

Google Cloud

Prof. Dr. Lenuța Alboai
lalboai@gmail.com



Contents

- Google in Cloud
 - ...services
 - Tools (GWT, GAS,...)
 - Features
 - Architectural aspects
 - Conclusions



Google in Cloud | ...pasi

- ? Google manages the largest server farms in the world

Initial

- All the content offered was available; it relied only on the support provided by advertising


(Services: Google Maps, Google Finance, Google Voice)

- Mantra: *“it’s free to the consumer”* 😊

- 15GB of storage per account
- Google Apps – free hosting for the e-mail server (with your own domain name), Google Talk, Google Calendar, Google Drive, Google Sites <- Rajen Sheth
- Google Apps for Education
- Following Apple's model: Apps Market-
www.google.com/enterprise/marketplace
- For developers: free versions of services, e.g. Google App Engine™
<http://code.google.com/appengine>



Google in Cloud



1 Billion+


Users

Google in Cloud | ...pasi

– Google App for Business

Google Cloud Platform Customers

Discover why the world's most innovative organizations are choosing Google Cloud Platform



Get your business
on Google for free

WHO'S USING CLOUD PLATFORM?

Over 4 million applications are built on Cloud Platform

GCP FOR STARTUPS PACKAGES

Google in Cloud | ...pasi

– Google App for Business/ **G Suite**

=>
many users
Business
(prices in 2016)

Google Apps

\$5

per user per month

or \$50 per user per year plus tax

Get started

Google Apps for Work includes:

- Business email addresses (name@yourcompany.com)
- Video and voice calls
- Integrated online calendars
- 30GB of online storage for file syncing and sharing
- Online text documents, spreadsheets and slides
- Easy to create project sites
- Security and admin controls
- 24/7 phone and email support

Google Apps
with unlimited storage and Vault

\$10

per user per month

or \$120 per user per year plus tax

Get started

Everything in Google Apps for Work plus:

- Unlimited Storage (or 1TB per user if fewer than 5 users)
- Advanced admin controls for Drive
- Audit and reporting insights for Drive content and sharing
- Google Vault for eDiscovery covering emails, chats, docs and files
- Easily search and export to different formats
- Archive all emails sent by your company
- Set message retention policies
- Place and enforce litigation holds on inboxes

Get your business
on Google for free

Google | SaaS


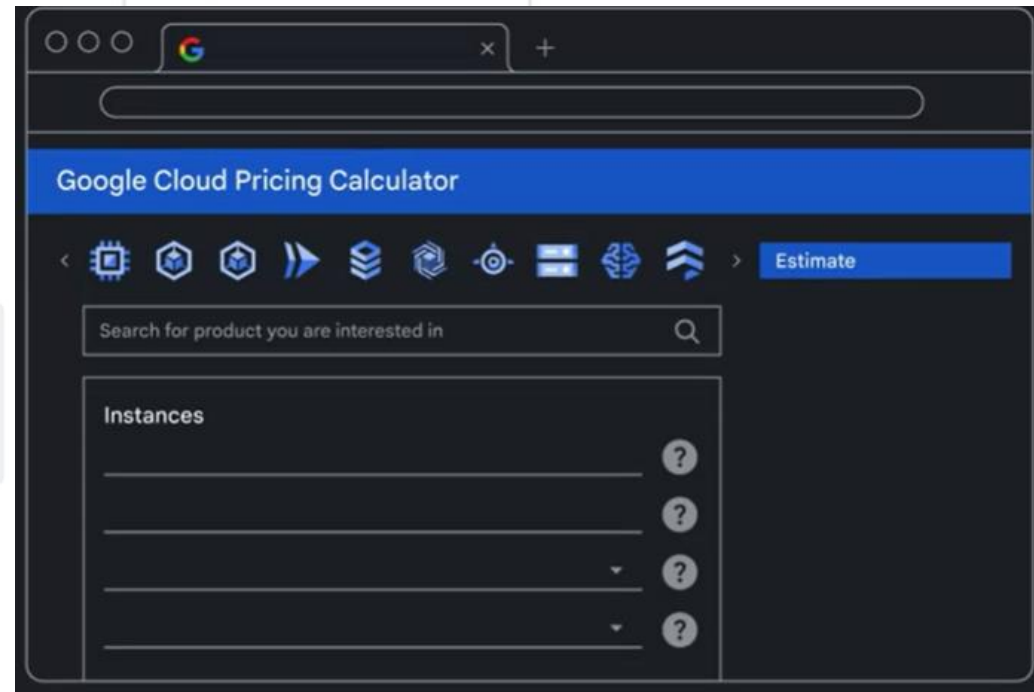
Google Workspace – 2022 - <https://workspace.google.com/pricing.html>

MOST POPULAR			
Business Starter	Business Standard	Business Plus	Enterprise
\$6 USD /user/month	\$12 USD /user/month	\$18 USD /user/month	Contact sales for pricing
Get started	Get started	Get started	Contact sales
<ul style="list-style-type: none">✓ Custom and secure business email✓ 100 participant video meetings✓ 30 GB cloud storage per user✓ Security and management controls✓ Standard Support	<ul style="list-style-type: none">✓ Custom and secure business email✓ 150 participant video meetings + recording✓ 2 TB cloud storage per user✓ Security and management controls✓ Standard Support (paid upgrade to Enhanced)	<ul style="list-style-type: none">✓ Custom and secure business email + eDiscovery, retention✓ 500 participant video meetings + recording, attendance tracking✓ 5 TB cloud storage per user✓ Enhanced security and management controls, including Vault and	<ul style="list-style-type: none">✓ Custom and secure business email + eDiscovery, retention, S/MIME encryption✓ 500 participant video meetings + recording, attendance tracking, noise cancellation, in-domain live streaming✓ As much storage as you need*


Google Cloud



Per-second billing




Budgets



Alerts



Reports



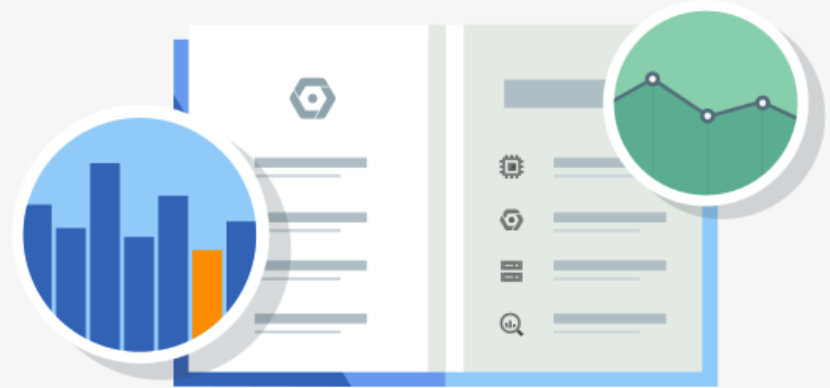
Quotas

Google



Cloud Price Leader

Google Cloud Platform gives you the best price to performance. Your cloud shouldn't break the bank and compromised performance should never be the only tradeoff.



 Boot up in

35 seconds

 Archive Restore


milliseconds

Milliseconds matter

Google Compute Engine instances boot up in 35 seconds on average. Coldline delivers millisecond data availability for archive restore – other public clouds can take up to 5 hours¹. Our Local SSDs offer 680,000 IOPS of sustained read performance – some other systems don't reach half of that IOPS. BigQuery can scan up to 35 billion rows, 20 TB of data, in seconds. Price and performance: we have both.

[<https://cloud.google.com/pricing/price-leader>]

Google



Do more for less

Pricing innovations like rightsizing and sustained-use discounts help deliver an average savings of 35% for many compute workloads. Our storage prices average 21% less than AWS for online storage workloads.¹

Get started for free


Contact sales



*21% sustained-use discounts and 18% rightsizing recommendations

[<https://cloud.google.com/pricing/price-leader>]

Google & Environmental impact



Google's data centers
were the first to achieve
ISO 14001 certification



[<https://cloud.google.com>]

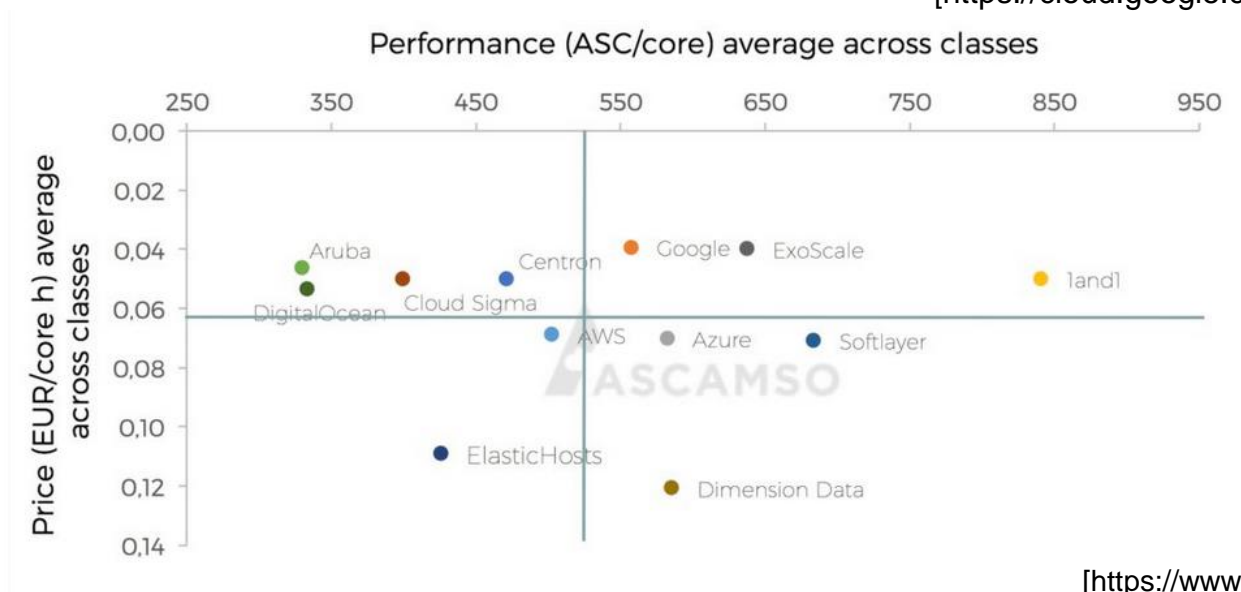
Google & Environmental impact

Data center efficiency through DeepMind machine learning

Google has applied artificial intelligence to optimize power usage in Google data centers. The up to **40% reduction in electricity per month** needed for cooling is considered a phenomenal step forward in the industry. On average, a Google data center **uses 50% less energy** than a typical data center.



[<https://cloud.google.com/pricing/price-leader>]



[<https://www.ascamsso.com/905-2/>]

price performance comparison of providers with European datacenters

Google

Google Web Toolkit – GWT

- <http://code.google.com/webtoolkit/gettingstarted.html>
- Set of open-source tools that allows the development of complex Web applications
- *“The GWT SDK contains the Java API libraries, compiler, and development server. It lets you write client-side applications in Java and deploy them as JavaScript.”*
-
- ***“GWT is a development toolkit for building and optimizing complex browser-based applications. Its goal is to enable productive development of high-performance web applications without the developer having to be an expert in browser quirks, XMLHttpRequest, and JavaScript. GWT is used by many products at Google, including AdWords, AdSense, Flights, Hotel Finder, Offers, Wallet, Blogger. It’s open source, completely free, and used by thousands of developers around the world.”***
-

Google

Google Apps Script - <https://developers.google.com/apps-script/>

- *“A cloud-based JavaScript platform that lets you integrate with and automate tasks across Google products.”*
- It offers the possibility of automating the business process, by means of scripts that add more functionality compared to the suite of already existing applications
 - Linking Google applications with *third party* services: a script can send emails and an invitation to a list from a MySQL database
 - Creating specialized functions for spreadsheets: performing complex analyzes on data from Google spreadsheets
 - Building an attractive user interface: an internal application in a company can benefit from an interface built with GAS

“With Apps Script, you can create add-ons for Google Sheets, Docs, or Forms, automate your workflow, integrate with external APIs, and more. “

11 Google apps, 1 platform in the cloud



Increase the power of your favorite Google apps – like [Calendar](#), [Docs](#), [Drive](#), [Gmail](#), and [Sheets](#).

Apps Script lets you [do more with Google](#). All on a JavaScript platform in the cloud.

Google Cloud

Hosting + Compute



App Engine



Compute Engine

Storage



Cloud Storage



Cloud Datastore



Cloud SQL

Big Data



BigQuery

Services



Cloud Endpoints



Translate API



Prediction API



2015 - <https://cloud.google.com/>

Google Cloud

Compute



App Engine



Compute Engine



Container Engine

Storage



Cloud Storage



Cloud Datastore



Cloud SQL



Cloud Bigtable

Big Data



BigQuery



Cloud Dataflow



Cloud Dataproc



Cloud Pub/Sub

Services



Cloud Endpoints



Translate API



Prediction API



2016 - <https://cloud.google.com/>

Google Cloud



Compute



Storage & Databases



Networking



Big Data



Machine Learning



Identity & Security



Management Tools



Developer Tools



2017 - <https://cloud.google.com/>

Google Cloud

Compute	Storage & Databases	Networking
Compute Run VMs on Google's infrastructure	Cloud Storage Object storage with global edge-caching	Cloud Virtual Network Managed networking for GCP resources
App Engine PaaS for apps and backends	Cloud SQL Fully-managed MySQL and PostgreSQL database service	Cloud Load Balancing High performance, scalable load balancing
Container Engine Run containers on GCP	Cloud Bigtable Fully managed NoSQL database service	Cloud CDN Content delivery on Google's global network
Container Registry Private container image storage	Cloud Spanner ^{BETA} Mission-critical, relational database service	Cloud Interconnect Connect directly to GCP's network edge
Cloud Functions ^{BETA} Serverless environment to build and connect cloud services	Cloud Datastore NoSQL database for non-relational data	Cloud DNS Reliable, resilient, low-latency DNS serving
	Persistent Disk Block storage for VM instances	



Google Cloud

The image shows a screenshot of the Google Cloud console navigation menu. The menu is organized into several main categories, each with a sub-menu of services. The categories and their services are:

- Compute**
 - Compute Engine: Run VMs on Google's infrastructure
 - App Engine: PaaS for apps and backends
 - Kubernetes Engine: Run containers on GCP
 - Cloud Functions ^{BETA}: Serverless environment to build and connect cloud services
- Storage & Databases**
 - Cloud Storage: Object storage with global edge-caching
 - Cloud SQL: Fully-managed MySQL and PostgreSQL database service
 - Cloud Bigtable: Fully managed NoSQL database service
 - Cloud Spanner: Mission-critical, relational database service
 - Cloud Datastore: NoSQL database for non-relational data
 - Persistent Disk: Block storage for VM instances
 - Data Transfer
 - Google Transfer Appliance: Securely migrate large volumes of data to Google Cloud Platform
- Networking**
 - Virtual Private Cloud (VPC): VPC networking for GCP resources
 - Cloud Load Balancing: High performance, scalable load balancing
 - Cloud CDN: Content delivery on Google's global network
 - Cloud Interconnect: Connect directly to GCP's network edge
 - Cloud DNS: Reliable, resilient, low-latency DNS serving
 - Network Service Tiers ^{ALPHA}: Optimize your network for performance or cost
 - Cloud AI
 - Cloud AutoML ^{Alpha}: Train high quality custom machine learning models with minimum effort and machine learning expertise
- Big Data**
 - BigQuery: Fully managed large-scale data warehouse
 - Cloud Dataflow: Real-time batch and stream data processing
 - Cloud Dataproc: Managed Spark and Hadoop service
- Security & Identity**









2018 - <https://cloud.google.com/> 19

Google Cloud

GOOGLE CLOUD PLATFORM

- **Featured products**
- **AI and machine learning**
- **API management**
- **Cloud Services Platform**
- **Compute**
- **Data analytics**
- **Databases**
- **Developer tools**
- **Internet of Things (IoT)**
- **Management tools**
- **Media**
- **Migration**
- **Networking**
- **Security**
- **Storage**

 Compute >	 Storage & Databases >	 Networking >
Compute Engine Run VMs on Google's infrastructure >	Cloud Storage Object storage with global edge-caching >	Virtual Private Cloud (VPC) VPC networking for GCP resources >
App Engine PaaS for apps and backends >	Cloud SQL Fully-managed MySQL and PostgreSQL database service >	Cloud Load Balancing High performance, scalable load balancing >
Kubernetes Engine Run containers on GCP >	Cloud Bigtable Fully managed NoSQL database service >	Cloud CDN Content delivery on Google's global network >
Cloud Functions ^{BETA} Serverless environment to build and connect cloud services >	Cloud Spanner Mission-critical, relational database service >	Cloud Interconnect Connect directly to GCP's network edge >
 Big Data >	Cloud Datastore NoSQL database for non-relational data >	Cloud DNS Reliable, resilient, low-latency DNS serving >
BigQuery Fully managed large-scale data warehouse >	Persistent Disk Block storage for VM instances >	Network Service Tiers ^{ALPHA} Optimize your network for performance or cost >
Cloud Dataflow Real-time batch and stream data processing >	 Data Transfer >	 Cloud AI
Cloud Dataproc Managed Spark and Hadoop service >	Google Transfer Appliance Securely migrate large volumes of data to Google Cloud Platform >	Cloud AutoML ^{Alpha} Train high quality custom machine with minimum effort and machine



Google Cloud

Google Cloud Platform

Grow your business with our secure storage, powerful compute, and integrated data analytics products.

[Get started for free](#)

[See pricing](#)

Featured products

Compute Engine

Scalable, high-performance VMs.

Cloud Run

Run stateless containers on a fully managed environment or on Anthos.

