beside it as the landing zone, and loading files from a bucket is a standard way in. A team that would elsewhere provision and tune a warehouse cluster can start with a bucket, a dataset and a query.

The project boundary is the second reason, because so much attaches to it. Billing, quotas, enabled APIs and IAM all meet at the project, so one project per application and environment draws a clean line: production and test get separate permissions and separate lines on the bill. Shutting a project down stops its billing and every resource in it at once. Its owners can restore it for 30 days, though some services, Cloud Storage among them, may delete their data sooner. Above the projects, an organisation policy set once on a folder can forbid whole classes of risky setting, such as default networks or downloadable service account keys, in every project beneath it. That is how a small team keeps a dozen projects consistent without checking each one by hand.

What it looks like in practice

Three portfolio entries show Google Cloud as the landing zone for a data warehouse, the offsite store for backups and one end of a private network.

In the data work, six ETL/ELT pipelines fed BigQuery from a warehouse behind an IPsec VPN. Curated datasets and tables were exported, compressed with gzip and transferred to a Google Cloud Storage bucket, and the upload and schema mapping into BigQuery were automated with the bq CLI and Python ingestion scripts. Once the Cloud Storage hop and the BigQuery load ran on their own, the whole job took under 20 minutes instead of 3+ hours by hand, BigQuery was refreshed daily rather than weekly, and the Python APIs reading from it gained 30% in performance.

For backups, a hundred critical data backups were automated, with Barman handling the PostgreSQL side and Google Cloud holding the offsite copies. Bash and Python orchestrated and validated the whole regime, with retention policies to keep the copies current and integrity checks to confirm that each one was good rather than merely present. The backups then ran automatically and could be verified across every database.

On the network side, eight WireGuard and IPsec tunnels between Google Cloud and on-premises Linux systems were managed and troubleshot as one set. Their configuration was standardised, their health was monitored, and the recurring routing and key-exchange problems were fixed at the root rather than with a restart, after which all eight tunnels ran securely and reliably.

Where it goes wrong

Google Cloud’s sharpest edges are its conveniences: settings that make the first afternoon easy and are still in place a year later, long after anyone remembers choosing them.

The default network stays open. Unless an organisation policy prevents it, every new project starts with a default network that has a subnet in every region and four pre-populated firewall rules. Three of them accept traffic from 0.0.0.0/0, which is the whole internet: SSH on tcp:22, Remote Desktop on tcp:3389, and ICMP. The fourth allows all TCP, UDP and ICMP traffic from 10.128.0.0/9 inside the network. Any virtual machine given an external address on that network is therefore open to SSH attempts from anywhere. The `compute.skipDefaultNetworkCreation` constraint keeps new projects from getting one, and a custom network with rules that name their sources replaces it.

A service account key ends up in a repository. A user-managed key is a JSON file that lets whoever holds it act as the service account, and by default it stays valid until someone deletes it. Such files spread easily: into CI secrets, into container images and, by mistake, into repositories. Google's guidance is to choose a safer alternative wherever one exists. Workloads on Google Cloud use the service account attached to the resource they run on, and workloads elsewhere, such as a CI system or another cloud, use Workload Identity Federation, which trades the outside system's own credentials for short-lived Google tokens. Organisations created on or after 3 May 2024 have key creation and key upload disabled by default through organisation policy constraints, and older ones can set the same constraints.

A small query reads a large table. On the on-demand model BigQuery bills for the bytes a query reads, and on a table that is not clustered a `LIMIT` clause leaves that number unchanged: `SELECT * ... LIMIT 10` returns ten rows and bills every byte of every column. Select only the columns needed, partition large tables by date so that queries filtering on it skip old partitions, check the estimate from a dry run before running anything new, and set maximum bytes billed. A query whose estimate exceeds that limit fails before it runs, without incurring a charge.

Cold storage holds warm data. Cloud Storage offers Nearline, Coldline and Archive classes that cost less per gigabyte stored than Standard, in exchange for retrieval fees and minimum storage durations of 30, 90 and 365 days. An object deleted, replaced or moved before its minimum is charged as if it had been stored for the whole period. Backups that rotate daily, or scratch files a pipeline writes and removes within the hour, can therefore cost more in Archive than they would in Standard. The colder classes suit data that is written once and read rarely, and a lifecycle rule can move objects into them as they age.

Further reading

- Google Cloud Documentation — Google's index of guides, references and tutorials for every Google Cloud product.
- About resource hierarchy — how organisations, folders and projects nest, and how roles and policies flow down them.
- VPC firewall rules — how firewall rules are evaluated, including the rules the default network starts with.
- Best practices for managing service account keys — how keys leak, how to restrict them and what to use instead.

- Estimate and control costs — BigQuery's guidance on dry runs, maximum bytes billed and why `LIMIT` saves little.
- Storage classes — the classes side by side, with their minimum durations, retrieval fees and availability.

Glossary

Choose for your actual constraints

Bring the deployment path, the operational risks and the people on call. We can scope a change with checks, a recovery route and clear ownership.

See what it costs Get in touch →

Make the terminology practical

Ask how "What is Google Cloud?" applies to your organisation. A concrete example from your current system is enough to start the discussion.

Explore our services

welcome@engineer.company +45 50 10 61 84 Copenhagen, Denmark

Take this page with you

Contact card

Carefully Crafted With Love by D.

CVR 45969703 · Tagensvej 105C, 3. th, 2200 København N

Copyright © 2025 – present · Engineer ApS · All rights reserved.

Printed from https://engineer.company/glossary/google-cloud/