



EUROPYTHON 2011
FLORENCE, JUNE 20-26



What is... Google App Engine?

Wesley Chun

Developer Advocate


<http://code.google.com/appengine>

@wescpy / @app_engine


About the Speaker

- Software engineer by profession
 - Currently at Google (cloud products)
- Course instructor: teaching Python since 1998
 - Private Corporate Training & Public Courses
- Community volunteer
 - User groups: BayPIGgies and SF Python Meetup
 - Other: Tutor mailing list; Python conferences
- Author/co-author (books, articles, blog, etc.)
 - *Core Python Programming* ([2009,]2007, 2001)
 - *Python Fundamentals* LiveLessons DVD (2009)
 - *Python Web Development with Django* (2009)

I Teach



I Write



I Code



Cloud Computing

What is it (besides being buzzword-compliant)?

What is Cloud Computing?

"Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

National Institute of Standards and Technology (NIST) Forum May 2010
http://csrc.nist.gov/groups/SNS/cloud-computing/forum-workshop_may2010.html

Cloud Computing concept... it's been around longer than you think

"In 1984, John Gage and Bill Joy of Sun Microsystems were credited with saying, 'The network is the computer,' one of the most memorable slogans of the information age. This concept, based on the tenet of making computing resources available to all users irrespective of location, as long as they are connected to the network, forms the basis of the Internet as we know it, and foretold the advent of cloud computing."

Garry Metcalf, Analysys Mason, Sep 2009
<http://www.analysysmason.com/about-us/news/newsletter/Previous-news-articles/Now-the-network-really-is-the-computer>

The Benefits of Cloud Computing

- **Reduced Cost**
 - Incremental payment based on usage
- **Elastic Storage**
 - Data storage can easily grow with business
- **High Automation**
 - Platform updates/patches updated automatically
- **Flexibility**
 - Resources can be changed on-demand
- **More Mobility**
 - Universal Accessibility
- **Shift towards Innovation**
 - Less maintenance, can focus on product innovation

Bottom line: More focus on business!

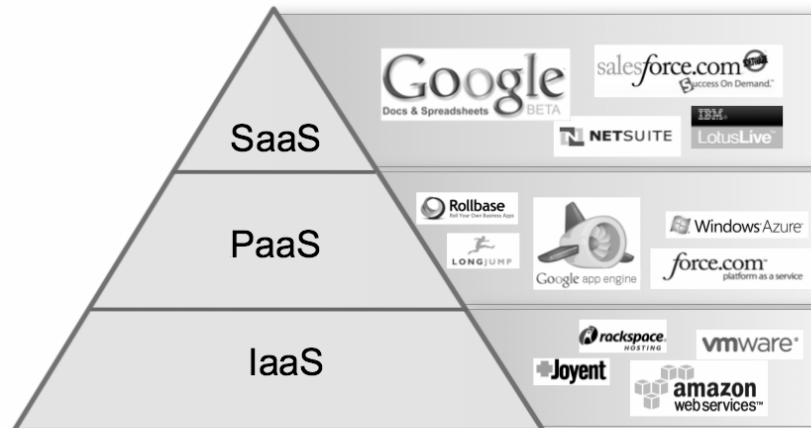
The shift towards Cloud Computing

"By 2012, 20% of Global 2000 enterprises will be using public cloud services, up from under 5% in 2009."

Yefim Natis - Gartner Cloud Computing Analyst, AADI Dec 2009


The Cloud Pyramid

Cloud Computing Defined



Source: Gartner AADI Summit Dec 2009

What is App Engine?




- App Engine is a platform
- You build & test your app
- Then upload your app to Google
- App Engine runs everything
- No need to worry about machines, network, storage, scalability, etc.

Challenges building web apps



What keeps you up at night? DIY Hosting means hidden costs

- Idle capacity
- Software patches & upgrades
- License fees
- Lots of maintenance
- Traffic & utilization forecasting
- Upgrades





Easy to start

Easy to scale

Easy to maintain

We do the dirty work...