Anthos

Modernize existing apps and build new apps rapidly in hybrid and multi-cloud environments.

Vision AI

Derive insights from images, text, and more with AutoML Vision and Vision API.

Cloud Storage

Object storage with global edge-caching.

Cloud SQL

MySQL, PostgreSQL, and SQL Server database services.

BigQuery

A fully managed, highly scalable data warehouse with built-in ML.

Security key enforcement

Enforce the use of security keys to help prevent account takeovers.

Google Cloud Platform

Featured products

AI and Machine Learning

API Management

Compute

Containers

Data Analytics

Databases

Developer Tools

Healthcare and Life Sciences

Hybrid and Multi-cloud

Internet of Things (IoT)

Management Tools

Media and Gaming

Migration

Networking

Operations

Security and Identity

Serverless Computing

Storage



Google Cloud



Compute



Storage



Big Data



Machine
Learning




Application
Services

Google Cloud

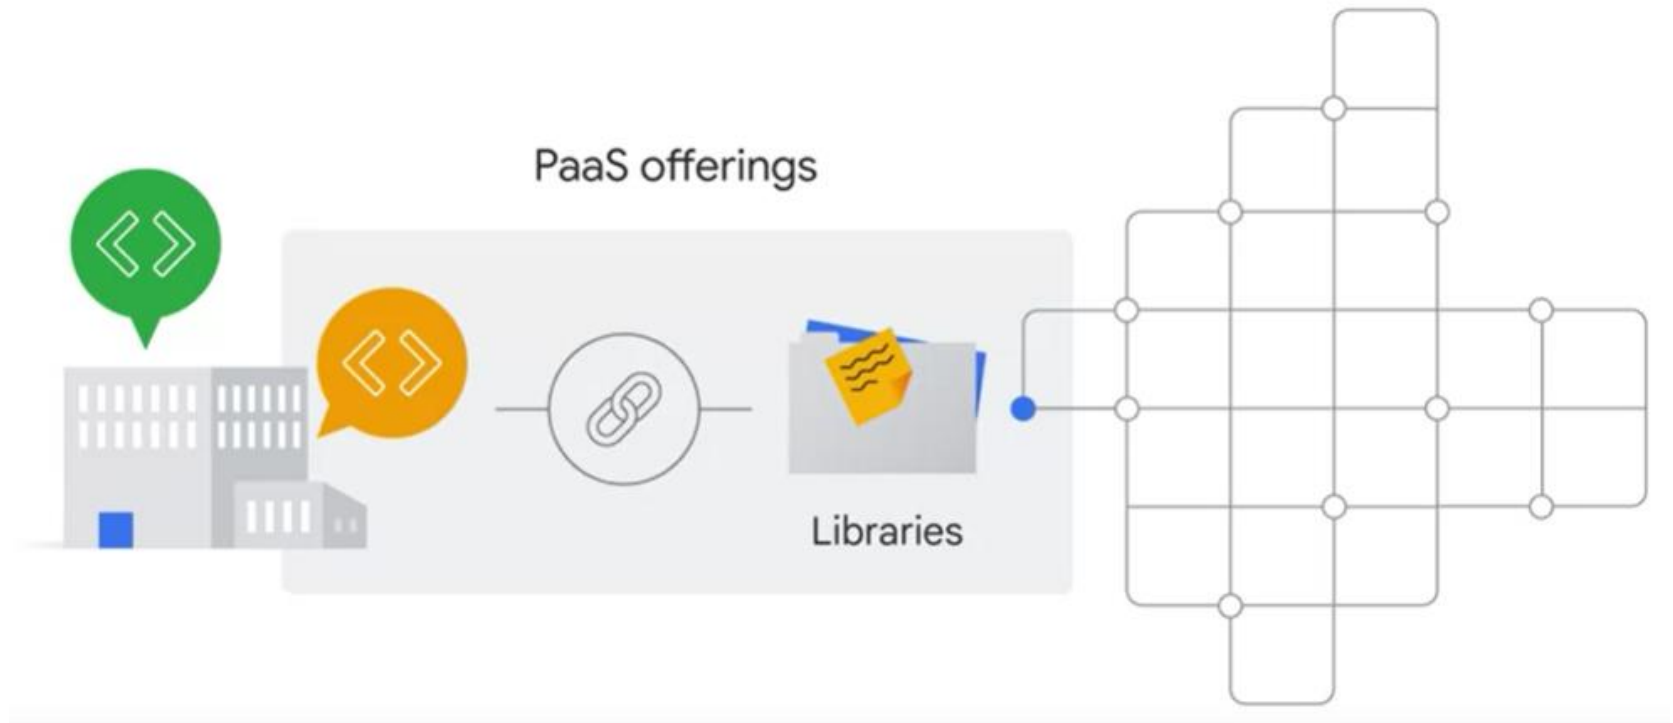
IaaS offerings provide:

- ✓ Raw compute
- ✓ Storage
- ✓ Network capabilities



<https://cloud.google.com/docs/overview/>

Google Cloud



<https://cloud.google.com/docs/overview/>

Google Cloud Platform

 35

REGIONS

 106

ZONES

 176

NETWORK EDGE LOCATIONS

AVAILABLE IN
 200+

COUNTRIES AND TERRITORIES

COMING SOON! Google Cloud will continue expanding into the following regions: Doha (Qatar), Turin (Italy), Berlin (Germany), Dammam (Kingdom of Saudi Arabia), Querétaro (Mexico), Malaysia, Thailand, New Zealand, Greece, Norway, South Africa, Austria and Sweden.



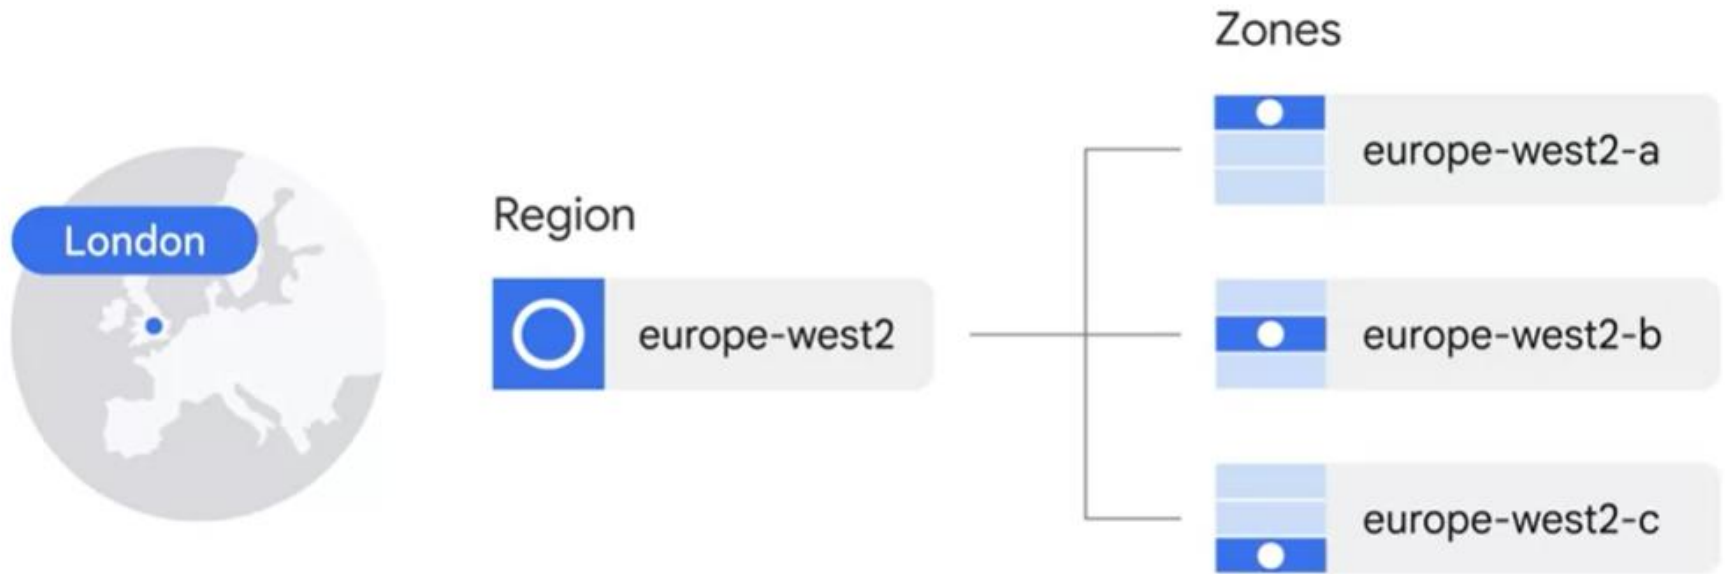
Products available by location

Deploy resources in specific zones, regions and multi-regions.



2023- <https://cloud.google.com/about/locations>

Google Cloud Platform



2023- <https://cloud.google.com>

Google Cloud | Network

- ✓ Highest possible throughput
- ✓ Lowest possible latencies
- ✓ 100+ content caching nodes worldwide
- ✓ High demand content is cached for quicker access




2023- <https://cloud.google.com/about/locations>

Google Cloud Platform | *Compute*



Compute Engine

Run large-scale workloads on virtual machines hosted on Google's infrastructure. Choose a VM that fits your needs and gain the performance of Google's worldwide fiber network.

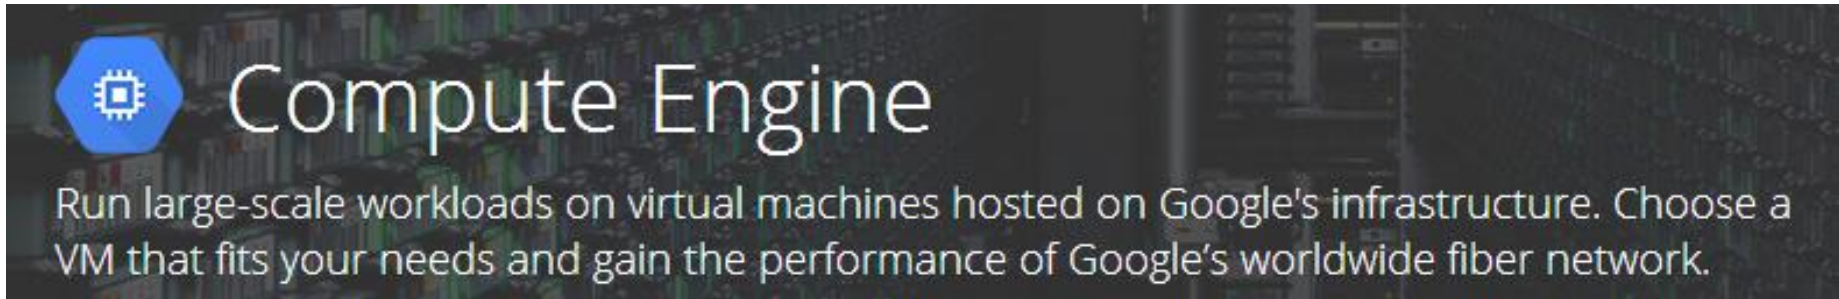
 Compute >
Compute Engine > <small>Run VMs on Google's infrastructure</small>
App Engine > <small>PaaS for apps and backends</small>
Kubernetes Engine > <small>Run containers on GCP</small>
Cloud Functions ^{BETA} > <small>Serverless environment to build and connect cloud services</small>

- <https://cloud.google.com/compute/>

Facilitati:

- *High-performance virtual machines* (Debian, CentOS, <-> instante: microVM -> large)
- *Comunication: Google's private global fiber network*
- *pay-per-use* – Obs. “*Google bills in seconds-level increments ...*”
- Automatic and easy management (API RESTful,
 1. `// CREATE INSTANCE WITH 4 vCPUs and 5 GB MEMORY`
 2. `gcloud compute instances create my-vm --custom-cpu 4 --custom-memory 5`
- **Security:** Certifications for Google Compute Engine: ISO 27001, SSAE-16, SOC 1, SOC 2, si SOC 3

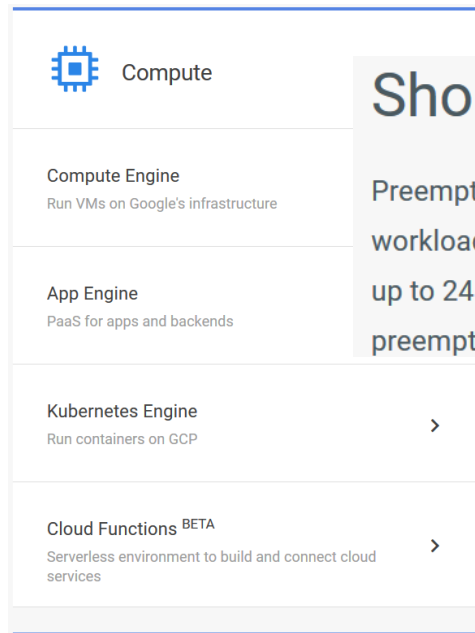
Google Cloud Platform | *Compute*



Compute Engine

Run large-scale workloads on virtual machines hosted on Google's infrastructure. Choose a VM that fits your needs and gain the performance of Google's worldwide fiber network.

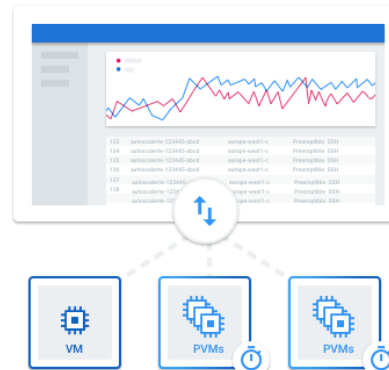
- **Spot VMs** - <https://cloud.google.com/spot-vms/>



- Compute**
- Compute Engine
Run VMs on Google's infrastructure
- App Engine
PaaS for apps and backends
- Kubernetes Engine >
Run containers on GCP
- Cloud Functions ^{BETA} >
Serverless environment to build and connect cloud services

Short Lived, Low Cost VMs

Preemptible VMs are highly affordable, short-lived compute instances suitable for batch jobs and fault-tolerant workloads. Preemptible VMs offer the same machine types and options as regular compute instances and last for up to 24 hours. If your applications are fault-tolerant and can withstand possible instance preemptions, then preemptible instances can reduce your **Google Compute Engine** costs significantly.



```
$ gcloud compute instances create my-vm  
--provisioning-model=SPOT --zone us-centra
```

Preemptible VMs are up to 80% cheaper than regular instances.

Google Cloud Platform | *Compute*



Compute Engine

Run large-scale workloads on virtual machines hosted on Google's infrastructure. Choose a VM that fits your needs and gain the performance of Google's worldwide fiber network.



Compute



- “All machine series support preemptible VMs, with the exception of the M2 machine series.”

- General-purpose –best price-performance ratio for a variety of workloads.
- Compute-optimized –highest performance per core on Compute Engine and optimized for compute-intensive workloads.
- Memory-optimized –ideal for memory-intensive workloads, offering more memory per core than other machine families, with up to 12 TB of memory.
- Accelerator-optimized –ideal for massively parallelized Compute Unified Device Architecture (CUDA) compute workloads, such as machine learning (ML) and high performance computing (HPC). This family is the best option for workloads that require GPUs.

