
Google
Developers

313: Google Compute Engine

Technical Details

Joe Beda
Staff Software Engineer, Tech Lead

Google™ 12

Introducing Google Compute Engine

Infrastructure as a Service at Google

- Compute
- Network
- Storage
- Tools

Introducing Google Compute Engine

Google's Advantages

- Scale
- Speed
- Global footprint
- Integrated platform

Guiding Principles

What to expect from Google Compute Engine

- Secure
- Open and Flexible
- Consistent
- Proven
- Enables an ecosystem

Hello Google Compute Engine

A quick demo

Evan Anderson, Tech Lead, Networking

The Architecture

Moving parts and how they fit together

System Components

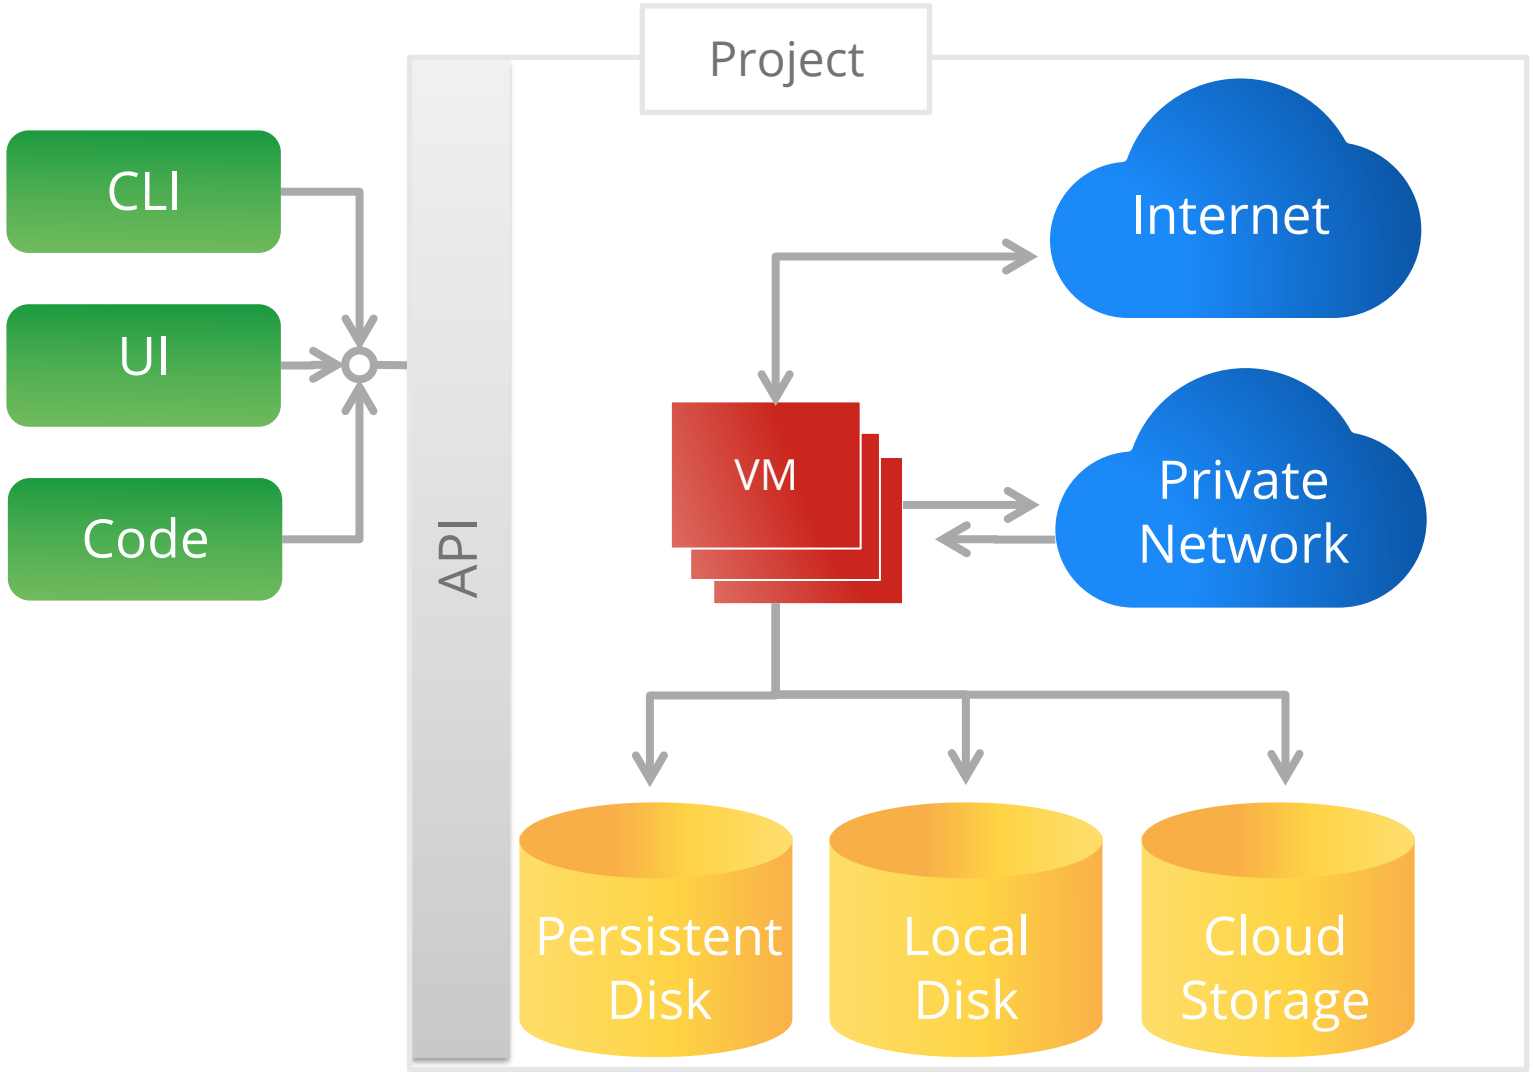

API Basics

- **JSON over HTTP**, REST-inspired
- Main Resources (Nouns)
  - Projects
  - Instances
  - Networks and Firewalls
  - Disks and Snapshots
  - Zones
- Actions (Verbs):
  - GET
  - POST (create) and DELETE
  - Custom ‘verbs’ for updates
- Auth via OAuth2

