

Introduction to Google Cloud Platform



Agenda

- What is Google Cloud Platform (GCP)
- GCP services & benefits
- Why GCP ?
- What is Google Compute Engine (GCE)
- How to create a VM using GCE
- Overview of Google Cloud Shell & gcloud CLI tool
- Demo, Discussion, Q&A session

Cloud Players

vmware®


rackspace 


DigitalOcean

 terremark®

 **amazon**
web services



 Joyent®



Google Cloud Platform

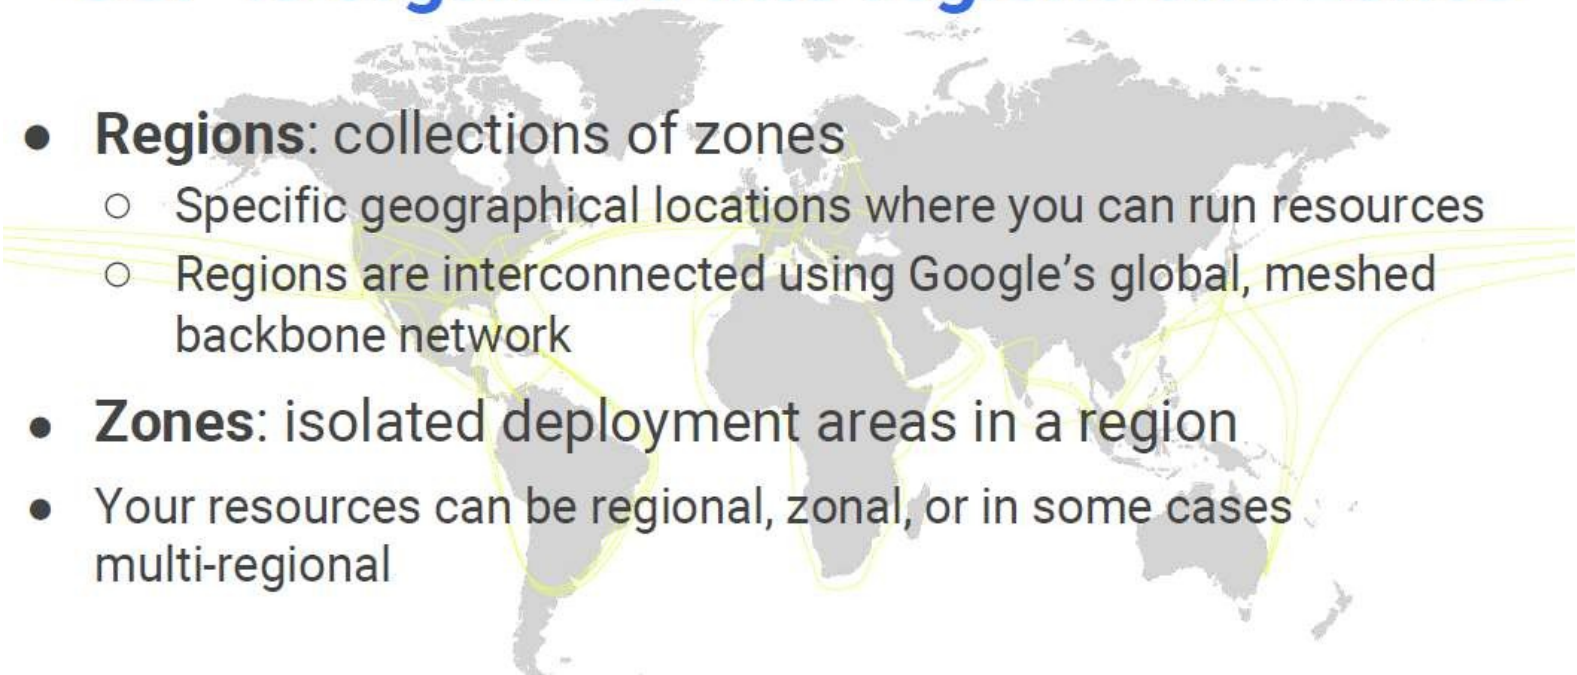
 Microsoft Azure

What is GCP ?