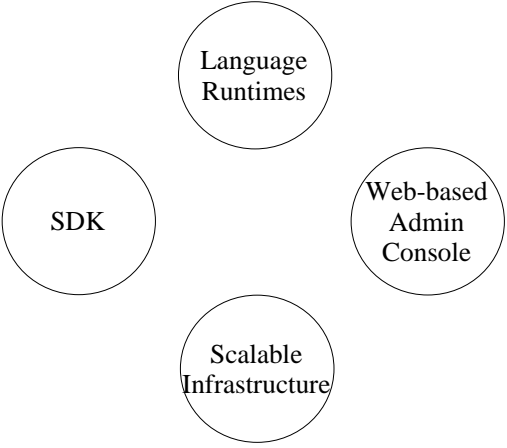
Google App Engine

“We wear pagers so you don’t have to”


The Components

of Google App Engine


Components



Components




Components



Alternative API/Runtimes also available (more later)

Language runtimes



Duke, the Java mascot
Copyright © Sun Microsystems Inc., all rights reserved.

What are programming languages?

Various dialects that let humans give instructions to computers



- Some are easier than others
- Some are special-purpose
- Adults: Java, C/C++, Python, Ruby, PHP
- Kids: Scratch, Alice, Python, BASIC



The king of ease-of-use




- Extremely rapid development
- Very low barrier of entry
- Simple yet robust syntax
- Rich library of packages/modules
- App Engine's first language API

The king of enterprise development



- Ubiquitous in Enterprise computing
- Adheres to Java servlet standard
- Rich library of packages/modules
- Eclipse Plug-in support
- Alternative language support

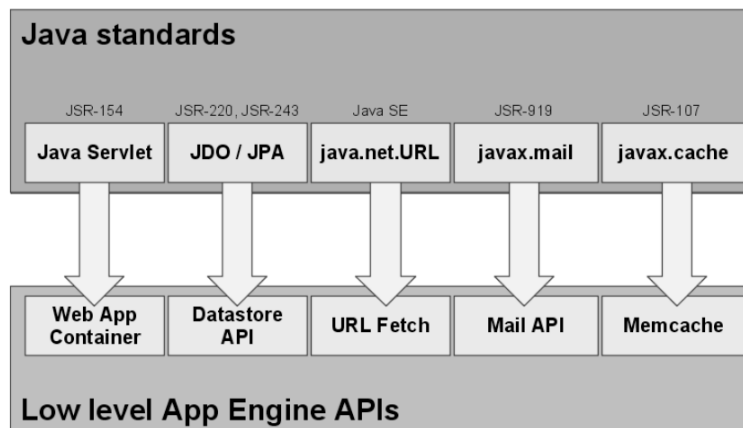
The new kid on the block



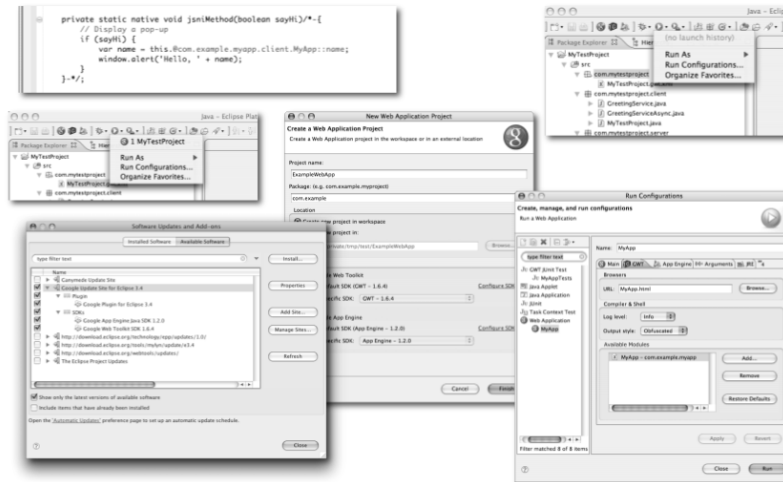
- Best of both worlds
- The complexity & power of a statically-type language
- Mixed with the world dynamically-typed languages
- Flexible alternative to Python & Java

Java compliant Servlet Standard

Ensuring Portability



Google Plugin for Eclipse



Extended Language support through JVM

- Java
- Scala
- JRuby (Ruby)
- Groovy
- Quercus (PHP)
- Rhino (JavaScript)
- Jython (Python)



