



Jon Weissman  
Department of CS&E  
University of Minnesota

Open Cirrus October 2011

# Key Trends

- Client technology
  - devices: smart phones, tablets, sensors
- Big data
  - distributed and dynamic
- Privacy/trust
  - small circles
- Multiple DCs
  - global services

# Minnesota Cloud Research

- Eye towards cloud evolution
- Projects
  - Nebula
  - Mobile cloud
  - DMapReduce
  - Proxy cloud
  - Active cloud storage
  - Virtualization

# Big Data Trend


- Big data is distributed
  - earth science: weather data, seismic data
  - life science: GenBank, PubMed
  - health science: GoogleEarth + CDC pandemic data
  - web 2.0: user multimedia blogs
  - “everyone is a sensor”
- Cost in moving data to the cloud

# Big Data Trend: Nebulas


- Process data “close by”
  - fully and/or on-route to the central cloud
  - cost: save time and money
  - privacy (think: patient records, local google doc)
- Close by
  - network distance
  - trusted peers

# Example: Dispersed-Data-Intensive Services

- Data is geographically distributed
  - Costly, inefficient to move to central location




# Example Instance: Blog Analysis



# Nebula

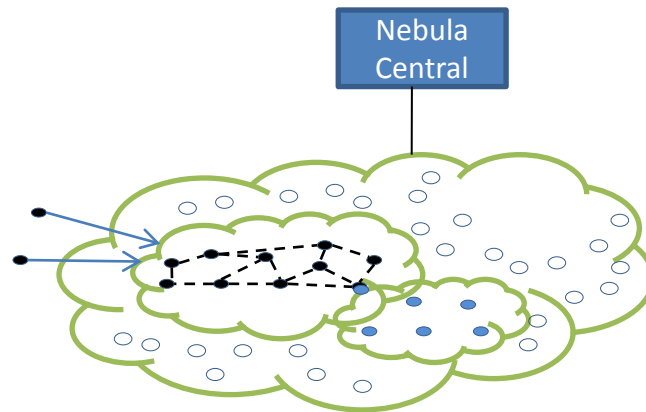
- Decentralized, less-managed cloud
  - dispersed storage/compute resources
  - low user cost






# Nebula: A New Cloud Model

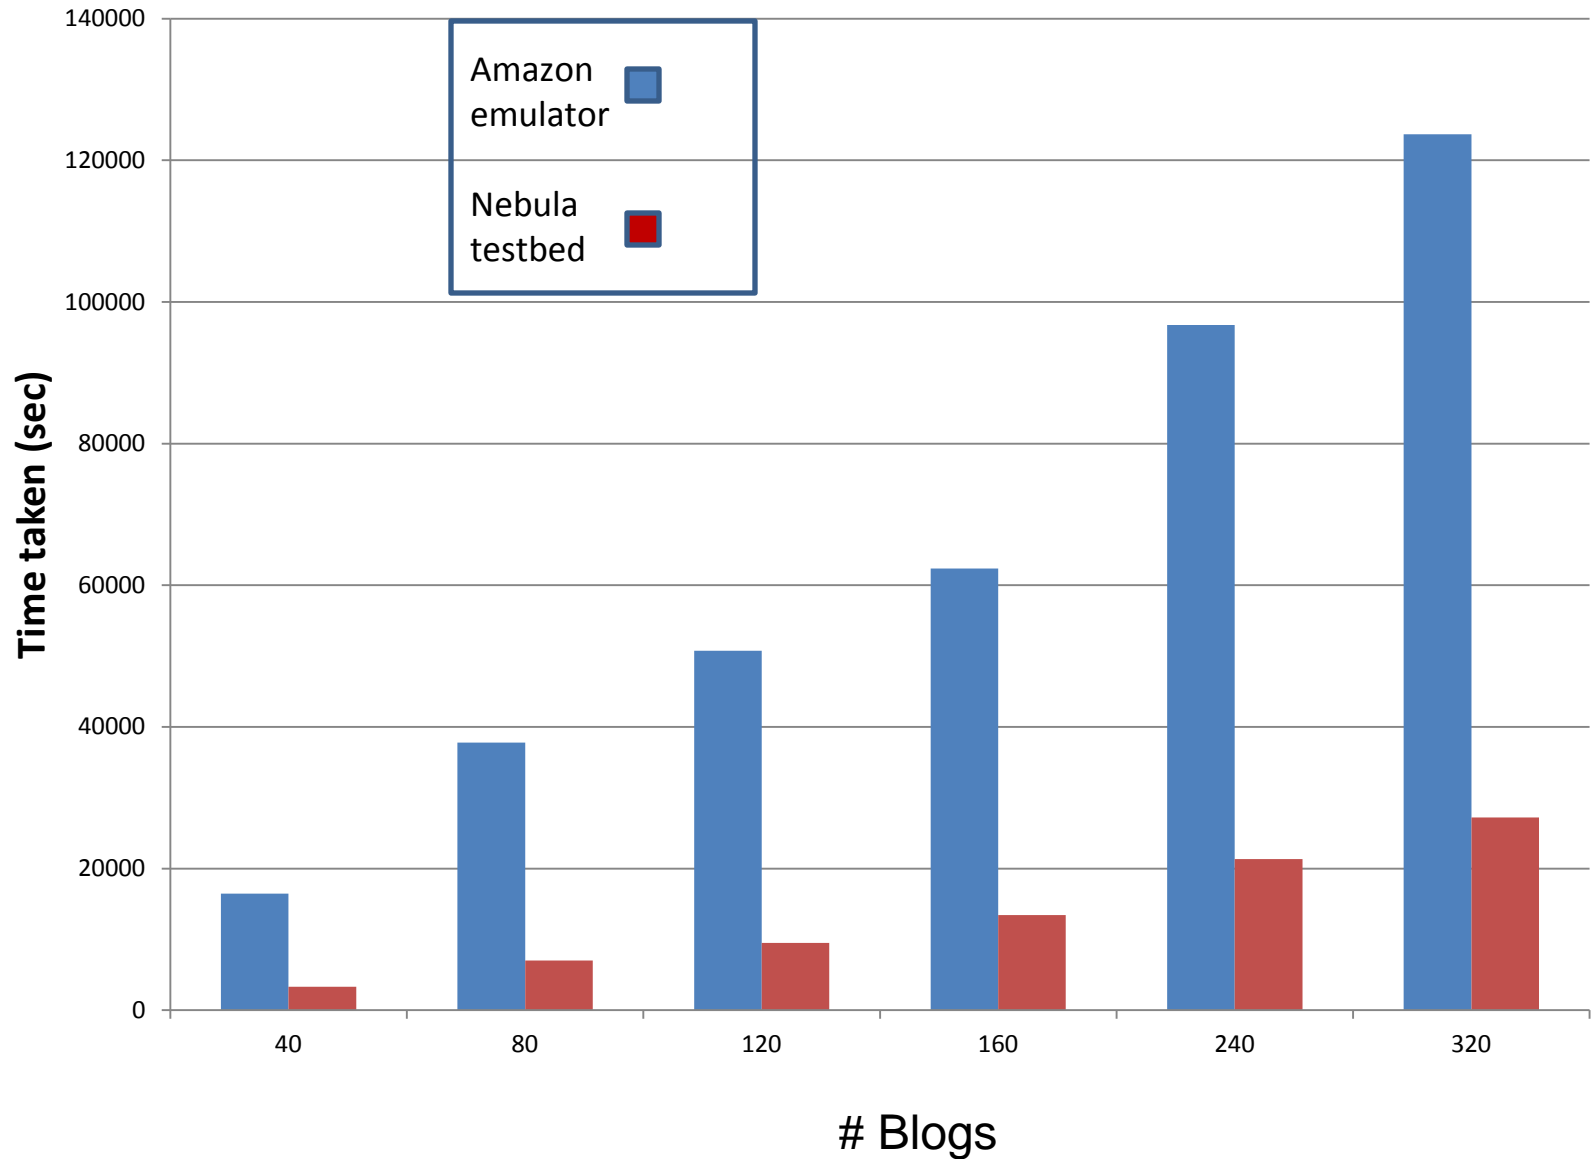
- Make the cloud more “distributed”
  - exploit the rich collection of edge computers
  - volunteers (P2P, @home), commercial (CDNs)
  - enormous computing potential, network diversity
  - lower latency: “on demand”, native client sandboxing