- ★ **Google Cloud Platform** enables developers to build, test and deploy applications on Google's highly-scalable and reliable infrastructure.
- ★ **Google Cloud Platform** is a set of modular cloud-based services that allow you to create anything from simple websites to complex applications.

GCP is organized into regions and zones

- **Regions:** collections of zones
 - Specific geographical locations where you can run resources
 - Regions are interconnected using Google's global, meshed backbone network
 - **Zones:** isolated deployment areas in a region
 - Your resources can be regional, zonal, or in some cases multi-regional
- 



Google Network

More than a collection of data centers



Why Google Cloud Platform?

For the past 15 years, Google has been building out the fastest, most powerful, highest quality cloud infrastructure on the planet.

Why Google Cloud Platform?

“[Google's] ability to build, organize, and operate a huge network of servers and fiber-optic cables with an efficiency and speed that rocks physics on its heels. This is what makes Google Google: its physical network, its thousands of fiber miles, and those many thousands of servers that, in aggregate, add up to the mother of all clouds.”

- Wired



Why Google Cloud Platform?

When you build on Google Cloud Platform, you're building on Google's fast, scalable, and highly reliable infrastructure. You can innovate and iterate faster, since you don't have to worry about the underlying operational aspects of your business.





Why GCP ?

#1 Run on Google's Infrastructure

Build on the same infrastructure that allows Google to return billions of search results in milliseconds, serve **1 billion hours** of YouTube video per **day** and provide storage for **1.4 billion** Gmail users.