[<https://cloud.google.com/compute/docs/machine-types>] 30

Google Cloud Platform | *Compute*

Google Cloud Functions

Event-driven serverless compute platform

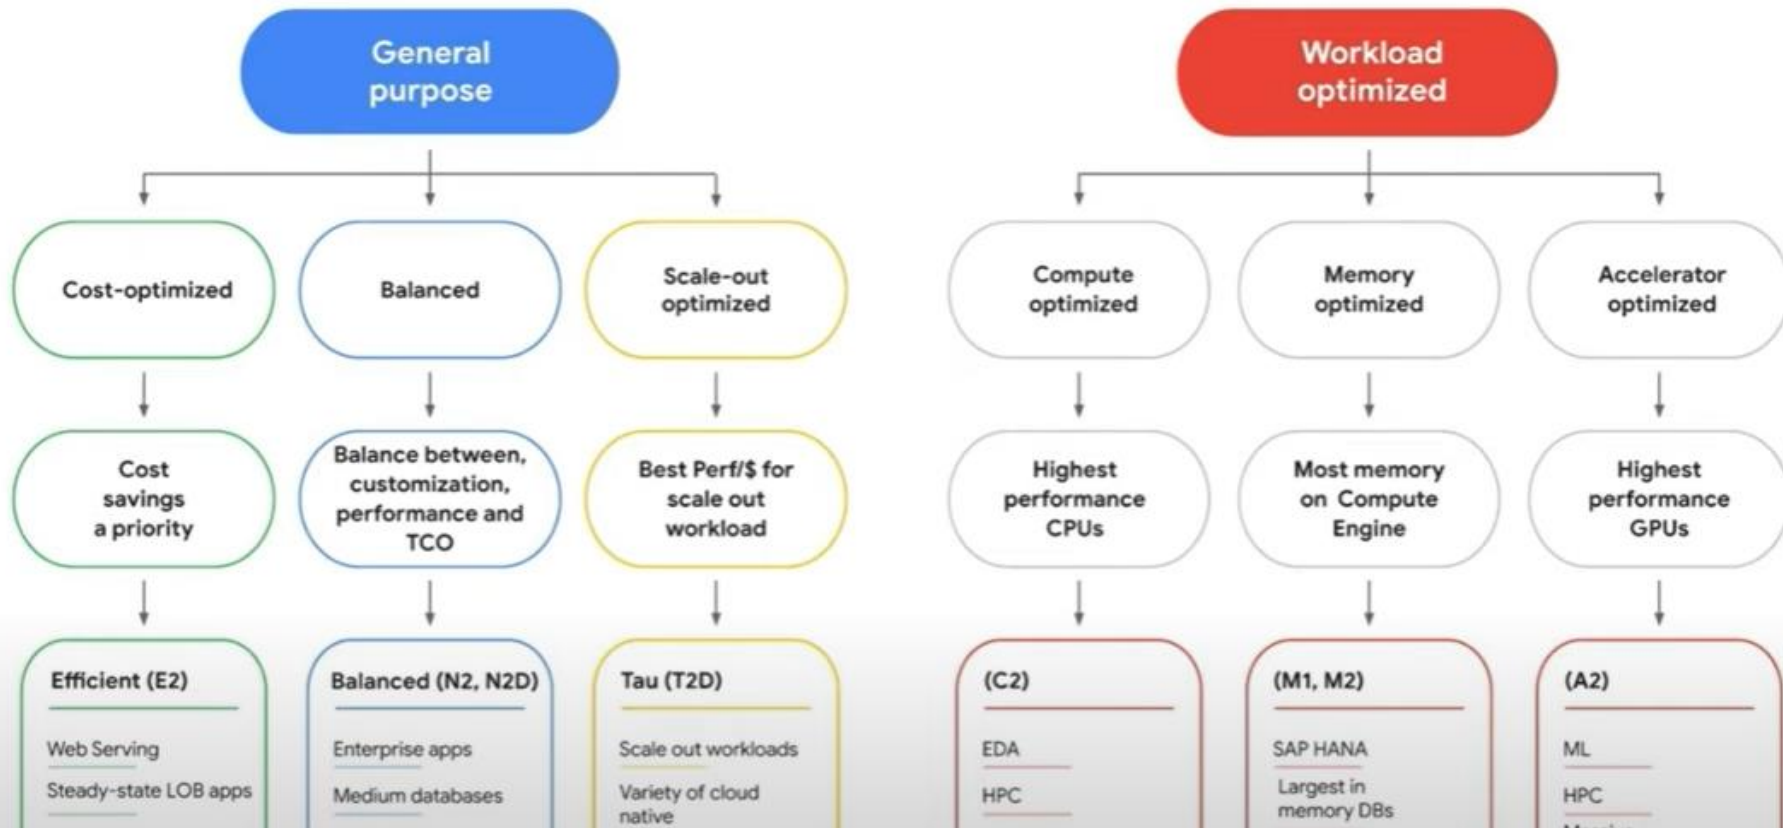
– *Serverless applications*



- ✓ Simplest way to run your code in the cloud
- ✓ Automatically scales, highly available and fault tolerant
- ✓ No servers to provision, manage, patch or update
- ✓ Pay only while your code runs
- ✓ Connects and extends cloud services

Cloud Functions lets application developers spin up code on demand in response to events originating from anywhere. Treat all Google and third-party cloud services as building blocks, connect and extend them with code, and build applications that scale from zero to planet-scale—without provisioning or managing a single server.

Google Cloud Platform | *Compute*



Compute Engine

<https://www.youtube.com/watch?v=YPWv-Ln9IPw&t=178s>

Google Cloud Platform | Storage



Cloud Storage

Use a powerful, simple and cost effective object storage service. With global edge-caching, your users have fast access to your app's data from any location.

- <https://cloud.google.com/storage/>
- Standard Storage – useful for data storage that requires a high level of availability and performance (frequent access, mobile applications, video streaming)
- Archive Storage – useful for data storage that requires low latency
- Nearline Storage – useful for storing data accessed less than once a month (e.g. backup)
- Coldline Storage – for long-lived data storage but which is accessed less frequently (once a year at most, disaster-recovery)

Google Cloud Platform | Storage



Cloud Storage

Use a powerful, simple and cost effective object storage service. With global edge-caching, your users have fast access to your app's data from any location.

- ✓ Unlimited storage with no minimum object size
- ✓ Worldwide accessibility and worldwide **storage locations**
- ✓ Low latency
- ✓ High durability (99.999999999% annual durability)
- ✓ **Geo-redundancy** if the data is stored in a multi-region or dual-region



BLOG POST

Twitter chooses Google Cloud for its flexibility in storage and compute.

[<https://www.youtube.com/watch?v=wN0s3LlsH6k>]

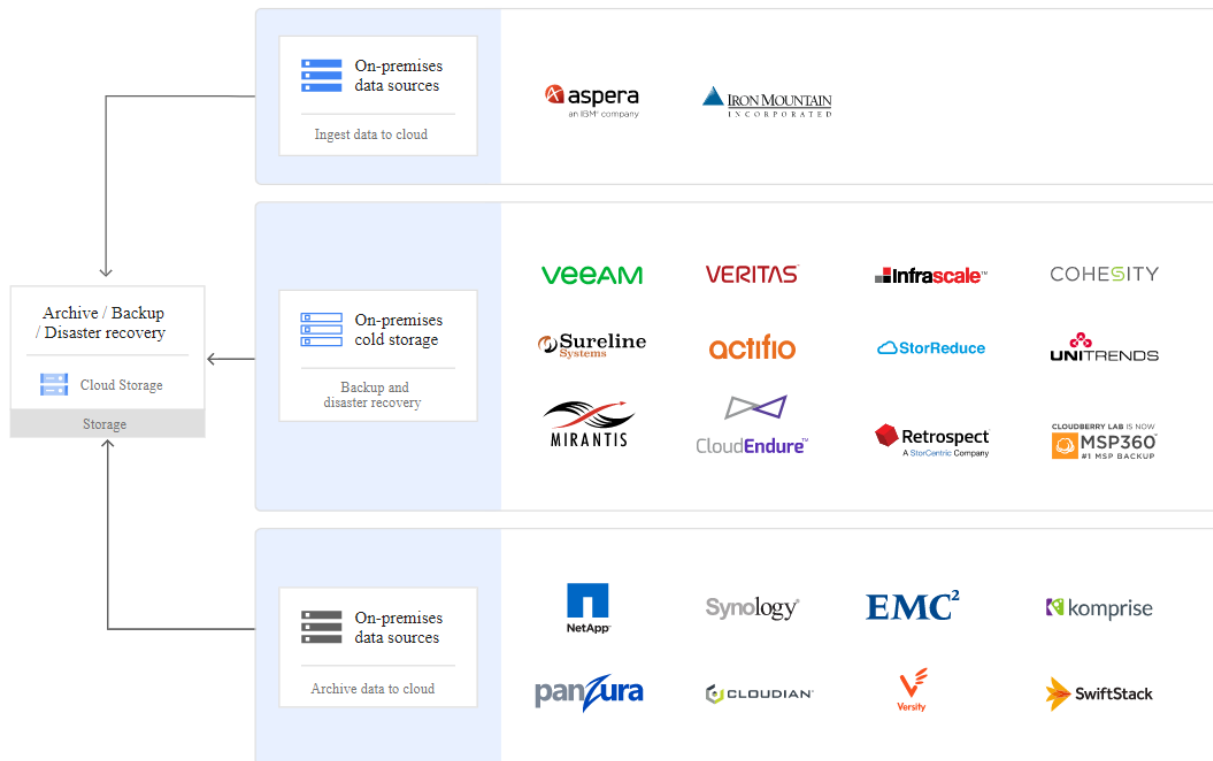
Google Cloud Platform | Storage




Cloud Storage

Use a powerful, simple and cost effective object storage service. With global edge-caching, your users have fast access to your app's data from any location.

Backups and archives



Google Cloud Platform | Storage



Cloud Datastore

Use a managed, NoSQL, schemaless database for storing non-relational data. Cloud Datastore automatically scales as you need it and supports transactions as well as robust, SQL-like queries.

INDUSTRY NEWS
Forrester names Google a leader in The Forrester Wave™: Big Data NoSQL, Q1 2019. Get your complimentary report.


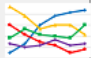


- <https://cloud.google.com/datastore/> - *Datastore is a highly scalable NoSQL database for your web and mobile applications.*
- Facilities:
 - Schemaless storage service for document-oriented databases
 - Allows automatic replication and sharding => availability and consistency
 - Cloud Datastore Development Kit – allows the local development of web or mobile applications

US (multi-region) ▾

	FREE <small>limit per day</small>	PRICE <small>above free limit (per unit)</small>	Price Unit
Stored data	1 GB storage	\$0.18	GB/Month
Entity Reads	50,000	\$0.06	per 100,000 entities
Entity Writes	20,000	\$0.18	per 100,000 entities
Entity Deletes	20,000	\$0.02	per 100,000 entities
Small Operations	50,000	Free	-

Google Cloud Platform | Storage

- **Firestore** - <https://cloud.google.com/firestore> - next generation of Datastore

Editorial information provided by DB-Engines		
Name	Google Cloud Datastore X	Google Cloud Firestore X
Description	Automatically scaling NoSQL Database as a Service (DBaaS) on the Google Cloud Platform	Cloud Firestore is an auto-scaling document database for storing, syncing, and querying data for mobile and web apps. It offers seamless integration with other Firebase and Google Cloud Platform products.
Primary database model	Document store	Document store
DB-Engines Ranking  Trend Chart 	Score 5.48 Rank #76 Overall #12 Document stores	Score 9.15 Rank #52 Overall #9 Document stores
Website	cloud.google.com/datastore	firebase.google.com/products/firestore
Technical documentation	cloud.google.com/datastore/docs	firebase.google.com/docs/firestore
Developer	Google	Google
Initial release	2008	2017
License 	commercial	commercial
Cloud-based only 	yes	yes

[<https://db-engines.com/en/system/Google+Cloud+Datastore%3BGoogle+Cloud+Firestore>]

Google Cloud Platform | Storage

DB-Engines Ranking

The DB-Engines Ranking ranks database management systems according to their popularity. The ranking is updated monthly.

Read more about the [method](#) of calculating the scores.



trend chart

52. ↑ 53. ↑ 54. [Google Cloud Firestore](#)

[Document](#)

9.15 +0.09 +1.70

388 systems in ranking, March 2022

Rank			DBMS	Database Model	Score		
Mar 2022	Feb 2022	Mar 2021			Mar 2022	Feb 2022	Mar 2021
1.	1.	1.	Oracle +	Relational, Multi-model ?	1251.32	-5.51	-70.42
2.	2.	2.	MySQL +	Relational, Multi-model ?	1198.23	-16.45	-56.59
3.	3.	3.	Microsoft SQL Server +	Relational, Multi-model ?	933.78	-15.27	-81.52
4.	4.	4.	PostgreSQL + ?	Relational, Multi-model ?	616.93	+7.54	+67.64
5.	5.	5.	MongoDB +	Document, Multi-model ?	485.66	-2.98	+23.27
6.	6.	↑ 7.	Redis +	Key-value, Multi-model ?	176.76	+0.96	+22.61
7.	7.	↓ 6.	IBM Db2	Relational, Multi-model ?	162.15	-0.73	+6.14
8.	8.	8.	Elasticsearch	Search engine, Multi-model ?	159.95	-2.35	+7.61
9.	9.	↑ 10.	Microsoft Access	Relational	135.43	+4.17	+17.29
10.	10.	↓ 9.	SQLite +	Relational	132.18	+3.81	+9.54
11.	11.	11.	Cassandra +	Wide column	122.14	-1.83	+8.51
12.	12.	12.	MariaDB +	Relational, Multi-model ?	108.31	+1.20	+13.85
13.	13.	13.	Splunk	Search engine	95.36	+4.55	+8.44
14.	↑ 15.	↑ 30.	Snowflake +	Relational	86.23	+3.05	+63.04
15.	↓ 14.	↑ 16.	Microsoft Azure SQL Database	Relational, Multi-model ?	84.68	-0.28	+13.79
16.	↑ 17.	↑ 17.	Amazon DynamoDB +	Multi-model ?	81.80	+1.45	+12.91
17.	↓ 16.	↓ 14.	Hive +	Relational	81.22	-0.66	+5.18

[<https://db-engines.com/en/ranking>]

Google Cloud Platform | Storage



Cloud SQL

Store and manage data using a fully-managed, relational MySQL database. Google handles replication, patch management and database management to ensure availability and performance.


- <https://cloud.google.com/sql>
- Facilities:
 - MySQL and PostgreSQL databases
 - Security (ISO/IEC 27001), performance, scalability and ease of use are ensured by applications that can run anywhere
 - Access: Web console or command line interface
 - *No Lock-in*

Gartner

Gartner recognizes Google Cloud as a Leader in the 2020 Magic Quadrant for Cloud Database Management Systems

Google Cloud Platform | Storage

- <https://cloud.google.com/bigtable/>
- *A fully managed, scalable NoSQL database service for large analytical and operational workloads with up to 99.999% availability.*
- *Consistent sub-10ms latency—handles millions of requests per second (optimized for heavy reads and writes)*
- It is used by: Search, Analytics, Maps, Gmail...
- It is designed to perform management in massive load contexts, ensuring low latency and high processing speed (=> a good option for applications such as IoT, user analytics, or financial data analysis)



Google Cloud Platform | Storage

PERSISTENT DISK

Reliable, high-performance block storage for virtual machine instances

- <https://cloud.google.com/persistent-disk/>

Google Persistent Disk is durable and high performance block storage for the Google Cloud Platform. Persistent Disk provides SSD and HDD storage which can be attached to instances running in either Google Compute Engine or Google Kubernetes Engine. Storage volumes can be transparently resized, quickly backed up, and offer the ability to support simultaneous readers.

Scale Without Interruption

You no longer have to worry about undersizing your block devices. Persistent Disk gives you unlimited flexibility by allowing you to resize your storage while it's in use by one or more virtual machines with no downtime.



Google Cloud Platform | Big Data



Google is named a Leader in Gartner 2020 Magic Quadrant for Cloud Database Management Systems.

Analyze Big Data in the cloud with BigQuery. Run fast, SQL-like queries against multi-terabyte datasets in seconds. Scalable and easy to use, BigQuery gives you real-time insights about your data.

- <https://cloud.google.com/bigquery/>
- *Serverless, highly scalable, and cost-effective multi-cloud data warehouse designed for business agility.*
- 1TB queries/month + 10Gb storage
- Queries can be executed asynchronously in the background; access to the history of queries and jobs is allowed

Resource	Pricing
Loading Data	Free
Exporting Data	Free
Storage	\$0.020 per GB / month ^{1,4}
Interactive Queries	\$5 per TB processed ^{2,3,4}
Batch Queries	\$5 per TB processed ^{2,3,4}
Streaming Inserts	\$0.01 per 100,000 rows.



BigQuery

Gartner

Gartner names Google a leader in the 2019 Magic Quadrant for Data Management Solutions for Analytics (DMSA). Get your complimentary report here.

Analyze Big Data in the cloud with BigQuery. Run fast, SQL-like queries against multi-terabyte datasets in seconds. Scalable and easy to use, BigQuery gives you real-time insights about your data.

Language correlation

Top 10 repos by total number of forks

Token occurrence per language

Search for the number of occurrences of a token across all commits per language:

Send request

```

1 SELECT repository_language, COUNT(*) AS cntlang
2 FROM [githubarchive:github.timeline]
3 WHERE repository_language != ''
4 AND payload_commit_msg != ''
5 AND PARSE_UTC_USEC(created_at) > PARSE_UTC_USEC('2014-03-09')
6 AND REGEXP_MATCH(payload_commit_msg,
7   r'(?i)\b(hello)\b')
8 GROUP BY repository_language
9 ORDER BY cntlang DESC

```

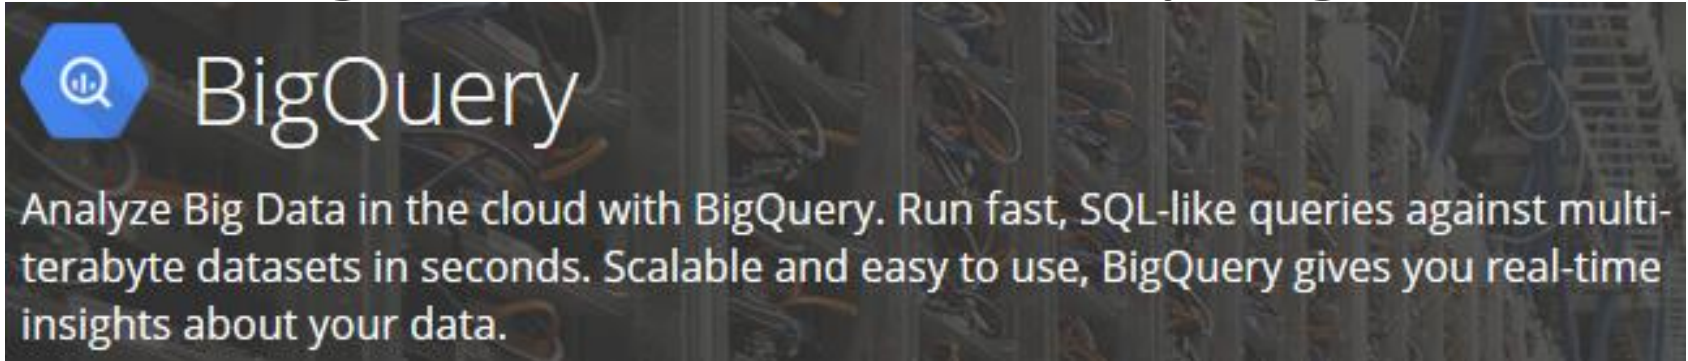
Query complete: 2.5s elapsed

[<https://cloud.google.com/bigquery/docs/running-queries>]

JavaScript	14989
Java	12000
CSS	9076
Ruby	8661

Examples: query the **BigQuery public dataset** `usa_names` to determine the most common names in the United States between the years 1910 and 2013, ...

