Networking For Big Data

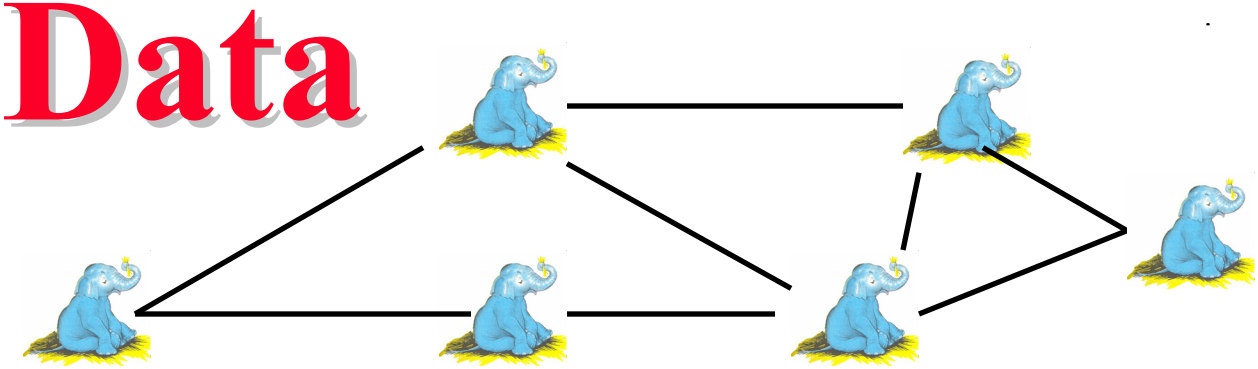

Raj Jain
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@cse.wustl.edu

IEEE CS Keynote at 19th Annual International Conference on Advanced Computing and Communications (ADCOM) 2013, Chennai, India, October 22, 2013

These slides and audio/video recordings of this talk are at:
http://www.cse.wustl.edu/~jain/talks/adcom13.htm

Overview

1. Why, What, and How of Big Data:
   It's all because of advances in networking
2. Recent Developments in Networking and their role in Big Data (Virtualization, SDN, NFV)
3. Networking needs Big Data

