

Programming the Cloud: the Internet as Platform



Gregor Hohpe
Software Engineer
www.EnterpriseIntegrationPatterns.com



© 2009 Google, Inc. All rights reserved.

Internet as a Platform: The Good



Falling cost of
storage and
computing
power




Ubiquitous
broadband
connectivity



Democratized
tools of
production



© 2009 Google, Inc. All rights reserved.



Internet as a Platform: The Bad

Architect's Dream



- Loosely coupled
- Extensible
- Standards-based
- Fault tolerant
- Unlimited computing power
- Ubiquitous


Developer's Nightmare



- NO Call Stack
- NO Transactions
- NO Promises
- NO Certainty
- NO Ordering Constraints



© 2009 Google, Inc. All rights reserved.



Internet as a Platform: The Ugly

ACID (before)


- Atomic
- Consistent
- Isolated
- Durable

ACID (today)

- Associative
- Commutative
- Idempotent
- Distributed



© 2009 Google, Inc. All rights reserved.




Starbucks Does not Use 2-Phase Commit Either

- Start making coffee before customer pays
- Reduces latency
- What happens if...

Customer rejects drink	➔	Remake drink Retry
Coffee maker breaks	➔	Refund money Compensation
Customer cannot pay	➔	Discard beverage Write-off



© 2009 Google, Inc. All rights reserved.




Programming the Cloud – The Google Way

- Fault tolerant distributed disk storage: *Google File System*
- Distributed shared memory: *Bigtable*
- New programming abstractions: *MapReduce*
- Domain Specific Languages: *Sawzall*




Google.stanford.edu (Circa 1997)



Current Rack Design

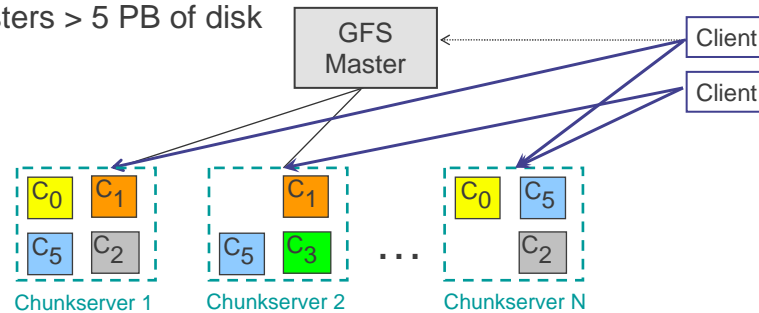


© 2009 Google, Inc. All rights reserved.



Fault Tolerant Distributed Disk Storage: GFS


- Data replicated 3 times. Upon failure, software re-replicates.
- Master: Manages file metadata. Chunk size 64 MB.
- Optimized for high-bandwidth sequential read / writes
- Clusters > 5 PB of disk



<http://research.google.com/archive/gfs-sosp2003.pdf>



© 2009 Google, Inc. All rights reserved.




Distributed Shared Memory: Bigtable

- Sparse, distributed, persistent, multidimensional, sorted
- Not a relational database (RDBMS): no schema, no joins, no foreign key constraints, no multi-row transactions
- Each row can have any number of columns, similar to a dictionary data structure for each row.
- Basic data types: string, counter, byte array
- Accessed by row key, column name, timestamp
- Data split into tablets for replication
- Largest cells are > 700TB

<http://research.google.com/archive/bigtable-osdi06.pdf>



© 2009 Google, Inc. All rights reserved.




Programming Abstraction: MapReduce

- Represent problems as Map and Reduce step (inspired by functional programming)
- Distribute data among many machines, execute same computation at each machine on its dataset
- Infrastructure manages parallel execution
- Open source implementation: Hadoop

```
map(in_key, data)
  → list(key, value)


reduce(key, list(values))
  → list(out_data)
```



<http://research.google.com/archive/mapreduce.html>



© 2009 Google, Inc. All rights reserved.





Parallel Log Processing: Sawzall

- Separate stateful aggregation from stateless record processing
- Commutative and associative aggregation enables parallelized execution and aggregation

```
while !eof(file) {
  line = readline(file);
  words = split(line);
  for each word in words:
    map[word]++;
}
```

```
map: table sum[string] of int;
line = input;
words = split(line);
for each word in words:
  emit map[word] <- 1;
```



<http://labs.google.com/papers/sawzall.html>




© 2009 Google, Inc. All rights reserved.



Google, the Cloud, and You!