Google Cloud Platform | Big Data



BigQuery

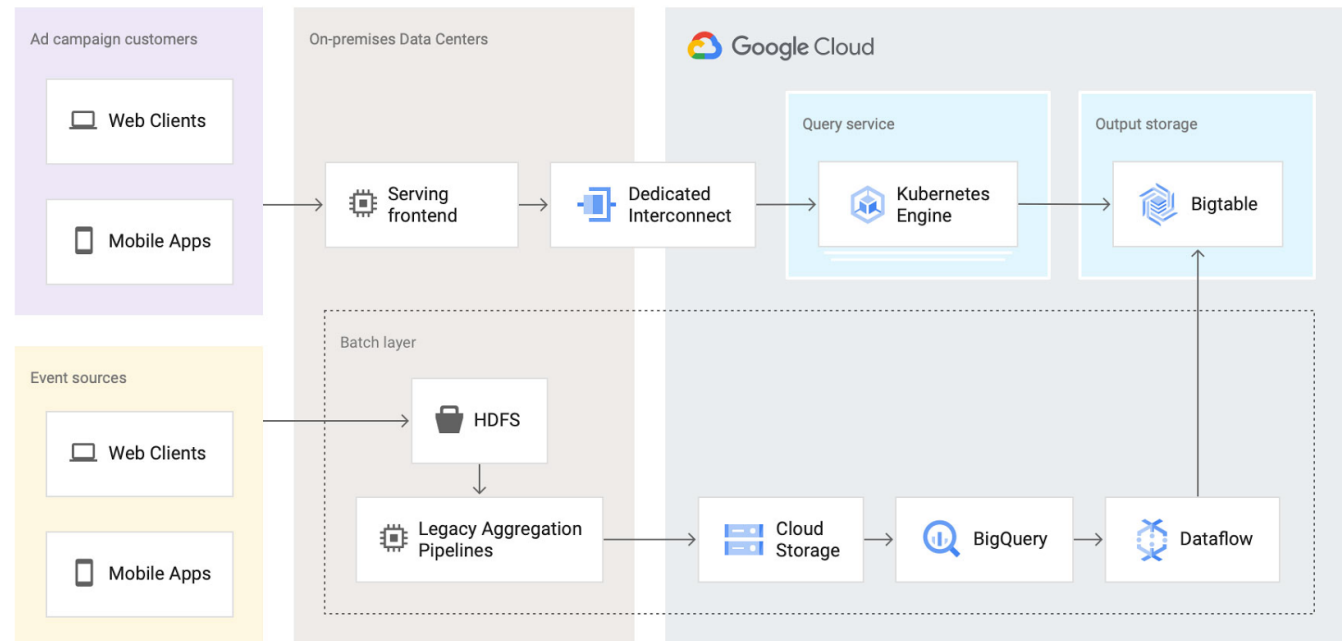
Analyze Big Data in the cloud with BigQuery. Run fast, SQL-like queries against multi-terabyte datasets in seconds. Scalable and easy to use, BigQuery gives you real-time insights about your data.



CASE STUDY

Twitter modernizes ad engagement platform; combines millions of metrics/second.

<https://cloud.google.com/blog/products/data-analytics/modernizing-twitters-ad-engagement-analytics-platform>

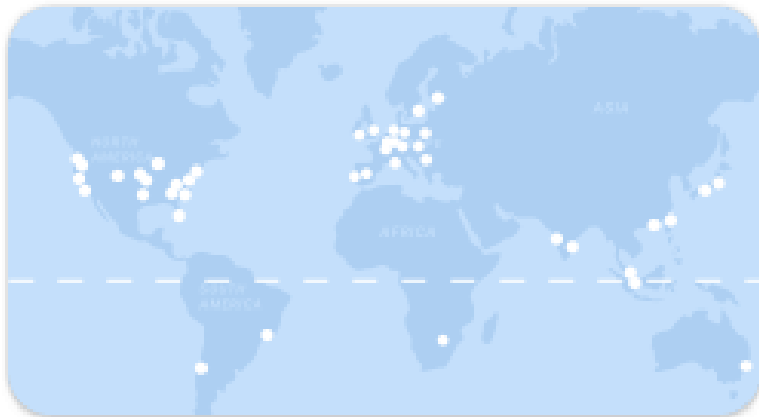


(prima iteratie)

Google Cloud Platform | Networking

- <https://cloud.google.com/cdn/>

Google Cloud CDN leverages Google's globally distributed edge points of presence to accelerate content delivery for websites and applications served out of Google Compute Engine and Google Cloud Storage. Cloud CDN lowers network latency, offloads origins, and reduces serving costs. Once you've set up HTTP(S) Load Balancing, simply enable Cloud CDN with a single checkbox.



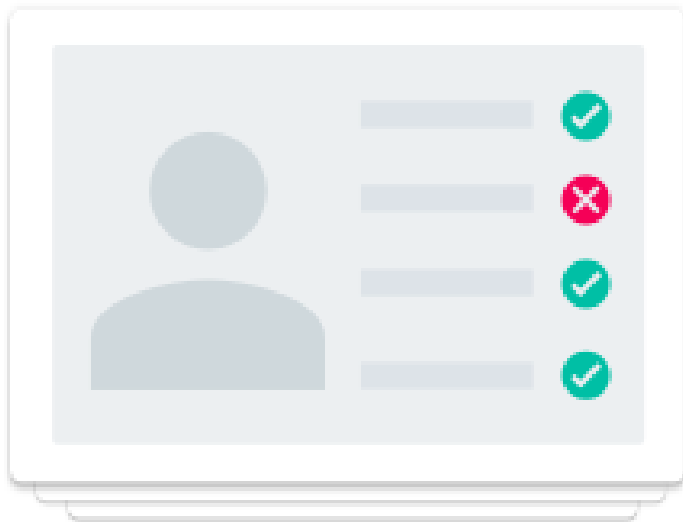
ITEM	PRICE (USD)
Cache egress	Charges vary based on location and usage: \$0.02/GB - \$0.08/GB for North America and Europe destinations \$0.04/GB - \$0.20/GB for other destinations worldwide
Cache fill	Charges vary based on location: \$0.04/GB - \$0.06/GB discounted pricing for in-region cache fills \$0.08/GB - \$0.15/GB for cross-region cache fills
HTTP(S) requests	\$0.0075 per 10,000 requests
Cache invalidation	\$0.005 per invalidation

With caches at more than 90 sites around the world,

Google Cloud Platform | Identity & Security

- <https://cloud.google.com/iam/> - *Fine-grained access control and visibility for centrally managing cloud resources.*

Google Cloud Identity & Access Management (IAM) lets administrators authorize who can take action on specific resources, giving you full control and visibility to manage cloud resources centrally. For established enterprises with complex organizational structures, hundreds of workgroups and potentially many more projects, Cloud IAM provides a unified view into security policy across your entire organization, with built-in auditing to ease compliance processes.

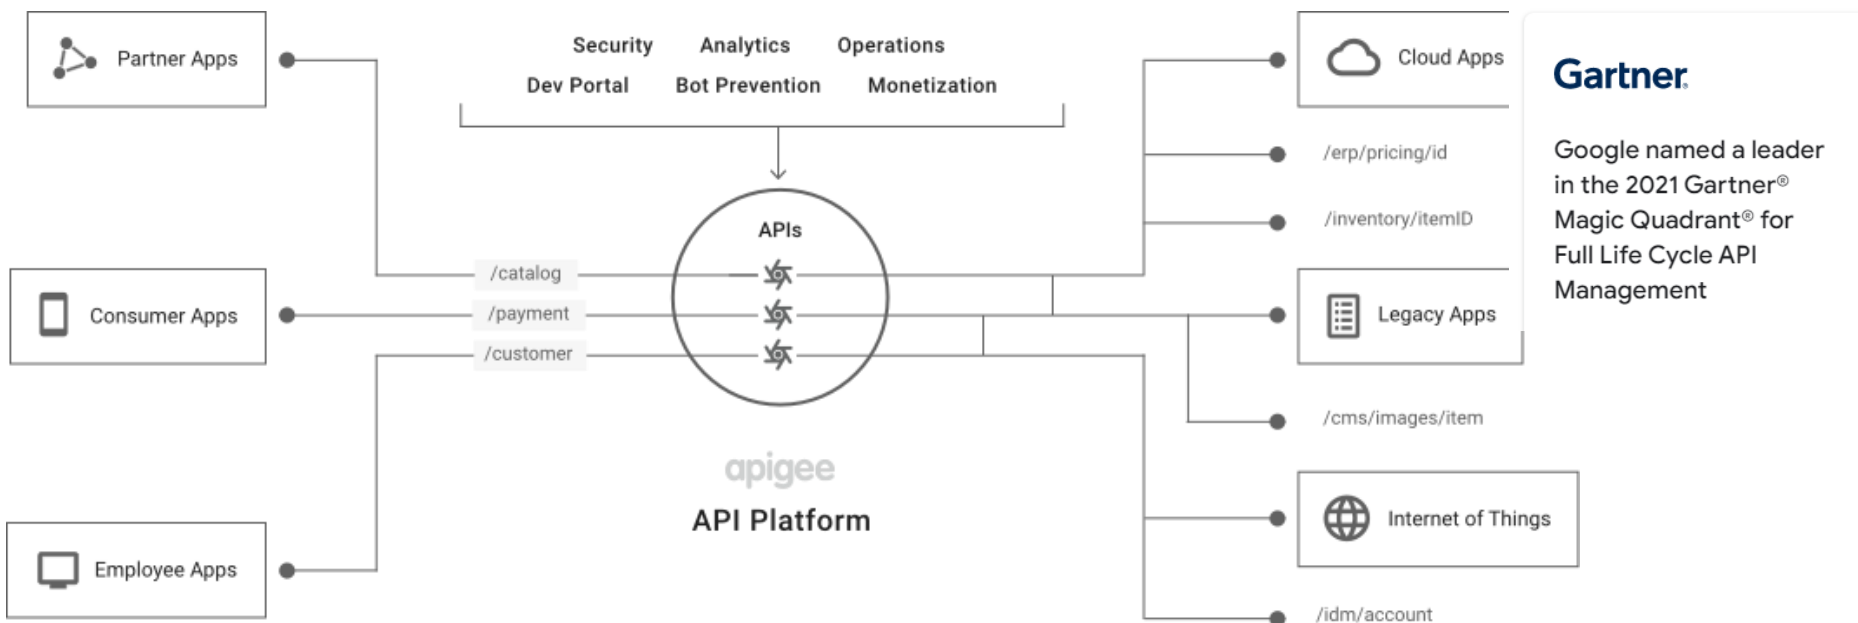


<https://cloud.google.com/iam/docs/quickstart>

Google Cloud Platform | API Platform and Ecosystems

- <https://cloud.google.com/apigee-api-management/>

Apigee is a full lifecycle API management platform that enables API providers to design, secure, deploy, monitor, and scale APIs. Apigee sits in-line with runtime API traffic and enforces a set of out-of-the-box API policies, including key validation, quota management, transformation, authorization, and access control. API providers use the customizable developer portal to enable developers to consume APIs easily and securely, and to measure API performance and usage.



Google Cloud Platform | Cloud AI

AI Building Blocks

AI building blocks

Easily infuse AI into applications with custom or pre-trained models.

AutoML

Custom machine learning model training and development.

Vision AI

Custom and pre-trained models to detect emotion, text, more.

Video AI

Video classification and recognition using machine learning.

Cloud Natural Language

Sentiment analysis and classification of unstructured text.

Cloud Translation

Language detection, translation, and glossary support.

Media Translation (beta)

Add dynamic audio translation directly to your content and applications.

Text-to-Speech

Speech synthesis in 220+ voices and 40+ languages.

Speech-to-Text

Speech recognition and transcription supporting 125 languages.

Dialogflow

Conversation applications and systems development suite for virtual agents.

AutoML Tables (beta)

Service for training ML models with structured data.

Cloud Inference API (alpha)

Quickly run large-scale correlations over typed time-series datasets.

Recommendations AI

Deliver highly personalized product recommendations at scale.

Google Cloud Platform | Cloud AI

CLOUD NATURAL LANGUAGE

Derive insights from unstructured text using Google machine learning

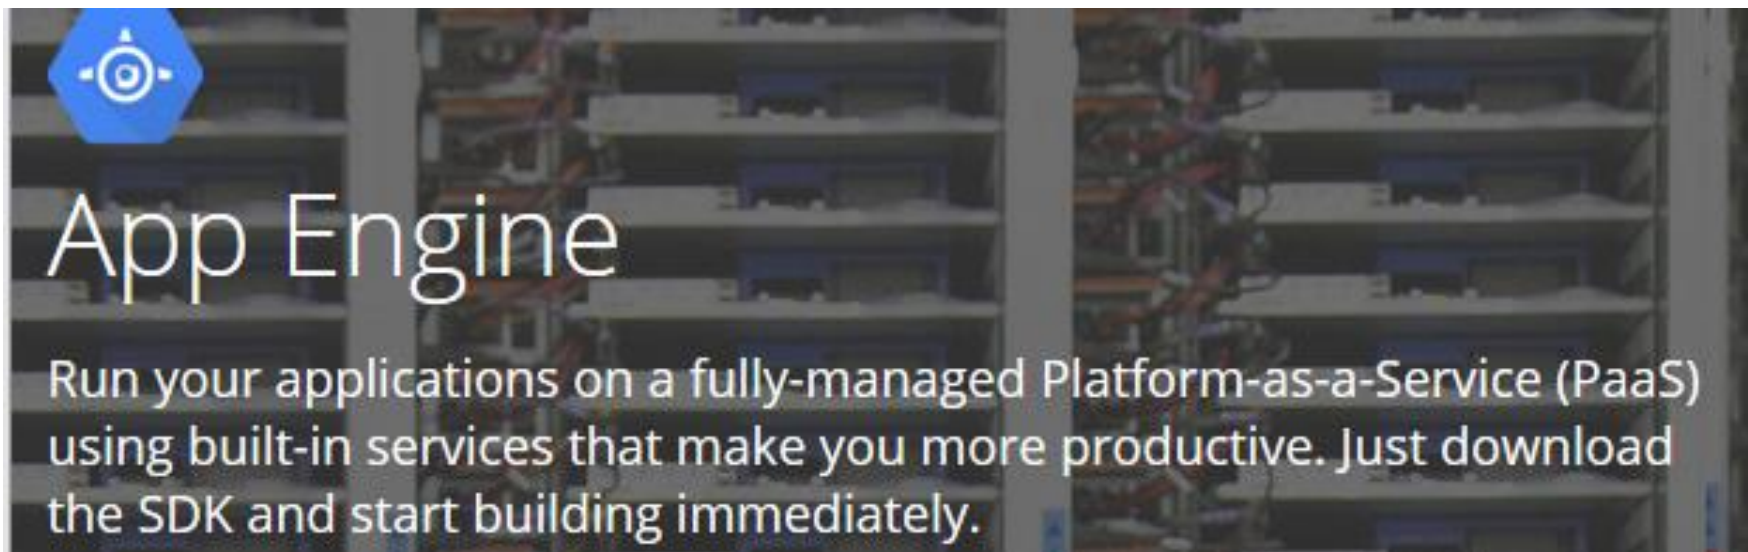
Powerful Text Analysis

Google Cloud Natural Language reveals the structure and meaning of text by offering powerful machine learning models in an easy to use REST API. You can use it to **extract information** about people, places, events and much more, mentioned in text documents, news articles or blog posts. You can use it to **understand sentiment** about your product on social media or **parse intent** from customer conversations happening in a call center or a messaging app. You can **analyze text uploaded in your request** or integrate with your document storage on Google Cloud Storage.