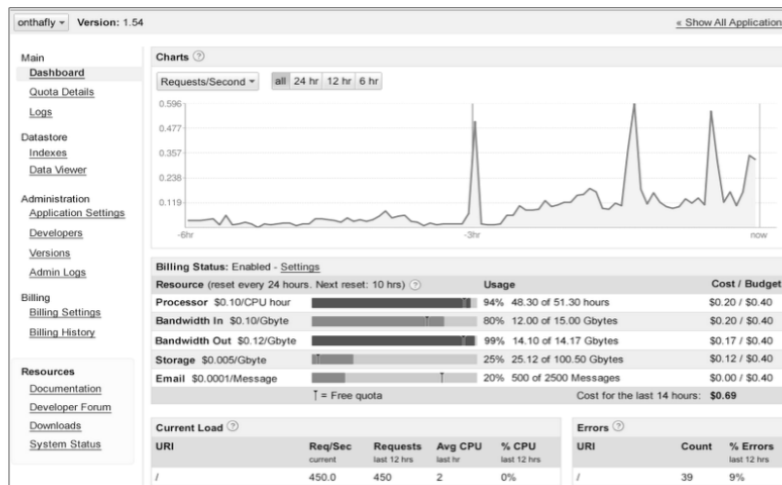
Duke, the Java mascot
Copyright © Sun Microsystems Inc., all rights reserved.

Components

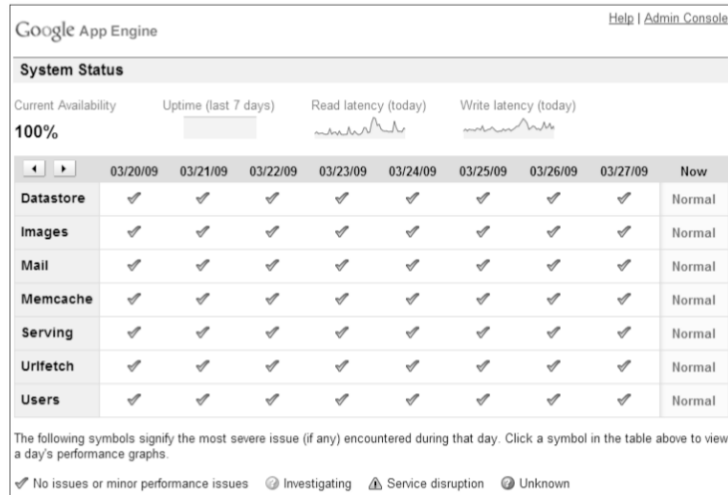


Application Monitoring

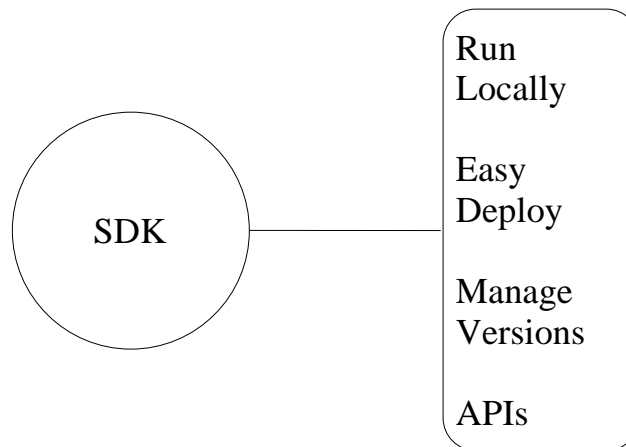
App Engine Dashboard



App Engine Health History




Components



Case Studies

Serving our Users

App Engine growth



2008
App Engine Launch
Python
Datastore
Memcache
logs export

2009
Batch write/read
Https
Status-
Dashboard

Java
DB Import
cron


Task Queues
XMPP
incoming email

2010
Blobstore
Appstats
cursors
Mapper


2011
Multitenancy
Instance Console
Always On
hi-perf imag
10 min tasks

Hi-Replication
Datastore
Channel API
Files API
Remote API
Prosp Search


How many _____ ? By the numbers



How many _____ ? By the numbers



How many _____ ? By the numbers



Developers who know that App Engine...