Gartner Hype Cycle 2013

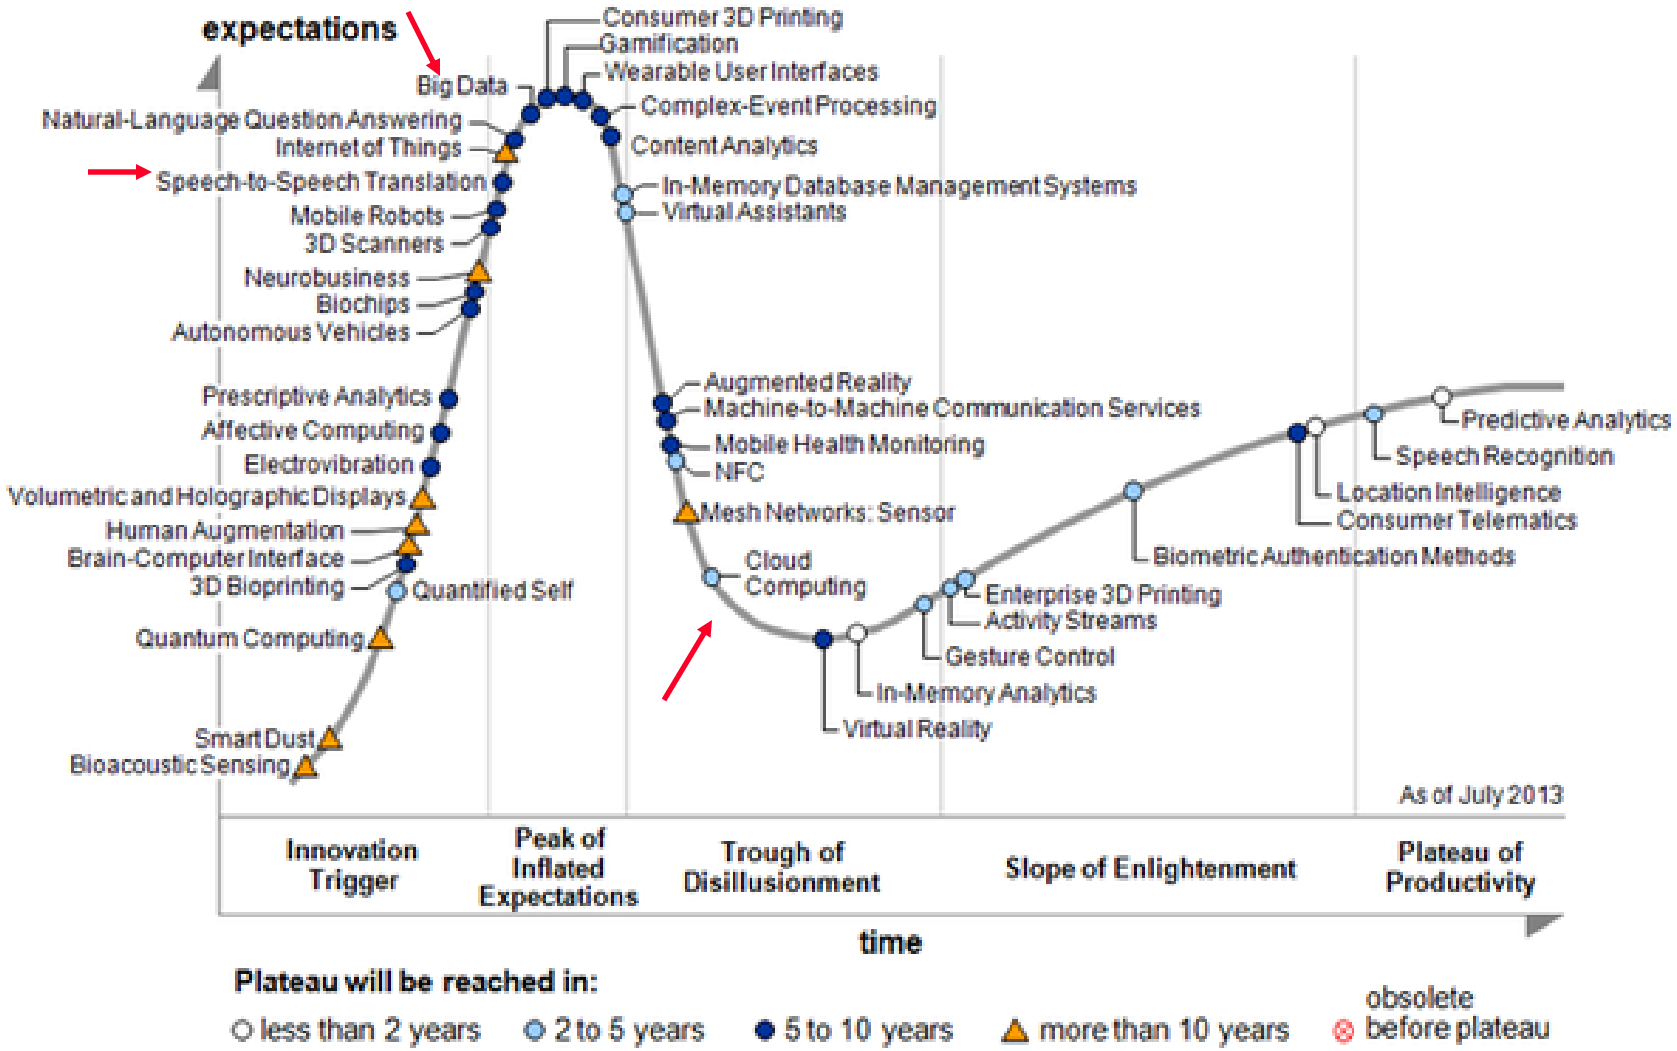

Ref: http://www.zdnet.com/gartners-2013-emerging-technologies-hype-cycle-focuses-on-humans-and-machines-7000019564/

What's Big?

- "Big data" is data larger than what you can handle
- "Big Data" first appeared as a problem in October 1997
- Sudden burst of activity in Q2 2012. Why? Let's ask Google…

What is Google Trends?

