Networking Architectures for Big-Data Applications

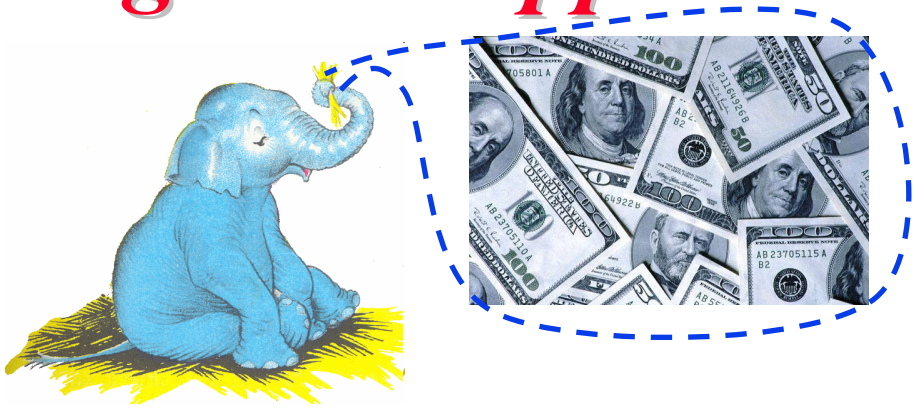

Raj Jain
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@cse.wustl.edu

Eighth Annual Microsoft Research Networking Summit
Woodinville, WA, June 18-19, 2012

These slides and audio/video recordings of this talk are at:
http://www.cse.wustl.edu/~jain/talks/bigdata.htm

Overview

1. Big Data: Why Now?
2. Database Solutions and Networking Bottlenecks
3. Intra-Cloud Solutions
4. Multi-Cloud Issue: OpenADN

Big Data: Why Now?

Search Database Analytics Storage Networking …

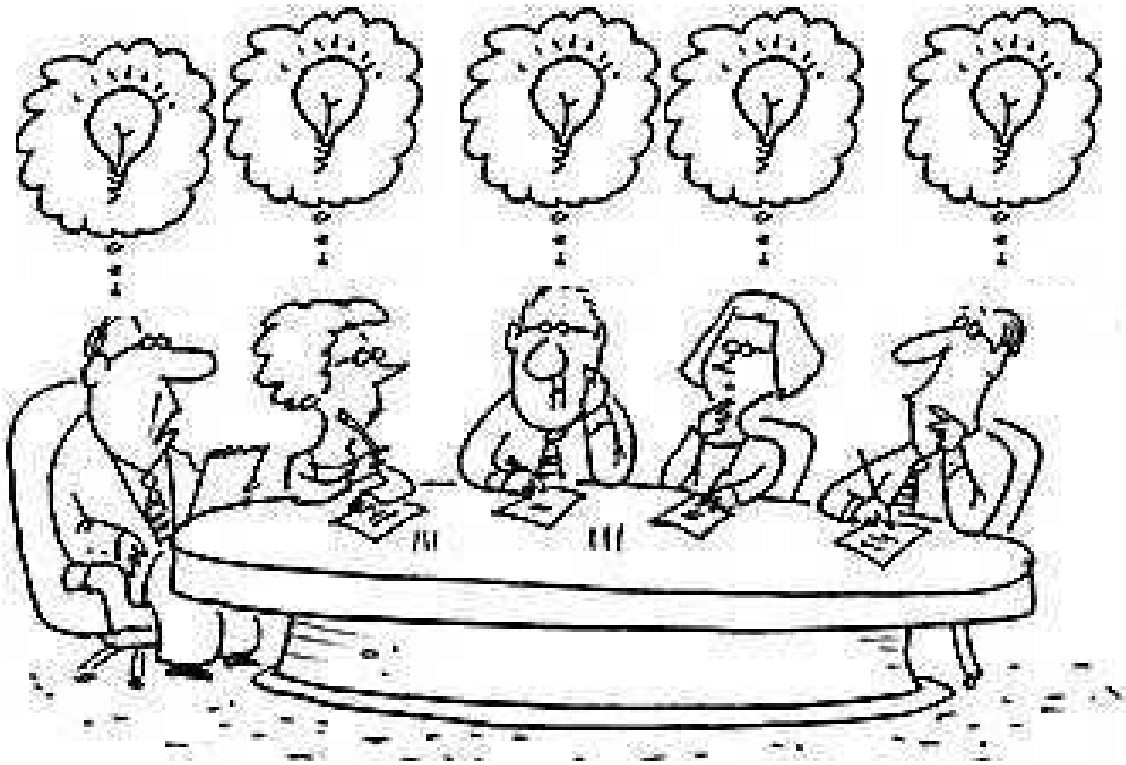

"Big Data" first appeared as a problem in October 1997

Big Research Funding in Q2 2012 ⇒ Big Academic Interest in all fields

Ref: http://whatsthebigdata.com/2012/06/06/a-very-short-history-of-big-data/

Samples of Recent News about Big Data

- NSF $80M, DoD $250M, DOE $25M
  http://gigaom.com/cloud/obamas-big-data-plans-lots-of-cash-and-lots-of-open-data/

