The Cloud and Google

International Supercomputing Conference '09

June 24th, 2009

Dr. Robin Williamson
Google Inc.

Outline

- The Cloud to Google
- Technology Stack
  - Hardware computing platform
  - Distributed systems software infrastructure
  - Products
- Development Platform and Tools
  - Google App Engine

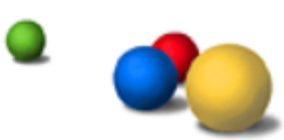

Berkeley View of Cloud Computing

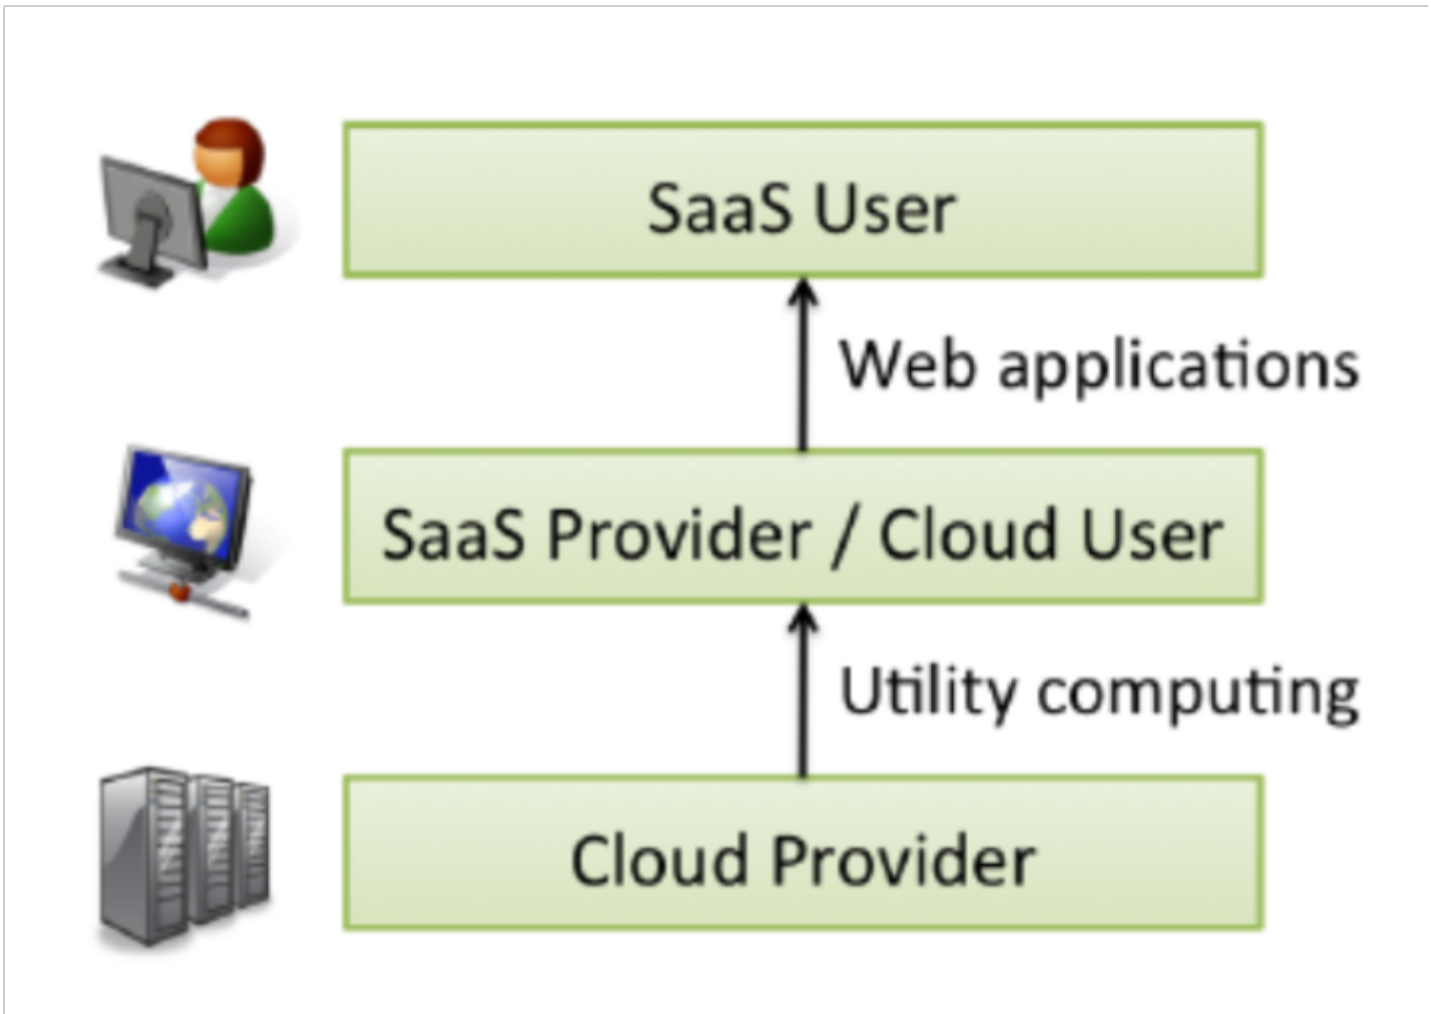

[*Above the Clouds: A Berkeley View of Cloud Computing, Ar* mbrust et al. 2009]