# Example: Blog Analysis



# Blog Results



# Current Status

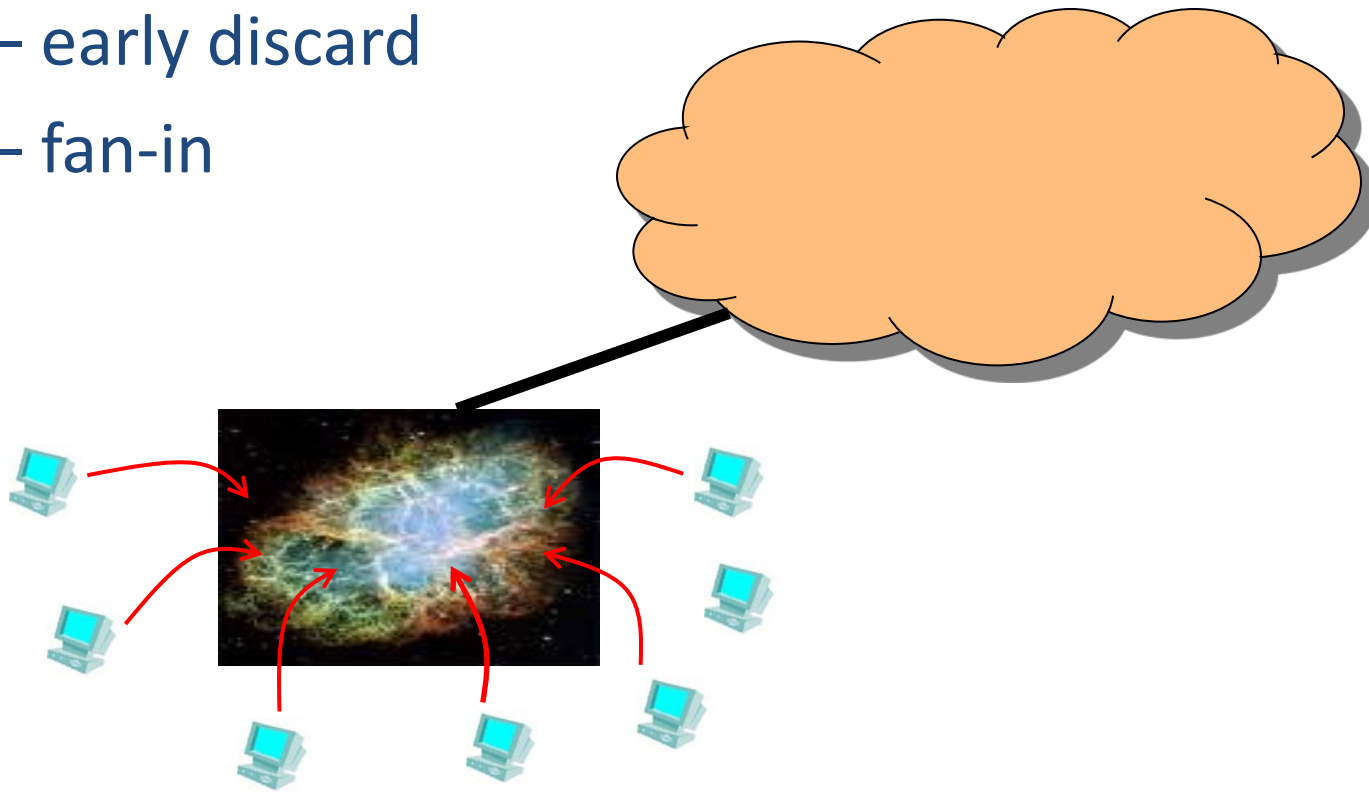
- Prototype running on PlanetLab
  - Chrome browser clients + native client
  - Distributed data-store service
  - Network dashboard service