- Global Network
- Redundancy
- Innovative Infrastructure

#2 Mix and Match Services

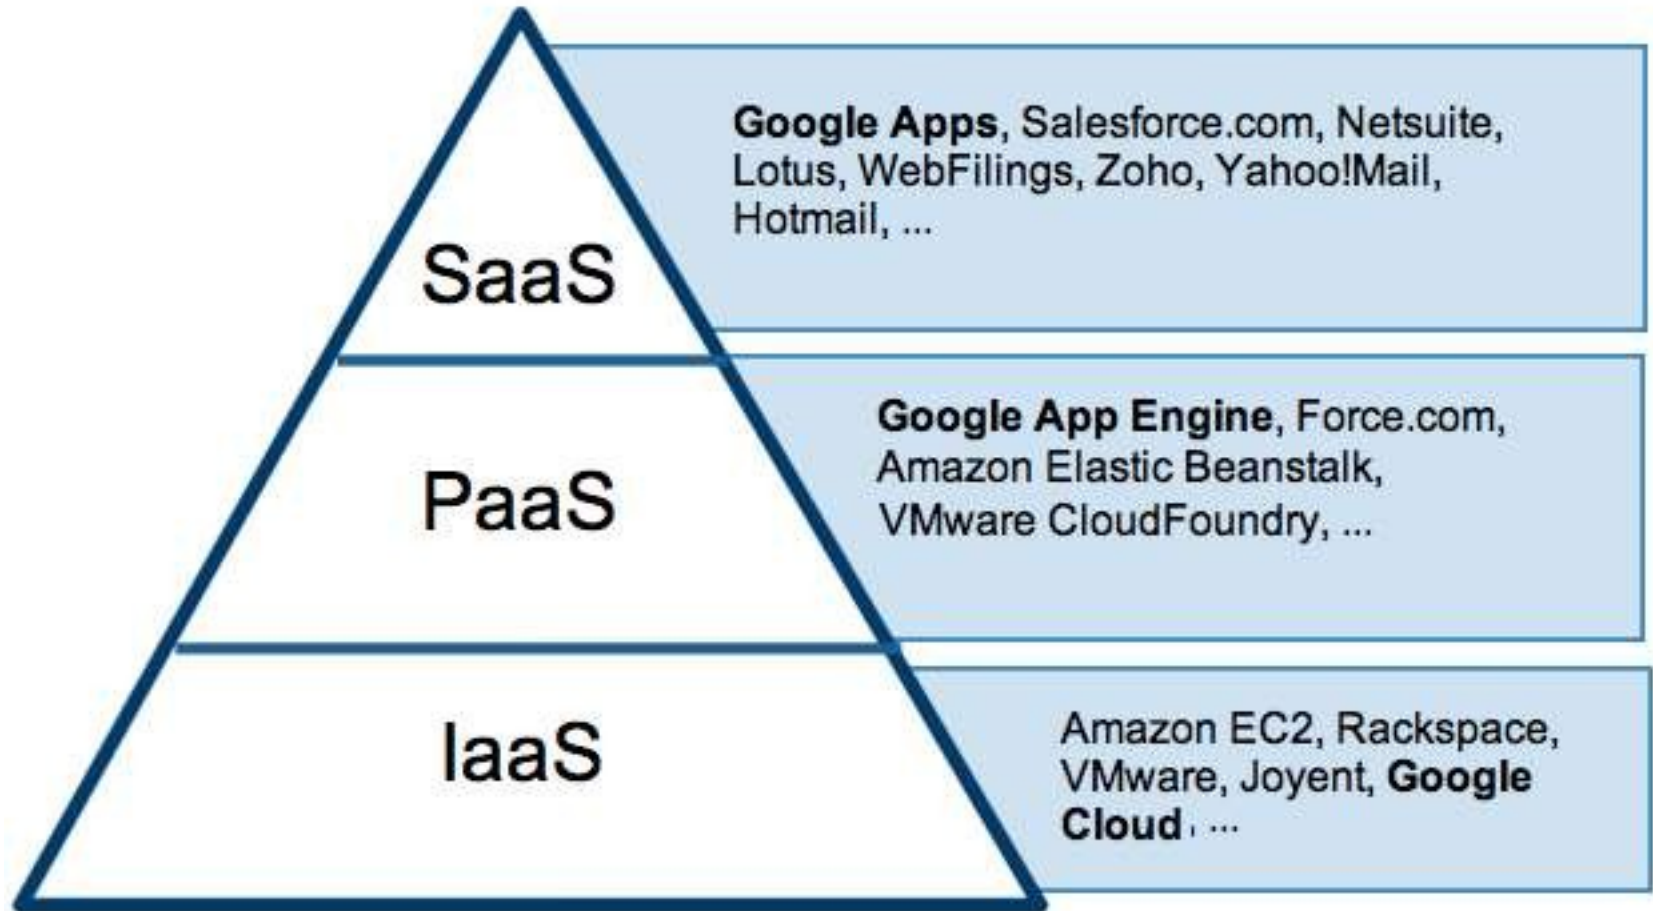
- ★ Virtual machines. Managed platform. Blob storage. Block storage. NoSQL datastore. MySQL database. Big Data analytics.
- ★ Google Cloud Platform has all the services your application architecture needs.
 - Compute
 - Storage
 - Services

#3 Performance you can count on

Google's compute infrastructure gives you consistent CPU, memory and disk performance. The network and edge cache serve responses rapidly to your users across the world.


- CPU, Memory and Disk
- Global Network
- Transparent maintenance

Cloud Computing service levels



Google Cloud Platform

Products and services of GCP can be broadly categorized as **Compute**, **Storage**, **Big-data** and **Machine Learning**



Google Cloud Platform

Compute



Compute Engine



Container Engine



App Engine



Cloud Functions

Storage



Bigtable



Cloud Storage



Cloud SQL



Cloud Spanner



Cloud Datastore

Big Data



BigQuery



Pub/Sub



Dataflow



Dataproc



Datalab

Machine Learning



Natural Language API



Vision API



Machine Learning



Speech API



Translate API



Compute Engine: Virtual machines hosted on Google's infrastructure - Infrastructure-as-a-Service