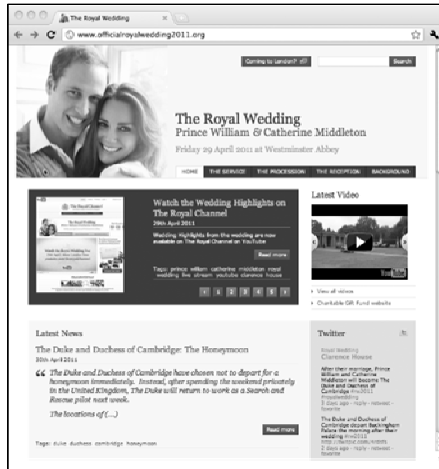
...scales for enterprise-targeted cloud apps



...scales for web and mobile gaming



...scales for event-based websites



Official Royal Wedding Website hosted on App Engine

On Wedding day, served:

- Up to 2000 requests per second
- 15 million pageviews
- 5.6 million visitors

<http://goo.gl/F1SGc>

...scales for social web and mobile apps

Social networking at scale

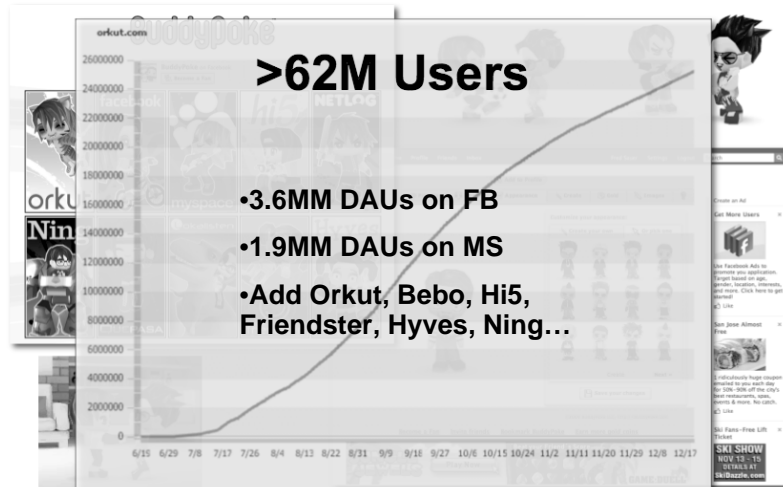
Buddypoke™



...and grows with you and your app

Social networking at scale

Buddypoke™



Gigya...


gigya Socialize



... scales

gigya Socialize

Flexible Scalability



"Although we typically host all our services in-house, on our own infrastructure, we felt that GAE would be a better fit for the live chat feature because of its unique traffic pattern, which is characterized by very low traffic most of the time with very high bursts during high profile events."

Raviv Pavel, Gigya VP of Research and Development

<http://googleappengine.blogspot.com/2010/02/scalability-means-flexibility.html#links>

Not all apps are web-based!!





- Need backoffice processing? Want to build your own?
- Go cloud with App Engine!
- No UI needed for app to talk to App Engine
- Great place for user info e.g., high scores, contacts, badges, etc.
- Better user interface: move user data off the phone

Features and Futures

Services and APIs now... and coming soon

Rich set of App Engine services/APIs



In Review: ~3 years of App Engine... an evolving platform

Apr 2008	Python launch
May 2008	Memcache API, Images API
Jul 2008	Logs export
Aug 2008	Batch write/delete
Oct 2008	HTTPS support
Dec 2008	Status dashboard, quota details
Feb 2009	Billing, Remote API, Larger HTTP request/response size limits (10MB)
Apr 2009	Java launch, Bulkloader (DB import), Cron jobs, SDC
May 2009	Key-only queries, Quota API
Jun 2009	Task queue API, Django 1.0 support
Sep 2009	XMPP API, Remote API shell, Django 1.1 support
Oct 2009	Incoming email
Dec 2009	Blobstore API
Feb 2010	Datastore cursors, Async URLfetch, App stats
Mar 2010	Denial-of-Service filtering, eventual consistency support
May 2010	OpenID, OAuth, App Engine for Business, new bulkloader
Aug 2010	Namespaces, increased quotas, high perf image serving
Oct 2010	Instances console, datastore admin & bulk entity deletes
Dec 2010	Channel API, 10-minute tasks & cron jobs, AlwaysOn & Warmup
Jan 2011	High Replication datastore, entity copy b/w apps, 10-minute URLfetch
May 2011	Backends, Pull Queues, Leaving Preview soon, Go launch
Jun 2011	Geolocation headers, Channel presence, WebP images, ProtoPRC, SDK HRD