Make the cloud more **accessible**



Make the client more **powerful**



Keep connectivity **pervasive**

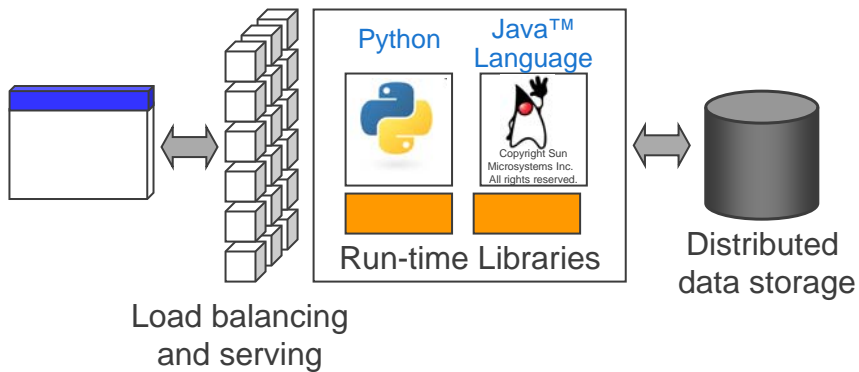


© 2009 Google, Inc. All rights reserved.



Google App Engine – Easy to Start, Easy to Scale

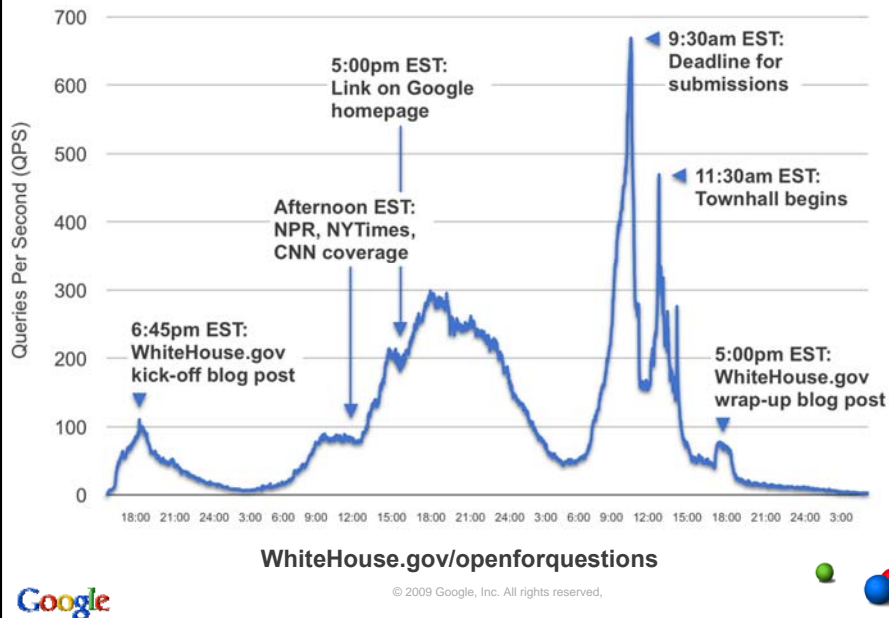
Your code running on Google infrastructure



© 2009 Google, Inc. All rights reserved.



Does it Scale?




Key Features

- Python or Java™ Language source code
- Develop locally, deploy to Cloud seamlessly
- Write once, scale automatically
- Local SDK & Eclipse Plugin
- Free quota of 500MB storage and ~5M pageviews / month



© 2009 Google, Inc. All rights reserved.




Programming & Run-time Model

- Responds to HTTP requests
- A programming platform, not “raw iron”
- API support for
 - User login and identity
 - Persistent state (on top of Bigtable, not RDBMS)
 - memcache
 - Mail, Images, URL Fetch
 - Django Templates / JSP



© 2009 Google, Inc. All rights reserved.



