



Google Cloud Platform Solutions for DevOps Engineers



Márton Kodok / @martonkodok
Google Developer Expert at **REEA.net** - Targu Mures

March 2019 - Vilnius, Lithuania

About me



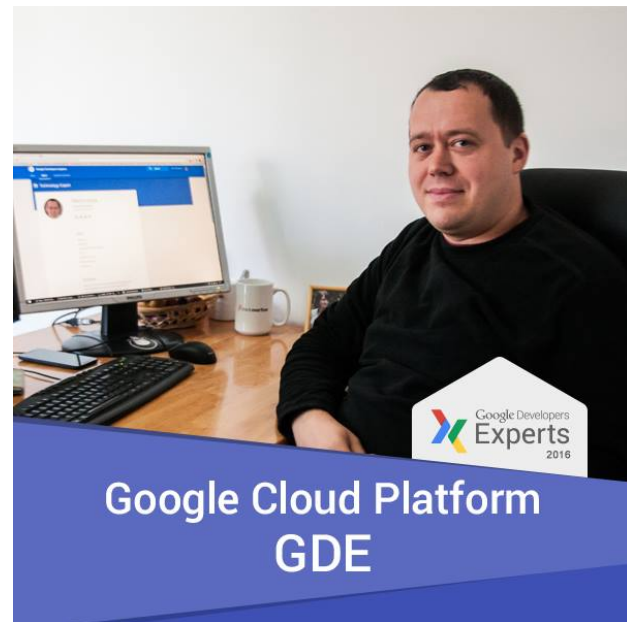
- Geek. Hiker. Do-er.
- Among the Top3 romanians on Stackoverflow 130k reputation
- Google Developer Expert on Cloud technologies
- Crafting Web/Mobile backends at **REEA.net**
- BigQuery/Redis and database engine expert
- Active in mentoring and IT community

StackOverflow: **pentium10**

GitHub: **pentium10**

Slideshare: **martonkodok**

Twitter: **@martonkodok**




Google Cloud Platform
GDE



1. Application development in the Cloud
2. **App Engine**: Scale your apps seamlessly from zero to planet scale
3. Cloud **Functions**: your gateway to GCP Services
4. Google **Stackdriver**: Metrics, logging, alerting are a universal right!
5. **BigQuery**: federated data access warehouse
6. Doing **hybrid cloud** mixing on premise with cloud
7. Practical use cases
8. Qwiklabs

What's so hard about traditional app development?



Build on the same infrastructure
that powers Google



Google Cloud Platform

Google Cloud Platform (GCP)



Compute



Compute Engine



App Engine



Kubernetes Engine



GPU



Cloud Functions



Container-Optimized OS

Big Data



BigQuery



Cloud Dataflow



Cloud Dataproc



Cloud Dataprep



Cloud Datalab



Cloud Pub/Sub



Genomics



Data Studio

Identity & Security



Cloud IAM



Cloud Resource Manager



Cloud Security Scanner



Key Management Service



BeyondCorp



Data Loss Prevention API



Identity-Aware Proxy



Security Key Enforcement

Internet of Things



Cloud IoT Core

Machine Learning



Cloud Machine Learning



Cloud Vision API



Cloud Speech API



Cloud Video Intelligence API



Cloud Natural Language API



Cloud Translation API



Cloud Jobs API



Advanced Solutions Lab

Storage & Databases



Cloud Storage



Cloud Bigtable



Cloud Datastore



Transfer Appliance



Cloud SQL



Cloud Spanner



Persistent Disk

Google Cloud Platform (GCP)



Management Tools



Stackdriver



Monitoring



Logging



Error Reporting



Trace



Debugger



Cloud Deployment Manager



Cloud APIs



Cloud Console



Cloud Shell



Cloud Mobile App



Cloud Billing API



Profiler

Networking



Virtual Private Cloud



Cloud Load Balancing



Cloud CDN



Cloud External IP Addresses



Cloud Firewall Rules



Cloud Router



Cloud Interconnect



Cloud DNS



Cloud Network



Cloud Routes



Cloud VPN



Dedicated Interconnect

Developer Tools



Cloud SDK



Cloud Deployment Manager



Cloud Source Repositories



Cloud Tools for Android Studio



Cloud Tools for IntelliJ



Container Builder



Cloud Tools for PowerShell



Cloud Tools for Visual Studio



Container Registry



Google Plug-in for Eclipse



Cloud Test Lab

Google sees serverless as



Programming model

- ✓ Focus on code
- ✓ Event-driven
- ✓ Stateless

Dev

Operational model

- ✓ Zero ops
- ✓ Automatic scaling
- ✓ Managed security

Ops

Billing model

- ✓ Pay for usage

\$



Serverless is more than a set of functions



Cloud Dataflow



BigQuery



Stackdriver



Cloud Functions



App Engine



Cloud Storage



Cloud PubSub



Cloud Tasks



Meet Serverless



Automatic scaling




No server management



Only pay for what you use



Event-driven



serverless data center depicted



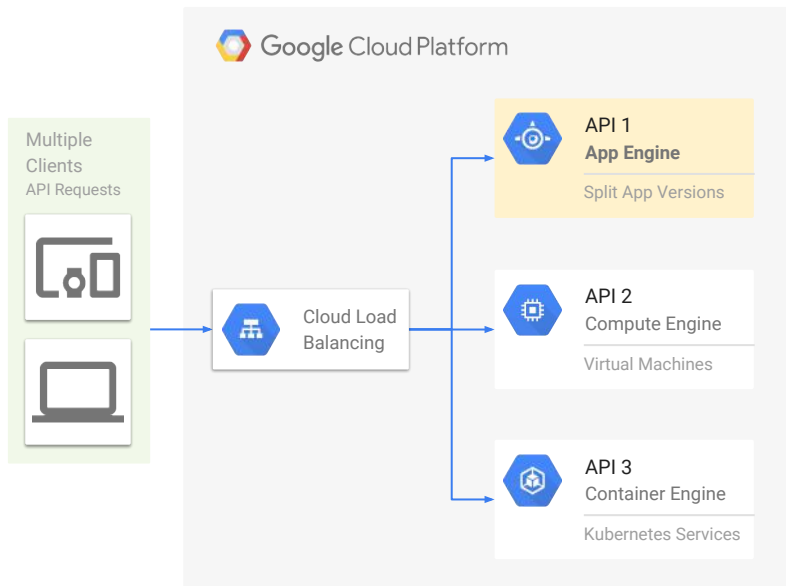
Serverless is about maximizing **elasticity**, **cost savings**, and **agility** of cloud computing.