Google

Where We Play

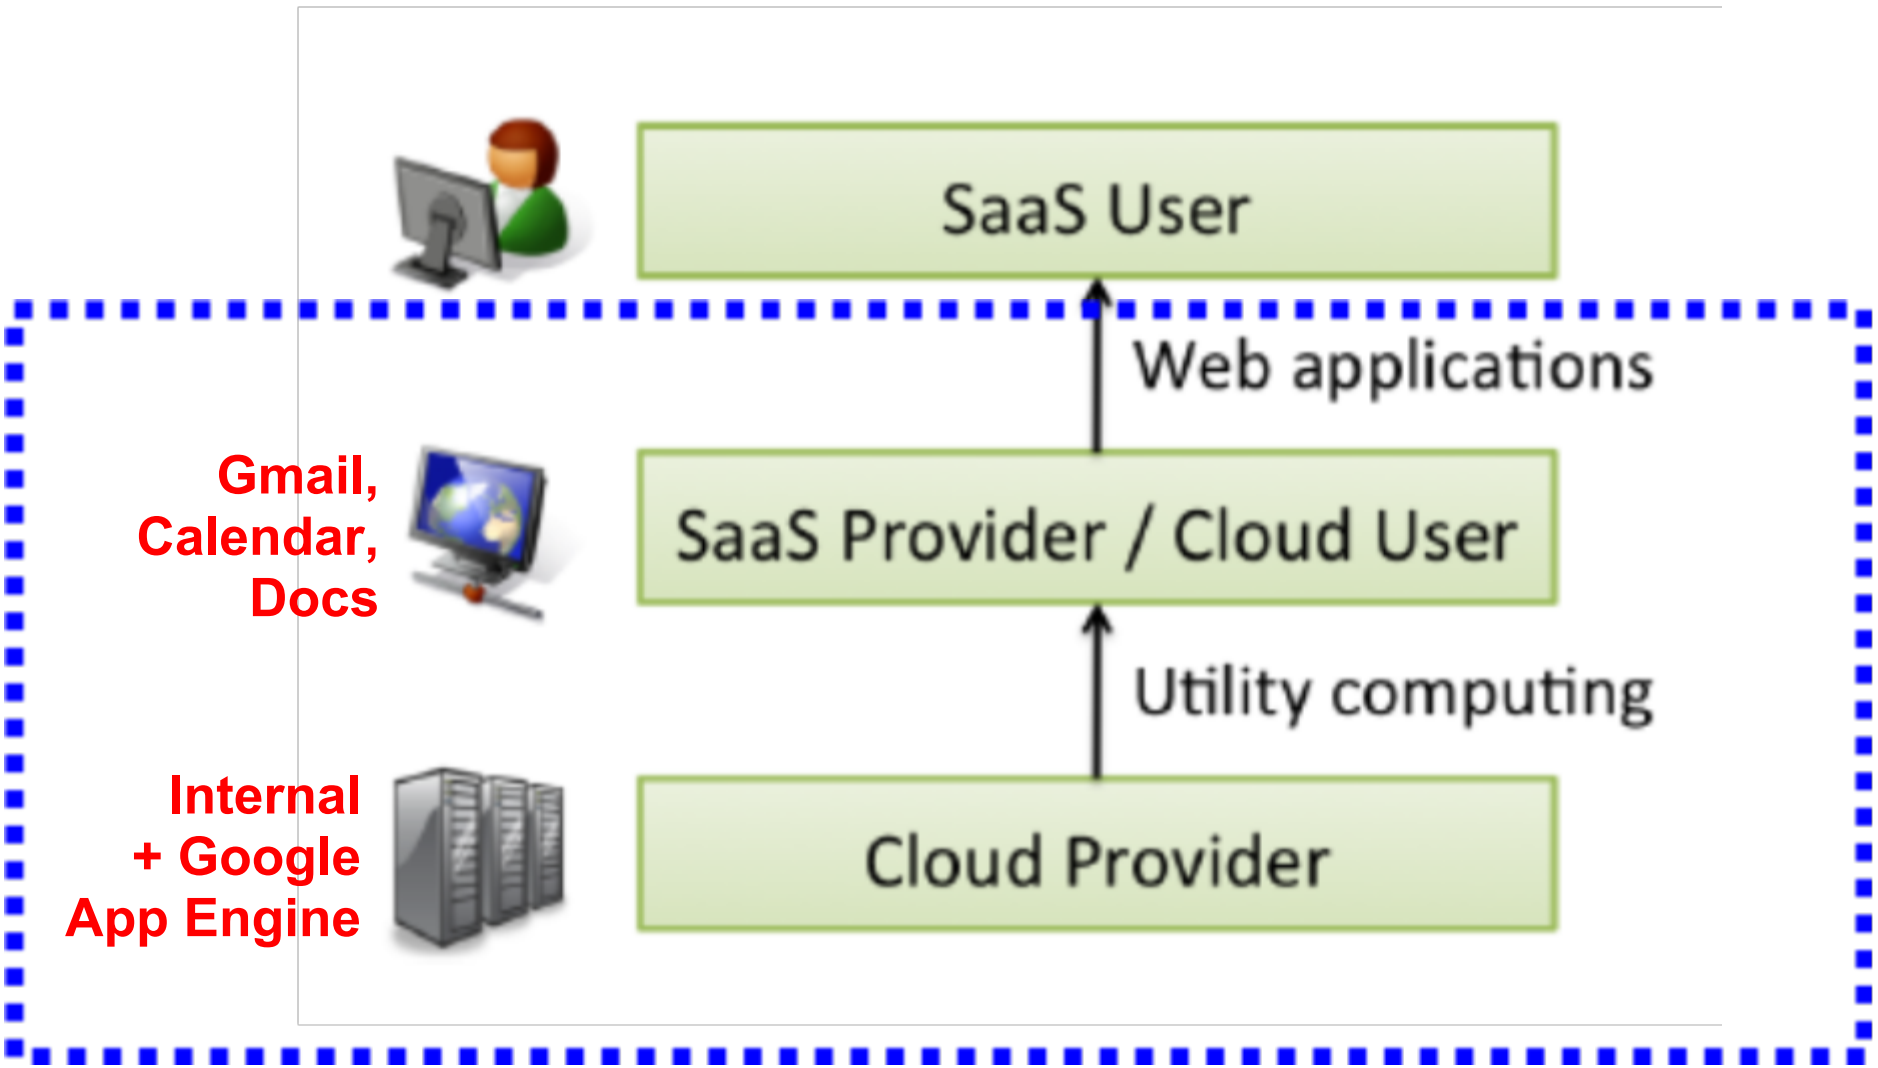

Google's Mission Statement

To organize the world's information and make it universally accessible and useful.

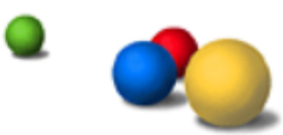

The *World's* Information

- There's *lots* of data—100s of TBs just for the Web
- Much of the data is common across large numbers of users
- Bandwidth is slow and expensive in the last mile (and improving more slowly than other dimensions)

- **Analyzing, transforming, querying all best done centrally (adjacent to the data)**