- 10gen Scores $42 Million in Big Data Funding
  http://cloudcomputing.sys-con.com/node/2286075 , May 30, 2012

- ACCEL PARTNERS LAUNCHES $100MM BIG DATA FUND
  http://www.accel.com/bigdata

- Big Data Funding Gets Bigger: IA Ventures Raised $105 M, Double Its Previous Fund
  http://betabeat.com/2012/02/big-data-funding-gets-bigger-ia-ventures-raises-105-m-double-its-previous-fund/

Magnitude of Data

- 2.5 Exa (=10^{18}) bytes of information created per day
  = 30k × US Library of Congress
- 9.57 Zetta (=10^{21}) bytes processed by servers in 2008
- One Zetta byte traffic on the Internet
- Solutions:
  - Database
  - Analytics
  - Storage
  - Networking
  - …

Ref: http://en.wikipedia.org/wiki/Big_data
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html

Database Solutions

1. **Share Nothing Architecture**

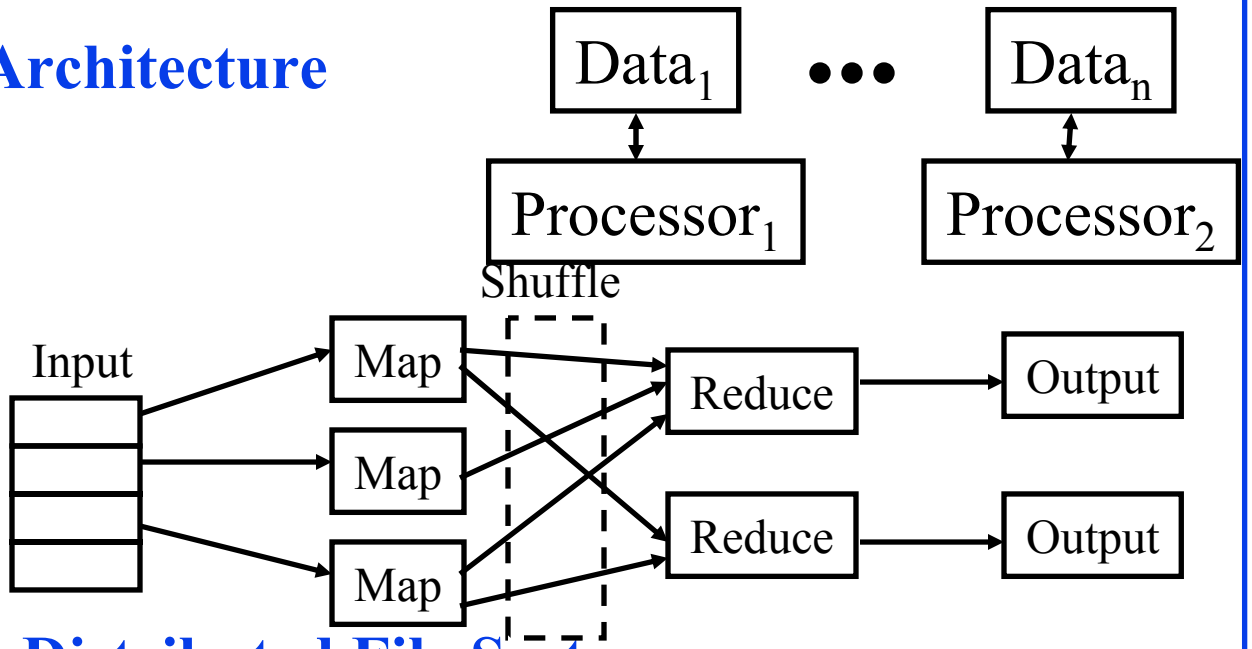

2. **Map-Reduce**

3. **Apache Hadoop Distributed File System**

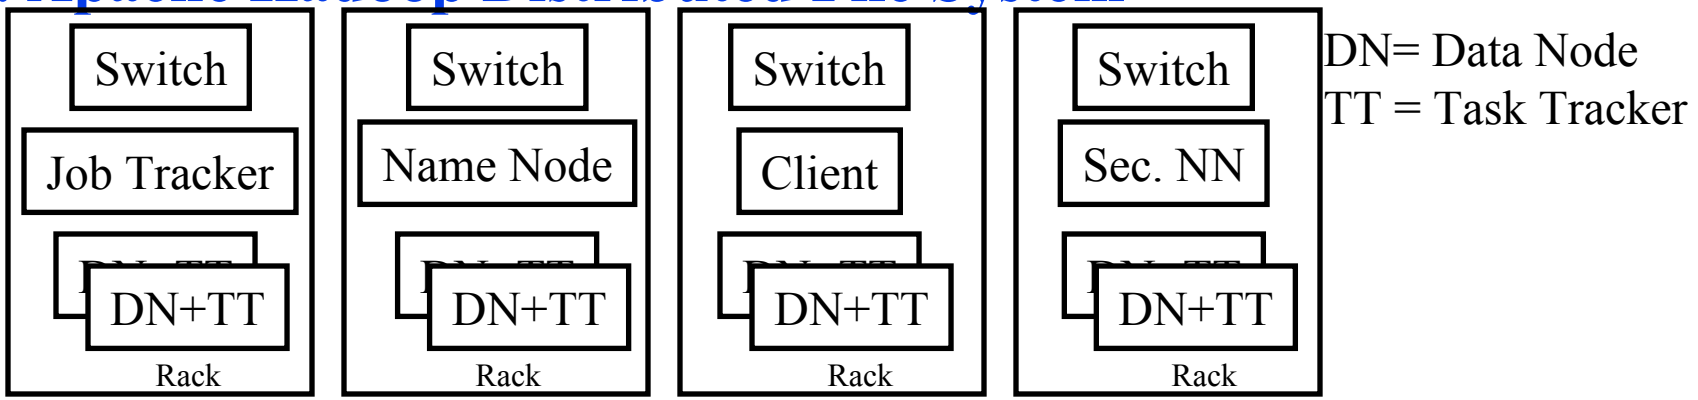

Networking Bottlenecks

1. Network Attached Storage (NAS):
   - Opposite of Share-nothing architecture
   - Networking Bottleneck ⇒ Directly attached storage
2. IP Subnets:

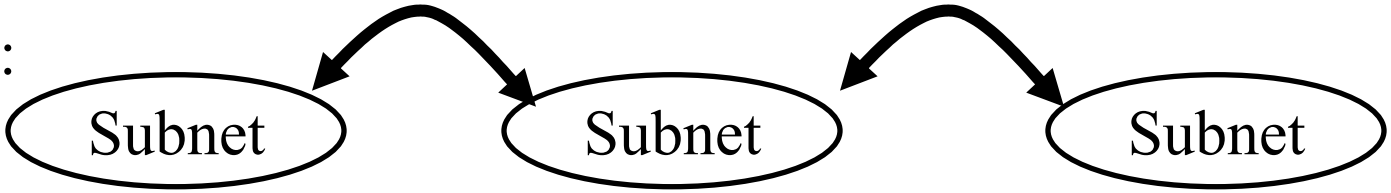

   - Routing much slower than switching
   - Moving VMs between subnets Addressing issues
     ⇒ Keep communication in a subnet
3. L2 Fabric: Multi-switch latency ⇒ Single Rack Traffic

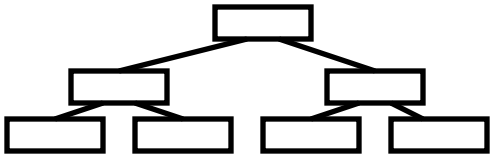

Intra-Cloud Solutions

1. **Flatter L2 Topologies**:

2. **Virtualization**: Ethernet over IP (VXLAN, TRILL)
   Multiple IP domains look like one L2 domain
   Does not solve the latency issue