Platforms

Triggered Code

App Engine - managed application platform



App Engine

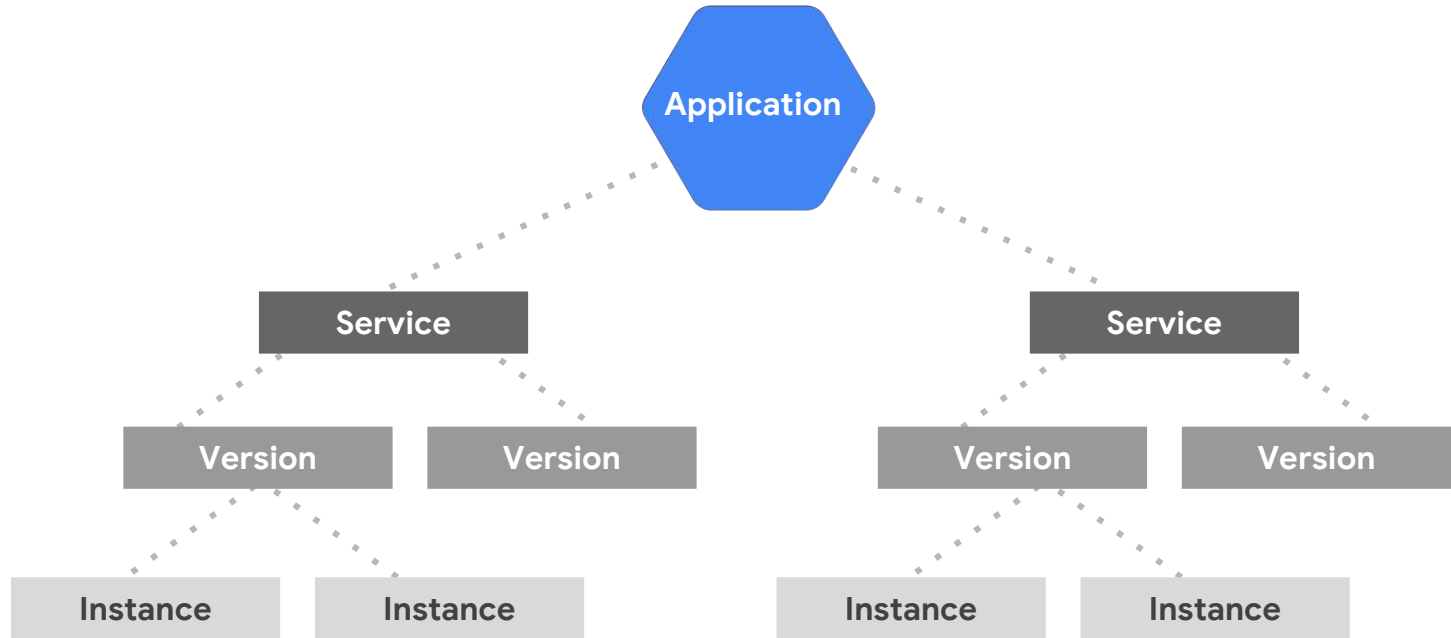
- Scale your applications seamlessly from zero to planet scale
- Automatically scales depending on your application traffic
- Traffic Splitting (app versions, A/B tests, incremental rollouts)

Best used:

- For HTTP services
- For existing applications



App Engine: Services and versions





App Engine: Traffic splitting

A/B testing and
Canary releasing
with a few clicks or
a single gcloud
command

App Engine

← Split traffic

You can split incoming traffic to different versions of your app. Traffic splitting is useful for slowly rolling out new versions or A/B testing different designs and features [Learn more](#)

Split traffic by

- IP address
- Cookie
- Random

Traffic allocation

3	will receive the remaining	10 %	×
2		90 %	×

+ Add version

Save Cancel



App Engine: Runtimes



Async message processing

Asynchronous task execution



Cloud Tasks

Message queue system
Guaranteed at-least-once delivery
Future Scheduling
Will only be dispatched once on dups

Best used:

- For HTTP services
- For App Engine handlers



Cloud PubSub

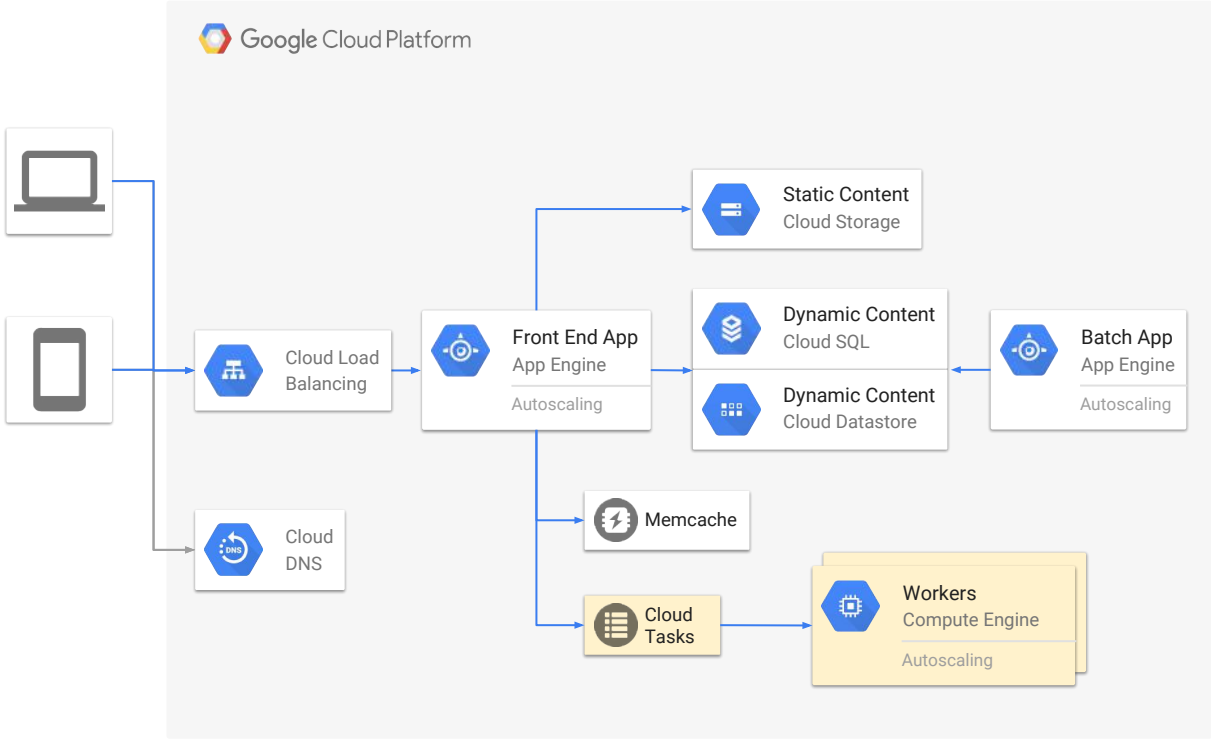
Deliver event data based on topics
Exactly-once processing
Build multi-cloud on premise, hybrid apps
Cross zone message replication