Management Console

onthally Version: 1.54 [Show All Applications](#)

Main

- [Dashboard](#)
- [Quota Details](#)
- [Logs](#)

Datastore

- [Indexes](#)
- [Data Viewer](#)

Administration

- [Application Settings](#)
- [Developers](#)
- [Versions](#)
- [Admin Logs](#)

Billing

- [Billing Settings](#)
- [Billing History](#)

Resources

- [Documentation](#)
- [Developer Forum](#)
- [Downloads](#)
- [System Status](#)

Charts

Requests/Second

Billing Status: Enabled - Settings

Resource (reset every 24 hours. Next reset: 10 hrs)	Usage	Cost / Budget
Processor \$0.10/CPU hour	94% 48.30 of 51.30 hours	\$0.20 / \$0.40
Bandwidth In \$0.10/Gbyte	80% 12.00 of 15.00 Gbytes	\$0.20 / \$0.40
Bandwidth Out \$0.12/Gbyte	99% 14.10 of 14.17 Gbytes	\$0.17 / \$0.40
Storage \$0.005/Gbyte	25% 25.12 of 100.60 Gbytes	\$0.12 / \$0.40
Email \$0.0001/Message	20% 500 of 2500 Messages	\$0.03 / \$0.40

↑ = Free quota Cost for the last 14 hours: \$0.69

Current Load

URI	Req/Sec current	Requests last 12 hrs	Avg CPU last hr	% CPU last 12 hrs
/	450.0	450	2	0%

Errors

URI	Count last 12 hrs	% Errors last 12 hrs
/	39	9%



© 2009 Google, Inc. All rights reserved.




Google App Engine for Java

Web App Container	Java Servlet	JSR-154
Datastore	JDO & JPA	JSR-220, JSR-243
HTTP	Java.net.URL	Java SE
Mail	Javax.mail	JSR-919
Memcache	Javax.cache	JSR-107

Java is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries



© 2009 Google, Inc. All rights reserved.




1 Year Anniversary: What's new?

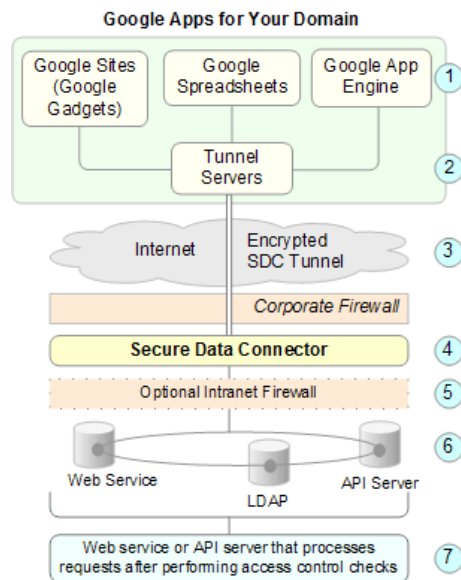
- 1. Purchase Additional Resources**
- 2. Cron Support**
- 3. Database Import**
- 4. Access to firewalled data:
Secure Data Connector**




© 2009 Google, Inc. All rights reserved.



Secure Data Connector



© 2009 Google, Inc. All rights reserved.




Programming the Cloud

- Programming the cloud requires different programming and run-time models
- Parallel execution, constraint-based programming instead of linear loops
- Distributed data storage instead of RDBMS
- Old ACID vs. New ACID
- The right tools and API's can make this a lot easier



© 2009 Google, Inc. All rights reserved.




Google and the Cloud

- Google Data API's
- Google App Engine
- Academic Cloud Computing Initiative (IBM & Google)
 - <http://code.google.com/edu/parallel>
- Developer community
 - <http://code.google.com/apis>
- Open Source
 - <http://code.google.com/opensource/>



© 2009 Google, Inc. All rights reserved.



日本語 Google Code サイト

The screenshot shows the Google Code website in Japanese. The main heading is "OpenSocial". Below it, there is a section titled "ソーシャル サイトへのトレンド" (Trends for Social Sites). The text discusses how social sites are becoming more popular and how OpenSocial provides APIs for developers to integrate with these sites. There is also a section titled "1つのAPIで多様なサイトに対応" (Supporting multiple sites with one API), which lists various social sites that are supported by the OpenSocial API, including Engage.com, Friendster, hi5, Hyves, imooon, LinkedIn, MySpace, Ning, Oracle, orkut, Plaxo, Salesforce.com, Six Apart, Tianji, Viadeo, and XING.

日本語ドキュメント

- Google Chart API
- Google Gadgets API
- Google Maps API
- OpenSocial
- YouTube Data API
- Google AJAX APIs
- Google AJAX Feed API
- Google AJAX Search API
- Google Gadgets API
- Google Themes API
- Google App Engine (partial)
- Google Visualization API
- KML (partial)
- iGoogle Gadget
- Orkut Developer Home

code.google.com/intl/ja/
www.google.co.jp/developer/prg/



© 2009 Google, Inc. All rights reserved.