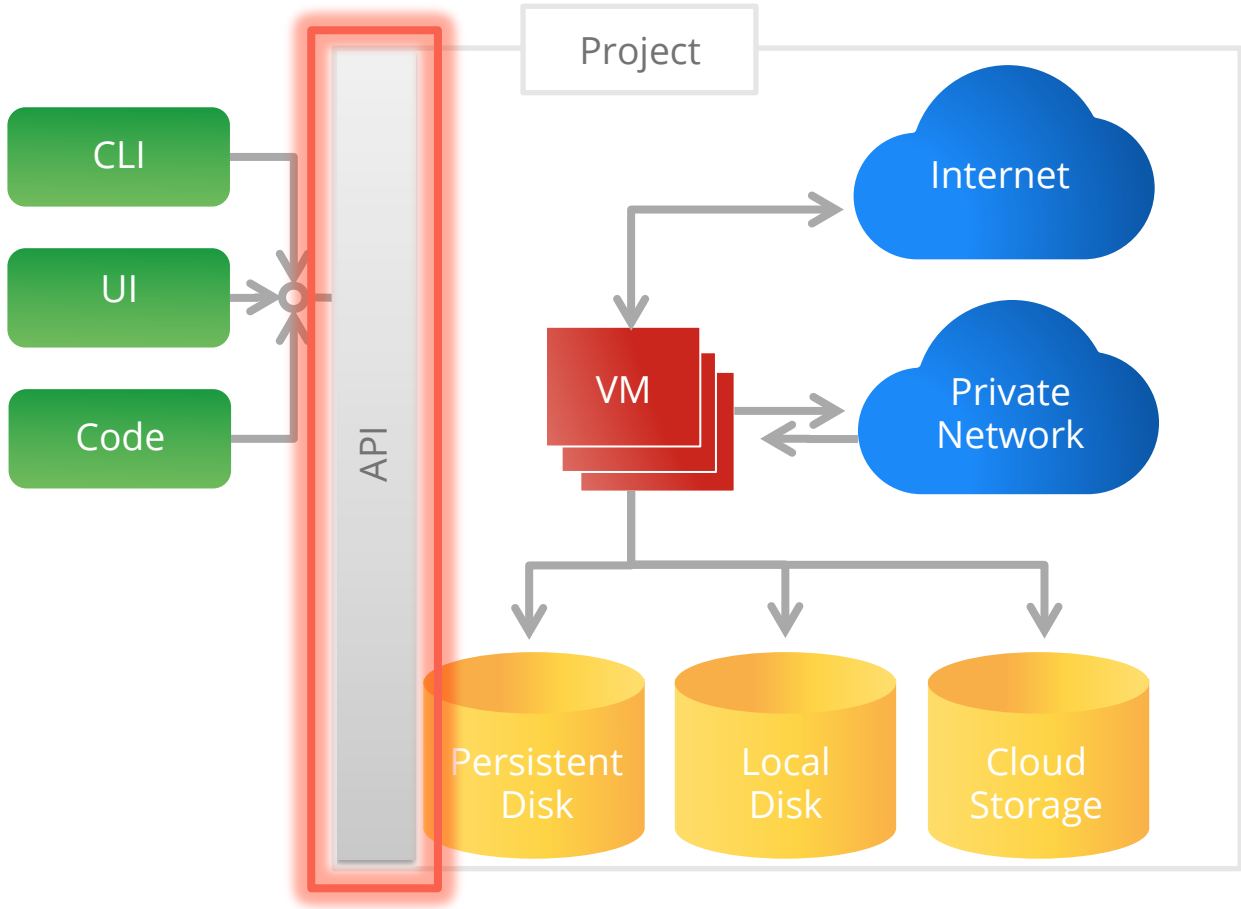

Clients and Libraries

- gcutil: command line utility
- Web UI: Built on GAE
- Libraries
- Partners and ecosystem

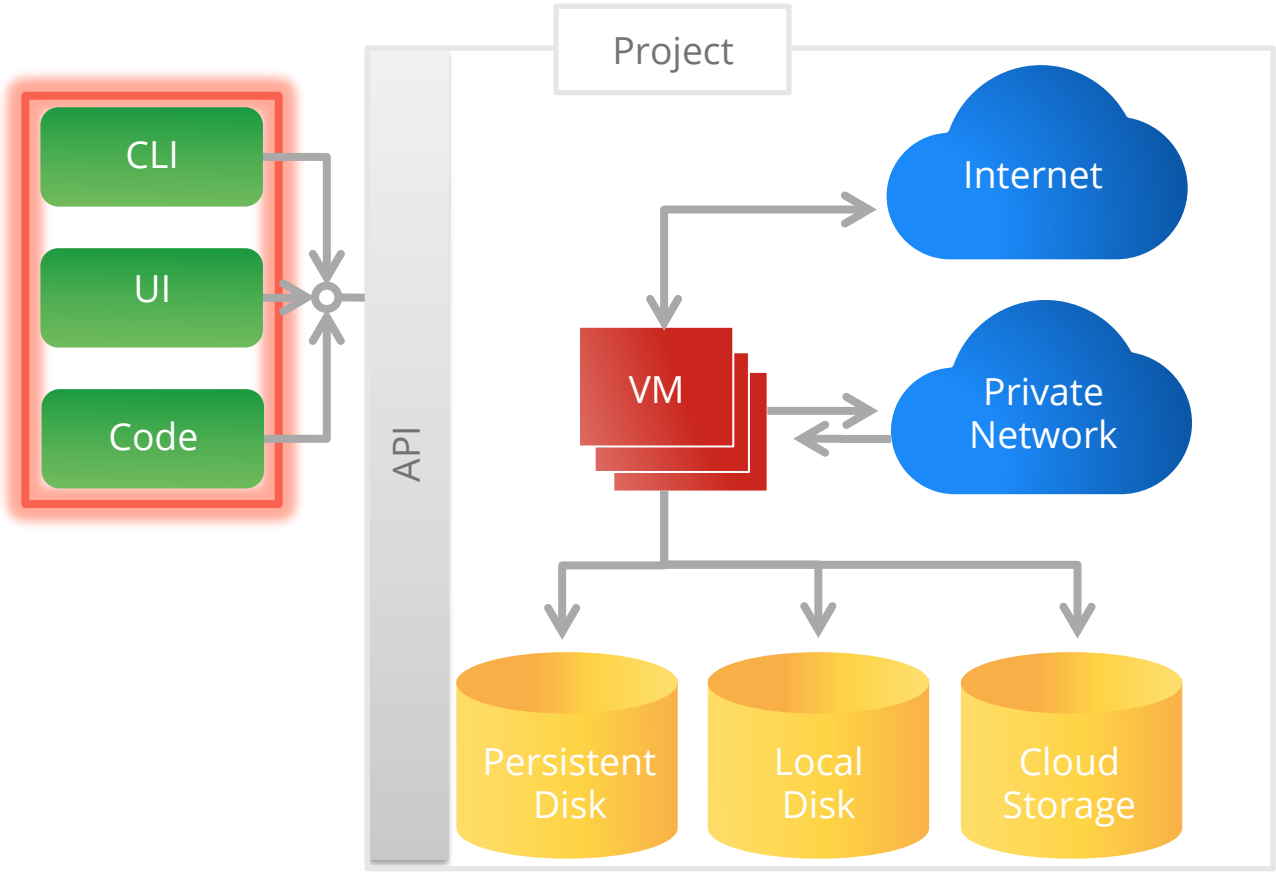

Projects

- Based on API Console projects
- Container for all resources
- Team membership
- Group ownership
- Billing