Google App Engine



2008 – launch of GAE - <http://www.youtube.com/watch?v=3Ztr-HhWX1c&feature=youtu.be>



[<https://developers.google.com/appengine/training/intro/whatisgae>]

Google Cloud

Google Cloud

2018 | The development environment offers

Standard Environment

About the Standard Environment

Python 2.7

Java 7

PHP 5.5

Go 1.6

Flexible Environment

About the Flexible Environment

Python 2.7, 3.6

Java 8

Node.js

Go 1.8, 1.9

Ruby

PHP 5.6, 7

.NET

Custom Runtimes

Known Issues

[2018 - <https://cloud.google.com/appengine/docs/>]

Google Cloud

Google Cloud

- 2019 | The development environment offers

Standard Environment

About the Standard Environment

The Standard Environment
Runtimes

Python 2.7, 3.7

Java 7, 8

Node.js 8, 10

PHP 5.5, 7.2

Go 1.9, 1.11

Flexible Environment

About the Flexible Environment

Python 2.7, 3.6

Java 8

Node.js

Go 1.9, 1.10, 1.11

Ruby

PHP 5.6, 7.0, 7.1, 7.2

.NET

Custom Runtimes

Known Issues

[2019 - <https://cloud.google.com/appengine/docs/>]

Google Cloud

Google Cloud

- 2020 | The development environment offers

- ▾ Standard Environment

- About the Standard Environment

- The Standard Environment Runtimes

- Long-term Support for Standard Runtimes

- Python 2.7, 3.7, 3.8

- Java 8, 11

- Node.js 10, 12

- PHP 5.5, 7.2, 7.3

- Ruby 2.5

- Go 1.11, 1.12, 1.13

- ▾ Flexible Environment

- About the Flexible Environment

- Python 2.7, 3.6

- Java 8

- Node.js

- Go 1.9, 1.10, 1.11

- Ruby

- PHP 5.6, 7.0, 7.1, 7.2

- .NET

- Custom Runtimes

[2020 - <https://cloud.google.com/appengine/docs/>]

Google Cloud

Google Cloud

- The development environment offers:
 - **App Engine Standard Environment - *Sandbox Environment*** –
cod + server Web+ *language runtime (modified)* a.i. the restrictions are respected *sandbox* => applications run in a secure environment, isolated from the hardware level, the operating system, and the physical location of the server
 - This limitation allows the distribution of web requests to several web servers (starting/stopping them) depending on the increase/decrease in requests
 - **App Engine Flexible Environment** –allows running App Engine applications using ***Google Compute Engine***
 - “*VM hosting environment offers more flexibility and provides more CPU and memory options*”.
 - There are no more restrictions from *sandbox runtimes*
[<https://cloud.google.com/appengine/docs/>]

Feature	Standard environment	Flexible environment
Instance startup time	Seconds	Minutes
Maximum request timeout	Depends on the runtime and type of scaling .	60 minutes
Background threads	Yes, with restrictions	Yes
Background processes	No	Yes
SSH debugging	No	Yes
Scaling	Manual, Basic, Automatic	Manual, Automatic
Scale to zero	Yes	No, minimum 1 instance
Writing to local disk	<ul style="list-style-type: none"> • Java 8, Java 11, Node.js, Python 3, PHP 7, Ruby, Go 1.11, and Go 1.12+ have read and write access to the <code>/tmp</code> directory. • Python 2.7 and PHP 5.5 don't have write access to the disk. 	Yes, ephemeral (disk initialized on each VM startup)
Modifying the runtime	No	Yes (through Dockerfile)
Deployment time	Seconds	Minutes
Automatic in-place security patches	Yes	Yes (excludes container image runtime)
Access to Google Cloud APIs & Services such as Cloud Storage , Cloud SQL , Memorystore , Tasks and others.	Yes	Yes
WebSockets	No Java 8, Python 2, and PHP 5 provide a proprietary Sockets API (beta), but the API is not available in newer standard runtimes.	Yes
Supports installing third-party binaries	<ul style="list-style-type: none"> • Yes for Java 8, Java 11, Node.js, Python 3, PHP 7, Ruby, Go 1.11, and Go 1.12+. • No for Python 2.7 and PHP 5.5. 	Yes
Location	North America, Asia Pacific, or Europe	North America, Asia Pacific, or Europe
Pricing	Based on instance hours	Based on usage of vCPU , memory , and persistent disks

[2022 - <https://cloud.google.com/appengine/docs/the-appengine-environments>]

Google Cloud

App Engine Flexible Environment

- Using Google Compute Engine Services is automatically ensured:
 - scaling and *load balancing*
 - checking the condition of the courts and optimal co-location with other services in the project
 - root access is provided to the Compute Engine VM instances (ssh access is disabled by default)
 - support for updates
 - it is allowed to modify the environment and the operating system by using *Dockerfiles*
- They are natively supported: *microservices, authorization, SQL and NoSQL databases, traffic splitting, logging, versioning, security scanning, and content delivery networks*

Google Cloud

App Engine Flexible Environment

– Runtimes

Google-supplied Dockerfile => standard runtime

- Provides native support for: Java 8 / Servlet 3.1 / Jetty 9, Python 2.7 and Python 3.6, Node.js, Ruby, PHP, .NET core, and Go
- *No* sandbox restrictions
- Queries can be made regarding the state of the VM, support is provided for services such as: Datastore, Memcache, Task Queues, Logging, Users

Docker image or a modified Dockerfile (open source) => custom runtime

- It is used if you want to write code in other languages
- Components such as interpreters or application servers can be configured

Google Cloud

- App Engine Flexible Environment

<https://cloud.google.com/appengine/docs/flexible/>

GO

JAVA 8

PHP 5 / 7

PYTHON 2.7 / 3.6

.NET

NODE.JS

RUBY

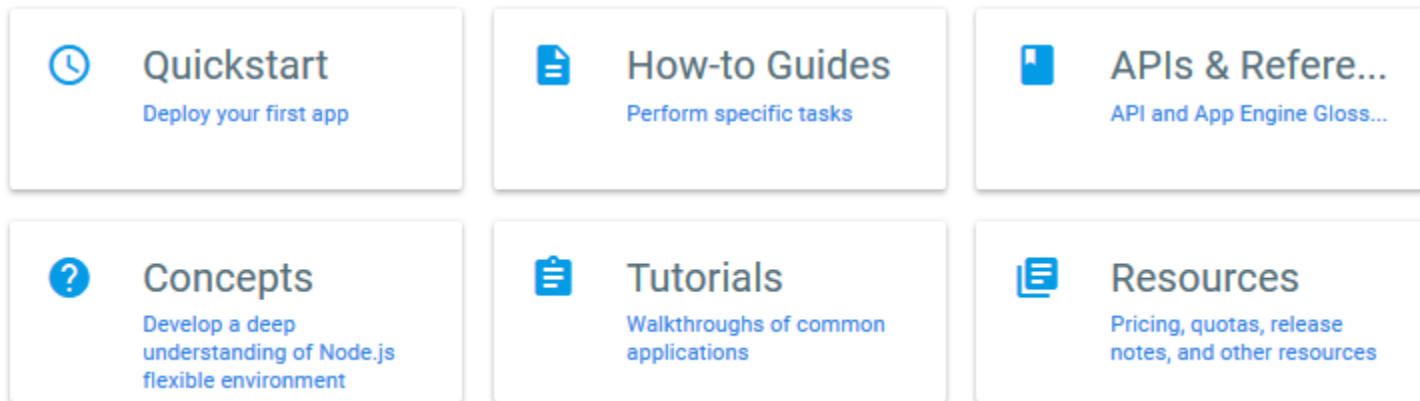
CUSTOM RUNTIMES

Google Cloud

- App Engine Flexible Environment - NodeJS

<https://cloud.google.com/appengine/docs/flexible/nodejs/>

The App Engine flexible environment is based on Google Compute Engine and automatically scales your app up and down while balancing the load.



Standard vs Flexible Environment

When to choose the standard environment

Application instances run in a [sandbox](#), using the runtime environment of a supported language listed below.

Applications that need to deal with rapid scaling.

The standard environment is optimal for applications with the following characteristics:

- Source code is written in **specific versions of the supported programming languages**:
 - Python 2.7, Python 3.7, Python 3.8, Python 3.9
 - Java 8, Java 11
 - Node.js 10, Node.js 12, Node.js 14, Node.js 16
 - PHP 5.5, PHP 7.2, PHP 7.3, and PHP 7.4
 - Ruby 2.5, Ruby 2.6, and Ruby 2.7
 - Go 1.11, Go 1.12, Go 1.13, Go 1.14, Go 1.15, and Go 1.16
- Intended to **run for free or at very low cost**, where you pay only for what you need and when you need it. For example, your application can scale to 0 instances when there is no traffic.
- Experiences **sudden and extreme spikes of traffic** which require immediate scaling.

When to choose the flexible environment

Application instances run within Docker containers on Compute Engine virtual machines (VM).

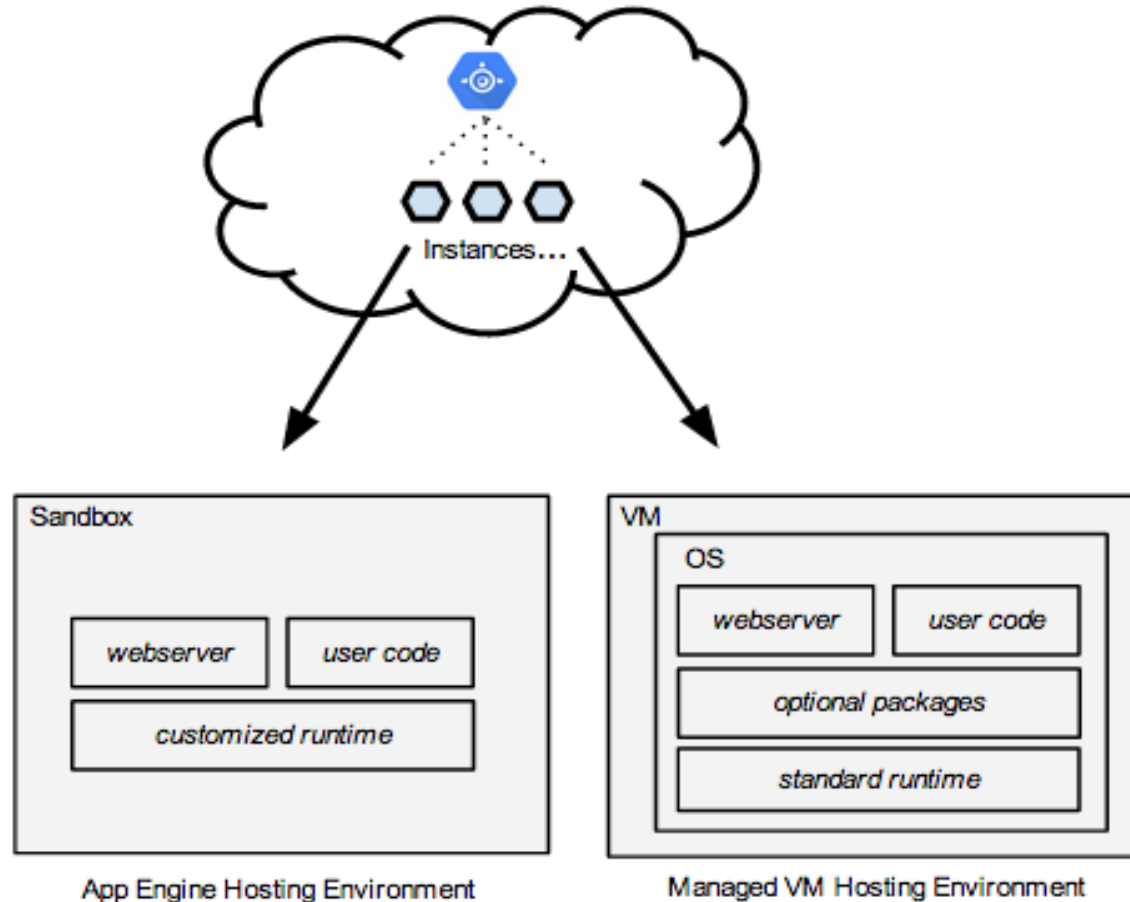
Applications that receive consistent traffic, experience regular traffic fluctuations, or meet the parameters for scaling up and down gradually.

The flexible environment is optimal for applications with the following characteristics:

- Source code that is written in a version of any of the supported programming languages: **Python, Java, Node.js, Go, Ruby, PHP, or .NET**
- Runs in a Docker container that includes a custom runtime or source code written in **other programming languages**.
- Uses or depends on frameworks that include **native code**.
- Accesses the resources or services of your Google Cloud project that reside in the **Compute Engine network**.

Google Cloud

- An application can contain modules that run in different hosting environments, e.g. it can be used for frontend *sandbox* and for processing ***App Engine Flexible Environment*** can be used



Google App Engine | Standard Environment



“Some users confuse Google App Engine with Amazon's EC2 service. The problem is that this is an apples to oranges comparison. Both operate at different cloud service levels, and each have their strengths and minuses. With App Engine, you only need to worry about your application and let Google take care of hosting and running it for you. With EC2, you're responsible for the app, but also its database server, web server, operating system, monitoring, load-balancing, upgrades, etc. This is the reason why typically, the costs for IaaS services run lower than that of PaaS services because with PaaS, you're "outsourcing" more work/responsibility. Cost estimates usually clouded by not considering the administration overhead when managing the infrastructure yourself. A better "apples-to-apples" comparison would be EC2 to the [Google Compute Engine](#) IaaS service.”

[<https://developers.google.com/appengine/training/intro/whatisgae>]

Google App Engine | Standard Environment



App Engine Standard Environment

Characteristics

- It allows the running (hosting) of Web applications using the Google infrastructure
 - *“App Engine Does One Thing Well”* [Rossum, Google]
- Do not understand "rent a piece of a server", because the application is not hosted on a single server
- Conceptually Google App Engine is at the PaaS level
- The control? Google...
 - No worries about infrastructure, load balancing, storage management

Google App Engine | Standard Environment



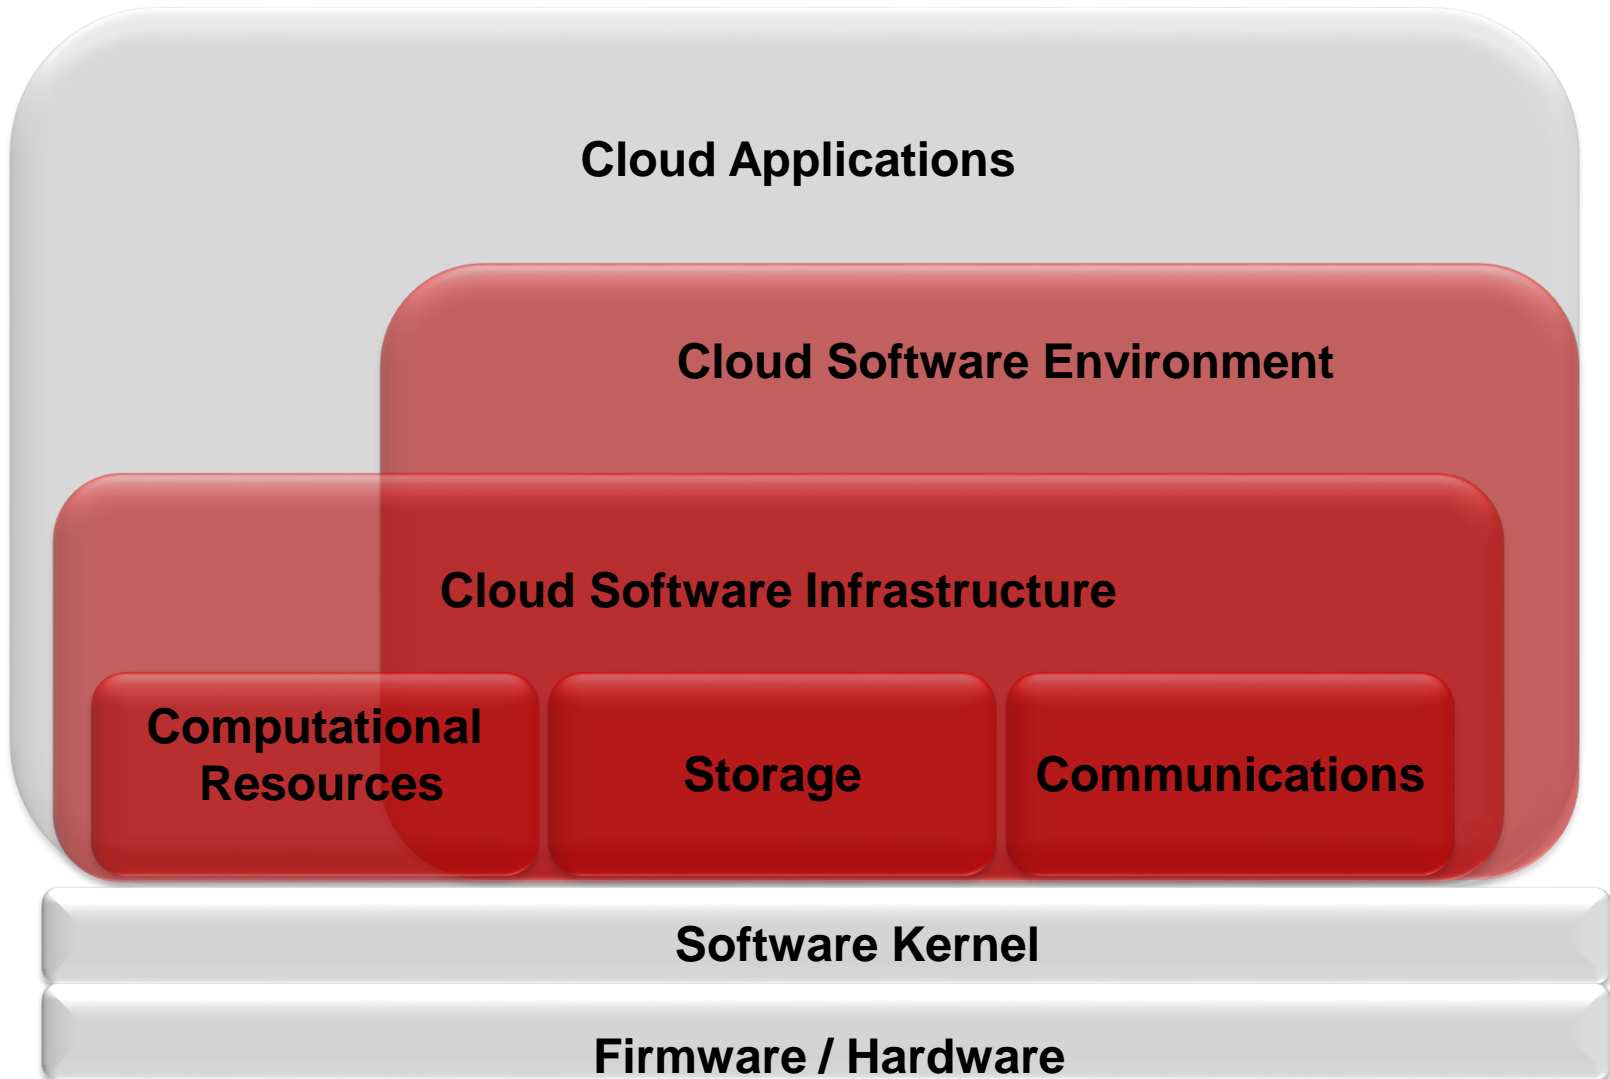
Characteristics

- GAE supports applications written in different languages (see the previous slides)
- Developers have access to storage technologies such as GFS (Google File System), Bigtable (storage system for unstructured data),...
- *“With Google App Engine, developers can write Web applications based on the same building blocks that Google uses,”* Kevin Gibbs, Google’s technical lead for the project, wrote in *The Official Google Blog*. *“Google App Engine packages those building blocks and provides access to scalable infrastructure that we hope will make it easier for developers to scale their applications automatically as they grow.”*