# Nebula Going Forward

- Organize Nebulas
  - around trusted peers
  - social groups
  - communities of interest
  - local resources
- Nebula + commercial cloud
  - “use the edge opportunistically”

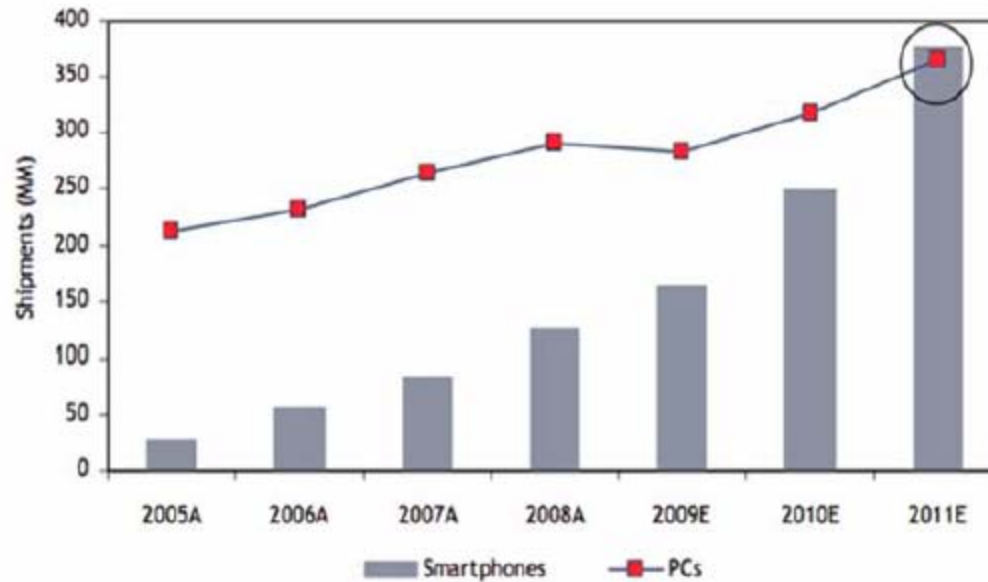
# Nebula Going Forward

- Hybrid paradigms
  - early discard
  - fan-in



# Mobility Trend: Mobile Cloud

*PCs = ( Desktops + Laptops + PDAs )*



Source: RBC Capital Markets estimates

# Mobility Trend: Mobile Cloud

- Mobile users/applications: s-phones, tablets
  - resource limited: power, CPU, memory
  - applications are becoming ↑ sophisticated
- Improve mobile user experience
  - performance, reliability, fidelity
  - tap into the cloud **dynamically** based on current resource state, preferences, interests, privacy