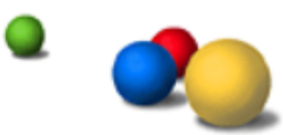

Query Serving Infrastructure

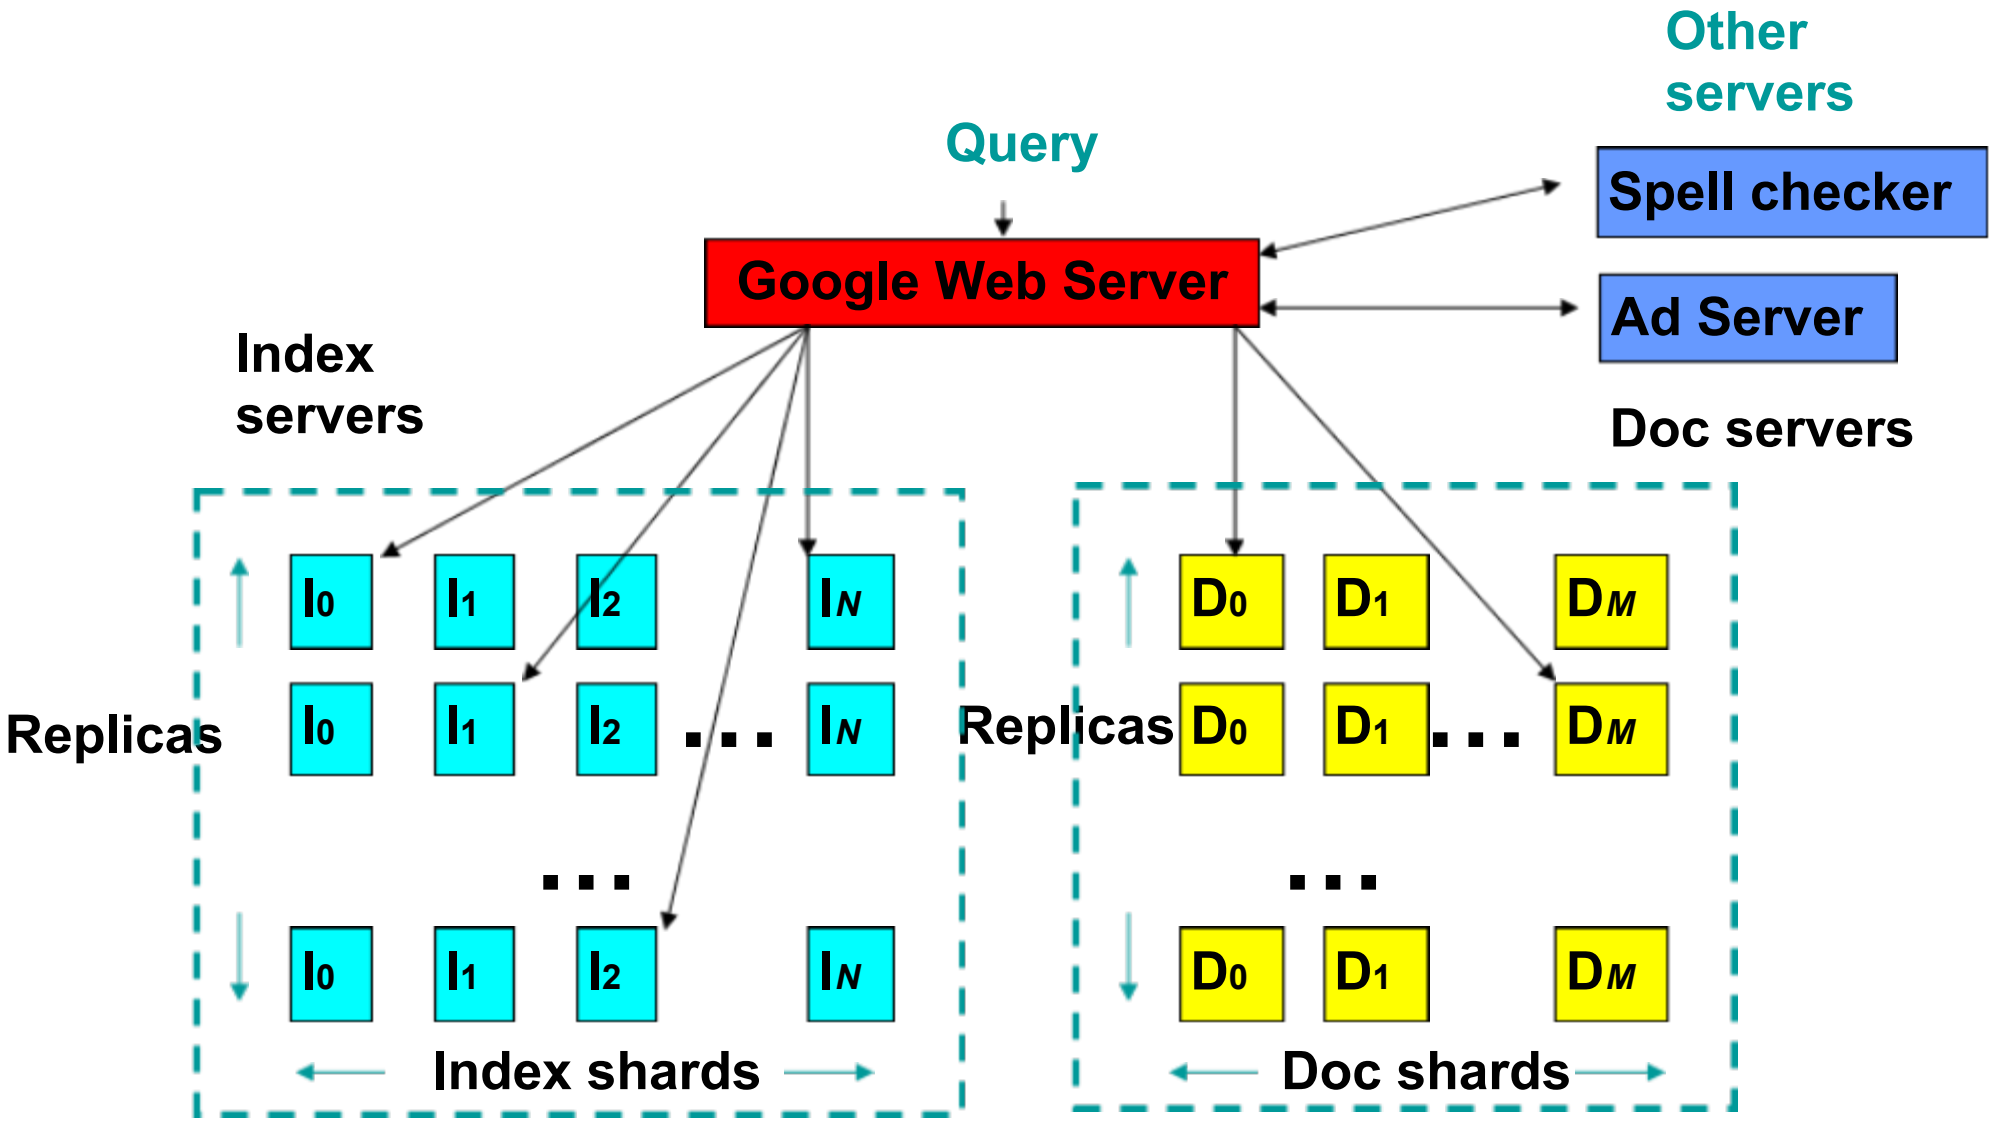

Elapsed time: 250ms, machines involved: 1000+

Outline

- The Cloud to Google
- Technology Stack
  - Hardware computing platform
  - Distributed systems software infrastructure
  - Products
- Development Platform and Tools
  - Google App Engine