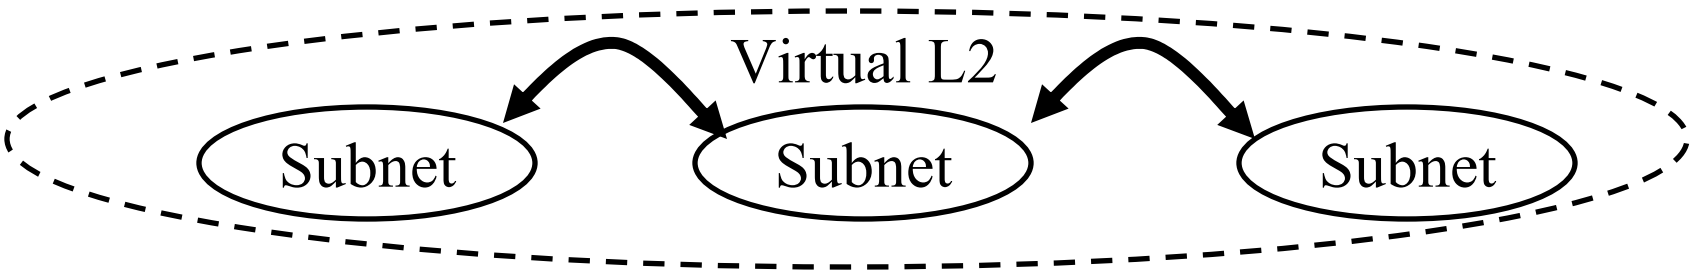

3. **Optimal Placement**: E.g., Camdoop uses reducers in the core
4. **Location Based Naming**: As in Tashi from CMU
5. **Programmable Networks**: Virtual topology similar to a single rack even though a the systems are physically in different racks

Really Big Data: Multi-Cloud Issue

- **Example:**

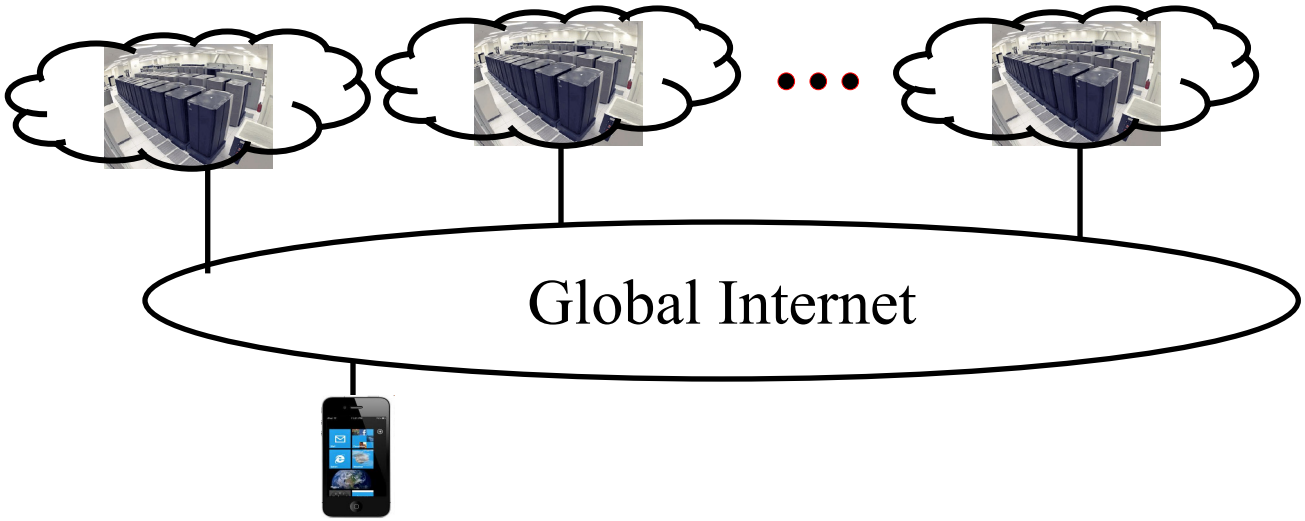

- Siri (Somewhat Intelligent Response Interpreter)
  Needs to consult global databases
  Where should the name tracker, data tracker, task tracker reside?

Our Solution: OpenADN

- Open Application Delivery Networking Platform.
  Platform = Clients, Servers, Switches, and Middle-boxes
  Servers = Name nodes, Task trackers, Data nodes, Clients
- Allows global Internet to appear as one datacenter
  Message-specific sequence of middle boxes and routes