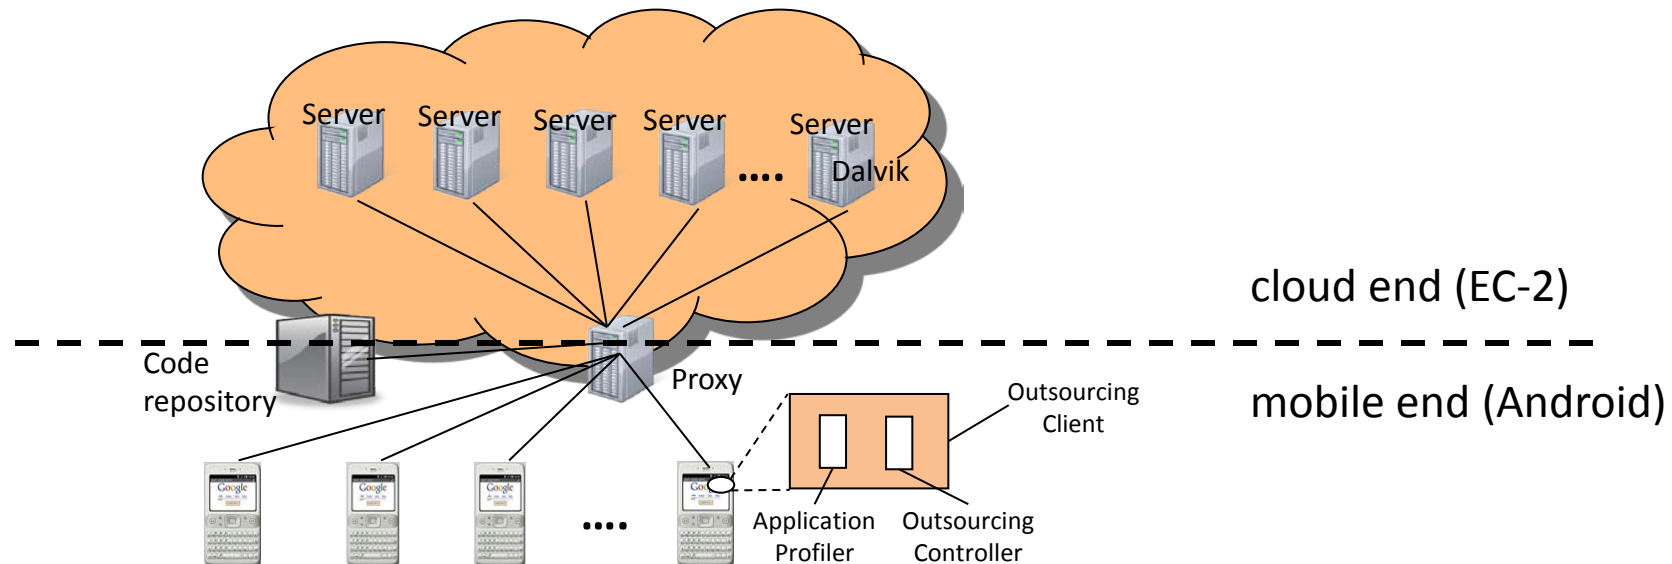


# Cloud Mobile Opportunity

- Dynamic outsourcing
  - move computation, data to the cloud dynamically
- User context
  - exploit user behavior to pre-fetch, pre-compute, cache
- Multi-user sharing
  - discover implicit cloud sharing based on interests, social ties

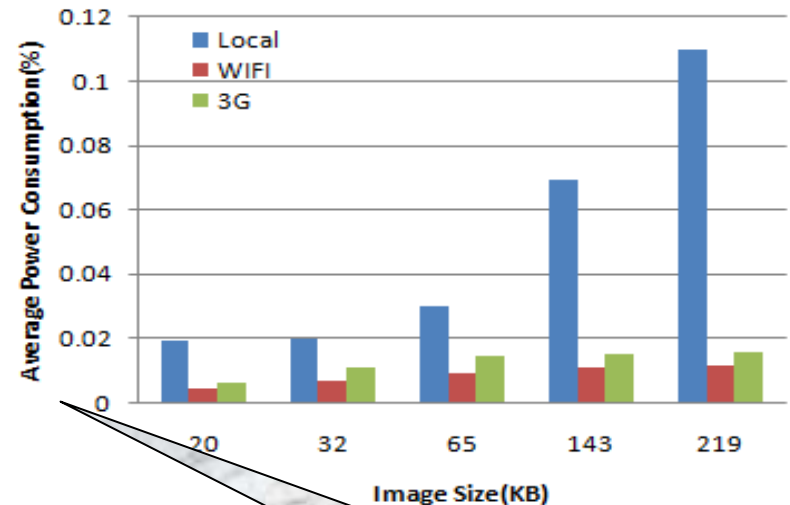
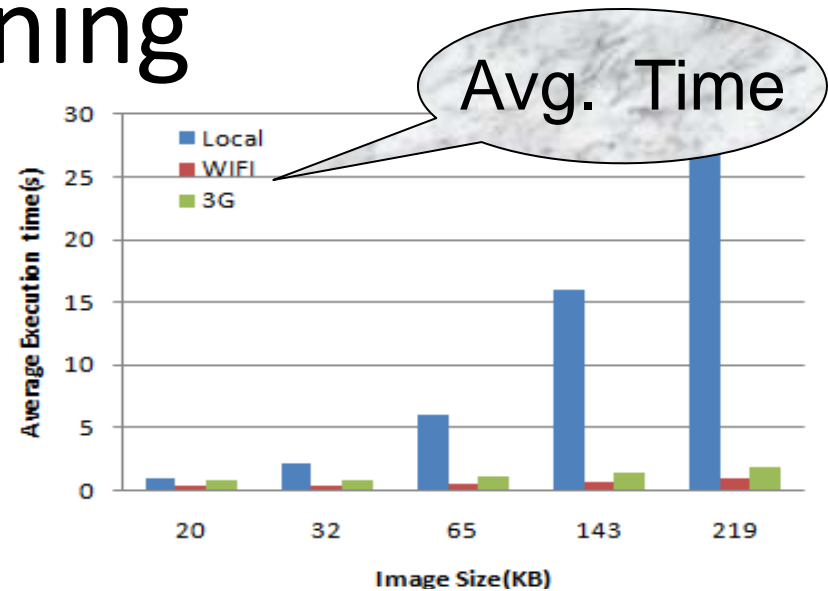
# Outsourcing

- Partitioning across (Mobile <-> Cloud)
  - local data capture + cloud processing
  - images/video, speech, digital design, aug. reality



# Experimental Results -Image Sharpening


- Response time
  - both WIFI & 3G
  - up to 27× speedup
- Power consumption
  - save up to 9× times




# Current Work

- Optimizations
  - User patterns => speculation, aggregation, data reduction
- Cloud-side
  - Cross-user patterns => sharing: VM provisioning, data placement, data re-use