App Engine Roadmap

- App Engine out-of-preview/official support
- SSL access on non-appspot.com domains
- Improved datastore import/export backup/restore
- MapReduce
- Full-text Search over Datastore
- Python 2.7
- Improved monitoring and alerting
- Raise request/response size limits for some APIs
- See code.google.com/appengine/docs/roadmap.html

Google Apps Integration

App Engine apps in your Apps domain

Google Apps + your apps

Our Google Apps




Your custom applications



Google's scalable serving architecture

Google Apps integration

☆ ▶



Google Apps

☆ ▶

Add your app to your Apps domain

Google app engine skirais@gmail.com | My Account | Help | Sign out

Application: axolotl100 No version deployed! [Show All Applications](#)

[Main](#)
[Dashboard](#)
[Quota Details](#)
[Logs](#)
[Cron Jobs](#)
[Task Queues](#)
[Datastore](#)

Add Domain

Domain Name:

Example: foo.com or bar.org

[Add Domain...](#) [Cancel](#)

Note: You must sign up for Google Apps to register this domain or prove that you already own it.

[Sign up for Google Apps »](#) | [Learn more](#)

Google Google Apps for [example.com](#) - Standard [See what's new!](#) [Inbox](#) [Calendar](#) [Help](#) [Sign out](#)

[Dashboard](#) [Users and groups](#) [Domain settings](#) [Advanced tools](#) [Support](#) [Service settings](#)

You have requested that the 'axolotl100' service be added to your domain

Please be careful when adding non-Google services to your domain, and make certain you know and trust the developer or originator of the service. Google cannot be held liable for any bad things that might happen as a result of adding this service to your domain.

Please accept the Google App Engine terms and conditions to continue

GOOGLE APP ENGINE TERMS AND CONDITIONS

Use of the Google App Engine is governed by the terms of service available at <http://code.google.com/appengine/terms.html>. By adding an application developed using the Google App Engine service (the "service") to your Google Apps domain, you understand and agree to the following:

Service is developed by a third party, not by Google. The Service provides you the ability to incorporate third-party applications into your use of Google Apps. These applications may not be developed by Google. You must independently evaluate

[Printable Version](#)

Agreement ☒ I accept. Continue to add this service. [Learn more](#)

By selecting "I accept, Continue to add this service," you are agreeing to the Google App Engine terms and conditions.

You can create custom web addresses for this service at any time on the service settings page.

[Activate this service](#) [Cancel](#)

Your app now part of your Apps suite

The screenshot shows the Google Apps dashboard for the domain **marzianicolai.com**. At the top, there's a navigation bar with links for Dashboard, Users and groups, Domain settings, Advanced tools, Support, and Service settings. Below this, a summary box for the domain shows 8 users and a link to create new users. A dropdown menu is open from the 'Service settings' link, showing a list of services: Email, Chat, Calendar, Docs, Sites, Mobile, my-app-id (App Engine), Code Reviews, Moderator, and Short Links. An arrow points to the 'my-app-id (App Engine)' option. Below the dropdown, the 'Service settings' section lists various services and their status: Email (Active), Chat (Active), Docs (Active), Mobile (Active), and Moderator (Active). The 'my-app-id' service is highlighted with a red box, showing it is 'App Engine - Active' with links to the appspot.com and the domain-specific URL.

Getting Started

<http://bit.ly/qcodelabs> (codelabs)
<http://code.google.com/appengine> (download, docs, etc.)
<http://code.google.com/p/googleappengine> (issues, wiki, etc.)
<http://code.google.com/eclipse> (Eclipse plug-in)
<http://appengine.google.com> (login, app management)
<http://googleappengine.blogspot.com> (blog)

App Engine online resources

<http://code.google.com/appengine> –
downloads, docs, forums, FAQ, etc.