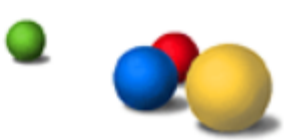

Technology Stack

- Innovation at all layers…

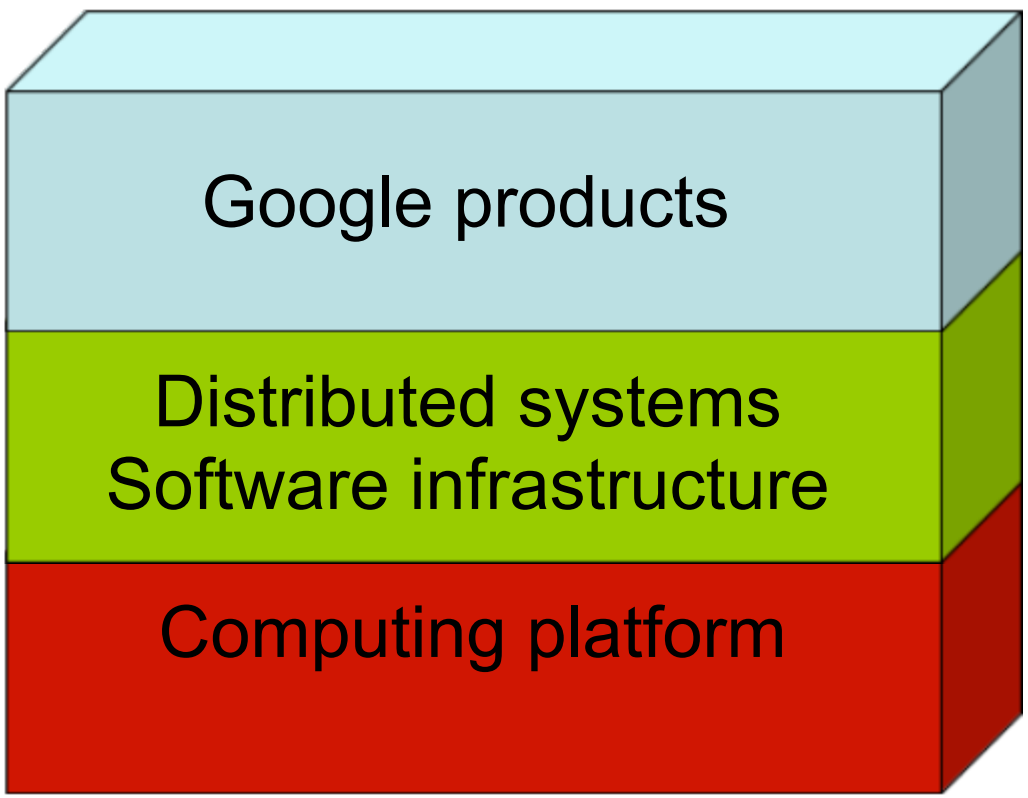

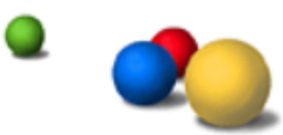

Computing Platform

- Single-threaded performance matters less
- Moore's law manifests as more cores
- The computer is the datacenter

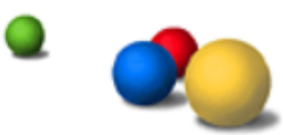

Higher-level Programming Abstractions

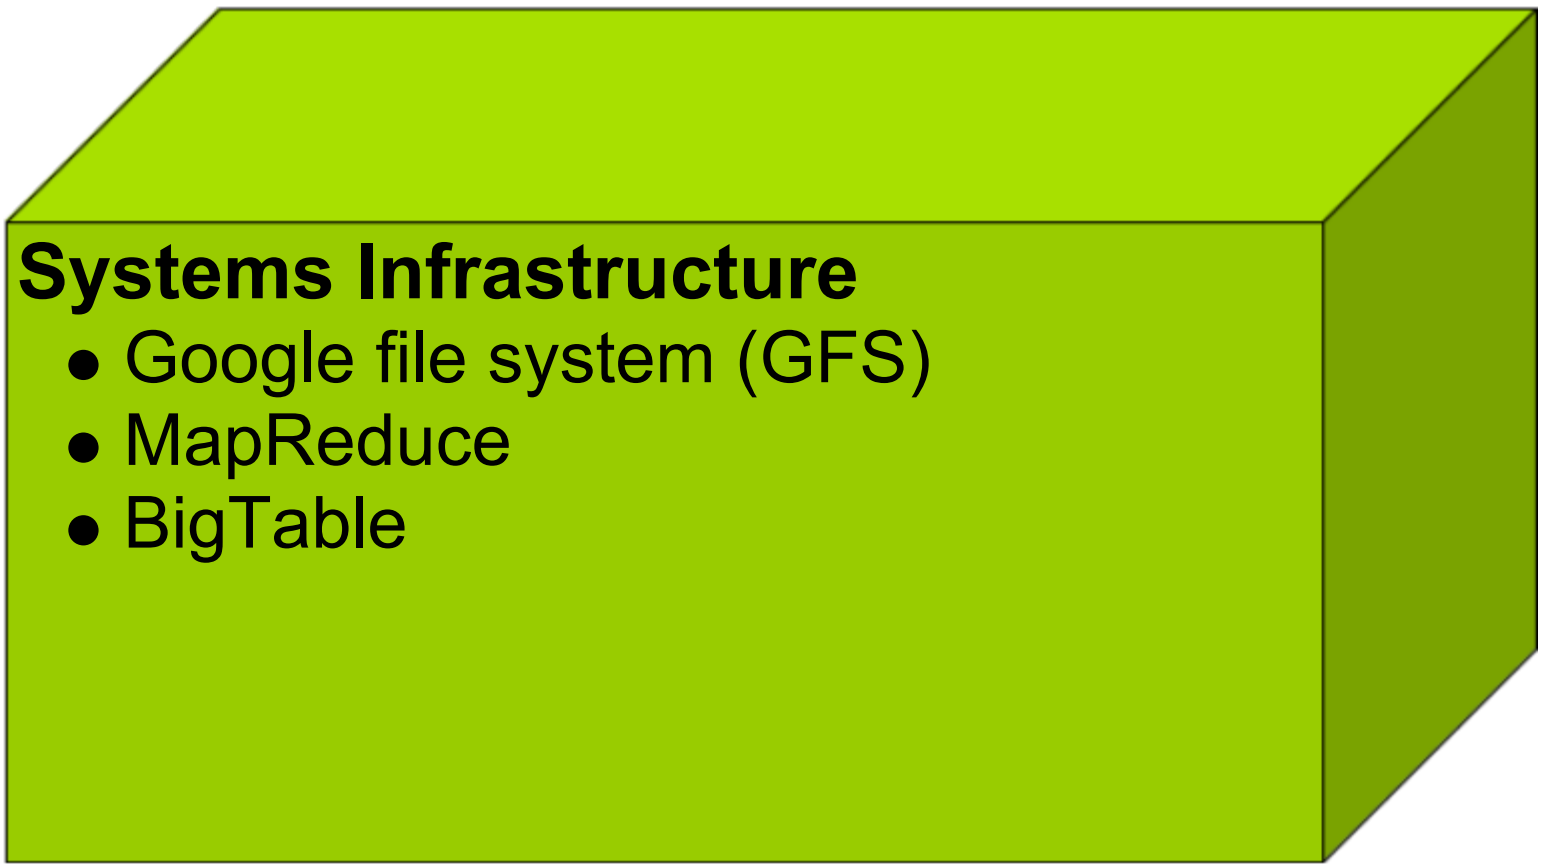

Google

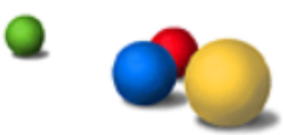

System Infrastructure

- **Google File System (GFS):**
  - Fault tolerant distributed disk storage
  - Optimized for high-bandwidth sequential read/writes
- **BigTable:**
  - A large-scale storage system for semi-structured data
  - Database-like model, but data stored on thousands of machines..
- **MapReduce:**
  - A programming model and library to simplify large-scale computations on large clusters
  - Distribute data among many machines, execute same computation at each machine on its dataset