# Big Data Trend + Multi-DCs: DMapReduce



# Wide-Area MapReduce

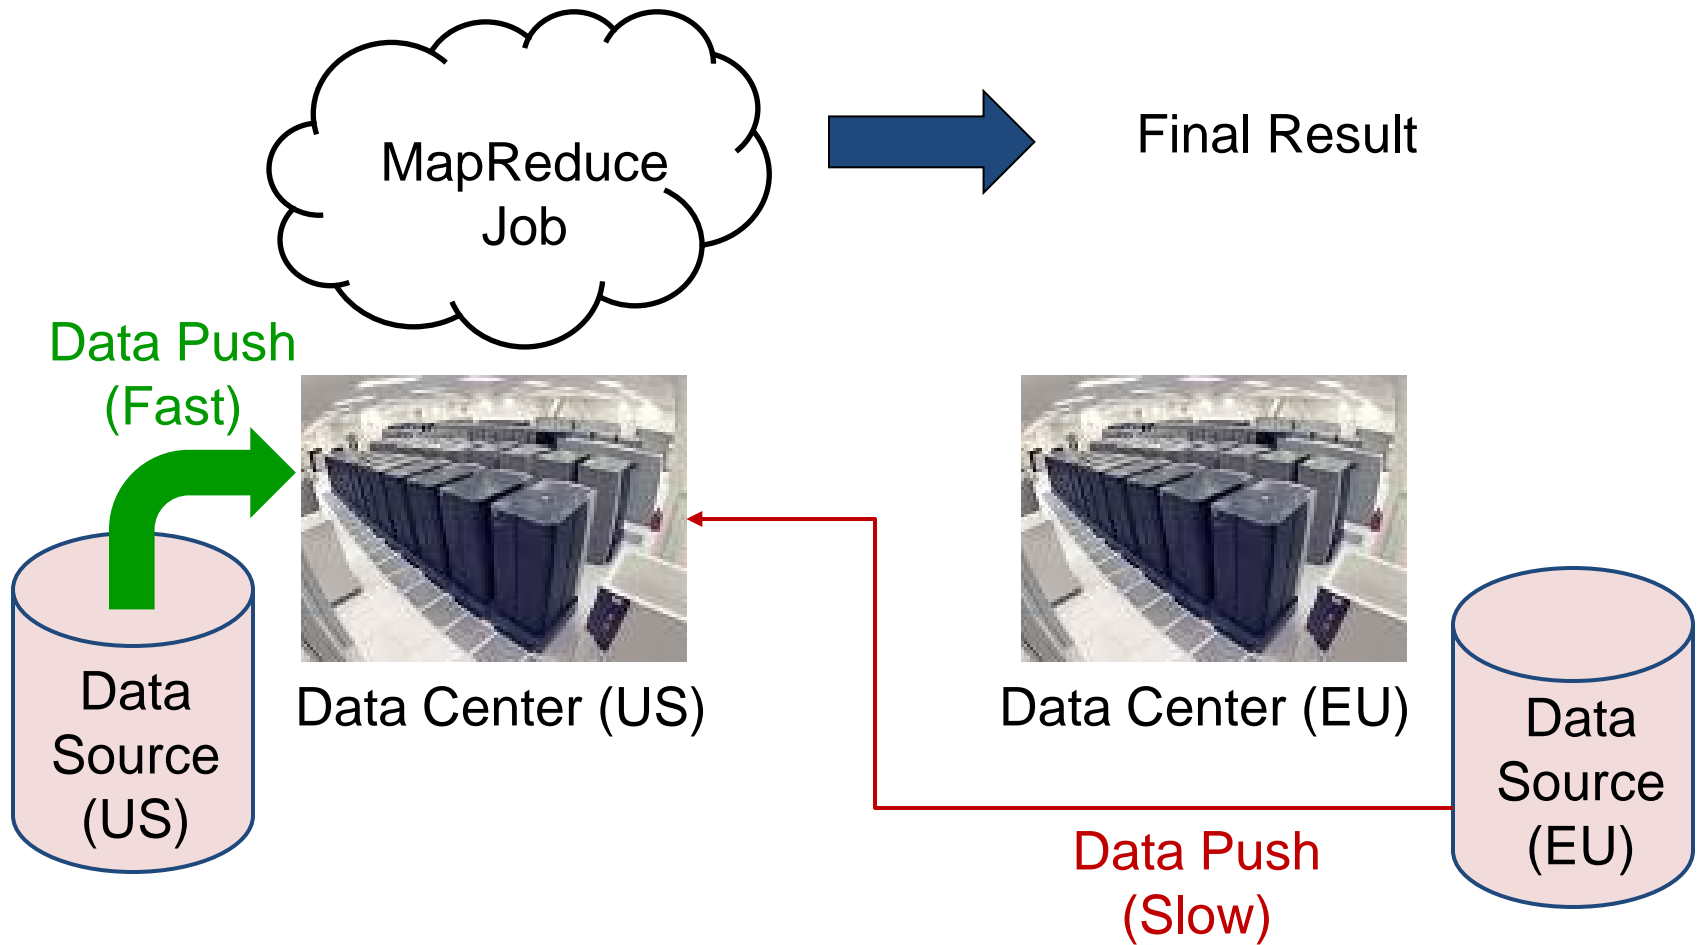


Big data is distributed

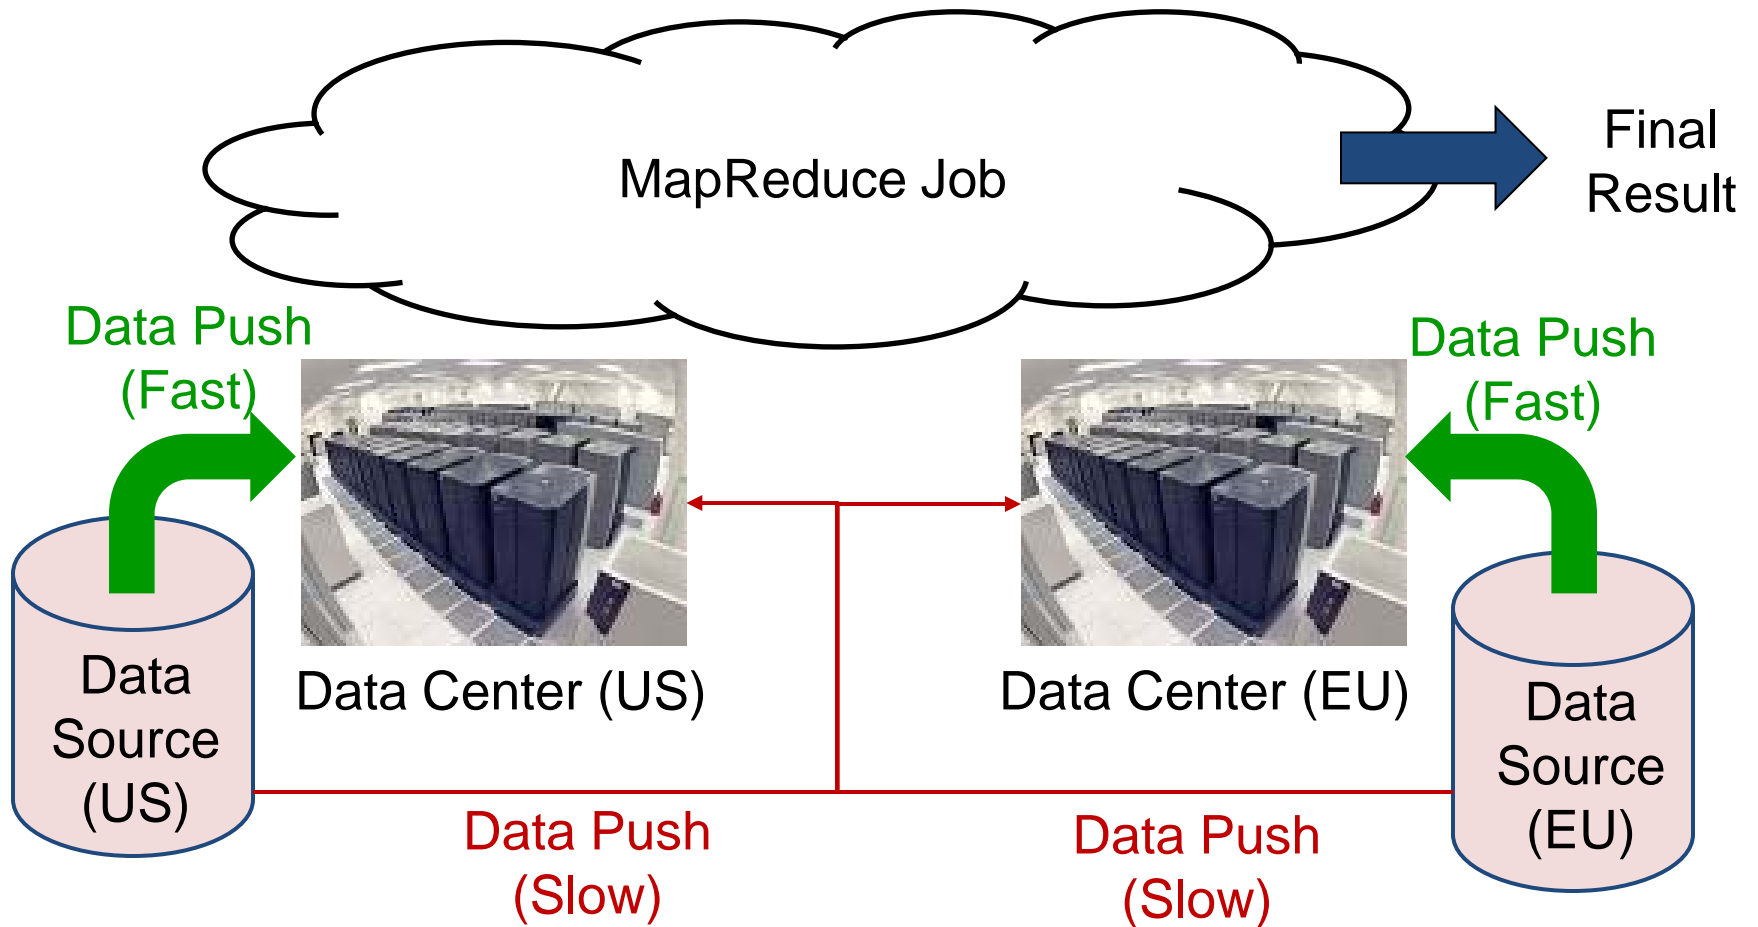
- earth science: weather data/seismic data
- life science: GenBank/PubMed
- health science: GoogleEarth /pandemic data
- web 2.0: user multimedia blogs