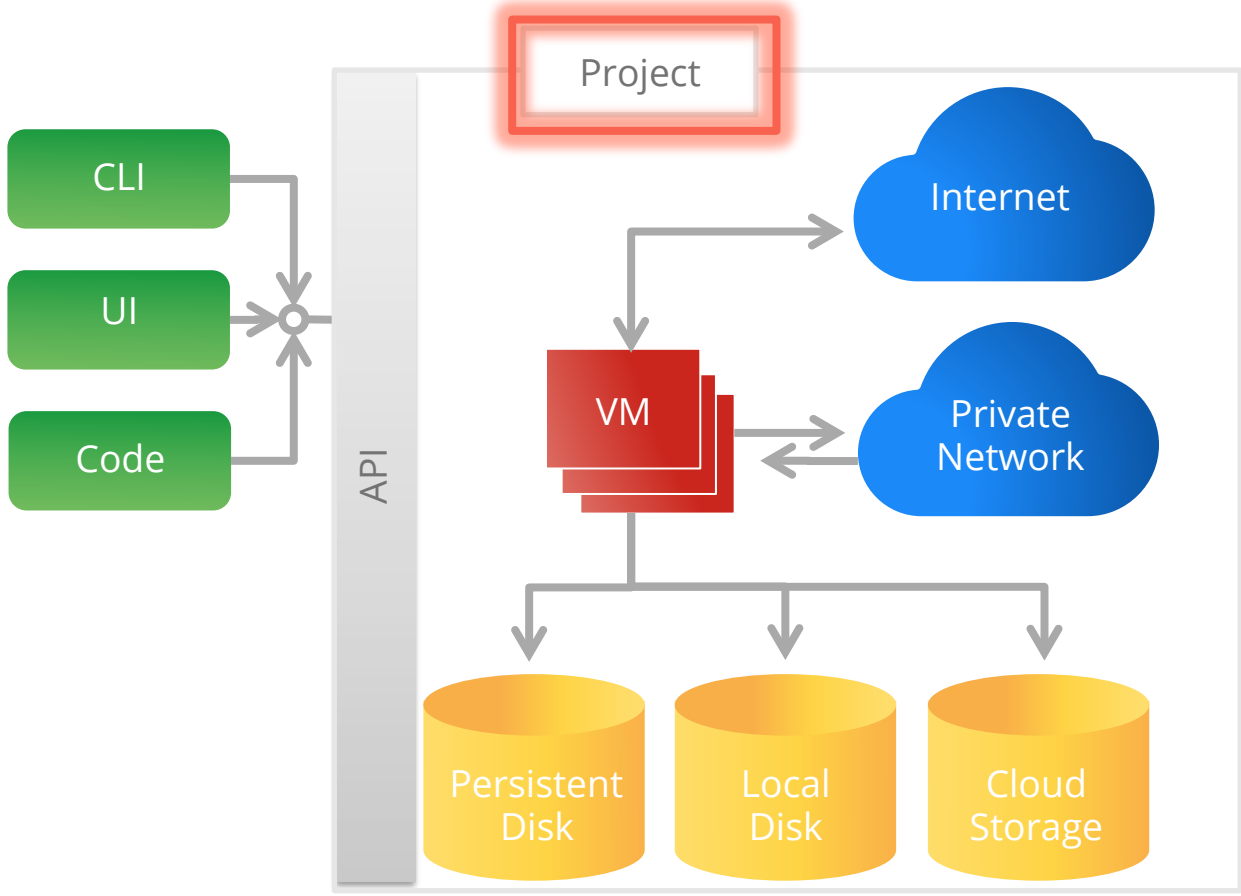

Instances

Linux Virtual Machines

- Root access, locked down kernel
- Stock Images: Ubuntu, CentOS
- Useful utilities preinstalled

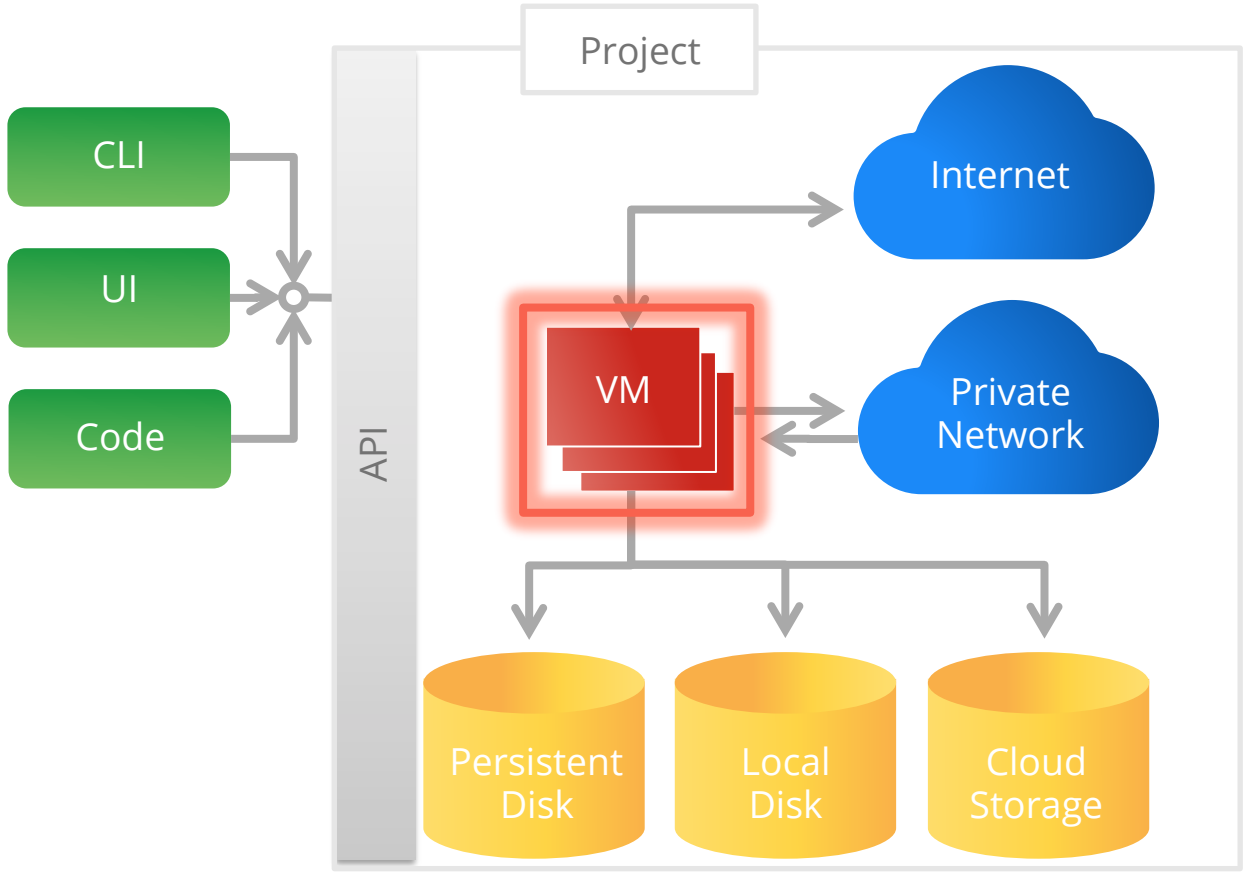

Instances

Machine Types

- Modern processor (Intel Sandy Bridge)
- 1, 2, 4 and 8 virtual CPUs
  - 1:1 virtual CPU to hyperthread
- 3.75GB RAM per virtual CPU
- Over 420GB ephemeral disk per CPU
  - Dedicated spindles on -4 and -8
- New Performance Metric
  - GCEU: Google Compute Engine Unit
  - 2.75 GCEUs per virtual CPU
- Smaller machine types coming soon