Best used:

- For large-scale ingestion of events, streams
- Topics, publish/subscribe patterns, IoT



Cloud Tasks - Message Queue systems



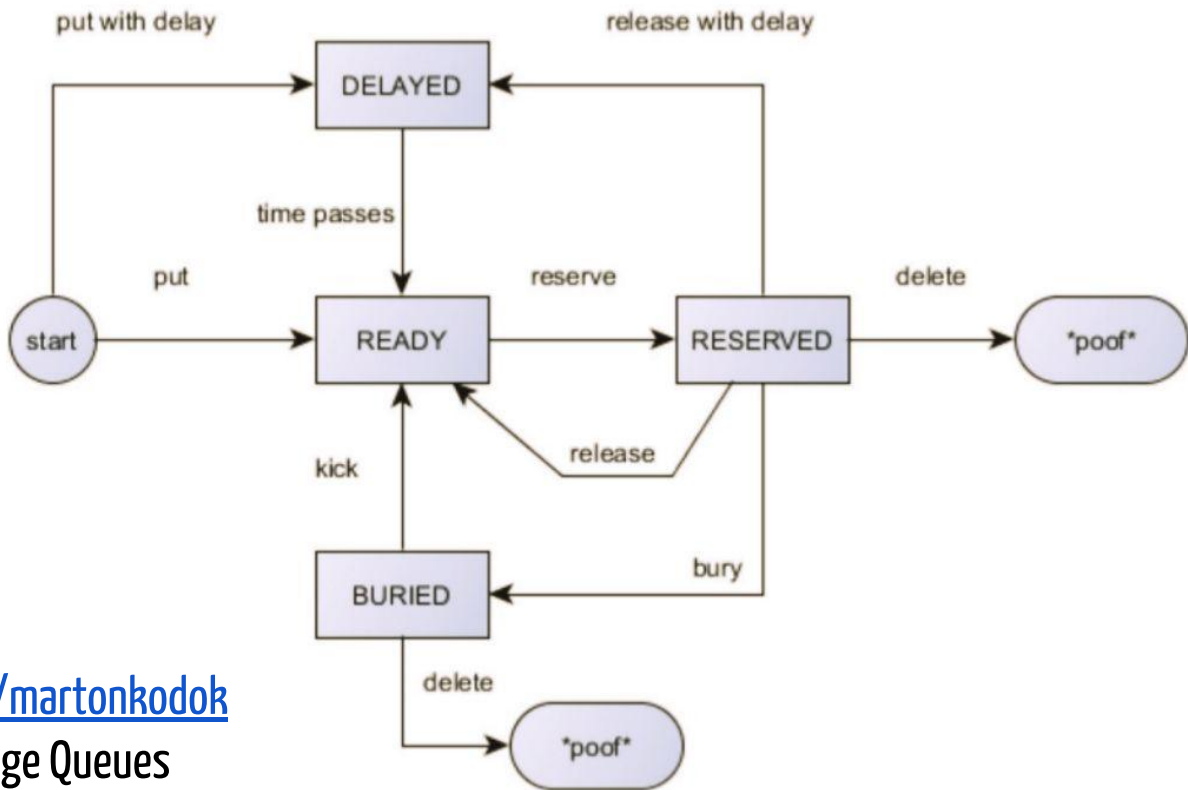
Task Queues

Push queues

- Put with delays

Pull queues

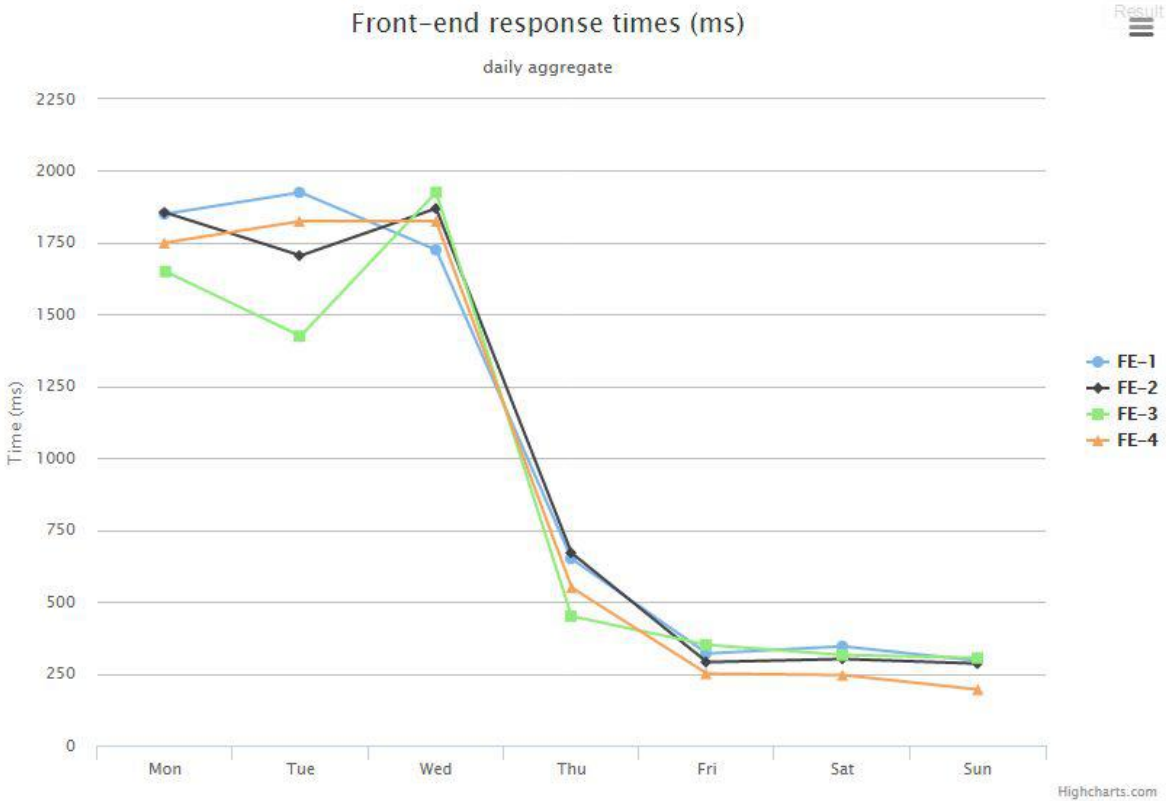
- Ability to “tag”
- Lease multiple eg: gameboard updated, game id as tag.