- A time series graph of number of searches on any term of your choice
- Includes geographical distribution of those searches
- Includes major news items
- Example: "Barack Obama"
  K: Nov 2008 election
  B: Nov 2012 election
  Popular in Africa

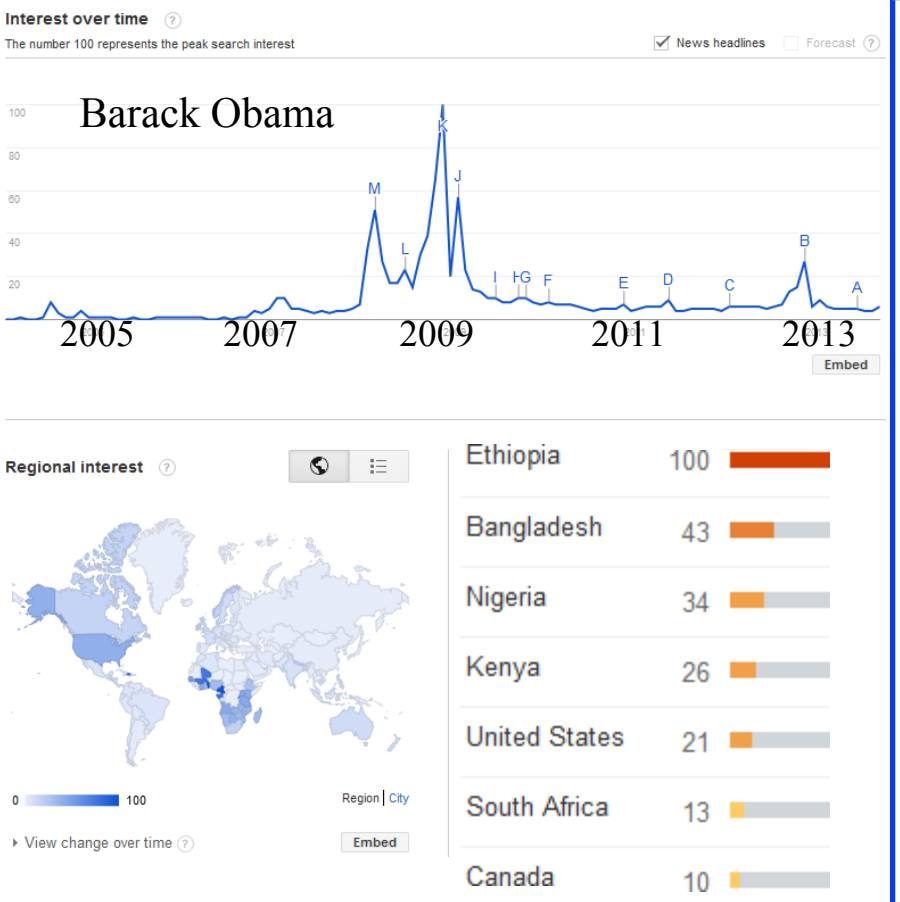

https://www.google.com/trends/explore#q=barack+obama

You are what you search!

- Internet users from countries with higher Gross Domestic Production (GDP) are more likely to search for future topics than about the past.
- Economic indicators are correlated to on-line behavior.

Google Trends reveals clues about the mentality of richer nations

When researchers looked up the Google searches of different countries, they …

by **Casey Johnston** - Apr 5 2012, 5:10pm CST

BUSINESS SCIENCE 14

Countries that focus on the future tend to have a higher per capita GDP, suggests a study published in *Nature's Scientific Reports*. Using Google Trends, a group of researchers compared the number of searches conducted for calendar years past to those for calendar years to come, and found that residents of countries with higher GDPs tend to search more about the year to come than the one past. Likewise, residents of countries with lower GDPs tend to search about the year past more than the one to come.