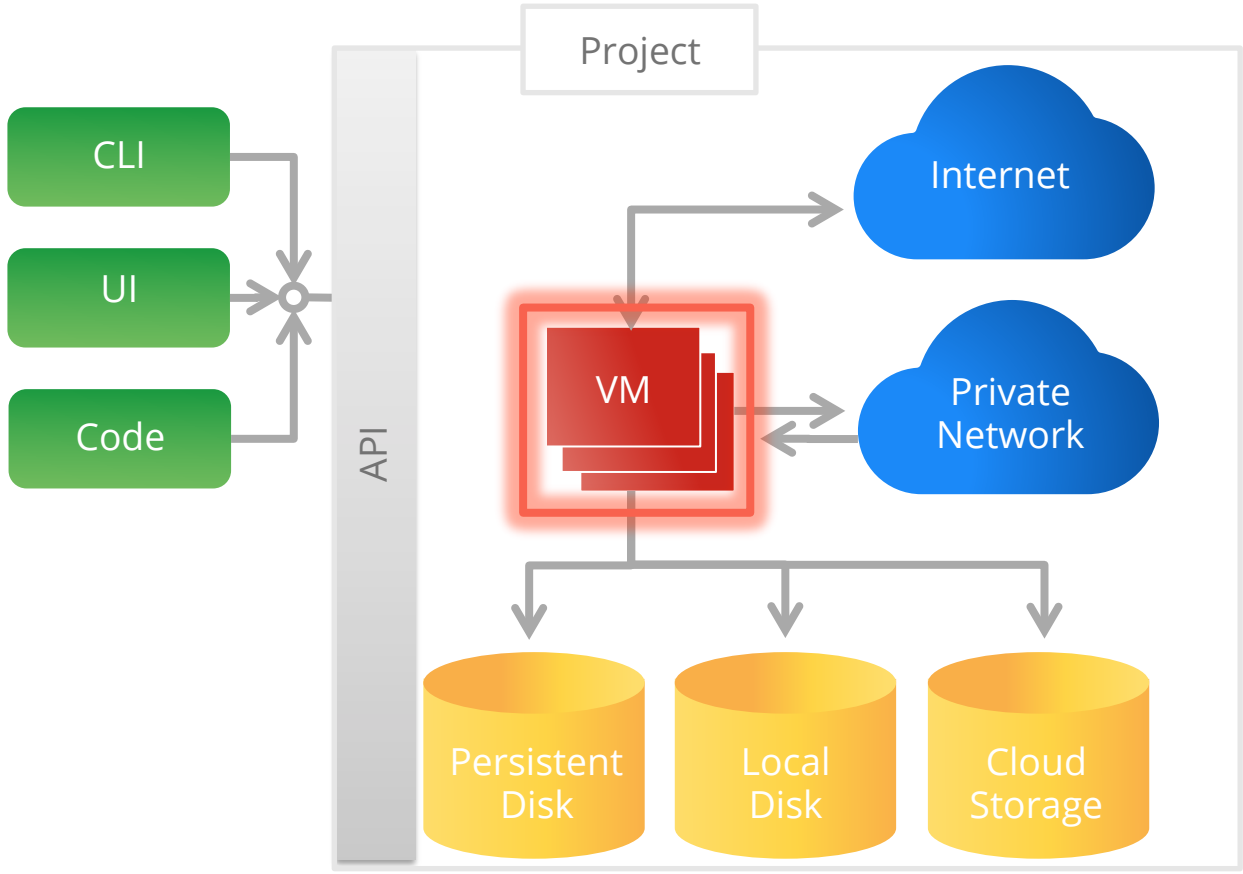

Instances

KVM Hypervisor + Linux cgroups

- Kernel Virtual Machines
  - Linux is the hypervisor
  - Virtualized, non-virtualized run side by side
  - Worked closely with Red Hat
- Linux cgroups
  - Resource isolation
  - Public linux feature driven by Google kernel engineering

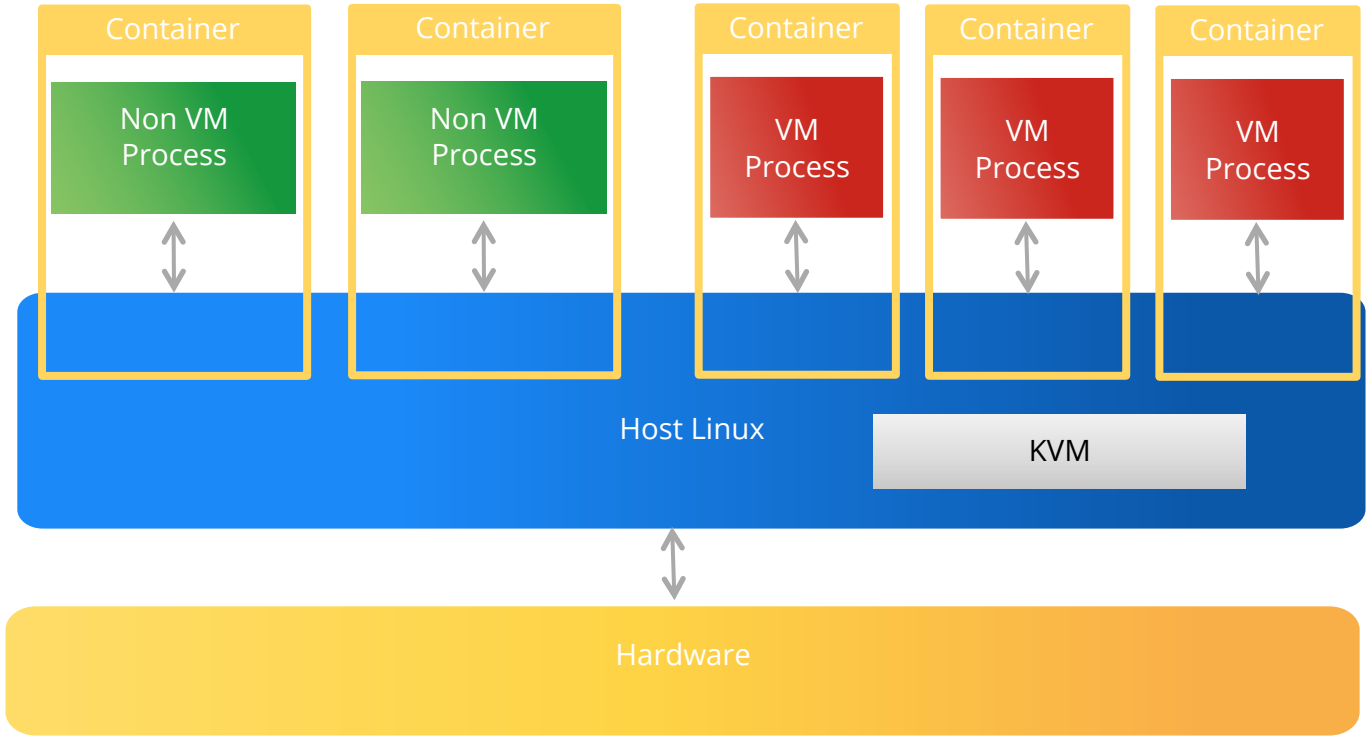

Networking

Private Virtual Network

- Isolated networks per project
- Private IPv4 space (RFC 1918)
- IP Level (Layer 3) network
- Flat across geographical regions
- Internal facing DNS
  - VM name = DNS name

