

T-Mobile Chair for Mobile Business & Multilateral Security

MOB2 - Google Exercise

Team 1

Benjamin Schäfer

Chris Schneider

Raphael Schöbel

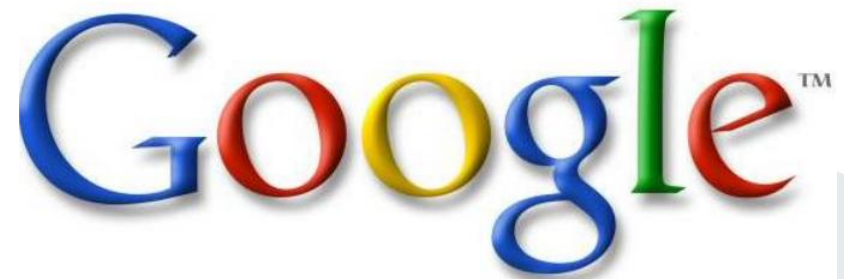
T-Mobile Chair for M-Business & Multilateral Security
J.W. Goethe Universität Frankfurt a. M.



Tasks Team 1

- 1) Describe the current operating model of Google and explain their objectives regarding the mobile internet
- 2) Give an overview of the telecommunication products and services which Google offers to the business and consumer customers
- 3) Give an overview of the hardware and software stack from Google

- Search advertising
- Display advertising
- Mobile advertising
- Tools for publishers
- Local Enterprise
- Enterprise Solutions



Google™

- \$1 billion in revenue in 2011
- “Mobile search is definitely going to surpass desktop search“
- Sustain market share of 97 % for Mobile searches
 - Voice Search
 - Goggles
 - Translate spoken conversations



Googles Telecommunication Products and Services



Search

- Google Websearch
- Google Alerts
- Google Insights

Advertising

- AdWords
- AdSense
- Google Analytics

Communication

- Google Mail
- Google Maps
- YouTube
- Google Chrome

For Developers

- Google Labs
- Google Base
- Google Reader

Consumer vs. Business Segment

CONSUMER


- Websearch
- Google Mail
- Google Toolbar
- Picasa

BUSINESS

- AdSense
- AdWords
- Google Analytics
- Google Insights

- The exact size of the data centers Google uses are unknown and facilities are not open to tours
- Google has numerous data centers around the world:
 - at least 12 significant Google data center installations are located in the United States
 - in Europe, the largest known centers are in Eemshaven and Groningen (NL) and Mons (B)
 - Google's Oceania Data Center is claimed to be located in Sydney (AUS)
- Strategy focuses on purchasing CPU generations that offer the best performance per dollar
- Network topology
 - DNS servers resolve `www.google.com` into multiple IP addresses which acts as the first level of load balancing and directs the client to different Google clusters
 - a cluster has thousands of servers and once the client has connected to the server additional load balancing is done to send the queries to the least loaded web server
 - this makes Google one of the largest and most complex content delivery networks

- Commodity-class x86 PCs running customized versions of Linux
- Hardware history of Google
 - 2003: one or more 80 GB hard disks per server
 - 2004: 2-4 GB of memory per machine
 - 2005: 200,000 servers
 - 2006: this number is claimed to be upwards of 450,000
 - 2009: ca. 16 GB RAM, 2 TB disk space per machine
- Power requirement for over 450,000 servers are about 20 megawatts which cost US\$ 2 million per month



Google's first production server rack (1999)

Google's server infrastructure is divided in several types, each assigned to a different purpose:

- **Google web servers**
 - coordinate the execution of queries sent by users, then format the result into an HTML page
- **Data-gathering servers**
 - are permanently dedicated to spidering the Web
- **Document servers**
 - return a summary for the searched document based on query words.
 - they can also fetch the complete document when asked; thus these servers need more disk space
- **Ad servers**
 - manage advertisements offered by services like AdWords and AdSense
- **Spelling servers**
 - make suggestions about the spelling of queries

- Most of the software stack was developed in-house
- Probably favored programming languages
 - C++
 - Java
 - Python
- Google Web Servers are custom Linux-based
- The software infrastructure includes special features
- Software development practices
 - focus on fault tolerance (when a system goes down, data is still available on other servers, which increases reliability)
 - queries are divided into sub-queries, where those sub-queries may be sent to different ducts in parallel