Top of the Stack

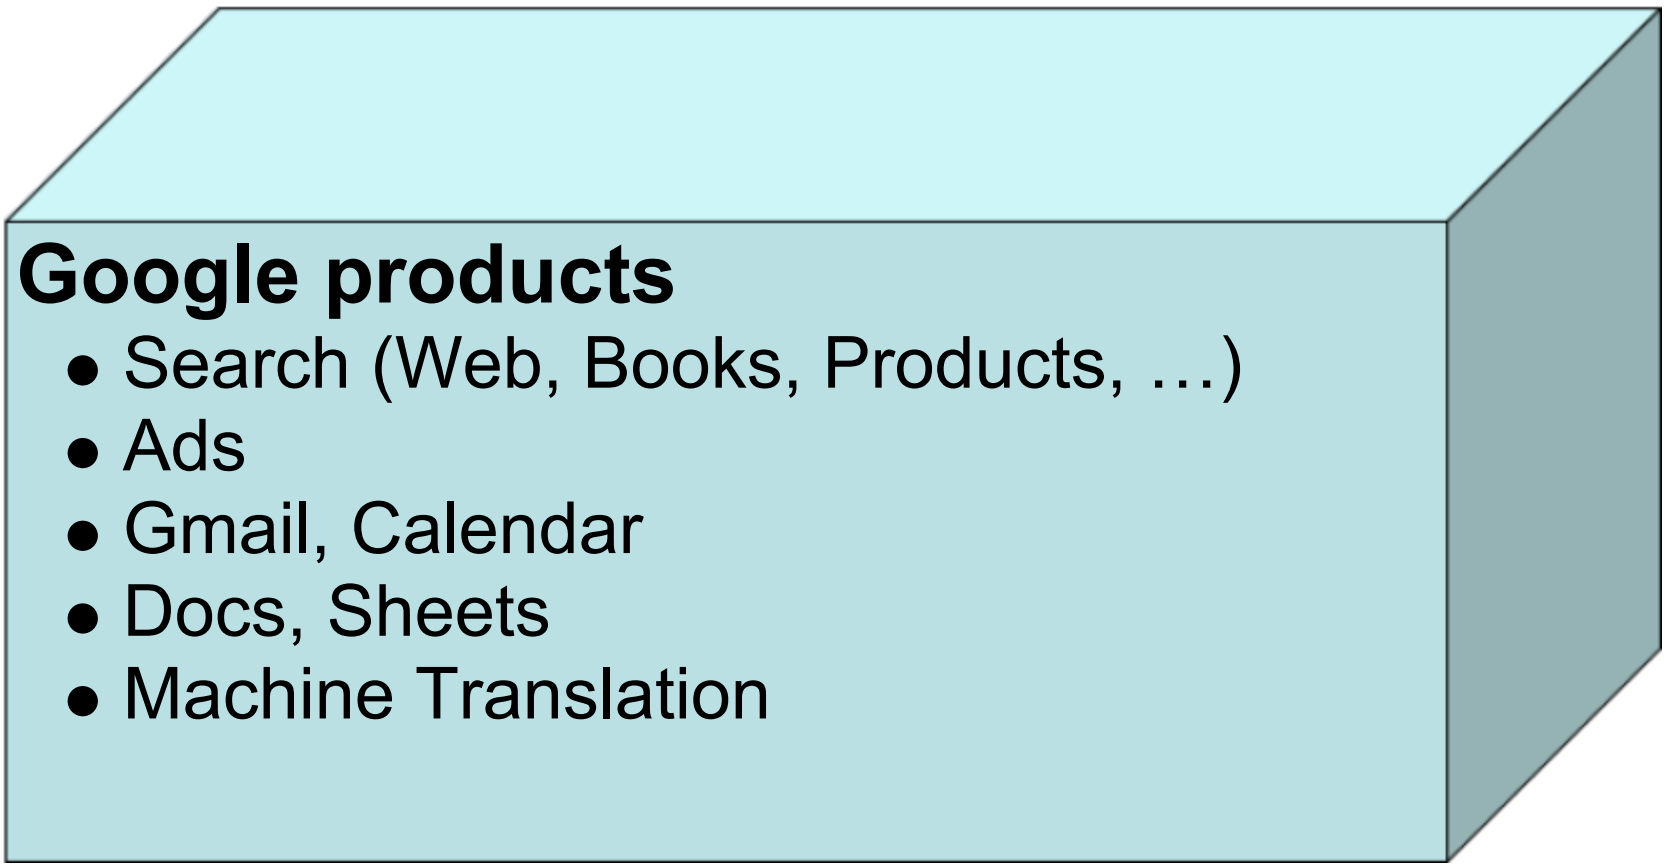

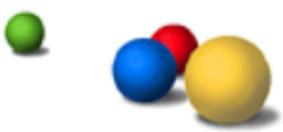

Outline

- The Cloud to Google
- Technology Stack
  - Hardware computing platform
  - Distributed systems software infrastructure
  - Products
- Development Platform and Tools
  - Google App Engine

