



Desired Properties in a Storage System  
(For building large-scale, geographically-distributed services)

Jeff Dean

Google Fellow

[jeff@google.com](mailto:jeff@google.com)

# Desired File System Characteristics

- Single global namespace
  - across many geographically dist. data centers
  - **data needs to be replicated** to multiple geographic regions for availability, reliability, and low-latency access
  - name for a piece of data is independent of its location(s)
    - `/user/jeff/gmail/2009/msg1376.subject`
- Large scale:
  - Deal well with many tiny files
  - Support  $\sim 10^{13}$  dirs,  $\sim 10^{15}$  files,  $\sim 10^{18}$  bytes of storage
  - Handle  $\sim 10^5$  to  $10^7$  machines, distributed in 100s to 1000s of locations around the world
  - Support direct access from  $\sim 10^9$  client machines (maybe?)

# Automated Management

- Users specify desired properties for data
  - “keep 5 copies of this data: 2 in U.S., 2 in Europe, 1 in Asia”
  - “map this kind of data into memory”
  - “99%ile latency to access this data should be  $\leq 50$  ms”
  - “never store this data in country X”
- Placement/replication decisions made automatically
  - based on hints, plus access statistics
  - while trying to minimize various costs (storage, bandwidth, access latency, etc.)
- Ability to attach computations to data
  - when data moves or is replicated, computation automatically moves, too

# Consistency and Sharing

- Support both strong-consistency and weak-consistency access modes
- Handle fine-grained sharing ( $\sim 10^9$  clients)
- Efficiently find and search all data accessible to a given user