Ref: C. Johnston, "Google Trends reveals clues about the mentality of richer nations," http://arstechnica.com/gadgets/2012/04/google-trends-reveals-clues-about-the-mentality-of-richer-nations/

Big Data: Google Trends

- Near the peak (As in Gartner's Hype cycle)
- Highest interest in India
- Followed by South Korea, Singapore, Taiwan, Hong Kong
- Knee on March 27, 2012 (Point G)

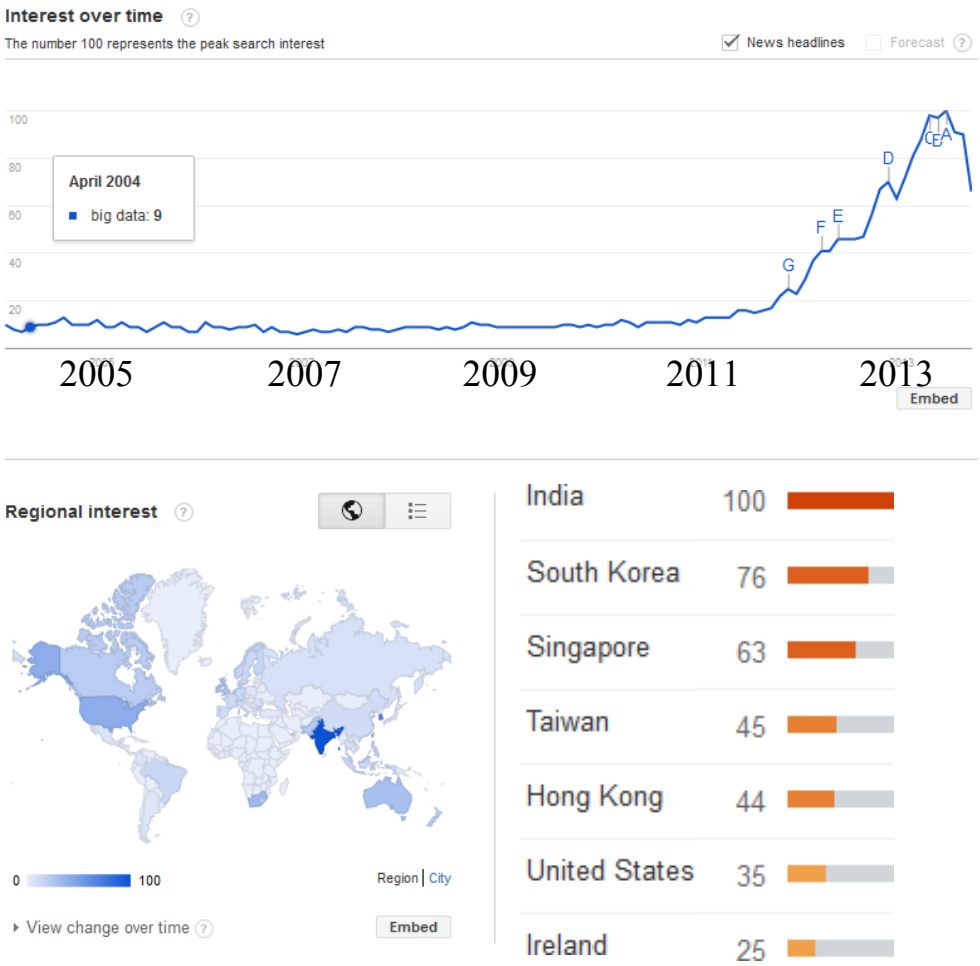

Snapshot: September 3, 2013

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/

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

- MongoDB Raises $150M, Now the Most Funded Big Data Startup
  http://siliconangle.com/blog/2013/10/04/mongodb-raises-150m-now-the-most-funded-big-data-startup/

- **Big Data Investment Keeps Climbing in 2013 – Funding Hits $1.28B Across 127 Deals**
  http://www.cbinsights.com/blog/trends/big-data-funding-venture-capital-2013

Ref:J. Kelly, "Big Data Start-up Funding by Vendor," Sep 30, 2013, http://wikibon.org/wiki/v/Big_Data_Start-up_Funding_by_Vendor

Networking is the Basis of Big Data

- Networking is the "plumbing" of computing
- Almost all areas of computing are network-based.
  - Distributed computing
  - Distributed databases
  - Distributed storage
  - Distributed Games

MapReduce

- Software framework to process massive amounts of unstructured data by distributing it over a large number of inexpensive processors
- **Map**: Takes a set of data and divides it for computation
- **Reduce**: Takes the output from Map outputs the result