Slides: <https://www.slideshare.net/martonkodok>

Title: Architectural Patterns - Message Queues

Reduce request latency



Highcharts.com

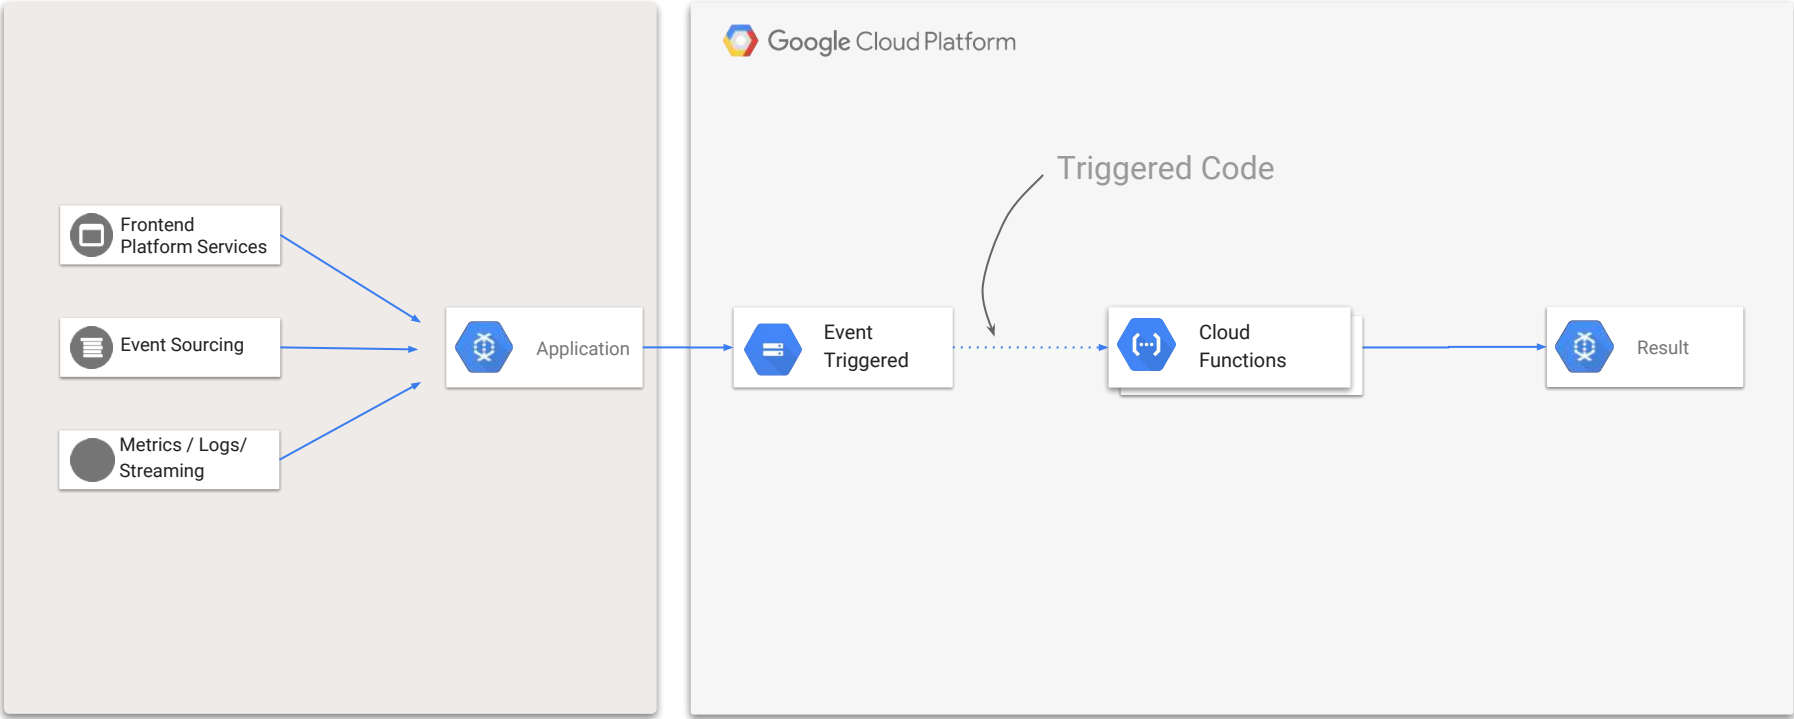




Cloud Functions



Cloud Functions





Unit: Function

Trigger: Events and HTTP

Best used:

- For Events & Async workloads
- For single-purpose microservices

Cloud Functions - Event-driven - FaaS

- Simplest way to run your code in the cloud - just deploy
- 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


(In alpha: Java)