App Engine: Deploy your code directly to a fully-managed platform - Platform-as-a-Service



Container Engine: Run Docker container cluster on Google Cloud Platform – Container-as-a-Service

Storage



Bigtable



Cloud Storage



Cloud SQL



Cloud Datastore



Cloud SQL: Full SQL support for an online transaction processing (OLTP) system



Cloud Datastore: Store highly structured objects and query with SQL-like statements



Cloud Storage: Store immutable blobs larger than 10 MB, such as large images or videos



Cloud BigTable: High-performance, extremely scalable NoSQL database, scales to billions of entries

Big-Data Services



BigQuery

Analytics data warehouse

Stream data at 100,000 rows per second



Dataflow

Stream and Batch processing of data

Unified programming model



Pub/Sub

Scalable & Reliable enterprise messaging
middleware



Dataproc

Managed Hadoop, Spark, Pig and Hive at
affordable pricing

Machine Learning



Vision
API

Recognizes objects in
images



Machine
Learning

Build you own Machine
Learning models



Speech
API

Recognizes speech

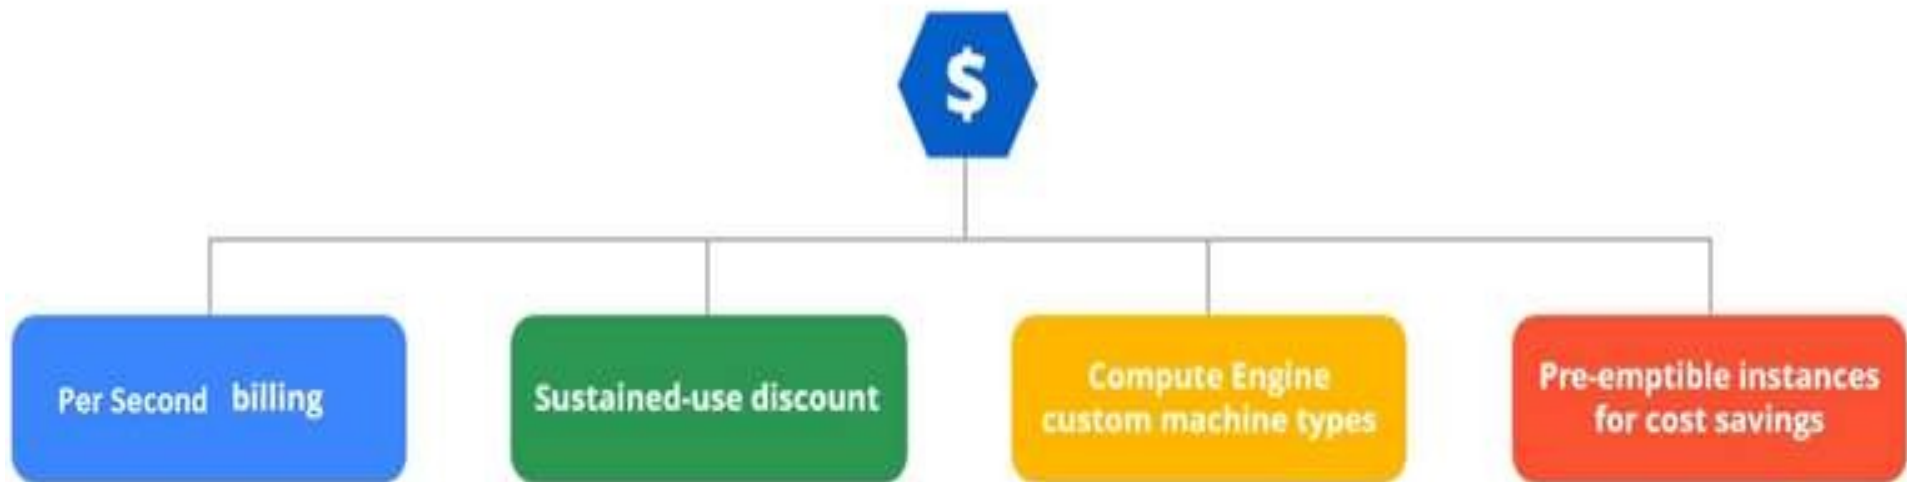


Translate
API

Translates different
languages

Pricing

Google Offers innovative and customer-friendly Pricing



Interacting with Google Cloud Platform



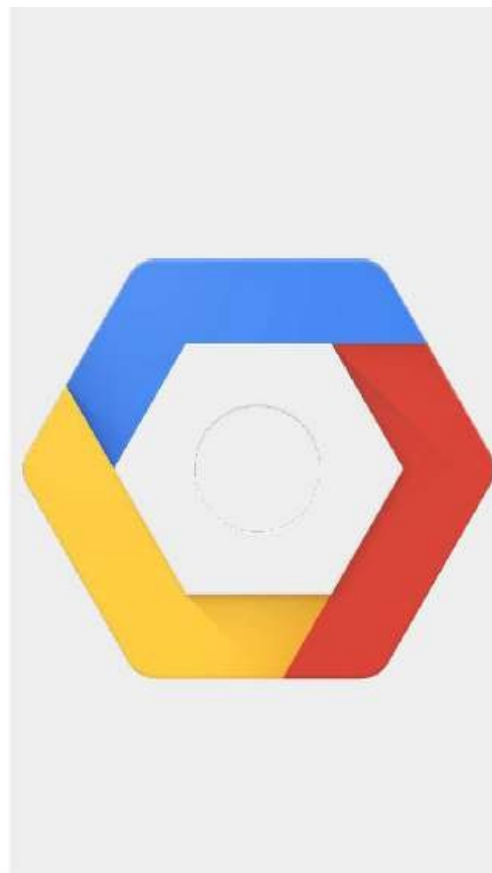
Projects organize resources

- All Google Cloud Platform services are associated with a project that is used to:
 - Track resource and quota usage
 - Enable billing
 - Manage permissions and credentials
 - Enable services and APIs



Google Cloud Platform Console

- Centralized console for all project data
- Developer tools
 - Cloud Source Repositories
 - Cloud Shell
 - Test Lab (mobile app testing)
- Access to product APIs