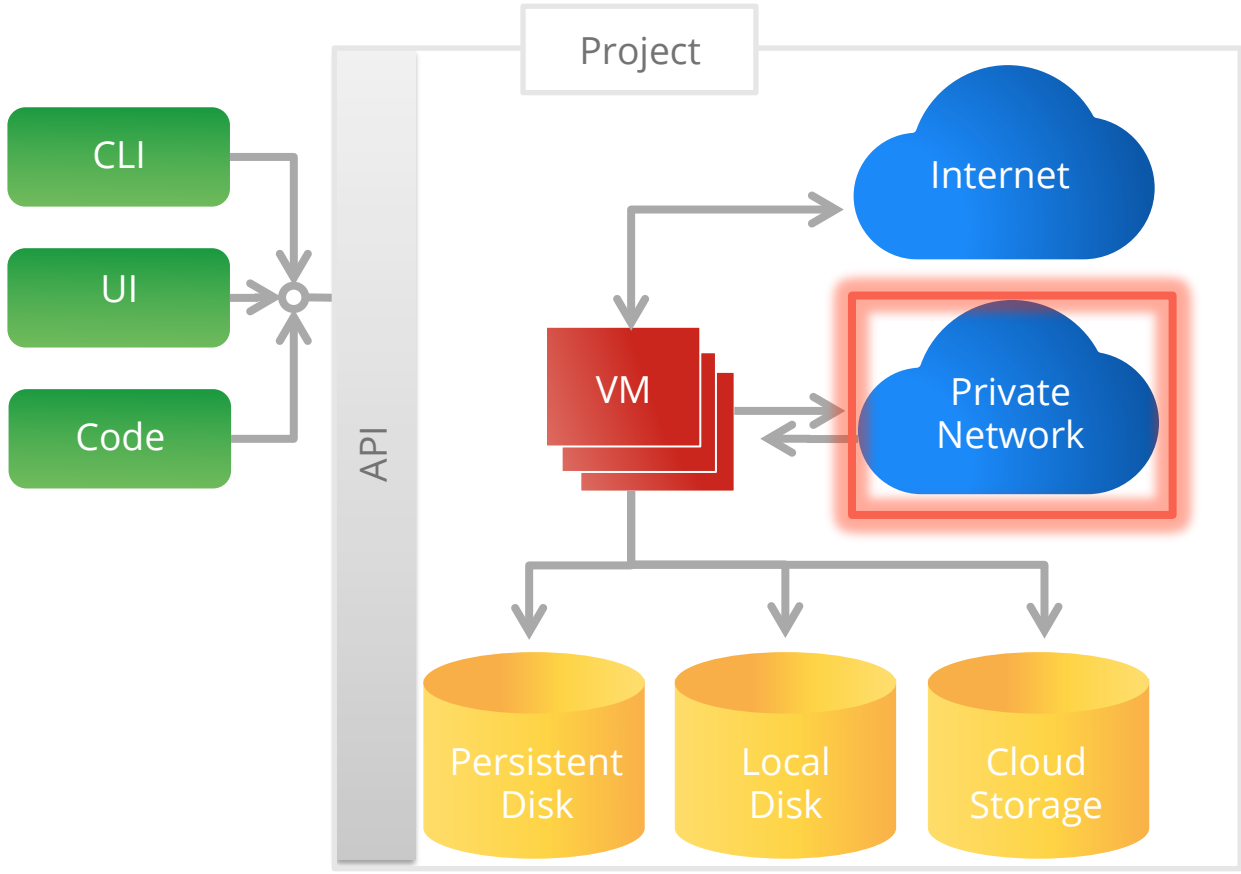

Networking

Internet Access

- External IPs:
  - Reserved, ephemeral, none
  - Not tied to region/zone
  - Dynamic attach/detach
- 1-to-1 NAT
- Built in firewall system
- Global network footprint
- Limitations
  - Outgoing SMTP blocked
  - UDP, TCP, ICMP only

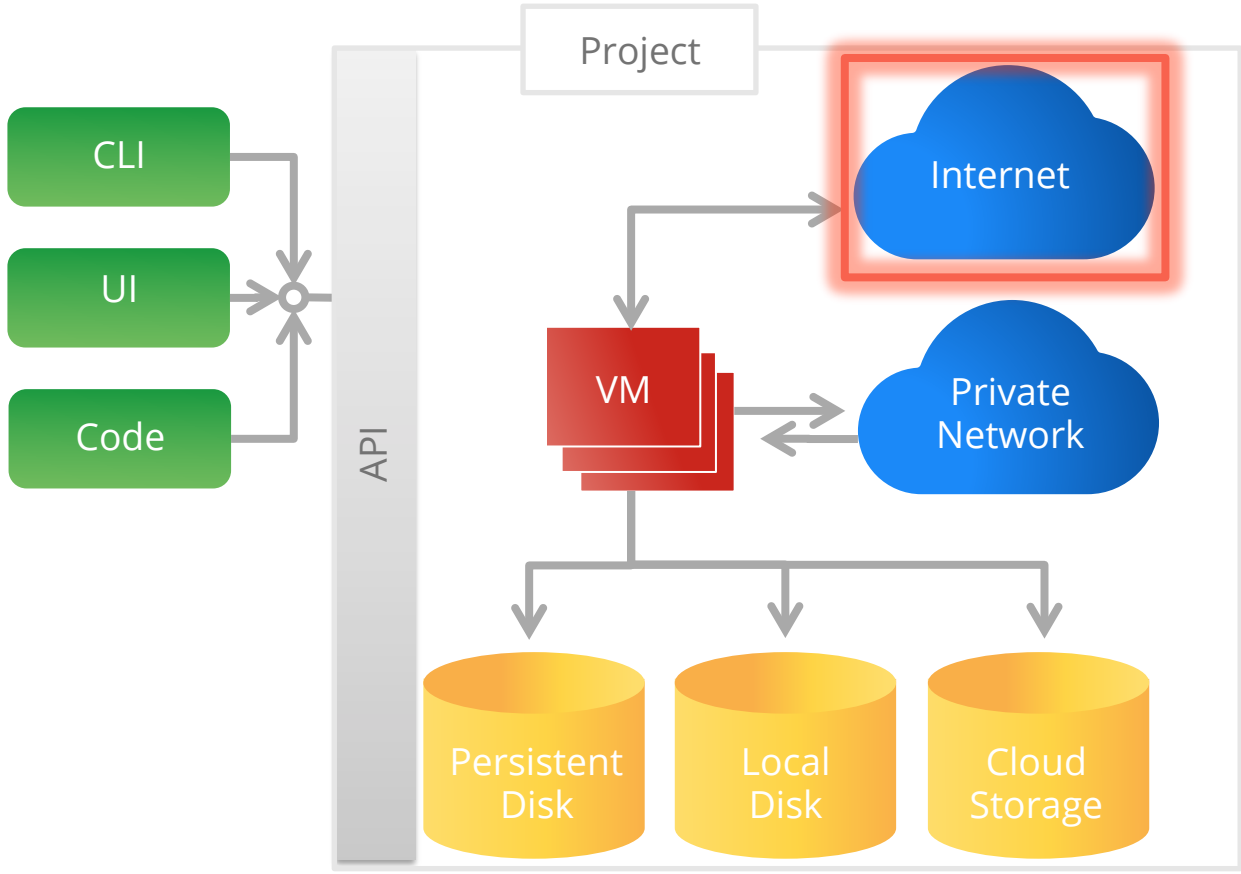

Storage

Persistent Disk

- **Fast, consistent performance**
- Provisioned via API
- Located in a zone
- R/W with single instance
- R/O with multiple instances
- Encrypted at rest