- Node 8.11
- Support for async/await




- Python 3.7.1
- Flask microframework



- Go 1.11
- Familiar building blocks




Functions: your gateway to GCP Services



Services



Applications



APIs



Access 20+ Google services from GCF



BigQuery



Cloud Pub/Sub



Cloud Storage



Cloud Bigtable



Cloud Datastore



Cloud Machine Learning



Cloud Vision API



Cloud Speech API



Cloud Spanner



Cloud Natural Language API



Cloud Translation API



Stackdriver



Firebase



Assistant



Dialogflow



Cloud Functions Trigger



Cloud Storage

Finalize/Create

Delete

Archive

Metadata Update

Compute Engine EAP

Instance

Instance Group

Autoscaler

Firewall

Set Labels

Cloud Pub/Sub

Topic



Cloud Functions Trigger



Firestore

Realtime DB change

Authentication

Remote config

Google Analytics log

BigQuery

Job create

Job complete

Cloud Storage

Create

Update

Write

Delete



Summary: When to use what



Cloud Functions

Smallest unit of computing

Event driven architecture

Connect & extend services

App Engine

Unit of computing is apps

HTTP request/response

Large scalable backends

Serverless add-on

Run functions, apps & containers on GKE

Full portability of your artifacts

Run on your own cluster



The unit of computing is yours to chose!



Slice and dice your application any way you think it makes sense to be **more agile**, to build **better apps** that **deliver value**.





How I Learned How to **Stop Worrying** and Dig Hosted Monitoring



What to do for monitoring?



- In-house or hosted?
- Modern vs proven?
- Specialized or general?
- Paid vs free?

Let's be honest:

- Not enough time
- Small environment
- More important work to do
- Everyone else is doing it





Intelligent monitoring and management for services, containers, applications, and infrastructure.

- Works with GCP, AWS, on prem
- Logging, tracing, alerts
- Collectd agent
- Cost per resource
- Easy point and click alerting policy
- Debugger, Error reporting, profiler *free
- Uptime checks *free

Google Stackdriver: Uptime check (offered free)




- Check Type: HTTP, TCP
- Resource Type: URL, Instance, App Engine, Load Balancer
- Global Locations
- Custom Headers (encrypted)
- Authentication
- Whitelisted source IPs
- Alerting policy: Email, SMS, Slack, PagerDuty, Hipchat, Campfire, Webhooks

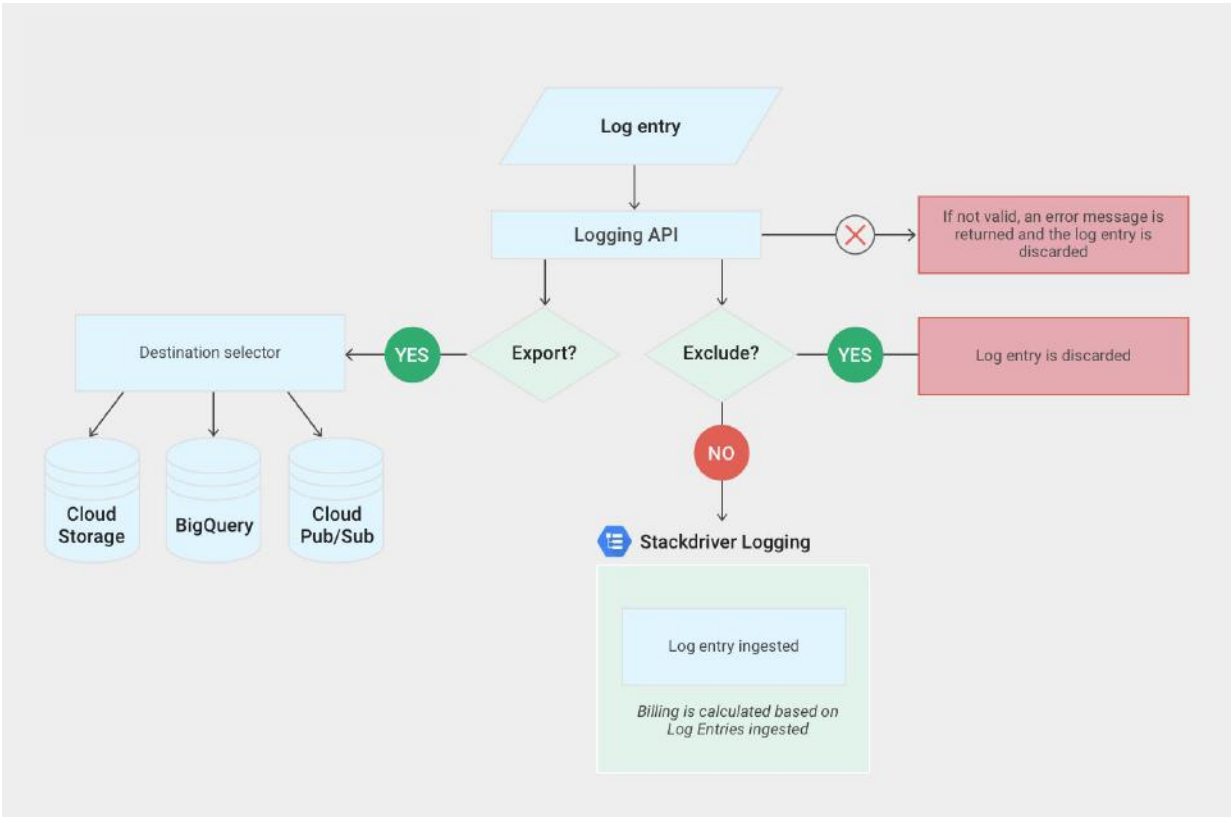
CHECKS	VIRGINIA	OREGON	IOWA	BELGIUM	SINGAPORE	SAO PAULO	POLICIES	ACTIONS
UptimeCheck for Licitatia								



Google Stackdriver: Uptime check



Life of a log





Metrics, logging, alerting are **a universal right!**

The Google BigQuery logo is a blue hexagon with a white magnifying glass icon inside, positioned to the left of the text 'Google BigQuery'.

Google BigQuery

What is BigQuery?