- Manage, create projects




Google Cloud SDK

- SDK includes CLI tools for Cloud Platform products and services
 - gcloud, gsutil (Cloud Storage), bq (BigQuery)
- Available as Docker image
- Available via Cloud Shell
 - Containerized version of Cloud SDK running on Compute Engine instance



Identity and Access Management



Who



can do what



on which resource

Google Compute Engine (GCE)

- High Performance Virtual Machines
- Power by Google's Global Network
- Pay for what you use – Really
- Fast and easy provisioning
- Compliance and Security
- Click to Deploy

Images

- Used to create boot disks for VM instances
- **Public images:**
 - provided and maintained by Google, open source communities, third party vendors
 - all projects have access and can use them
- **Custom images:**
 - Available only to your project
 - Create a custom image from boot disks and other images

Projects & Instances

- Each instance belongs to a project
- Projects can have any number of instances
- Projects can have upto 5 VPC (Virtual Private Networks)

Machine Types

Pre-defined

Custom

Machine Types

- Standard
- High-memory
- High-CPU
- Shared-core (small, non-resource intensive)
- Can attach GPU dies to most machine types




High Memory Machines

- More memory per vCPU as compared with regular machines
- Useful for tasks which require more memory as compared to processing
- 6.5 GB of RAM per core



High CPU Machines

- 
- More memory per vCPU as compared with regular machines

Custom Machines

- If none of the predefined machine types fit your workloads, use a custom machine type
- Save the cost of running on a machine which is more powerful than what you need
- Billed according to the **number of vCPUs** and the **amount of memory** used