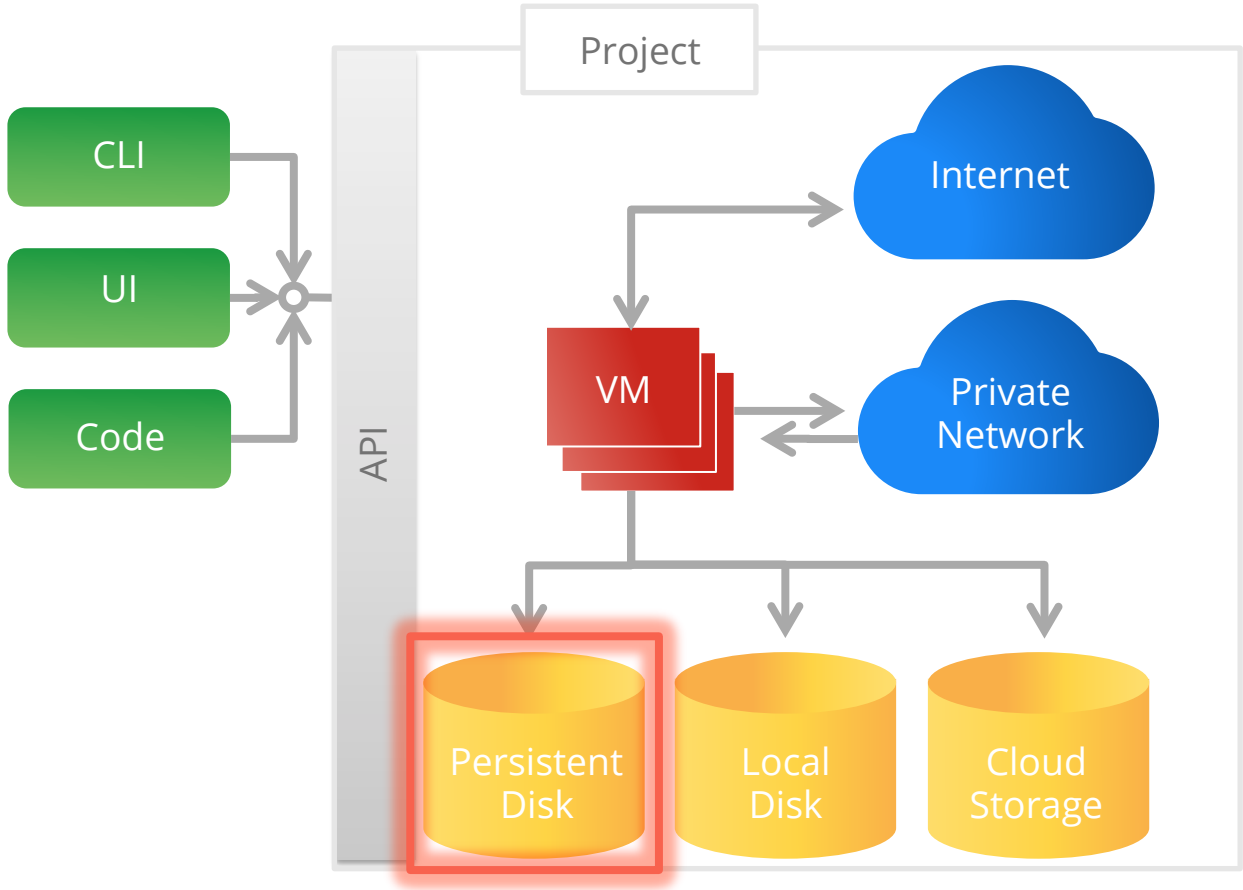

Storage

Ephemeral Disk

- Currently used for booting all instances
- Lives and dies with instance
- Large ‘extended’ devices
- On same physical machine
- Dedicated spindles (4 CPU and larger)
- Encrypted at rest

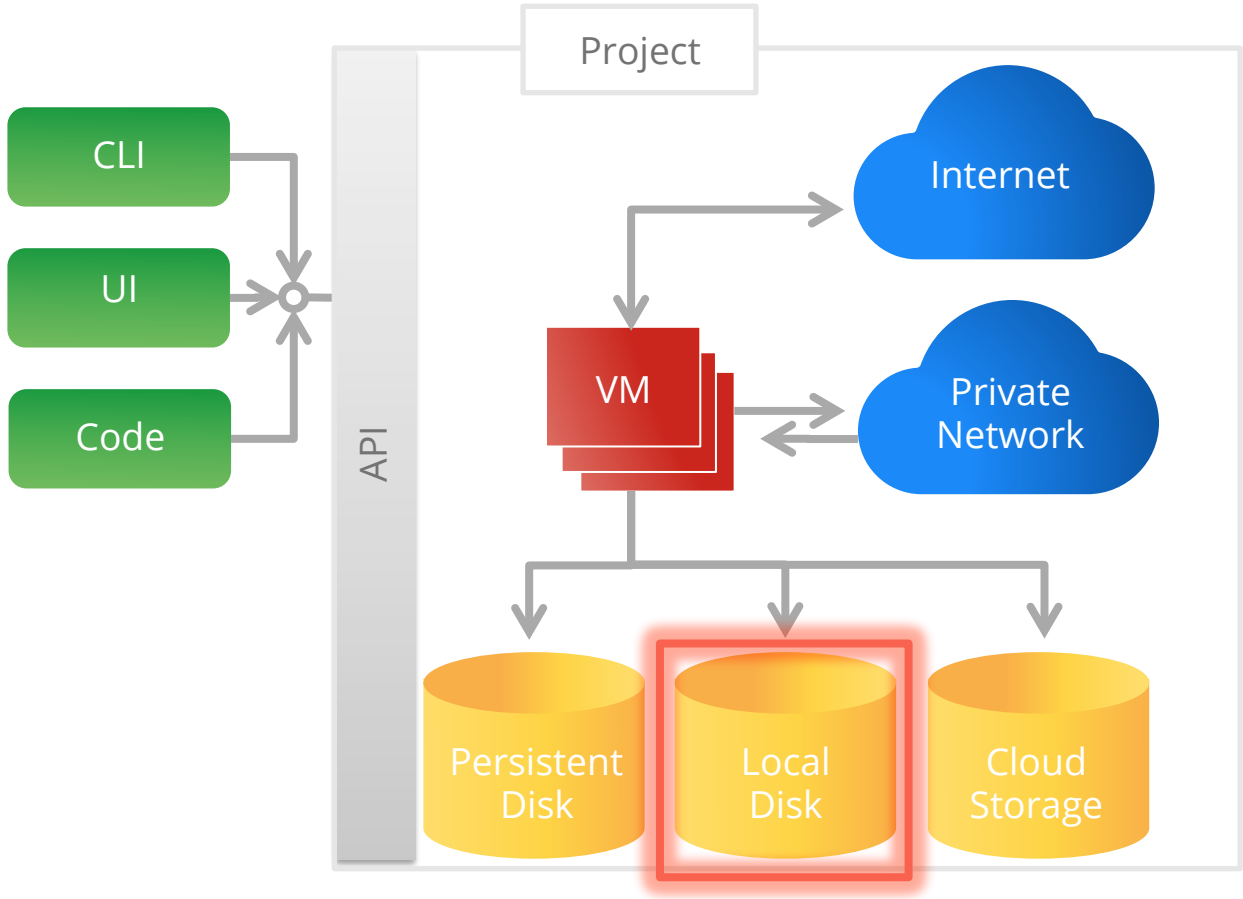

Storage

Google Cloud Storage

- Internet object store
- Global API based access
- Great for getting data in and out
- Frictionless access with service accounts