- Data in different data-centers
- Run MapReduce across them
- Data-flow spanning wide-area networks

# Option 1: Local MapReduce

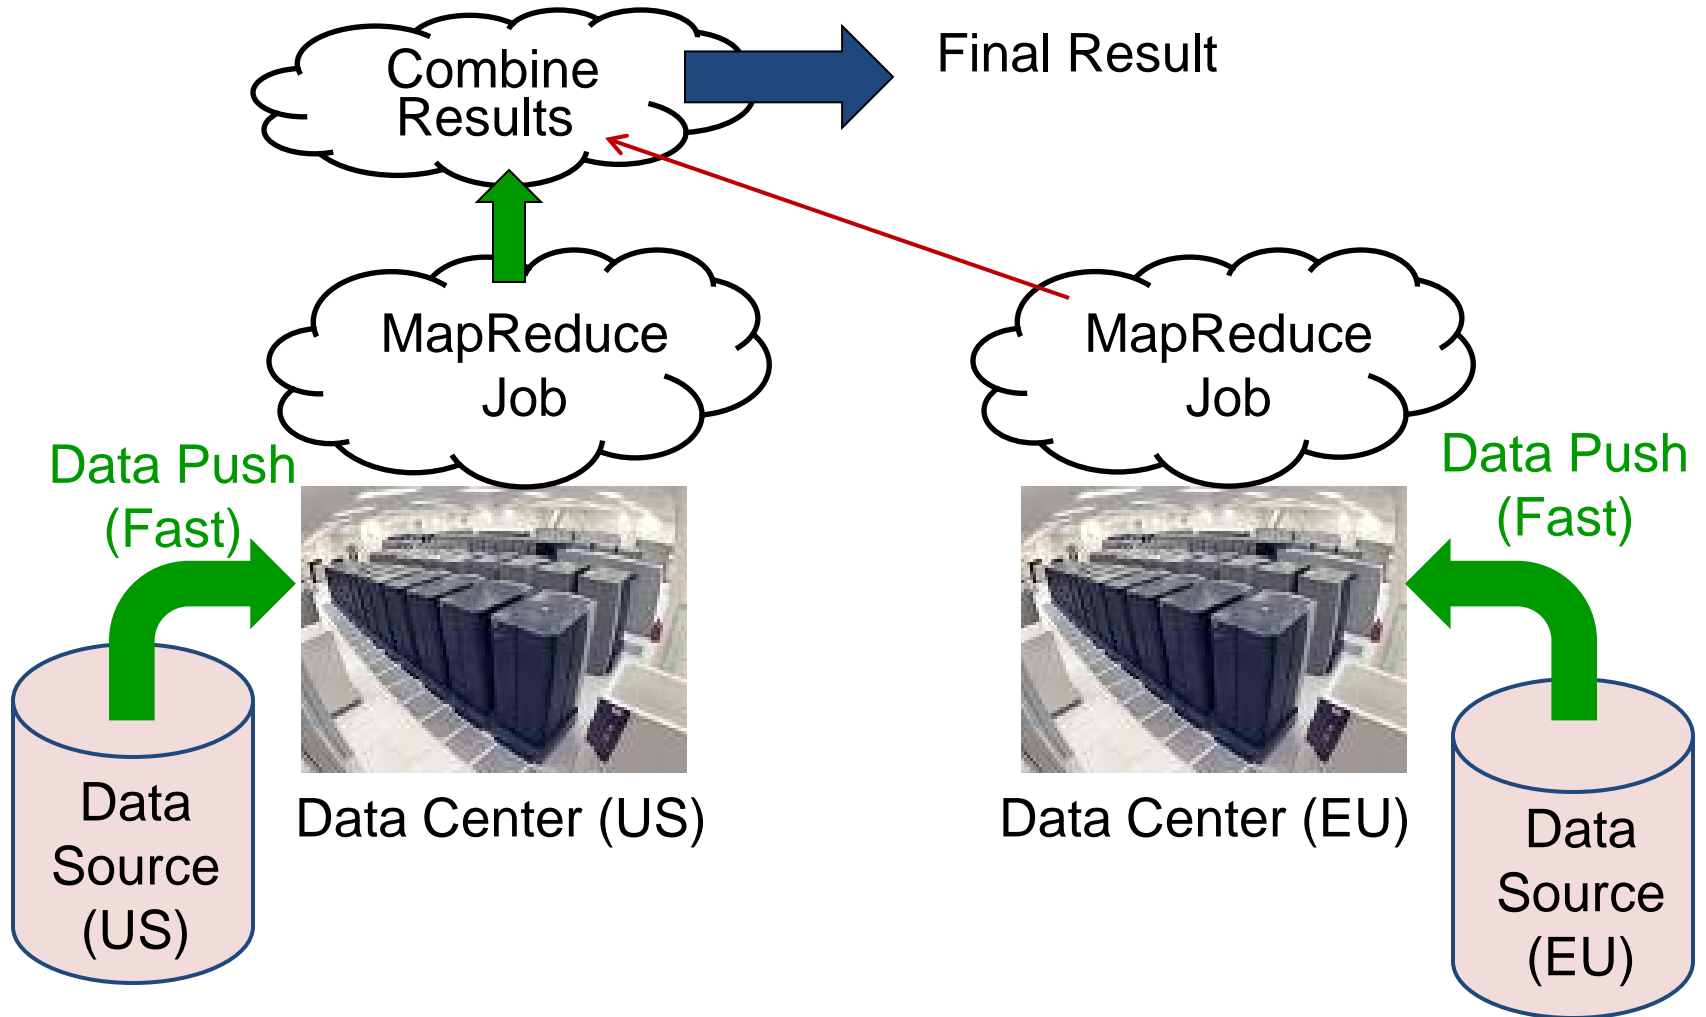


# Option 2: Global MapReduce






# Option 3: Distributed MapReduce




# PlanetLab: WordCount (text data)



- Distributed MR works best in presence of data aggregation

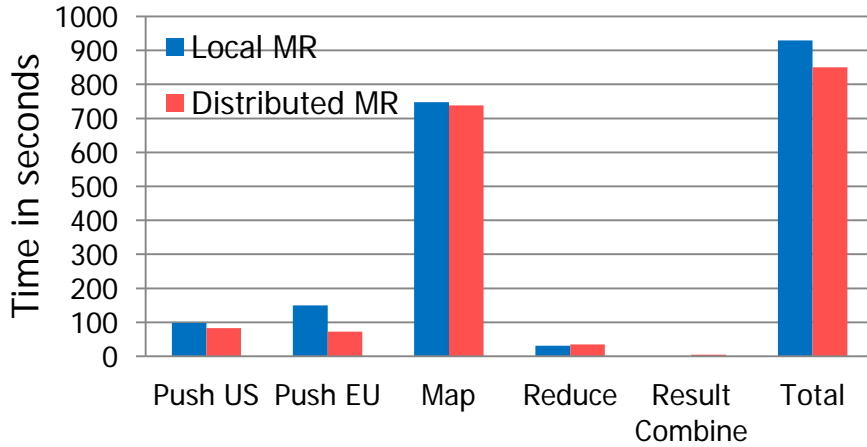
# PlanetLab: WordCount (random data)




- Local MR works best in presence of data ballooning