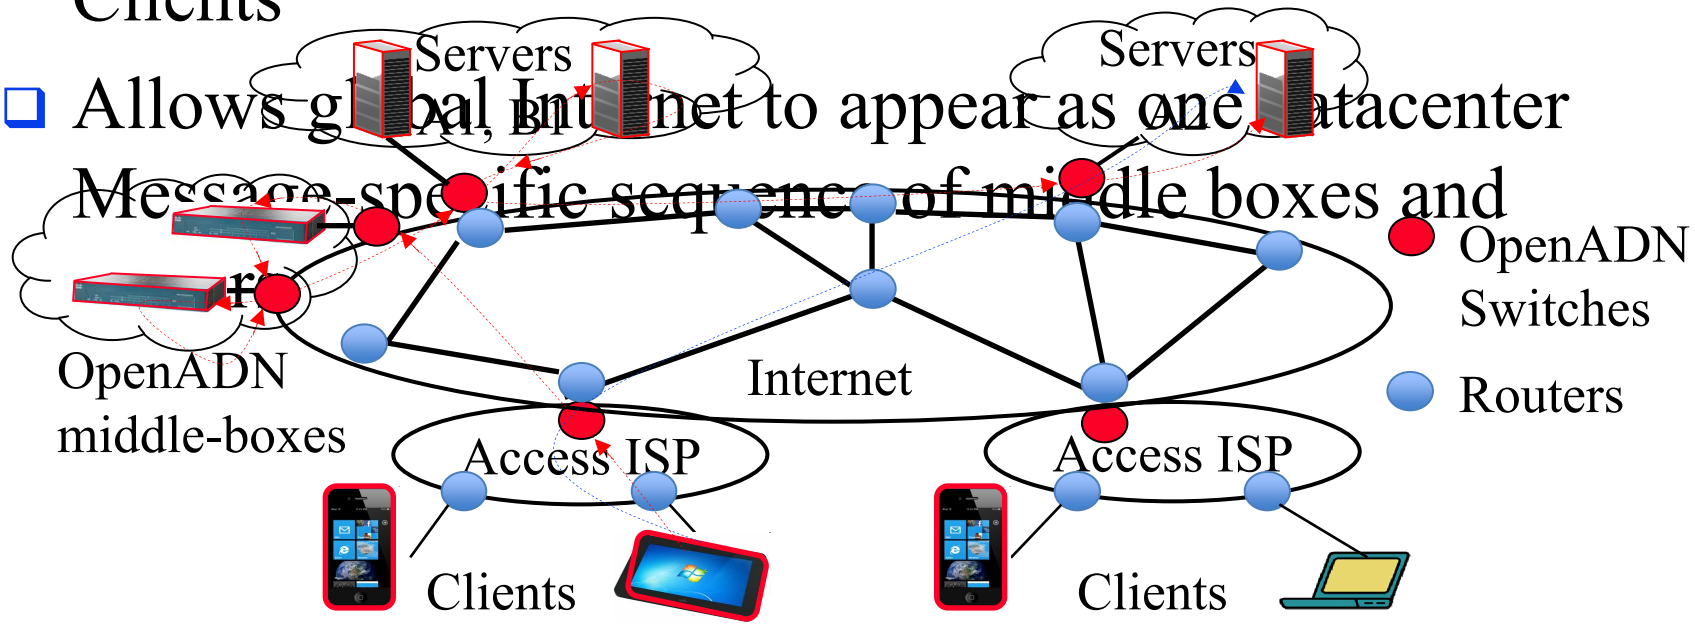

OpenADN Features

- Message-specific sequence of middle-boxes and Servers:
  Name nodes, Task trackers, Data nodes, Clients
- Middle Boxes:
  Load balancing, Fault tolerance, Intrusion Detection, …
  Control plane and data plane MBs
- Server mobility, User Mobility
- ISPs can offer proper routing and switching services without visibility into data
  CSPs (Cloud service providers) have visibility into data
- Message-specific policies

Key Features of OpenADN

1. Edge devices only.
   Core network can be current TCP/IP based, OpenFlow and SDN based
2. Coexistence (Backward compatibility):
   Old on New. New on Old
3. Incremental Deployment
4. Economic Incentive for first adopters
5. Resource owners (ISPs) keep complete control over their resources

Summary

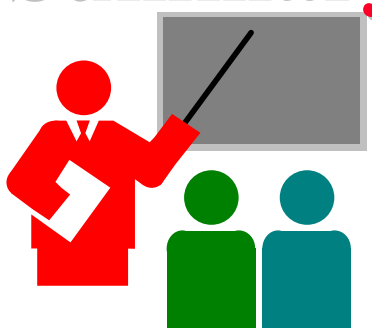

1. Big data = Big Funding ⇒ Big Interest
2. Need networking architectures that complement Shared Nothing Architecture, Map Reduce, Hadoop File Systems
3. NAS needs to be redesigned.
   Flat topologies are helpful.
   Virtualization needs to include latency issues
4. Programmable networks are potential solution
5. Solving multi-cloud is important for really big global data sets.

Thank You!