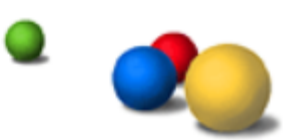

Google App Engine – Easy to Start, Easy to Scale

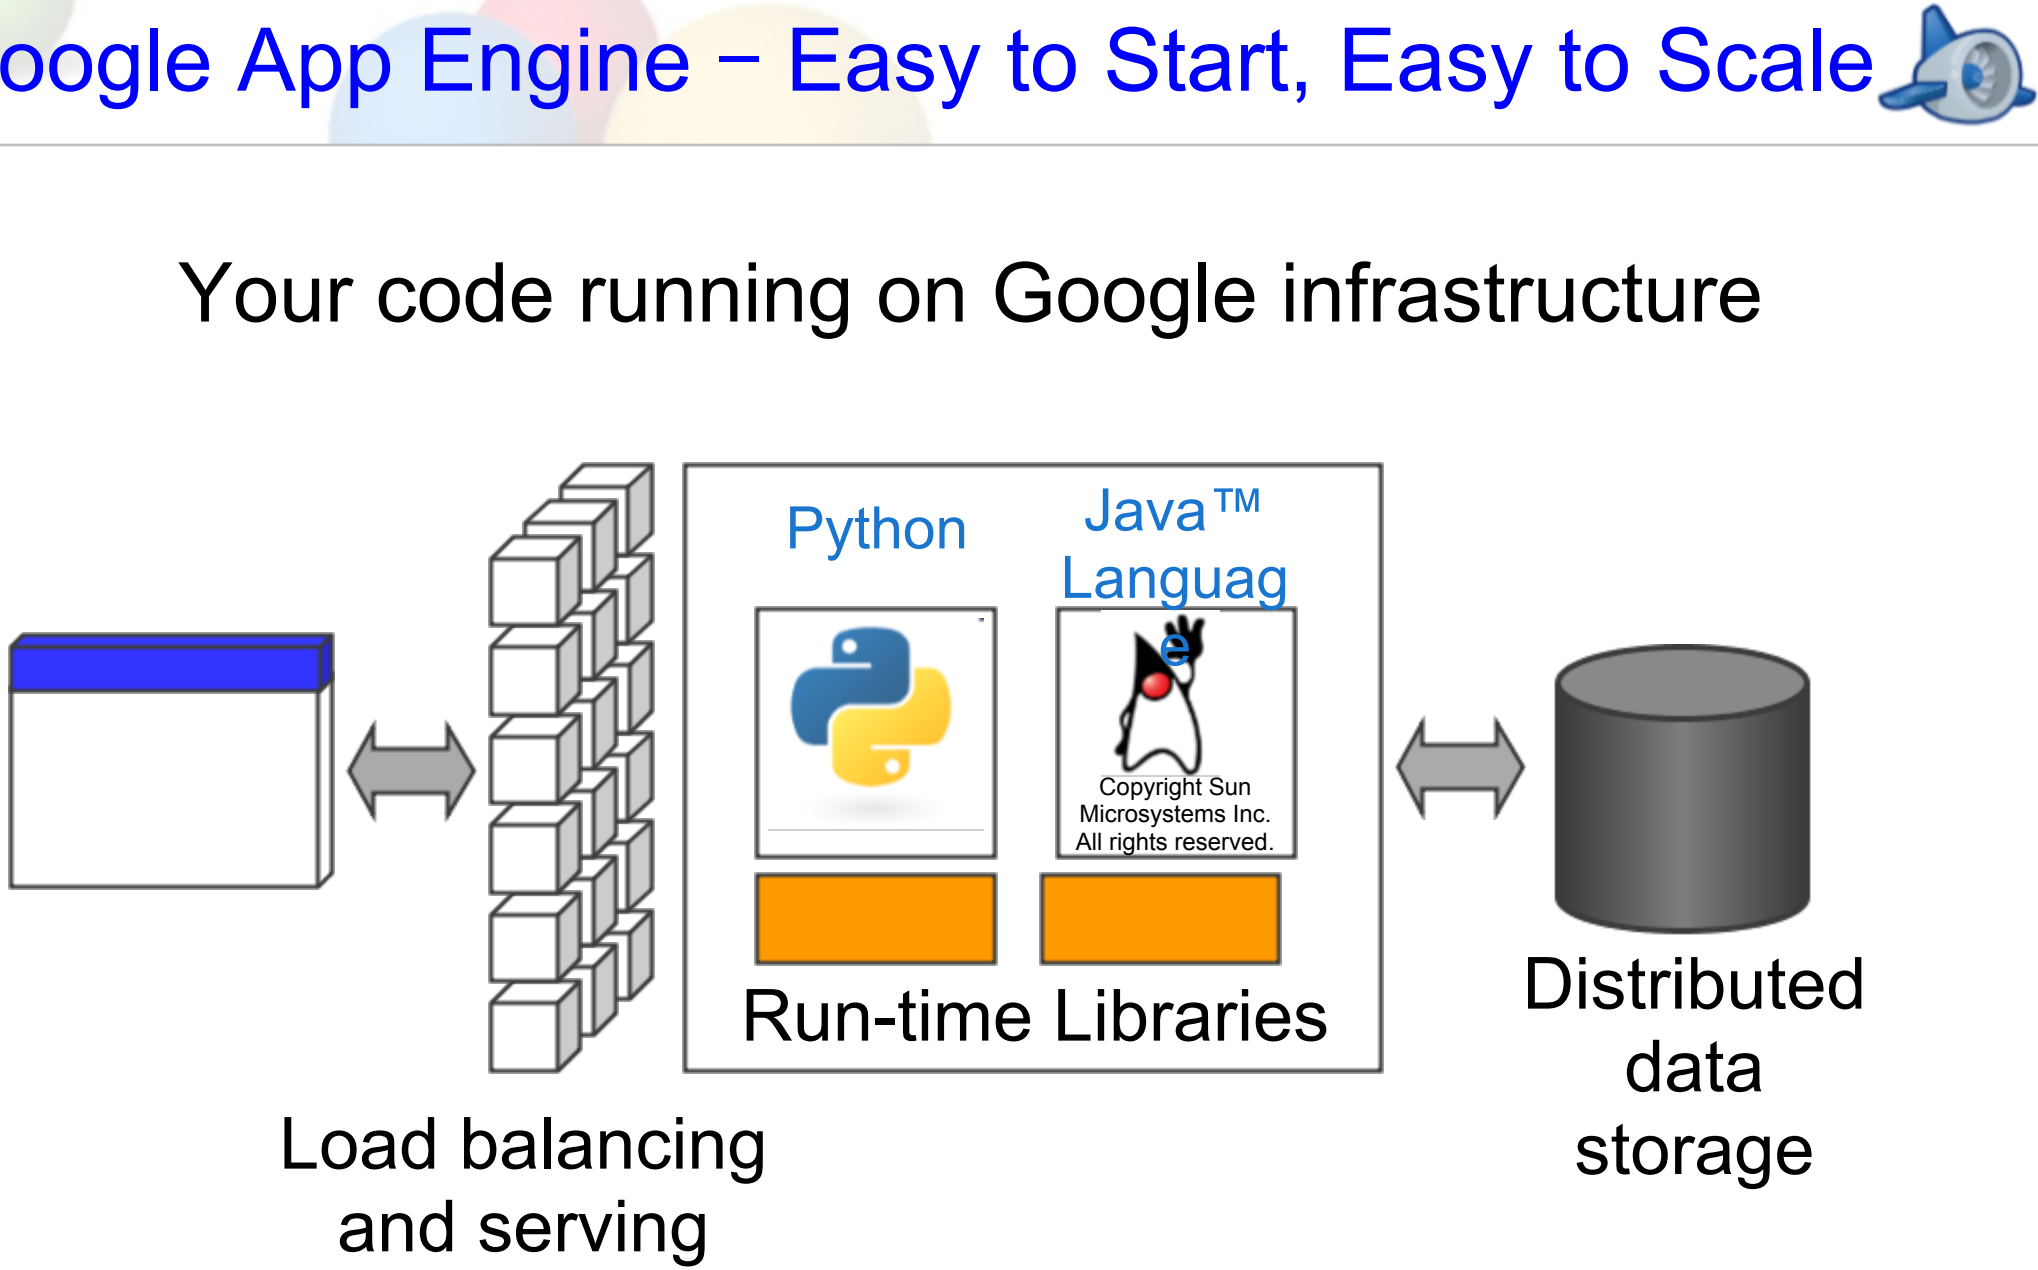

Key Features

- Python or Java™ Language source code
- Develop locally, deploy to Cloud seamlessly
- Write once, scale automatically
- Local SDK & Eclipse Plugin
- Free quota of 1 GB storage and ~5M pageviews / month

Google

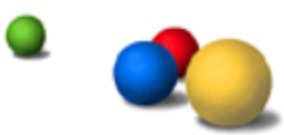

Programming & Run-time Model

- 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

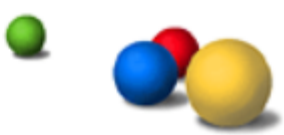

Does it Scale?