<http://appengine.google.com> –
managing your applications here

<http://googleappengine.blogspot.com>
– our blog

Also check <http://stackoverflow.com>

Hello World

Linux, MacOS, etc. command-line:

```
$ dev_appserver.py helloworld # run dev svr  
$ appcfg.py update helloworld # deploy live
```

Windows GUI (also avail for Mac):



Project Contents

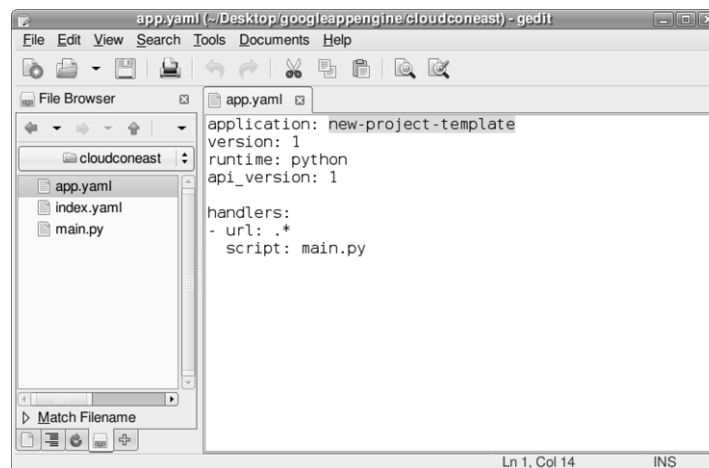


app.yaml – main configuration file

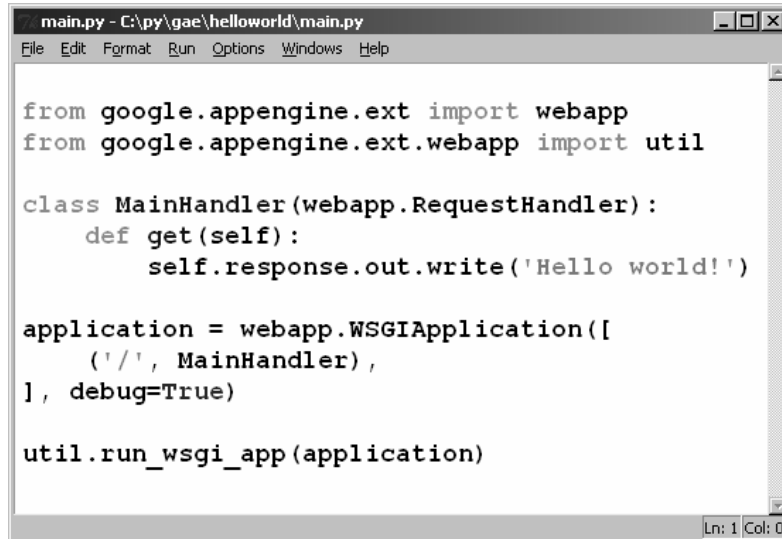
index.yaml – automatically generated to index your data

main.py – your main application "controller" code goes here

Hello World



Hello World

A screenshot of a Python IDE window titled 'main.py - C:\py\gae\helloworld\main.py'. The window has a menu bar with 'File', 'Edit', 'Format', 'Run', 'Options', 'Windows', and 'Help'. The code inside the window is as follows:

```
from google.appengine.ext import webapp
from google.appengine.ext.webapp import util

class MainHandler(webapp.RequestHandler):
    def get(self):
        self.response.out.write('Hello world!')

application = webapp.WSGIApplication([
    ('/', MainHandler),
], debug=True)

util.run_wsgi_app(application)
```

The status bar at the bottom right shows 'Ln: 1 Col: 0'.

Testing Your Install

```
$ cd google_appengine
~/Desktop/google_appengine/helloworld
$ dev_appserver.py helloworld
INFO      2009-03-04 17:51:22,354 __init__.py]
.
.
.
```

(optionally use the launcher for Macs & PCs)

Hello World



Registering

By default, your apps hosted at:
`http://APP-ID.appspot.com`

**Welcome to Google App Engine**



Run your web applications on Google's infrastructure.

Google App Engine enables developers to build web applications on the same scalable systems that power our own applications.

No assembly required.
Google App Engine exposes a fully-integrated [development environment](#).

It's easy to scale.
Google App Engine makes it easy to design scalable applications that grow from one to millions of users without infrastructure headaches.