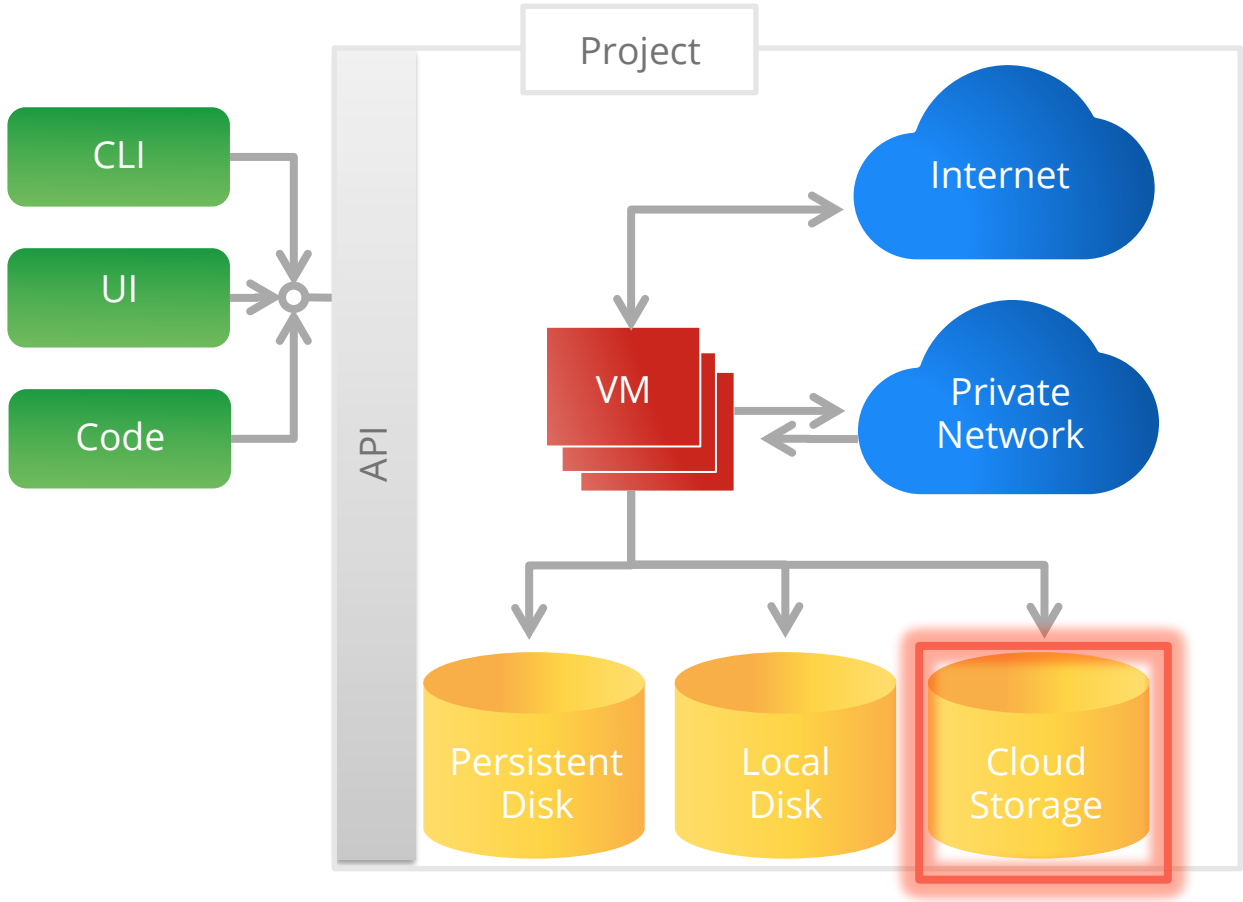

Locality

Managing Location and Availability

- Region: geography and routing
- Zone: fault isolation
- 3 Zones in limited preview, more coming

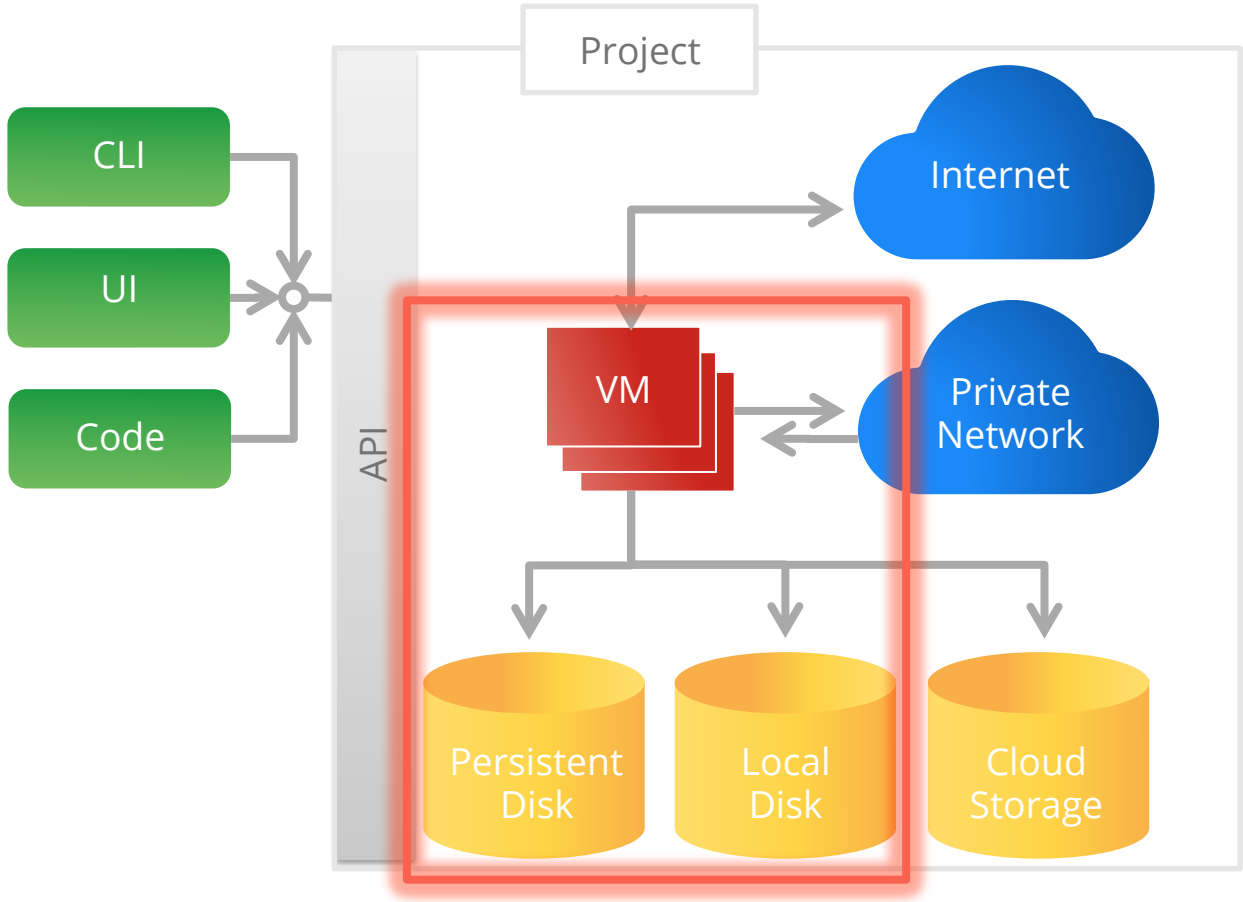

Invite Media on Compute Engine

Hamza Kaya
Software Engineer

invitemedia®

Running an Advertising Business on Google Compute Engine

- Automatically buy from multiple ad exchanges in real-time, through one interface
  - Need low latencies to multiple exchanges
  - Running high qps to multiple exchanges
- Built first on another IaaS cloud; ported to Google Compute Engine

invitemedia®

The port...

- Quick and easy
- Familiar API model
- 2 weeks of engineering time

invitemedia®

The results...

- Observed twice the computing power over original provider
  - Max QPS on a single 8 virtual CPU instance – from 350 qps to 650 qps
  - Half the number of 8 virtual CPU servers to manage – from 284 to 140
- While offering strong consistency of results
  - Connection Error Rate – from 5% to 0.5%
  - Deadline Exceeded Rate – from 11% to 6%