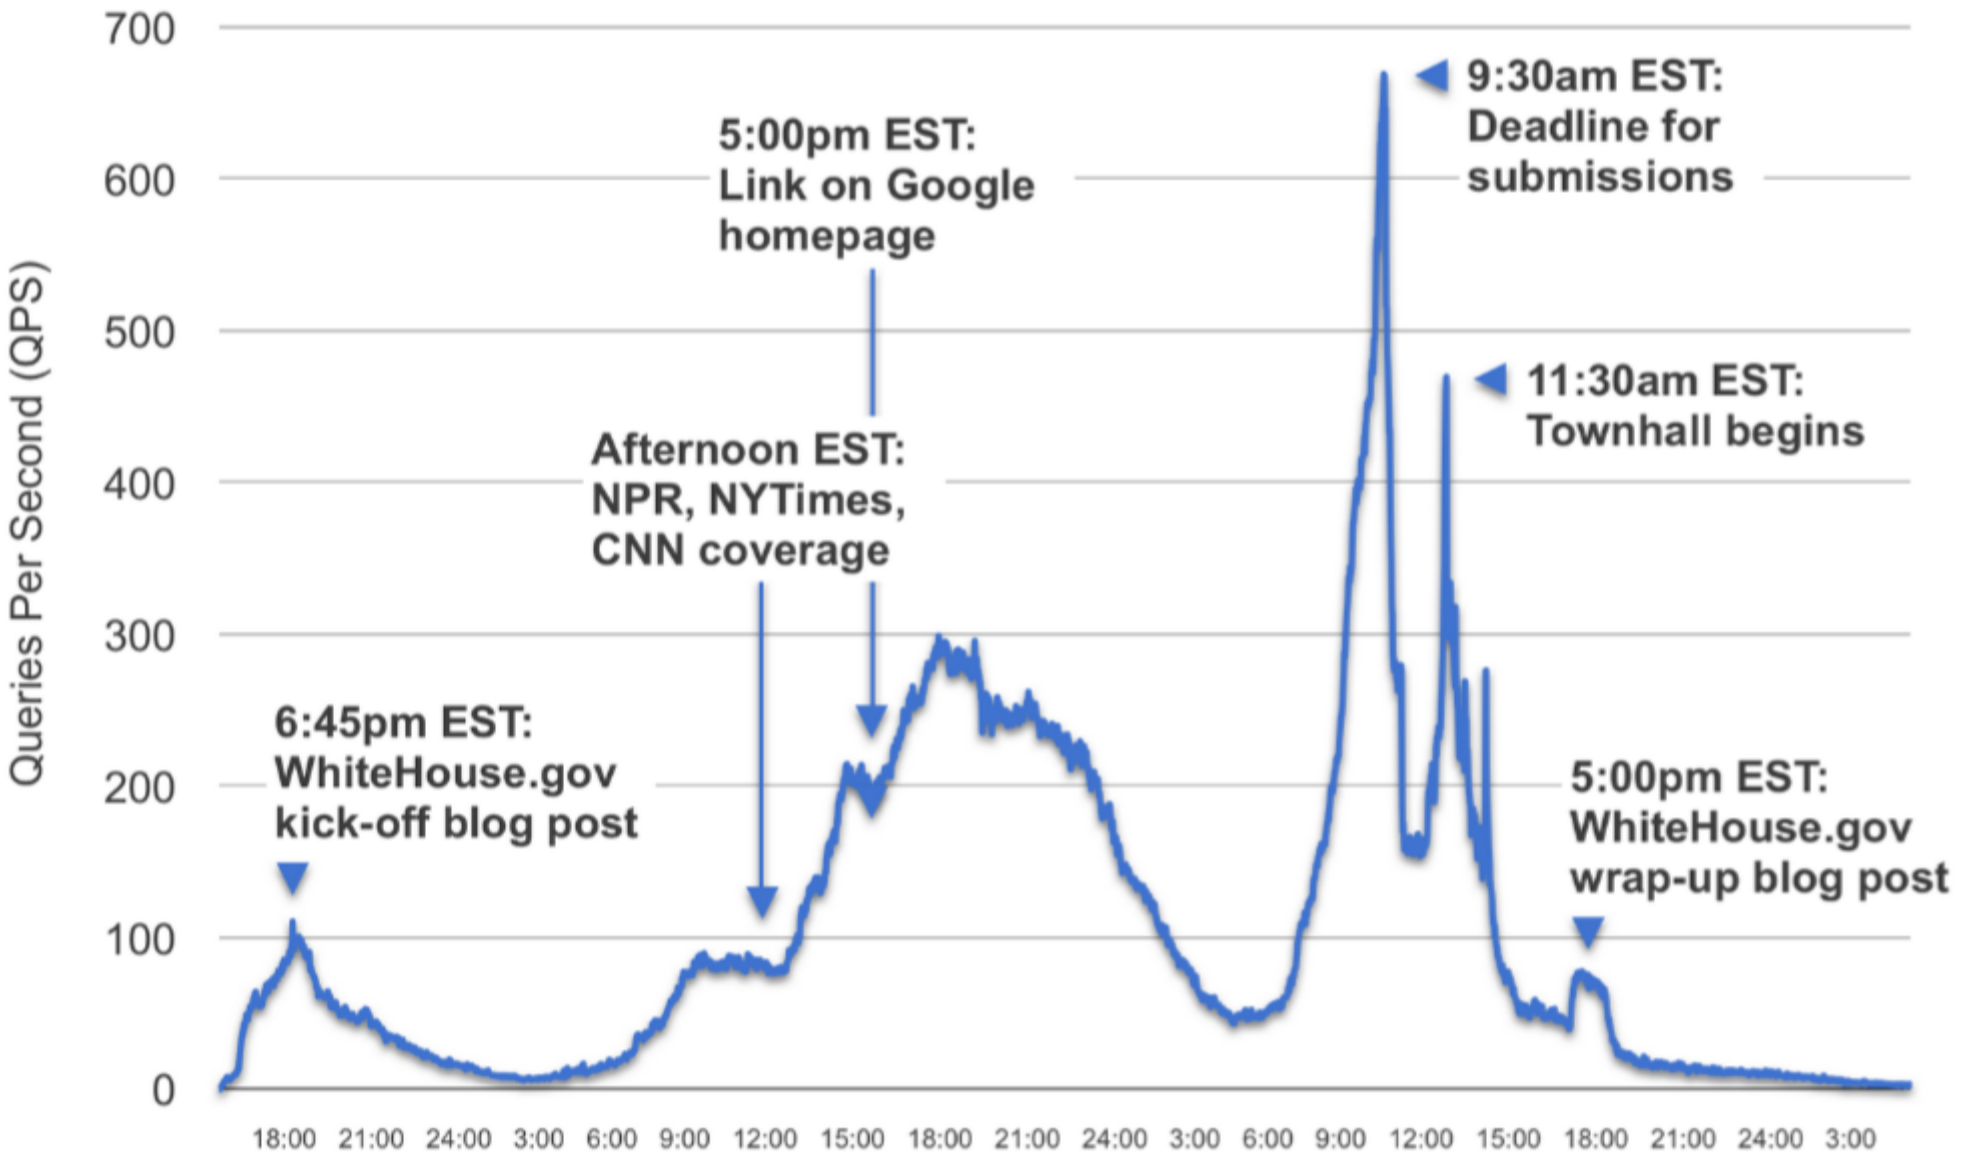

WhiteHouse.gov/openforquestions

Google