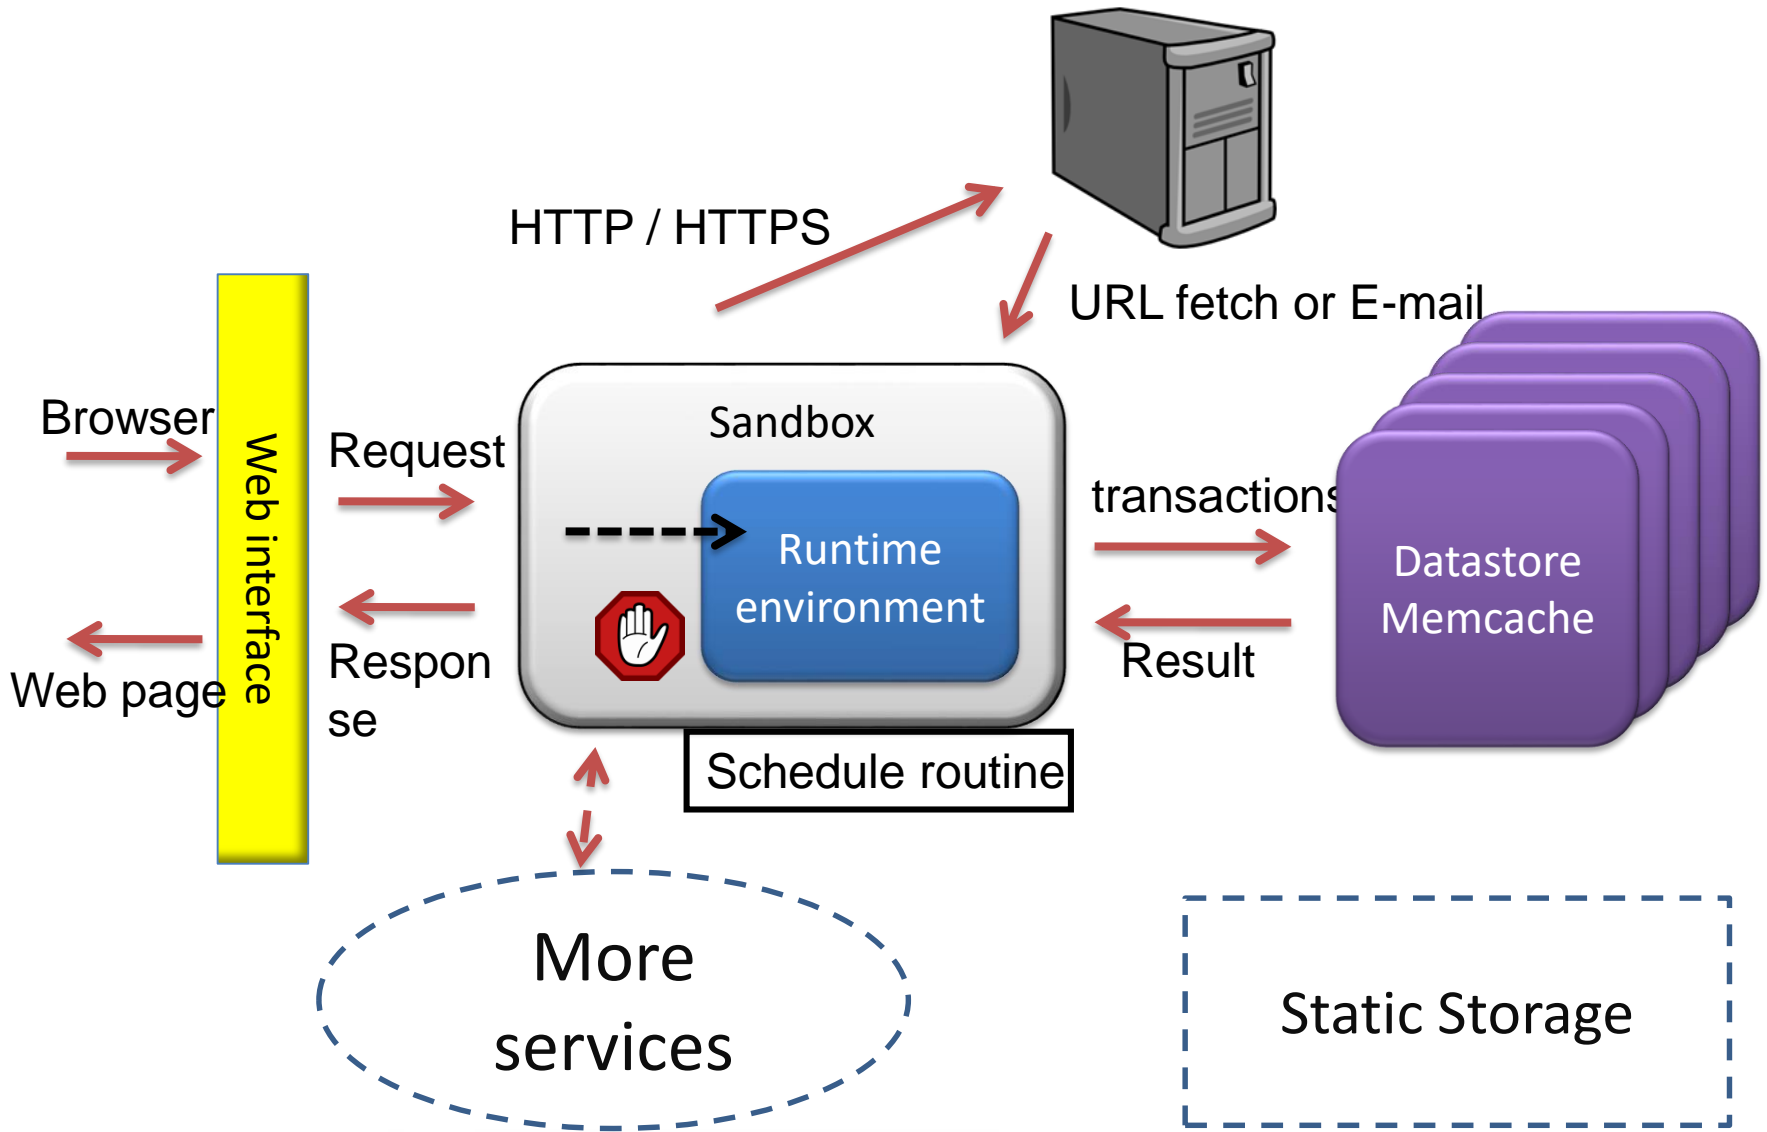
Google App Engine | Standard Environment



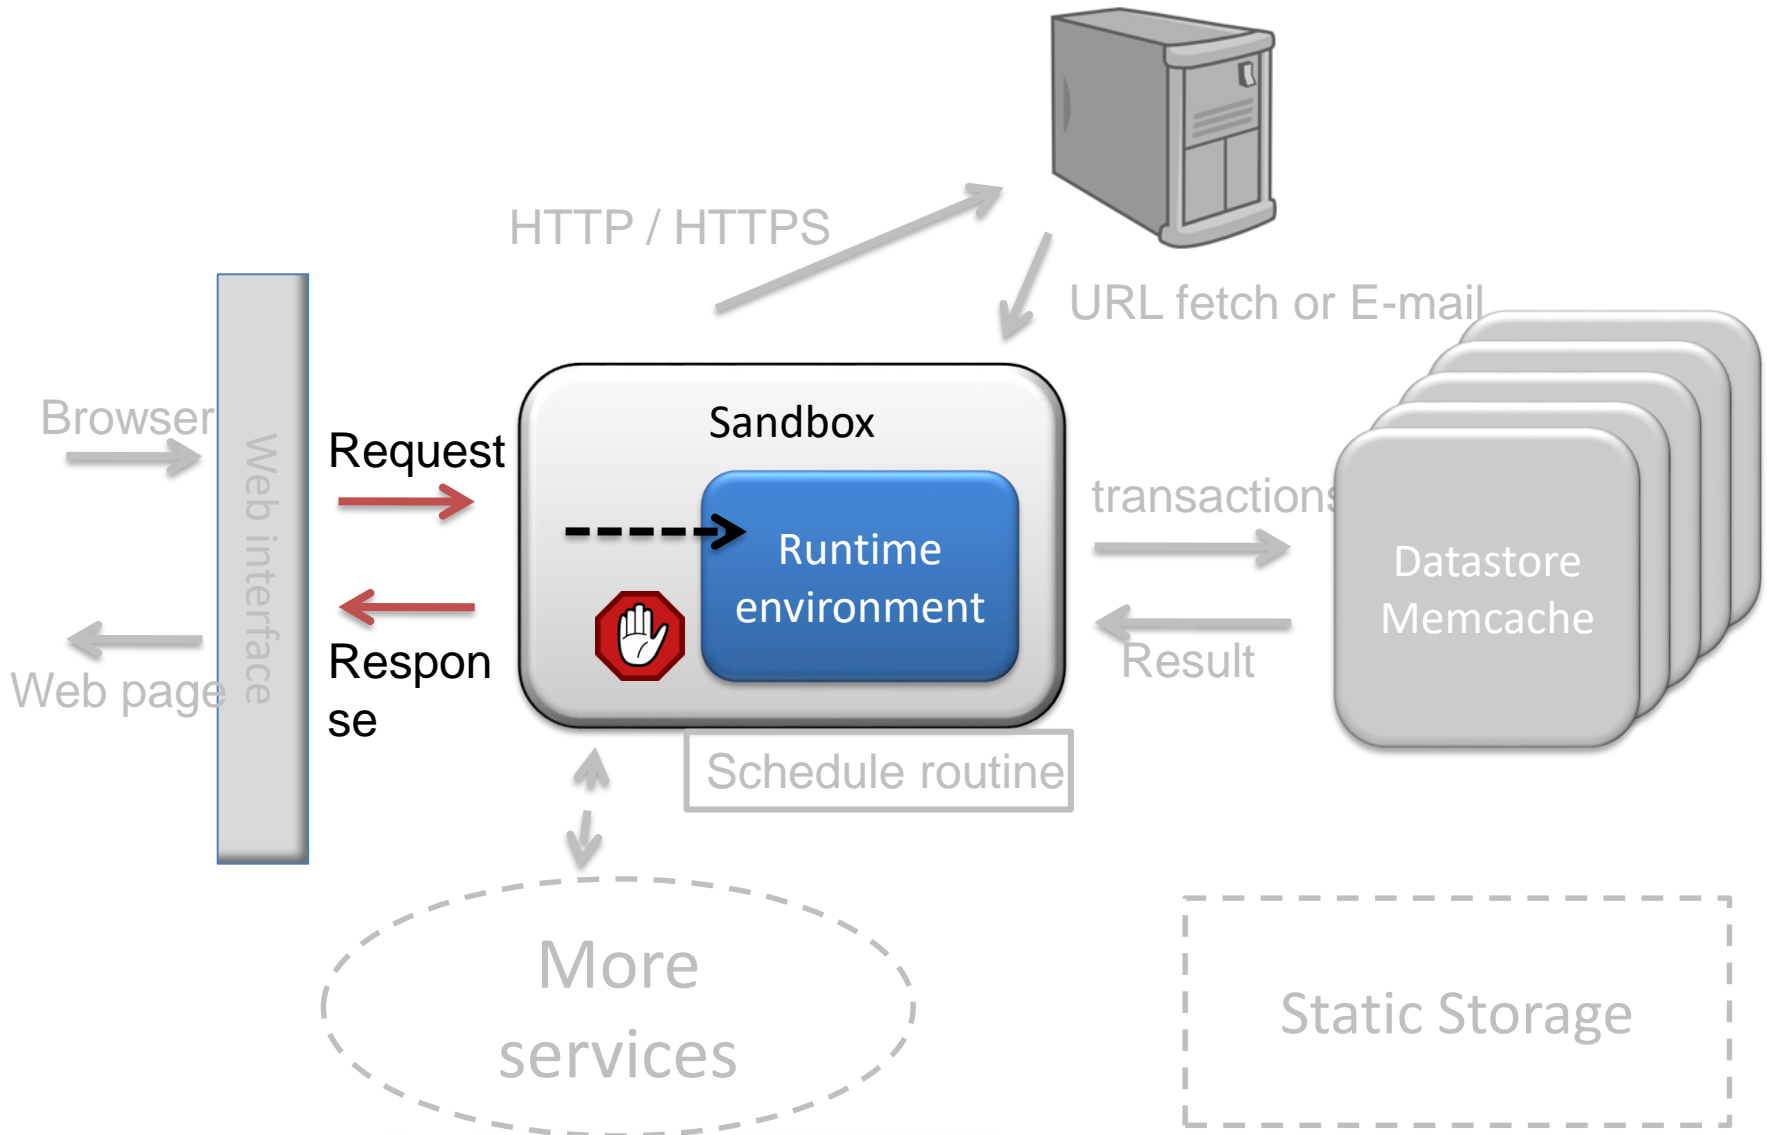
Google App Engine | Standard Environment



Google App Engine | Standard Environment



Google App Engine | Standard Environment



Google App Engine | Standard Environment



Application Standard Environment - Sandbox

- The applications run in a secure environment, and it is ensured limited access to the operating system
- It is independent of the hardware level, operating system, or physical location of the Web server
- Applications cannot write to the file system, they can only read what was loaded through the application code



Google App Engine | Standard Environment



Application Standard Environment

Mechanism:

- Upon receiving the request, AppEngine selects the server estimated to be the fastest, sends the request to the application and returns the response to the client
- Obs. The status between requests is not preserved => AppEngine ensures the same treatment for all requests, by distributing the traffic between servers
- Applications can only access their own files from the file system, and not those of other applications;
- Applications can see the environment variables set by AppEngine, but their manipulation is not necessarily persistent between requests
 - The application must use e.g. Datastore to ensure persistence between requests



Google App Engine | Standard Environment



Application Standard Environment

- The applications cannot access the network facilities at the hardware level, but it is allowed to perform network operations by using the services
- An application can access other computers only through email service or URL fetch
 - HTTPs requests using standard ports can have limitations regarding CPU usage, memory consumed per request
 - a request can be answered in up to 3600 seconds (60 min)
 - Obs. App Engine is optimized for applications that respond fast, therefore are some recommendations in the documentation regarding the timeout
 - If the application uses more processor cycles, App Engine will slow it down so as not to affect the performance of other applications running on the same machine

Google App Engine | Standard Environment



Application Standard Environment

- Obs. Python > 3 Environment ->... allows reading, writing and modifying *bytecode*
- **Python runtime**
 - A version of the Python interpreter is used
 - General mechanism: App Engine invokes a Python application using a CGI mechanism
 - Python frameworks like Django, Flask, uWSGI, Tornado work with App Engine

Google App Engine | Standard Environment



Application Standard Environment

– GAE Java runtime environment

- Java applications can be developed using the Java JVM
- Applications can be built using standard Java technologies
<https://cloud.google.com/appengine/docs/java/>

– Go Runtime

- It runs different versions of Go
- The SDK includes the Go compiler and standard libraries
- A GoAPI is provided for most App Engine services (Storage, Schedule, Communication,)
- Third-party libraries can be used, the condition is that they are implemented only in Go
- <https://cloud.google.com/appengine/docs/go/>

Google App Engine | Standard Environment



Application Standard Environment

- **PHP runtime environment**

The application is created using *sandboxed PHP environment*

- PHP runtime is the software stack responsible for installing your web service's code and its dependencies and running your service
- Frameworks: Laravel, Symfony, Slim

The App Engine application is integrated with Google Accounts for authentication

The PHP runtime contains a wrapper that allows easy access to the services offered by Google Cloud