Hadoop On Compute Engine

A real world sample application

Evan Anderson, Tech Lead, Networking

Hadoop On GCE

Application Architecture

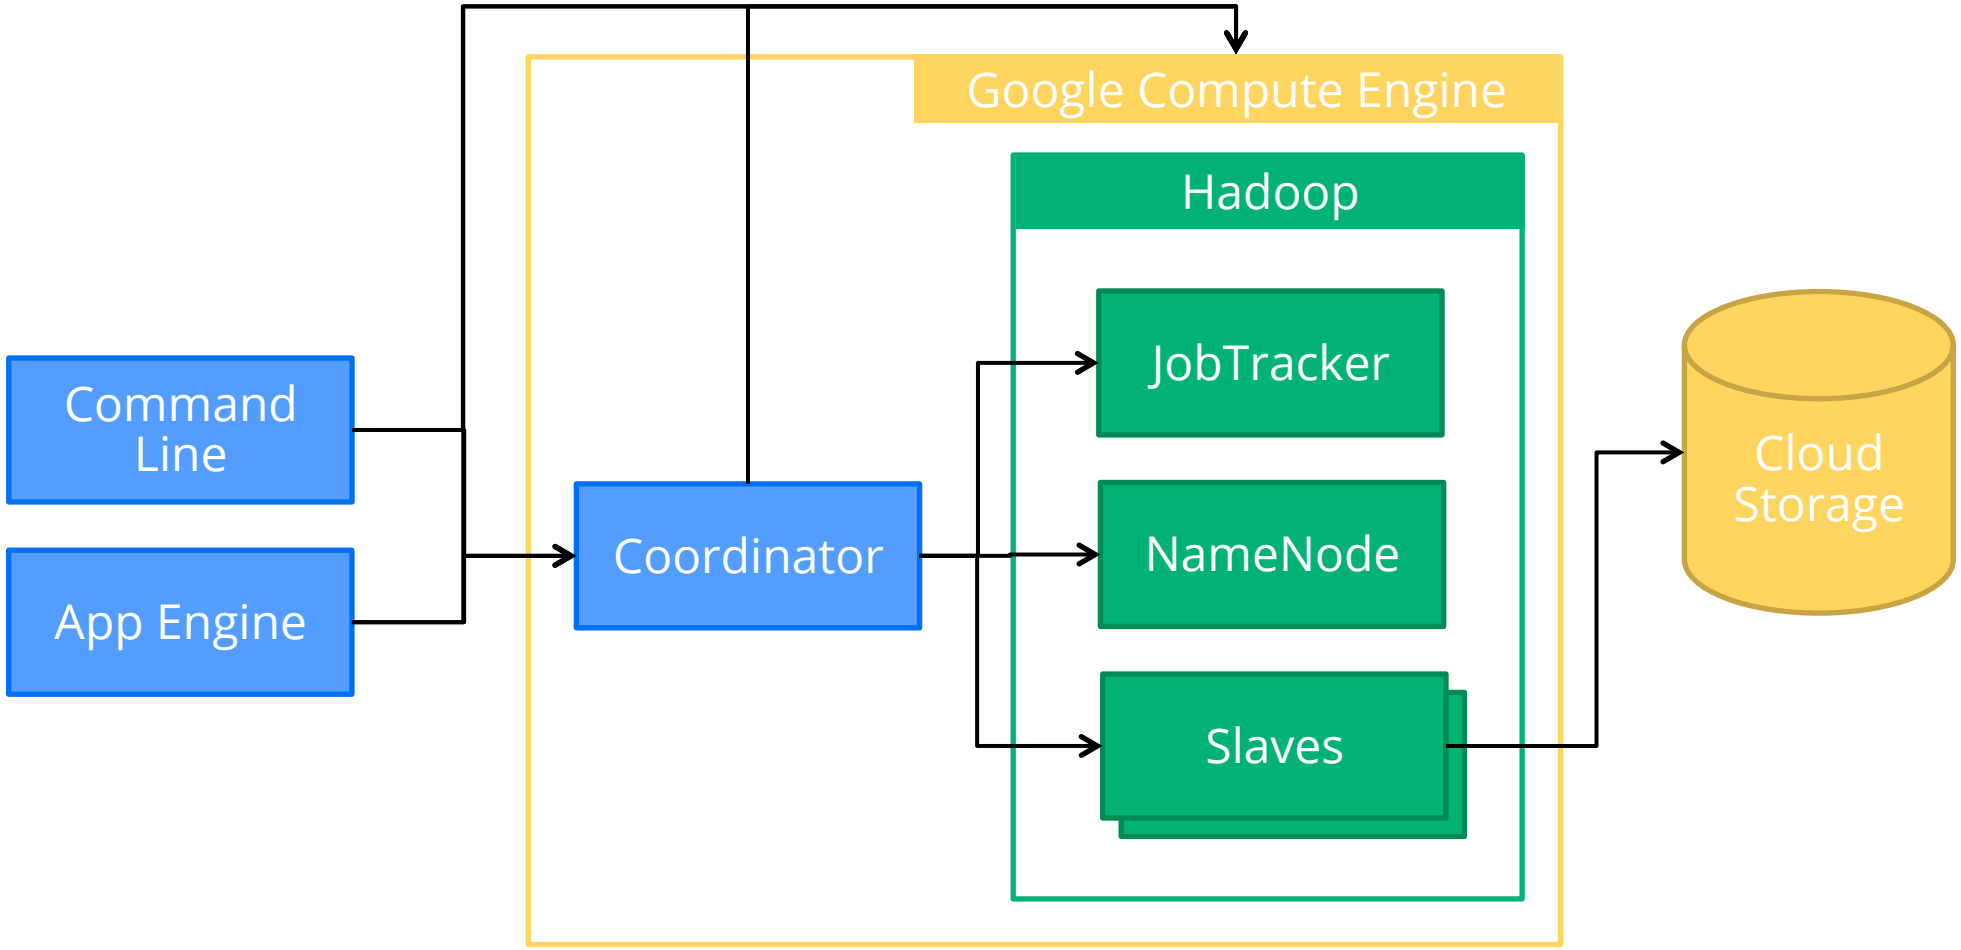

Exploring Compute Engine

Getting the most from Google Compute Engine

Service Accounts

Frictionless Access to Google APIs

- Synthetic identity for VMs and code
- Google Compute calling Google APIs
  - Examples: Cloud Storage, App Engine task queue API
- App Engine calling Compute Engine API
  - Use App Engine as ‘orchestrator’
  - Build your own customized dashboard and control logic

Service Accounts

Google Compute Engine Calling Google Cloud Storage

CLI

```
me@workstation$ gcutil addinstance sa-example --service_account_scopes=storage-rw
me@workstation$ gcutil ssh sa-example
[snip]
me@sa-example$ gsutil mb gs://unique-bucket-name
Creating gs://unique-bucket-name/...
```

No configuration or passwords required!

Instance Metadata

Parameters for VMs

- Dictionary of Key/Value pairs
- Set from the API, read from the Instance
- Accessible at metadata server (http://metadata/...)
- Useful for small amounts of configuration data
- Project level metadata inherited by instances.

Instance Metadata

CLI

```
me@workstation$ gcutil addinstance metadata-example \
  --metadata=role:master --metadata_from_file=config:config.txt
me@workstation$ gcutil ssh metadata-example
[...snip...]
me@metadata-example$ curl http://metadata/0.1/meta-data/attributes/role
master
me@metadata-example$ curl http://metadata/0.1/meta-data/attributes/config
[...file content...]
```

Start Up Scripts

Simple Bootstrapping

- Builds on metadata
- Equivalent to rc.local
- Example usage:
  - Install packages, start services
  - Use Google Cloud Storage to grab data, code and binaries
- Bootstrap other management infrastructures

Start Up Scripts

CLI

```
me@workstation$ cat render-stuff.sh
#! /bin/bash
apt-get install -y contextfree
cfdg -s 10000 /usr/share/doc/contextfree/examples/sierpinski.cfdg /tmp/out.png
gsutil cp -a public-read /tmp/out.png gs://contextfree-examples/sierpinski.png

me@workstation$ gcutil addinstance start-me-up \
  --metadata_from_file=startup-script:setup-my-instance.sh \
  --service_account_scopes=storage-rw
me@workstation$ gcutil ssh sa-example
[...snip...]
me@sa-example$ tail -f /var/log/google.log
```

Services, Not Servers

- Realities of a datacenter
  - Hardware and software fails
  - Build across zones
  - Scheduled maintenance:
    - Up to 2 weeks every 20 weeks, one zone at a time
    - Addressed in future versions
- Techniques
  - Ephemeral disk = cache
  - Start up scripts
  - Dynamic management
  - Automation

Limited Preview Program

Apply for access today!

- Apply for program at cloud.google.com
- Complimentary access to a quota of compute cores for a limited period
- SLA and support available to commercial customers

Thank You!

Please visit cloud.google.com

Come to 308 tomorrow: Google Compute Engine + Google App Engine

Google
Developers