5M Free PVs + Pay as you Go

Admin Interface

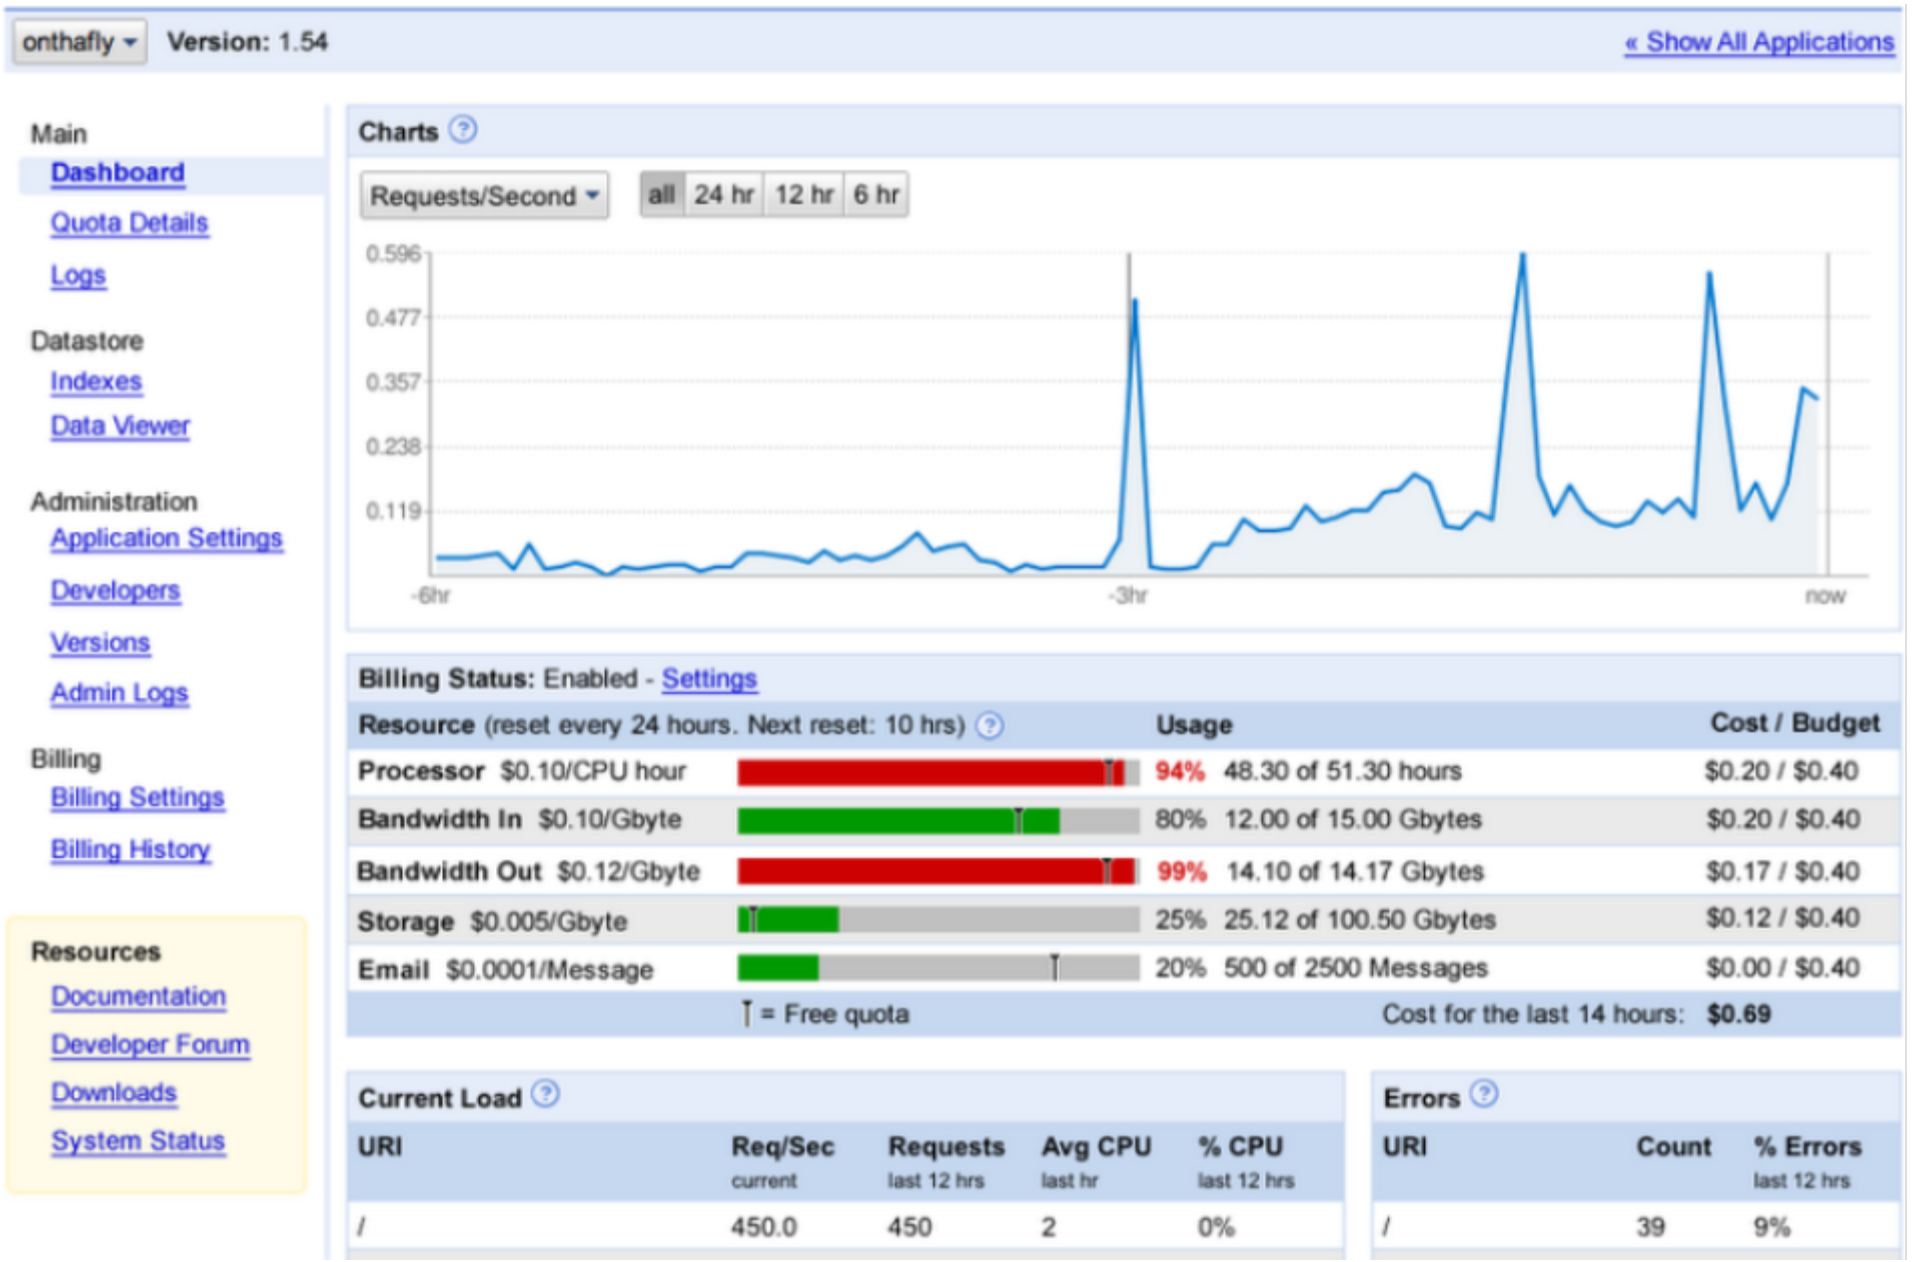

Thank you