Analytics-as-a-Service - Data Warehouse in the Cloud

Scales into Petabytes on Managed Infrastructure - load up to 5TB large files

SQL 2011 + Javascript UDF (User Defined Functions)

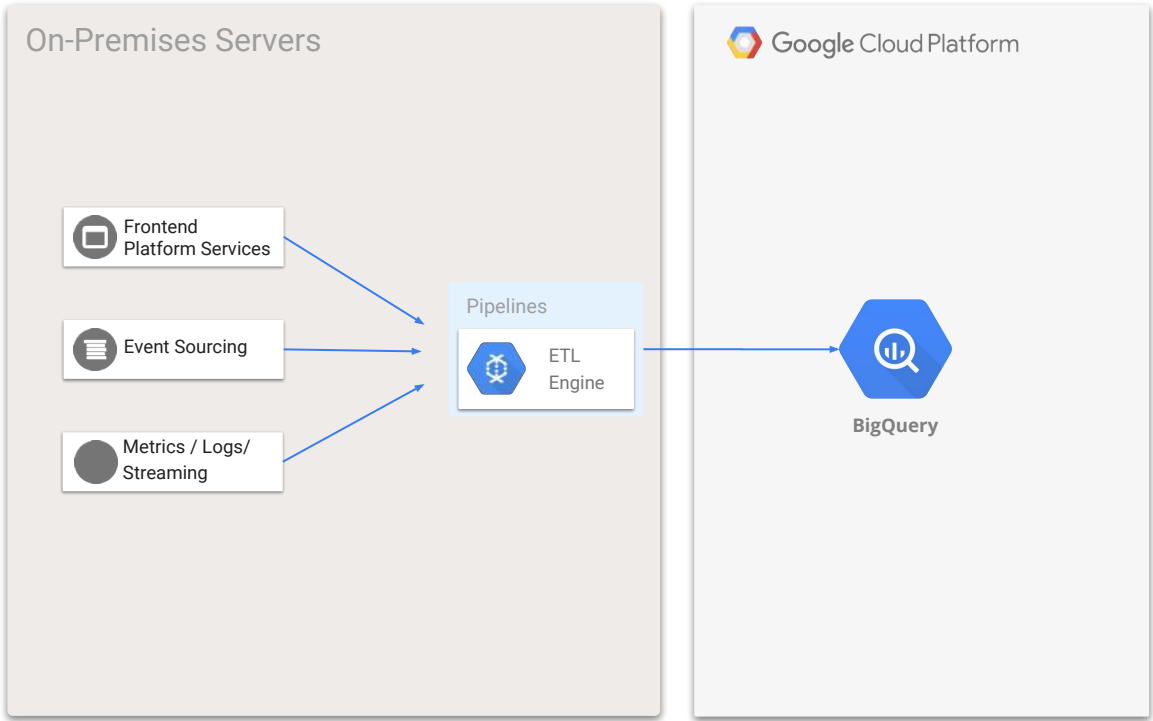
Familiar DB Structure (table, columns, views, struct, nested, JSON)

Integrates with Google Sheets + Cloud Storage + Pub/Sub connectors

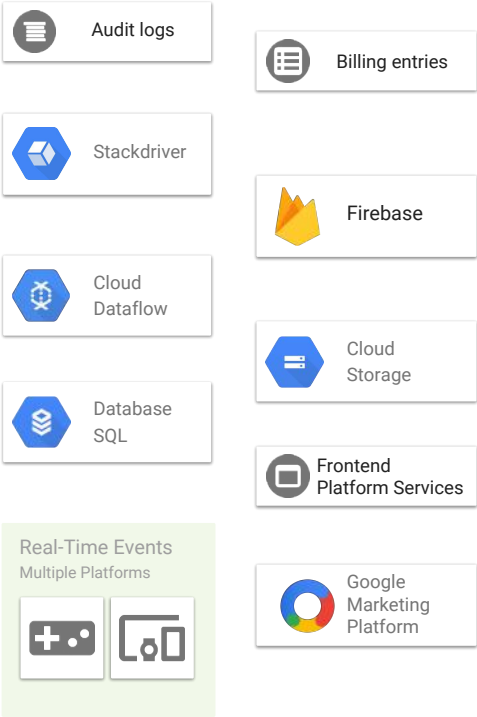
Decent pricing (storage: \$20/TB cold: \$10/TB, queries \$5/TB) *March 2019