<https://cloud.google.com/appengine/docs/standard/php-gen2/runtime>

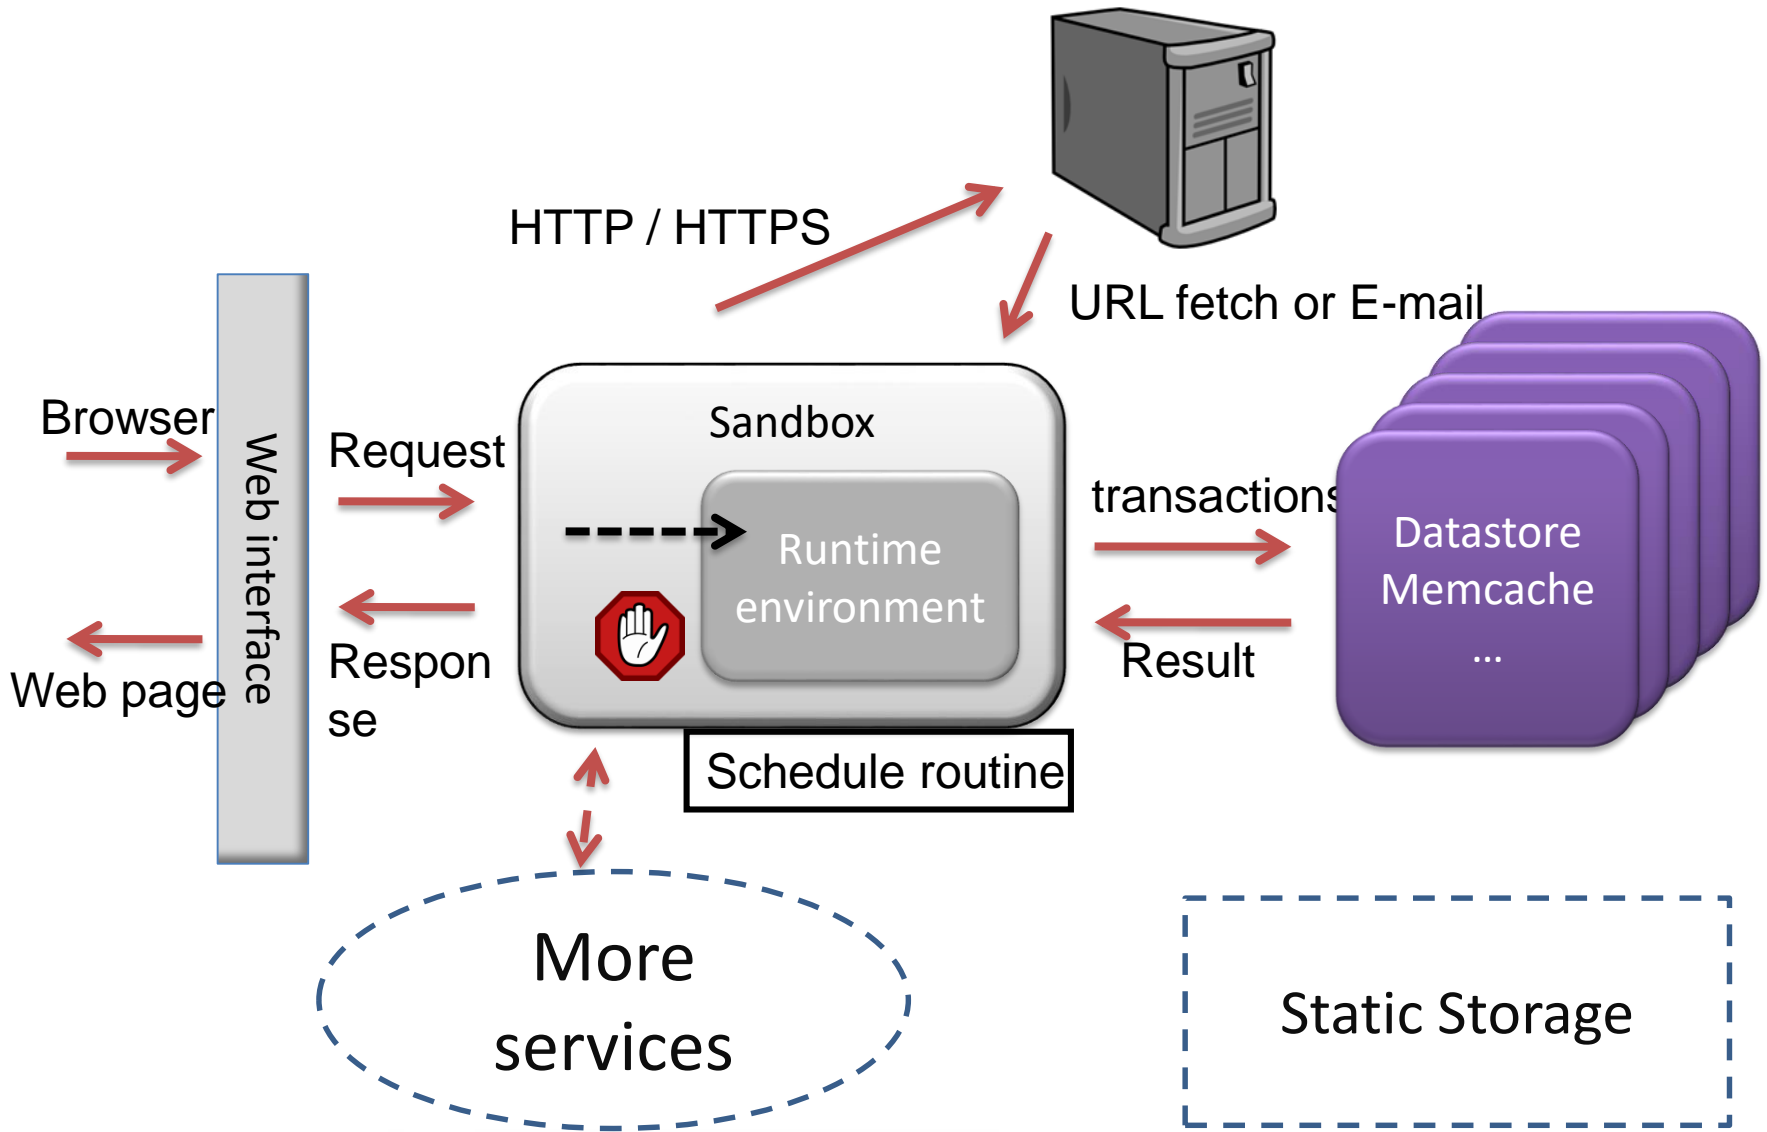
Google App Engine | Standard Environment



Application Standard Environment

- Aspects regarding the invocation of an application:
 - Mechanism: the request is routed to the server, the application is started, the response is created, the response is returned to the client
 - Each environment runs its own interpreter (JVM or Python) respecting the sandbox restrictions (e.g. trying to use any facility of the existing language or libraries, which access other resources than those allowed will lead to an exception)
 - Using a different server for each request
 - => advantage: scaling
 - => disadvantage: the operation of creating an instance for each request is time-consuming
- Solution:
- AppEngine mitigates start-up costs by keeping the application in the server's memory as much as possible and intelligently reusing servers
 - all servers have runtime environments preloaded before the request reaches the server

Google App Engine | Standard Environment



Google App Engine | Standard Environment



Storage Services

- **Static**
 - Static space (source files of Web services, configuration files, background images,...)
 - Blobstore
 - Contains files larger than 1MB (images, video or audio, etc.)
- **Dynamic**
 - Datastore
 - Service provided in the "dynamic provisioning" manner that supports dynamic data insertion/update/delete operations
 - Memcache
 - Used to increase the speed of datastore queries



Static File Servers

- Many websites have resources that do not undergo changes during normal operations on the site (e.g. images, CSS files, JavaScript code, HTML pages with static content)

They are called static files

Providing these resources does not involve application code => does not require application servers

AppEngine provides dedicated servers that provide this type of content - Static File Servers

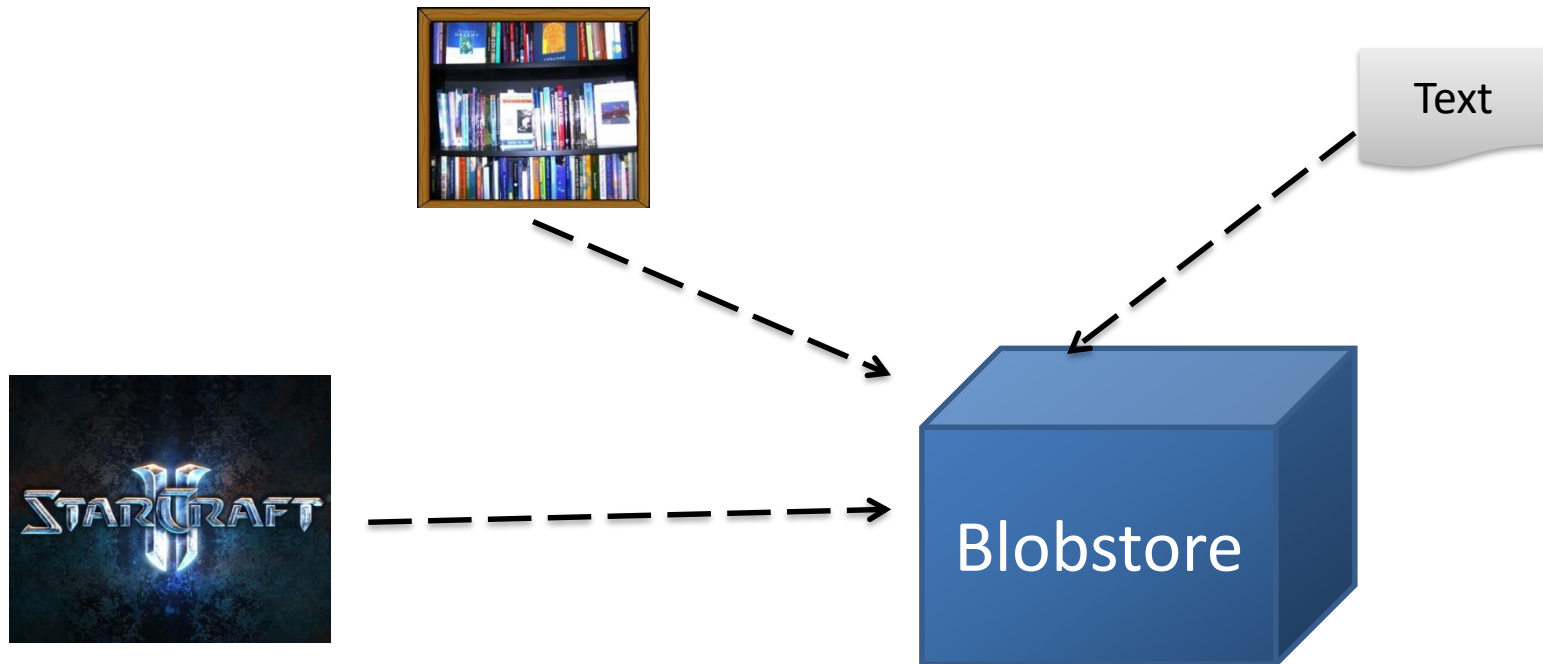
- For the end customer, such a static resource is similar to any other resource
- Developers can configure several aspects regarding how to provide static resources (URLs, content types, instructions for browsers to keep copies of files in the cache, etc.)



Storage Services

– Blobstore

- *Binary large objects* – which are larger than the size allowed for the objects in the datastore service (video, img,...)
- They are created by uploading files as a result of HTTPs requests





Storage Services

Datastore

- An AppEngine application stores data in one or more storage entities (*datastore entities* or *data objects*)
- *Entitatea* has associate *properties: name – value*

Status:= awake

Name:= Hamsa

photo



Age:= 1

Weight := 3 KG

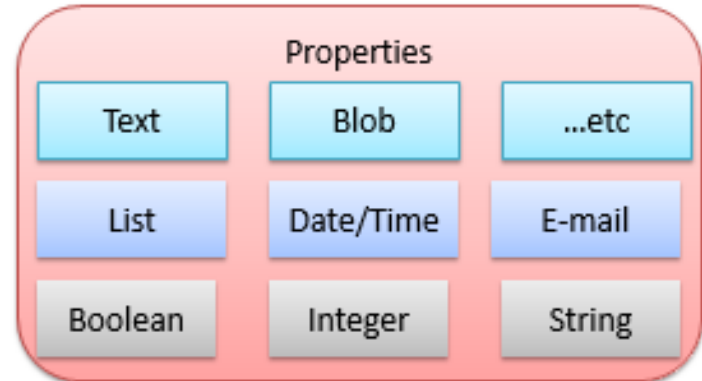
entity := Dog



Storage Services

Datastore

- Properties can have a specific type
- The constructor of the property can define
 - Name
 - Default value
 - A list of values...



```
from google.appengine.ext import db
class Dog(db.Model):
    name = db.StringProperty(default='dog')
    age = db.IntegerProperty(required=True)
    weight = db.IntegerProperty(indexed=False)
    status = db.StringProperty(choices = ['awake', 'eat', 'play'])
    photo = db.BlobProperty()
)
```



Storage Services

Datastore

- An entity has a **unique key (key)** that is provided by the application or App Engine (the developer decides)
 - This key is not an elementary property, but is an independent aspect of the entity
 - A key cannot be changed after the entity has been created
 - Knowing the key, queries can be made
- Knowing the type of the entity and the associated key can determine where the entity is stored in the entire collection of servers



Storage Services

Datastore

- Operations:
 - *put(key)* – upload or update
 - *delete(key)* – deleting an entity
 - *get(key)*, ...
 - ? Similarity to relational databases (... ? *entities* = rows from tables; ?*properties* = columns)
 - Differences:
 - An entity of a certain type can have different properties compared to another entity of the same type
 - An entity can have a property with the same name but of a different type compared to another entity of the same type
- => entities in the *datastore* are *schemaless* => increased flexibility & maintenance challenges



Interogari si indexari

- Datastore uses *indexes* for every query made by the application
- Examples of query types:
 - According to the properties of the entity, the list of entities ordered by the value of the property will be obtained
 - Filtering and sorting by keys





Queries and indexing

- Internal mechanism for query execution

Relational databases:

- the queries are planned and executed in real time

App Engine Standard Environment

- Each query has a correspondence index maintained in the *datastore*
- When the application performs a query, the datastore finds the index corresponding to that query, scans to the first row that matches the query, and returns the entity for each consecutive row in the index, repeating the process until the first row that does not match the query
- Obs. App Engine must know in advance what queries the application will make; it does not need to know a priori the values of the filters, but it must know the type of entity by which the query is made, the properties by which filtering or sorting is desired, etc.



Queries and indexing

- App Engine provides a set of indexes for simple queries based on what properties and entity types exist
 - For complex queries, an application must include specifications for indexes in its configuration
 - App Engine SDK helps to produce these configuration files (e.g. supervising the queries that the developer made during the development of the application - using the development server. When loading the application, the *datastore* will create indexes for each query of the application made during the development phase)
 - Obs. Indexes can also be configured manually
 - When the application creates new entities and updates the existing ones, the *datastore* updates each index
- => queries are fast



Transactions

- When an application has several clients that perform simultaneous read/write operations on the same data => the need to maintain data consistency
- When an application updates the properties of an entity, App Engine ensures the complete update or the update operation will give an error ⇔ the update of an entity is done through a transaction (Property: any transaction is atomic)
- When an application wants operations with several entities in a single transaction, the application will create an *entity group*, on which App Engine will supervise the execution of transactions
- App Engine uses *optimistic concurrency control* -> the application must try to complete a transaction several times before returning an error
- Reading an entity cannot fail due to concurrency, the application having access to the most stable state of the entity



Storage Services

- *Statefull services*

- **Blobstore**
- **Datastore**
- **Google Cloud SQL**
- **Memcache (memory cache)** – is a key-value storage service
 - The main advantage over the datastore: speed in storage and retrieval operations
 - It stores values in memory and not on disk
- It is not persistent: if a server fails, the memory is erased
 - It is used to cash the results of the most frequent queries or calculations

Several types are provided:

- **Shared memcache**
- **Dedicated memcache**

Google App Engine | Standard Environment



Communication Services

- ***Stateless service APIs***

- *URL Fetch*

- Allows App Engine applications to access other web resources (HTTP(S) requests for obtaining web pages, interacting with web services)
- Because *remote* servers can be slow to respond, URL Fetch supports *URL fetching* in the background, while applications can perform other operations
- Obs. Such an operation must exist only during the lifetime of the application
- The application can set a deadline for the *fetch* operation => the call will end if the answer is not received in the appropriate time

- *Mail/SendGrid/Mailgun*

- Applications can send and receive messages in the form of HTTP requests initiated by App Engine and sent to the application
- Example: notifying users, confirming user actions, validating contact information, etc.

Google App Engine | Standard Environment



Communication Services

- ***Stateless service APIs***

- *Sockets*

- Their behavior varies depending on the driving environment

- *Twilio*

- This is not a Google service but an independent cloud communications platform that provides APIs for voice and messaging services, including phone and VOIP calls. Twilio's services can be integrated into applications, including those hosted on Google App Engine, to enable such functionalities.

- *Twilio Client* allows making VOIP calls from any device (supports WebRTC)

- <https://www.twilio.com/>

- *Google Cloud Endpoints*

- It consists of tools and libraries that allow the generation of APIs based on an application, i.e. to facilitate access to data from other applications

Google App Engine | Standard Environment



Process Management Services

– *Task Queues*

The idea: a web application must provide the result as quickly as possible

- The problem: sometimes there are many operations to be done, which require more time, and which can be organized into tasks
- The solution: *task queues*
 - It allows the definition of tasks that can be executed in the background when the system resources allow this
- Types:
 - *Push Queues* - ensures that these tasks are automatically executed by the system that scales and performs the processing according to needs; *automate scaling* – the services must be completed in a maximum of 10 min; *manual scaling* – the services can run 24h
 - *Pull Queues*
 - » it offers better control over when tasks are executed, in a certain time interval, but requires increased effort in process management
 - » ensures the integration of the application with the *Task Queue REST API* - which allows the external management of existing tasks

Google App Engine | Standard Environment



Process Management Services

- *Migration Pull Queues to Pub/Sub*
- *It can be used Pub/Sub as a Task Queues pull queue: Subscriptions to a topic do not expire and can exist simultaneously for multiple workers. This means that a message can be processed by more than one worker, which is one of the primary use cases for Pub/Sub. To recreate your Task Queues pull queues as Pub/Sub pull subscriptions, create a topic for each worker and subscribe only the associated worker to the topic. This ensures that each message is processed by exactly one worker as in Task Queues.*

Task Queues workflow

You create the pull queue

You create and enqueue the task

The worker leases the task

The worker processes the task

The worker deletes the task from the queue

The lease expires

Pub/Sub workflow

You create the topic and subscribe your subscriber (i.e. worker) to the topic

You create the message and publish it to the topic

The subscriber pulls the message from the topic

The subscriber processes the message

The subscriber acknowledges the message

The topic deletes the message when all of its subscribers have acknowledged the message

Google App Engine | Standard Environment



Process Management Services

- *Migrarea Push Queues to Cloud Tasks*
- In Cloud Tasks, all queues operate as push queues
- Cloud Tasks is a fully managed service that allows you to manage the execution, dispatch and delivery of a large number of distributed tasks.
- You can asynchronously perform work outside of a user request.
- Your tasks can be executed on App Engine or any arbitrary HTTP endpoint.

[<https://cloud.google.com/appengine/docs/standard/python/taskqueue/push/migrating-push-queues>
<https://cloud.google.com/tasks/docs>]

Pricing

Cloud Tasks has its own [pricing](#). As with Task Queues, sending requests to your App Engine app with a task can cause your app to incur costs.

Quotas

The Cloud Tasks [quotas](#) are different from the quotas for Task Queues. Like with Task Queues, sending requests to your App Engine app from Cloud Tasks might impact your App Engine [request quotas](#).