Billing Model

- All machines types are charged for a minimum of 1 minute
- After 1 minute instances are charged in 1 second increments

Preemptible Instances

- Much much cheaper than regular Compute Engine instances
- But, might be terminated (preempted) at any time if Compute Engine needs the resources
- Use for fault-tolerant applications

Storage Options

- Each instance comes with a small root persistent disk containing the OS
- Add additional storage options
 - Persistent disks
 - Standard
 - SSD
 - Local SSDs
 - Cloud Storage

Persistent Disks

- Durable network storage devices that instances can access like physical disks in a desktop or a server
- Two types - Standard and SSD
- Standard Persistent - regular hard disks - cheap - OK for sequential access
- SSD Persistent - expensive - fast for random access



Local SSD

- Physically attached to the server that hosts your virtual machine instance
- Local SSDs have higher throughput and lower latency
- **The data that you store on a local SSD persists only until you stop or delete the instance**
- Very high IOPS and low latency



Cloud Storage Buckets

- use when latency and throughput are not a priority
- and
- when you must share data easily between multiple instances or zones.

Cloud Storage

- Create buckets to store data
- Buckets are globally unique
 - Name (globally unique)
 - Location
 - Storage Class

Cloud Storage Buckets

- Flexible, scalable, durable
- ~Infinite size possible
- Performance depends on storage class
 - Multi-regional
 - Regional
 - Nearline
 - Coldline

Bucket Storage Classes

- Multi-regional - frequent access from anywhere in the world
- Regional - frequent access from specific region
- Nearline - accessed once a month at max
- Coldline - accessed once a year at max

SSD or HDD Disks

- Use SSD unless skimping on cost
- SSD can be 20x faster on individual row reads
- More predictable throughput too (no disk seek variance)
- Don't even think about HDD unless storing > 10 TB and all batch queries
- The more random access, the stronger the case for SSD

Storage Options

	Standard persistent disks	SSD persistent disks	Local SSDs	Cloud Storage buckets
Storage type	Efficient and reliable block storage	Fast and reliable block storage	High-performance local block storage	Affordable object storage
Price per GB/month	\$0.040 - \$0.052	\$0.170 - \$0.221	\$0.218 - \$0.283	\$0.007 - \$0.026
Maximum space per instance	64 TB	64 TB	3 TB	Almost infinite

Compute Features

Feature	Amazon EC2	Compute Engine
Virtual machines	Instances	Instances
Machine images	Amazon Machine Image	Image
Temporary virtual machines	Spot instances	Preemptible VMs
Firewall	Security groups	Compute Engine firewall rules
Automatic instance scaling	Auto Scaling	Compute Engine autoscaler
Local attached disk	Ephemeral disk	Local SSD

Storage Options & Big-data

Block storage for compute VMs - persistent disks or SSDs

Immutable blobs like video/images - Cloud Storage

OLTP - Cloud SQL or Cloud Spanner

NoSQL Documents like HTML/XML - Datastore

NoSQL Key-values - BigTable (~HBase)

Getting data into Cloud Storage - Transfer service

Use cases

When you need

Storage for Compute, Block Storage

Storing media, Blob Storage

SQL Interface atop file data

Document database, NoSQL

Fast scanning, NoSQL

Transaction Processing (OLTP)

Analytics/Data Warehouse (OLAP)

Use

Persistent (hard disks), SSD

Cloud Storage

BigQuery

DataStore

BigTable

Cloud SQL, Cloud Spanner

BigQuery

Useful Links

- <https://cloud.google.com/free/>
- <https://cloud.google.com/products/>
- <https://cloud.google.com/docs/tutorials/>
- <https://cloud.google.com/getting-started/>
- <https://cloud.google.com/products/compute-engine/>
- <https://cloud.google.com/compute/docs/machine-types/>
- <https://cloud.google.com/security/>
- <https://status.cloud.google.com/>
- <https://google.qwiklabs.com/>