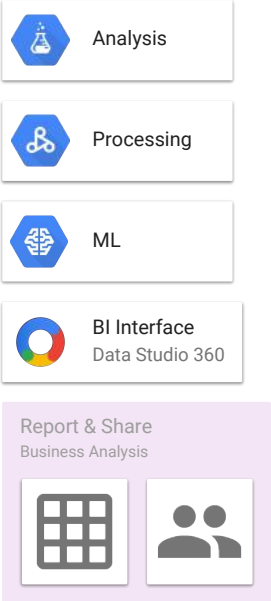
Architecting for The Cloud



BigQuery: federated data access warehouse



Application & Presentation

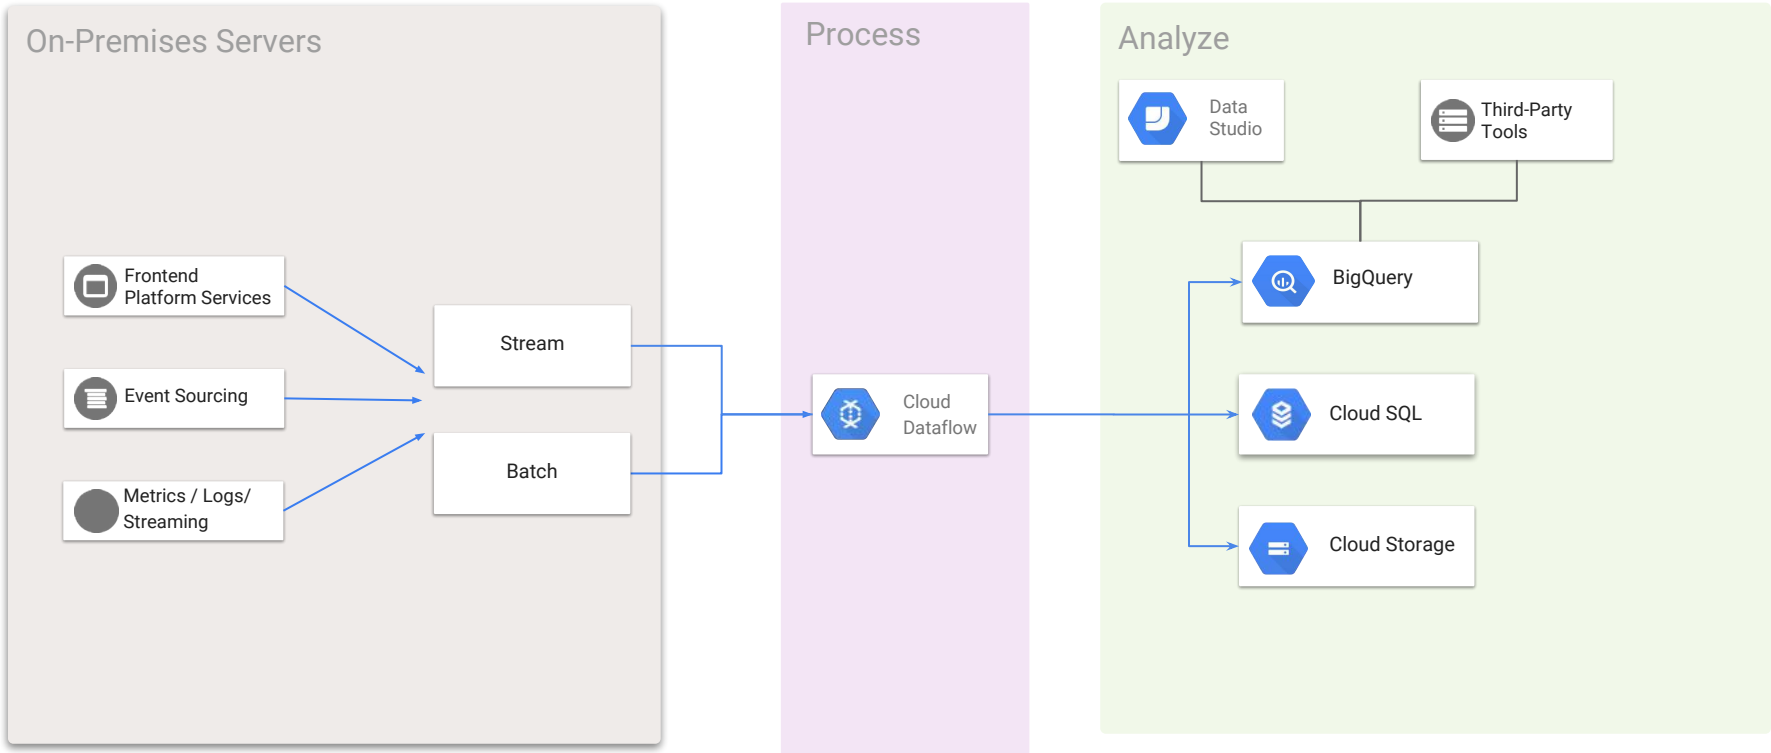




**Data needs to be processed in
multiple services.
How can we pipe to multiple places?**



Architecting for The Cloud

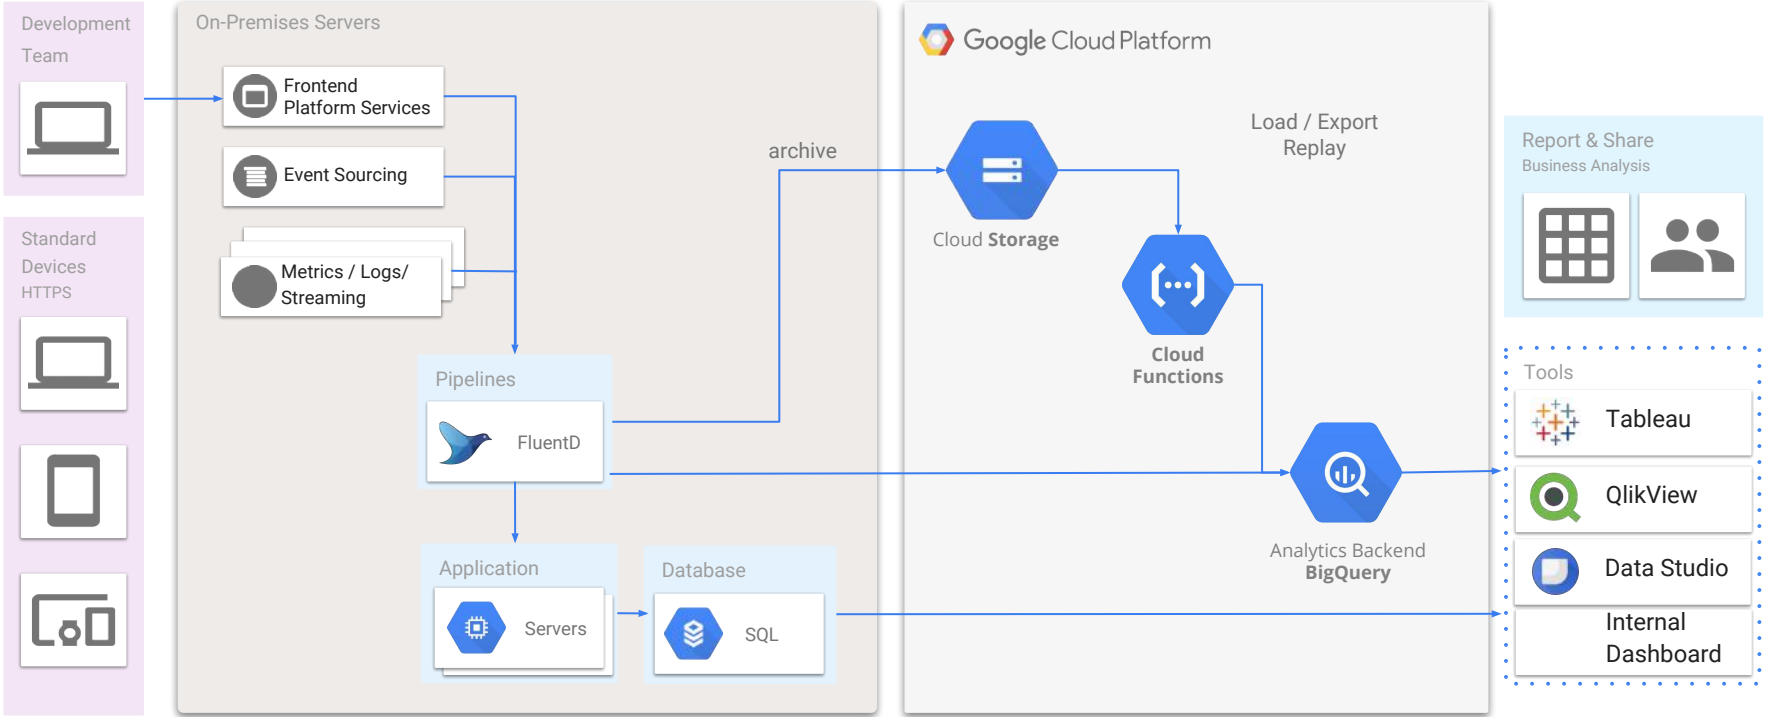




**We have our app outside of GCP.
How can we use the benefits of BigQuery?**



Data Pipeline Integration at REEA.net



The following slides will present a sample Fluentd configuration to:

1. **Transform** a record
2. **Copy** event to multiple outputs
3. **Store** event data in File (for backup/log purposes)
4. **Stream** to BigQuery (for immediate analyses)

```
<filter frontend.user.*>  
  @type record_transformer  
</filter>
```

```
<match frontend.user.*>  
  @type copy  
  <store>  
    @type forest  
    subtype file  
  </store>  
  <store>  
    @type bigquery  
  </store>
```

...

```
</match>
```

1

Filter plugin mutates incoming data. Add/modify/delete event data **transform attributes without a code deploy.**

2

The copy output plugin copies events to multiple outputs. File(s), multiple databases, DB engines.
Great to ship same event to multiple subsystems.

3

4

The Bigquery output plugin on the fly streams the event to the BigQuery warehouse. **No need to write integration.** Data is available immediately for querying.

Whenever needed other output plugins can be wired in: Kafka, Google Cloud Storage output plugin.

- 1 record_transformer
- 2 copy
- 3 file
- 4 BigQuery

```
<filter frontend.user.*>  
@type record_transformer  
enable_ruby  
remove_keys host  
<record>  
  bq {"insert_id":"${uid}","host":"${host}",  
      "created":"${time.to_i}"  
  avg ${record["total"] / record["count"]}  
</record>  
</filter>
```

syntax: Ruby, easy to use.

Great for:

- date transformation,
- quick normalizations,
- calculating something on the fly,
and store in clear log/analytics db
- renaming without code deploy.

- 1 record_transformer
- 2 copy
- 3 file
- 4 BigQuery

```
<match frontend.user.*>  
  @type copy  
  <store>  
    @type forest  
    subtype file  
    <template>  
      path /tank/storage/${tag}.*.log  
      time_slice_format %Y%m%d  
    </template>  
  </store>  
</match>
```

- 1 record_transformer 2 copy 3 file 4 BigQuery

```
<match frontend.user.*>
```

```
  @type bigquery
```

```
  method insert
```

```
  auth_method json_key
```

```
  json_key /etc/td-agent/keys/key-31da042be48c.json
```

```
  time_field timestamp
```

```
  time_slice_format %Y%m%d
```

```
  table user$%{time_slice}
```

```
  ignore_unknown_values
```

```
  schema_path /etc/td-agent/schema/user_login.json
```

```
</match>
```