Sign in with your
Google Account

Email:

Password:

☒ Remember me on this computer.

[I cannot access my account](#)

Don't have a Google Account?
[Create an account now](#)

Create an App Id

Google App Engine

joe.gregorio@gmail.com | [My Account](#) | [Help](#) | [Sign out](#)

My Applications

Application	Current Version
apex-learning-center	1
jcgregorio-test	1
just-overheard-it	1
project-playground	1
robaccia-test-app	1
sharded-counter-example	1
sparklines-bitworking	1
Create an Application	
You have 3 applications remaining.	

Modify 'app.yaml'

```
application: helloworld
version: 1
runtime: python
api_version: 1

handlers:
- url: .*
  script: main.py
```

Google

Upload the default application

```
$ cd google_appengine
~/Desktop/google_appengine/helloworld
$ appcfg.py update helloworld
Scanning files on local disk.
Initiating update.
Email: ...
```

Check it out



Thank You

Questions?

wesc+api@google.com
@wescpy / @app_engine


Recent+Upcoming Events

- Oct 18-20: Python course, San Francisco
 - <http://cyberwebconsulting.com>
- Jul 25-29 O'Reilly Open Source (OSCON), Portland
 - <http://oscon.com>
- Jul 11-13 ACM CSTA CS&IT Conference, NYC
 - <http://www.csitsymposium.org>
- Jun 20-25 EuroPython, Florence
 - <http://europepython.eu>
- May 8-10: Google I/O, San Francisco
 - <http://google.com/io>



Google cloud technologies

App Engine, Storage, Prediction, BigQuery



2

Google Storage

Store your data in a fast & reliable cloud!

Google Storage

- Cloud-based binary object store
 - Structured as buckets and objects
 - Many buckets, many objects, large objects
- You control your data
 - Private, shared, or public
 - Get your data back out at any time
- For developers
 - RESTful API
 - Many SDKs + tools
 - Integration with other Google services

Google Storage Benefits



High Performance and Scalability
backed by Google infrastructure

Flexible Authentication &
Sharing Models



Get Started Fast with
Google & 3rd Party Utilities

Google services using Google Storage

Pan^oramio
from Google

picnik

double
click
Partner Reporting

dl
Google
Data Liberation

google.org
Haiti Relief Imagery

You Tube[™]
Partner Reporting

Google
BigQuery

Google
Prediction API

Some current users

vmware[®]

syncplicity

QTECH

APPRIOR[™]

SnapABug



VIVU

Cloud
Sherpas

mēmeo[®]

widgetbox

theguardian

socialwok

XYLABS

Google Prediction API

Use machine learning in the cloud!!

Prediction API 101

Google's sophisticated machine learning algorithms

Available as an on-demand RESTful HTTP web service

Train a model offline/asynchronously

Predict results in real-time


















How does it work?

The Prediction API finds relevant **features** in the sample data during training.

The Prediction API later searches for those **features** during prediction.

"english"	The quick brown fox jumped over the lazy dog.
"english"	To err is human, but to really foul things up you need a computer.
"spanish"	No hay mal que por bien no venga.
"spanish"	La tercera es la vencida.
?	To be or not to be, that is the question.
?	La fe mueve montañas.

A virtually endless number of applications...

 Customer Sentiment	 Transaction Risk	 Species Identification	 Message Routing	 Diagnostics
 Churn Prediction	 Legal Docket Classification	 Suspicious Activity	 Work Roster Assignment	 Inappropriate Content
 Recommend Products	 Political Bias	 Uplift Marketing	 Email Filtering	 Career Counselling

... and many more ...

Three simple steps to use the Prediction API

1. Upload

Upload your training data to Google Storage

Use the API, gsutil or any compatible utility to upload your data to Google Storage

2. Train

Build a model from your data

`prediction/v1/train/{}
POST : a training request`

3. Predict

Make new predictions

`prediction/v1/query/{}
GET : model info
POST : a prediction request`

Prediction API Capabilities

Data

- Input Features: numeric or unstructured text
- Output: up to 100s of discrete categories

Training

- Many machine learning techniques
 - Automatically selected
 - Performed asynchronously

Access from many platforms:

- Web app from Google App Engine
- Apps Script (e.g. from Google Spreadsheet)
 - Desktop app