Google App Engine | Standard Environment



Process Management Services

- *Scheduled task (sau cron jobs)*
 - Allows the execution of tasks at fixed time intervals
 - Useful for periodic maintenance (update of some data from the cache every 10 minutes, ...) or sending notification messages every day
 - Free applications can have up to 20 such processes (100 in other cases)

Computation Services

- *Images*
 - It allows performing operations on images (resizing, cropping, flipping, ...)
- *MapReduce*
 - programming model for processing large amounts of data
 - It uses Datastore and TaskQueues

Google App Engine | Standard Environment




How we are thinking the application.... Microservices?

A monolithic application puts all its functionality into a single process...




... and scales by replicating the monolith on multiple servers



A microservices architecture puts each element of functionality into a separate service...



... and scales by distributing these services across servers, replicating as needed.



Google App Engine | Standard Environment

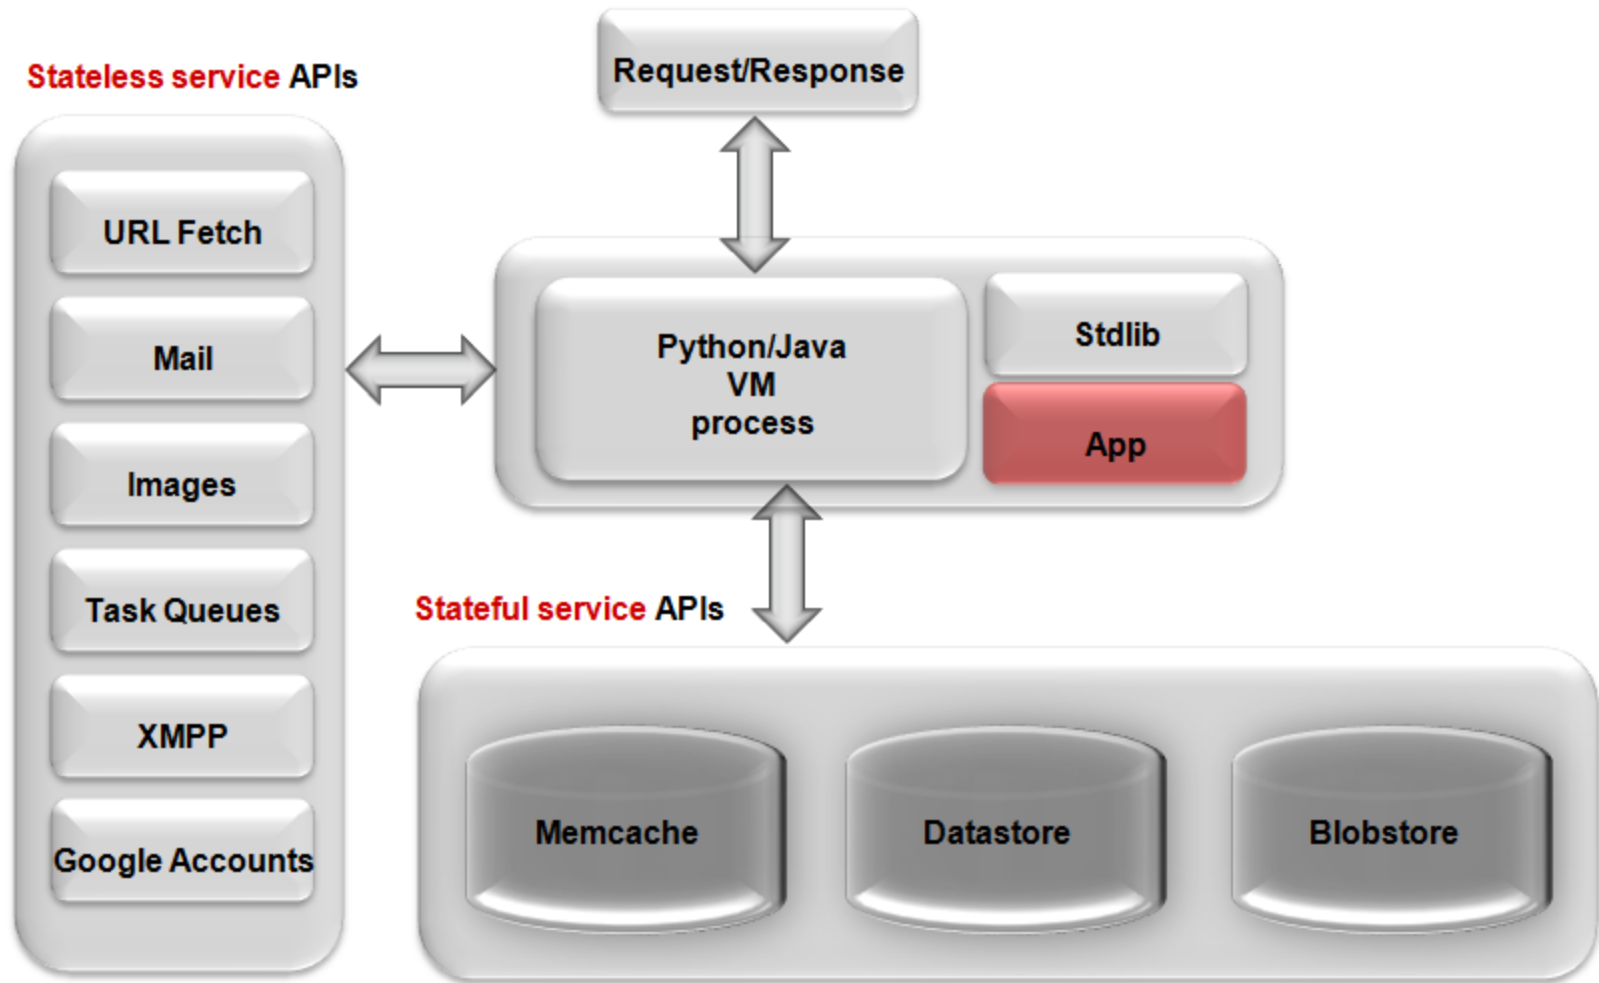


Figure. Possible way to use GAE services

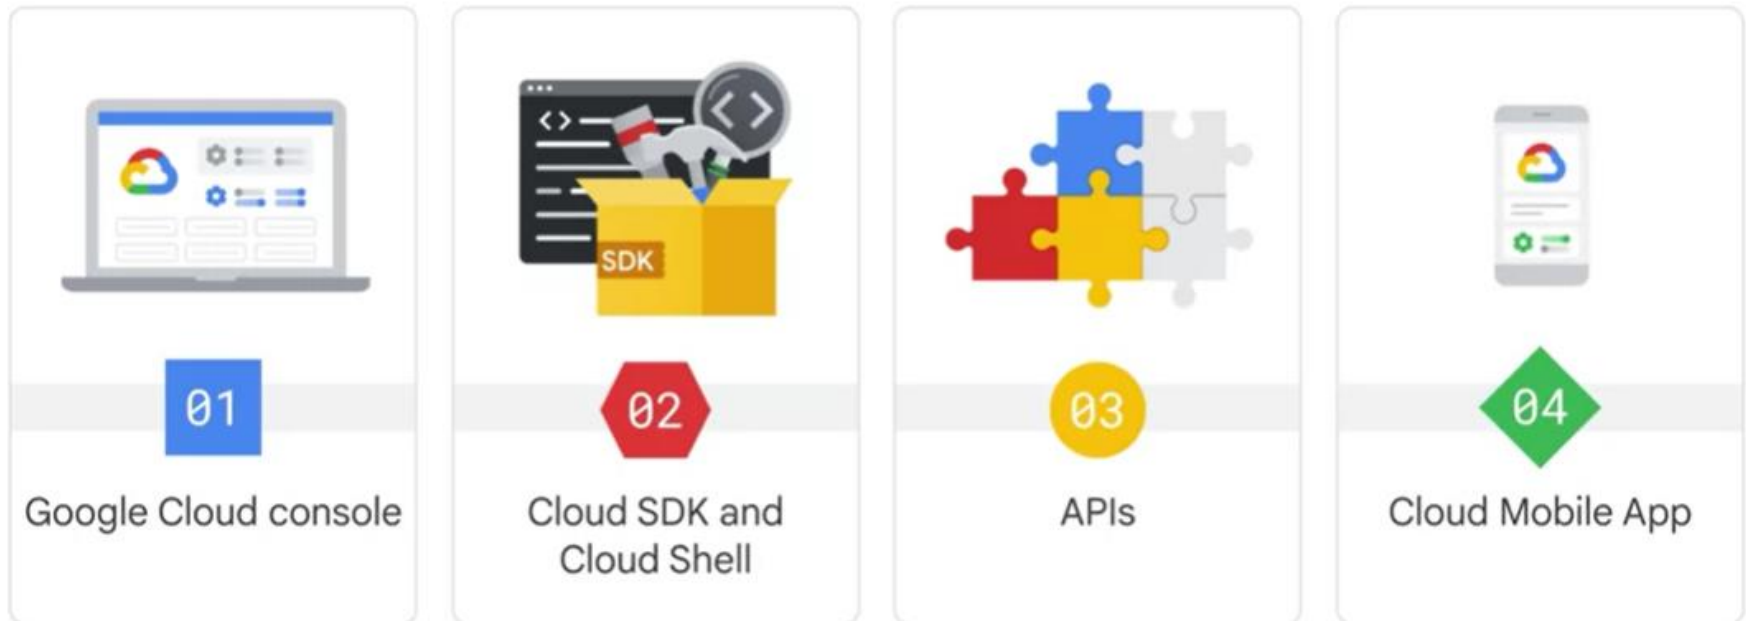


Conclusions

- allows the running of web, mobile applications, ...
- quick setup
- scalability
- security
- Everything is designed to scale
 - *Low-usage apps*: multiple applications per physical host
 - *High-usage apps*: multiple physical hosts per application
- Portability
 - **Kubernetes & Google Kubernetes Engine** – give the ability to mix and match microservices running across different clouds
 - **Operation Suite** – let customers monitor workloads across multiple cloud providers (see next Course)

Google Cloud

Interacting with Google Cloud:



Google App Engine



Costs:

- <http://code.google.com/appengine/docs/whatisgoogleappengine.html>

“Not only is creating an App Engine application easy, it's free! You can create an account and publish an application that people can use right away at no charge, and with no obligation. An application on a free account can use up to 1 GB of storage and up to 5 million page views a month. When you are ready for more, you can enable billing, set a maximum daily budget, and allocate your budget for each resource according to your needs. You can register up to 10 applications per developer account. ...”

<http://code.google.com/appengine/docs/billing.html>

Example: Costs for operations in datastore

March 2013

Operation	Cost
Write	\$0.10 per 100k operations
Read	\$0.07 per 100k operations
Small	\$0.01 per 100k operations

March 2014

Operation	Cost
Write	\$0.09 per 100k operations
Read	\$0.06 per 100k operations
Small	\$0.01 per 100k operations

March 2015

Operation	Cost
Read / Write	\$0.06 per 100,000 operations
Small	Free

Google App Engine



Resource	Free quota per day	Unit	Price beyond the free quota per unit
Stored data	1 GB	per GB per month	\$0.18
Entity reads	50,000	per 100K entities	\$0.06
Entity writes	20,000	per 100K entities	\$0.18
Entity deletes	20,000	per 100K entities	\$0.02
Small operations	Unlimited. Includes calls to allocate Cloud Datastore IDs, keys-only queries, and projection queries that do not use the distinct on clause. A keys-only query or a projection query that does not use the distinct on clause is counted as a single entity read for the query itself. The individual results are counted as small operations.		

[2017 - <https://cloud.google.com/appengine/pricing>]

Google App Engine



lowa (us-central1) ▾

Instance class	Cost per hour per instance
B1	\$0.05
B2	\$0.10
B4	\$0.20
B4_1G	\$0.30
B8	\$0.40
F1	\$0.05
F2	\$0.10
F4	\$0.20
F4_1G	\$0.30

Instance Class	Memory Limit	CPU Limit	Supported Scaling Types
F1 (default)	128 MB	600 MHz	automatic
F2	256 MB	1.2 GHz	automatic
F4	512 MB	2.4 GHz	automatic
F4_1G	1024 MB	2.4 GHz	automatic
B1	128 MB	600 MHz	manual, basic
B2 (default)	256 MB	1.2 GHz	manual, basic
B4	512 MB	2.4 GHz	manual, basic
B4_1G	1024 MB	2.4 GHz	manual, basic
B8	1024 MB	4.8 GHz	manual, basic

[2019 - <https://cloud.google.com/appengine/pricing>]

Google - Flexible Environment



Applications running in the [App Engine flexible environment](#) are deployed to virtual machine types that you specify. These virtual machine resources are billed on a per-second basis with a 1 minute minimum usage cost.

This table summarizes the hourly billing rates of the various computing resources:

Iowa (us-central1) ▾

Resource	Unit	Unit cost
vCPU	per core hour	\$0.0526
Memory	per GB hour	\$0.0071
Persistent disk	per GB per month	\$0.0400

[2019 - <https://cloud.google.com/appengine/pricing>]

Google - Flexible Environment



App Engine flexible environment pricing

App Engine does not provide free tier in the flexible environment.

Apps running in the flexible environment are deployed to virtual machine types that you specify. These virtual machine resources are billed on a per-second basis with a 1 minute minimum usage cost.

Billing for the memory resource includes the memory your app uses plus the memory that the runtime itself needs to run your app. This means your memory usage and costs can be higher than the [maximum memory you request for your app](#).

The following table summarizes the hourly billing rates of the various computing resources in the flexible environment.

Iowa (us-central1) ▾

Resource	Unit	Unit cost
vCPU	per core hour	\$0.0526
Memory	per GB hour	\$0.0071
Persistent disk	Priced as Compute Engine persistent disk , which is called "Storage PD Capacity" on your bill.	
Outgoing network traffic	Priced as Compute Engine internet egress .	
Incoming network traffic	Gigabytes	Free

[2022 - <https://cloud.google.com/appengine/pricing>]

Google - Flexible Environment



Google Cloud Datastore calls

Cloud Datastore operations are billed as follows:

Iowa (us-central1) ▾

Resource	Free quota per day	Unit	Price beyond the free quota per unit
Stored data	1 GB	per GB per month	\$0.18
Entity reads	50,000	per 100K entities	\$0.06
Entity writes	20,000	per 100K entities	\$0.18
Entity deletes	20,000	per 100K entities	\$0.02
Small operations	Unlimited. Includes calls to allocate Cloud Datastore IDs, keys-only queries, a that do not use the distinct on clause. A keys-only query or a projection query distinct on clause is counted as a single entity read for the query itself. The is counted as small operations.		

Search

Fees for use of the Search API are listed in the table below. Refer to the [Java](#) and [Python](#) documentation for a detailed description of each type of Search call.

Iowa (us-central1) ▾

Resource	Unit	Price per unit
Total storage (documents and indexes)	per GB per month	\$0.18
Queries	per 10K queries	\$0.50
Indexing searchable documents	per GB	\$2.00

[2019 - <https://cloud.google.com/appengine/pricing>]

Google - Flexible Environment



Iowa (us-central1) ▾

Resource	Unit	Unit cost (in US \$)
Outgoing network traffic - standard environment*	Gigabytes	\$0.12
Outgoing network traffic - flexible environment	Gigabytes	Google Compute Engine Network Rates
Incoming network traffic	Gigabytes	Free
Blobstore stored data**	Gigabytes per month	\$0.026
Dedicated memcache	Gigabytes per hour	\$0.06
Logs API	Gigabytes	\$0.12
Sending email, shared memcache, cron, APIs (Task Queues, Image, Files, Users)		No Additional Charge

[2019 - <https://cloud.google.com/appengine/pricing>]

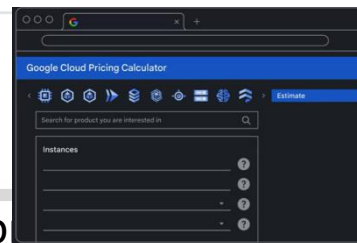
Google - Flexible Environment



Iowa (us-central1) ▾

Resource	Unit	Unit cost (in US \$)
Blobstore stored data*	Gigabytes per month	\$0.026
Dedicated memcache	Gigabytes per hour	\$0.06
Logs API	Gigabytes	\$0.12
Search API*: Total storage (documents and indexes)	per GB per month	\$0.18
Search API: Queries	per 10K queries	\$0.50
Search API: Indexing searchable documents	per GB	\$2.00
Sending email, shared memcache, cron, APIs (Task Queues, Image, Files, Users)		No Additional Charge

[2022 - <https://cloud.google.com/appengine/pricing>]



Future discussions...



First wave
Colocation



Second wave
Virtualized
data center



Third wave
Container-based
architecture

[<https://cloud.google.com>]



Bibliography

- <https://cloud.google.com/>
- Mark C. Chu-Carroll, Code in the Cloud, Programming Google App Engine, 2011
- Dan Sanderson, Programming Google App Engine, O'Reilly, 2010
- Implementing and Developing Cloud Computing Applications, DAVID E.Y. SARNA, CRC Press, Taylor&Francis Group, 2011
- Cloud Computing, Software Engineering Fundamentals, J. Heinzlreiter, W. Kurschl, www.fh-hagenberg.at
- [Rossum, 2008] Guido van Rossum, Google App Engine, Stanford EE380 Colloquium, Nov 5, 2008
- <http://code.google.com/appengine/>
- <http://www.python.org/download/>
- <https://www.ascamso.com/905-2/>
- <https://www.martinfowler.com/articles/microservices.html>
- Compute Engine: <https://www.youtube.com/watch?v=YPWv-Ln9IPw&t=178s>
- Cloud Storage: <https://www.youtube.com/watch?v=wNOs3LlsH6k>

Summary



- Google in Cloud
 - ...services
 - Tools (GWT, GAS,...)
 - Features
 - Architectural aspects
 - Conclusions

Google Cloud

Questions?