Connector uses:

- JSON key auth file
- JSON table schema

Pro features:

- streaming to Partitioned tables
- ignore unknown values
(not reflected in schema)



Our benefits using BigQuery

- SQL+JSON (big costs saving with partitioning/clustering)
- run raw ad-hoc queries (either by analysts/sales or Devs)
- ability to throw in / join all kind of data
- pricing model 1TB free every month
- no more throwing away-, expiring-, aggregating old data
- no provisioning/deploy
- no running out of resources
- no more focus on large scale execution plan



Qwiklabs - Find Your Topic and Skill Level

	Kubernetes	Machine Learning	Certification
Introductory For GCP newcomers 1 hour duration, 2 labs	Enroll in the Baseline: Infrastructure Quest . 1. Kubernetes Engine: Qwik Start (lab) (video) - Deploy a containerized application with Kubernetes Engine. 2. IoT: Qwik Start (lab) - Learn about GCP's Cloud IoT Core service, where almost any IoT thing is possible!	Enroll in the Baseline: Data, ML, AI Quest . 1. Cloud ML Engine: Qwik Start (lab) (video) - Train and deploy a TensorFlow model to Cloud ML Engine. 2. Cloud Natural Language API: Qwik Start (lab) - Perform sentiment analysis on a block of text. For fun, paste in the last email you sent!	Enroll in the Cloud Architecture Quest . 1. Stackdriver Qwik Start: Web UI (lab) - What's up with your VM? Learn to install monitoring and logging agents to collect info about your GCP resources. 2. Getting Started with Cloud KMS (lab) - Create your first KeyRing and CryptoKey with GCP.
Experienced Qwiklabs alumni 1 hour, 2 labs	Enroll in the Kubernetes Quest . 1. Kubernetes Engine: Qwik Start (lab) (video) - Deploy a containerized application with Kubernetes Engine. 2. Intro to Docker (lab) Are your containers running? Create, run, and debug containers, then pull and push images to and from Google Container Registry.	Enroll in the Machine Learning APIs Quest. 1. Cloud ML Engine: Qwik Start (lab) (video) - Train and deploy a TensorFlow model to Cloud ML Engine. 2. Detect Labels, Faces, and Landmarks in Images with the Cloud Vision API (lab) - How do you feel about self-driving cars? Learn what's behind the AI making the future possible.	Enroll in the Cloud Architecture Quest . 1. Stackdriver Qwik Start: Web UI (lab) - What's up with your VM? Learn to install monitoring and logging agents to collect info about your GCP resources. 2. Application Performance Management (APM) with Stackdriver (lab) - Dig deeper into Stackdriver with a problematic environment - can you identify the issues?



Thank you.

Slides available on: [slideshare.net/martonkodok](https://www.slideshare.net/martonkodok)

Reea.net - Integrated web solutions driven by creativity to deliver projects.