# EC-2 Results


## WordCount (Text)



## WordCount (Random)

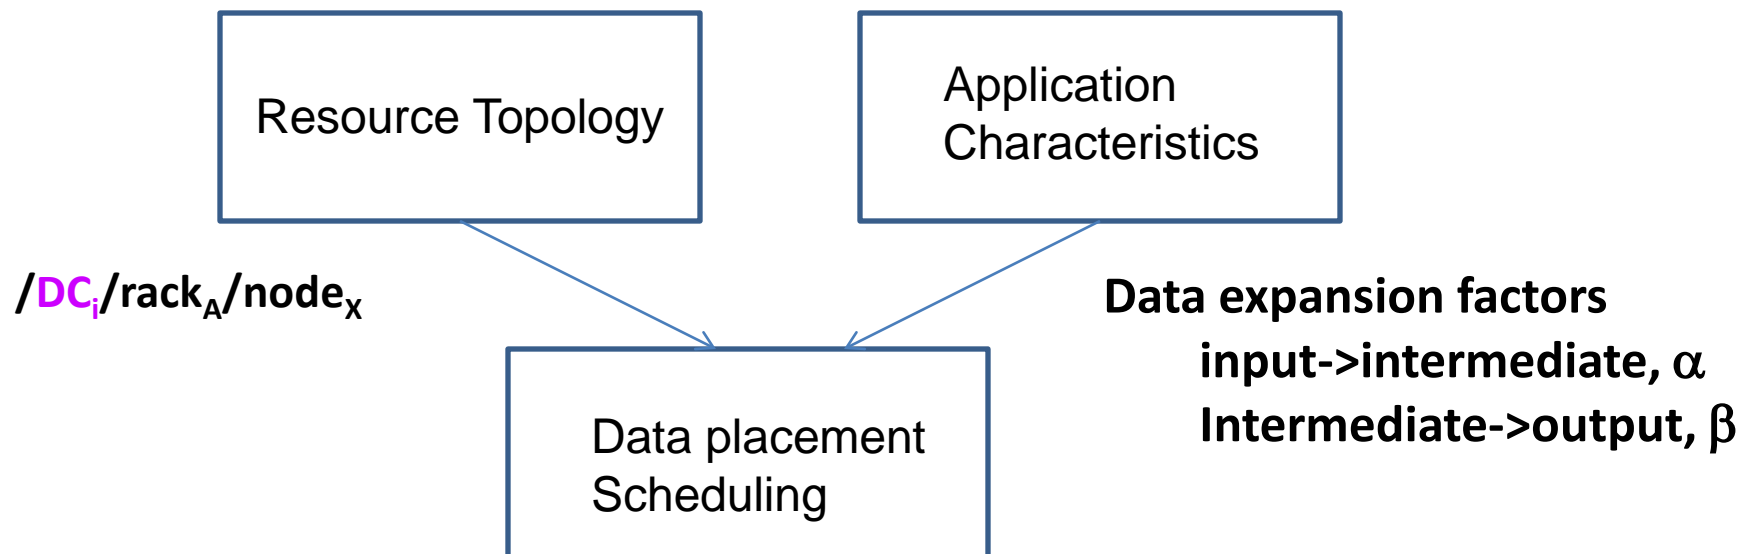


## Sort




# Intelligent Data Placement for Global Services

- HDFS push
  - local node, same rack, random rack



# Experimental Results (Word count)



# Summary

- Cloud Evolution
  - mobile users, big data, privacy/trust, global services
- Our Vision of the Cloud
  - locality of users, data, other clouds/data centers, user-centric behavior