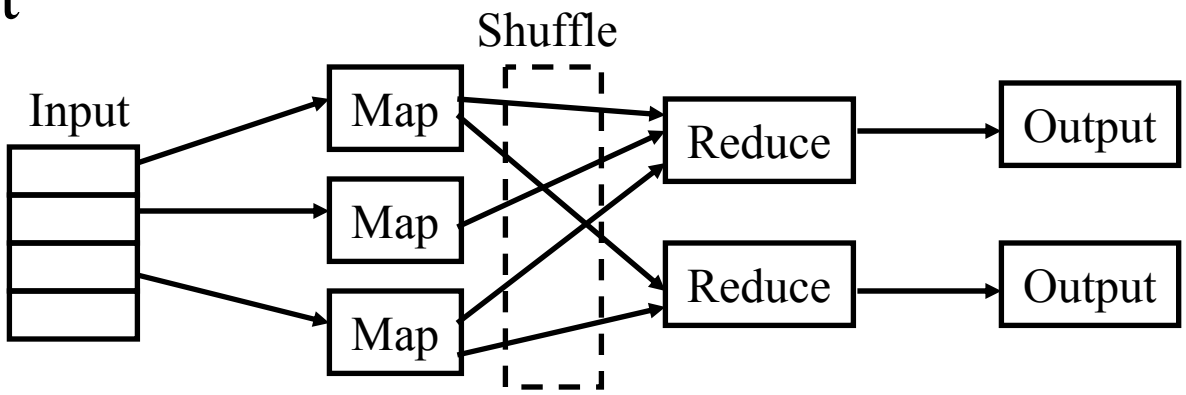

Ref: J. Dean and S. Ghemawat, "MapReduce: Simplified Data Processing on Large Clusters," OSDI 2004, http://research.google.com/archive/mapreduce-osdi04.pdf

Hadoop

- An open source implementation of MapReduce
- Named by Doug Cutting at Yahoo after his son's yellow plus elephant
- Hadoop File System (**HDFS**) requires data to be broken into blocks. Each block is stored on 2 or more data nodes on different racks.
- **Name node**: Manages the file system name space ⇒ keeps track of blocks on various **Data Nodes**.

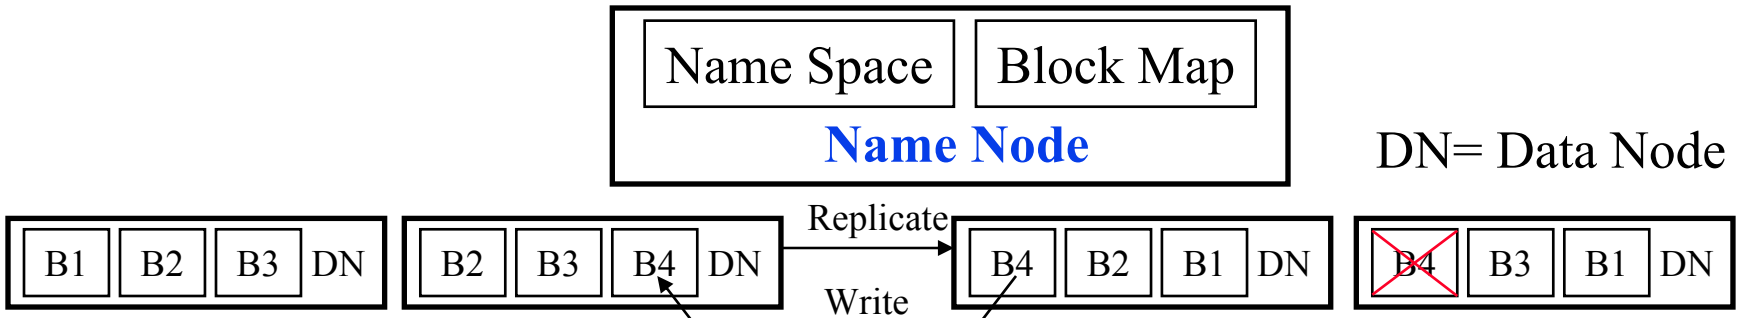

Hadoop (Cont)

- **Job Tracker**: Assigns MapReduce jobs to task tracker nodes that are close to the data (same rack)
- **Task Tracker**: Keep the work as close to the data as possible.

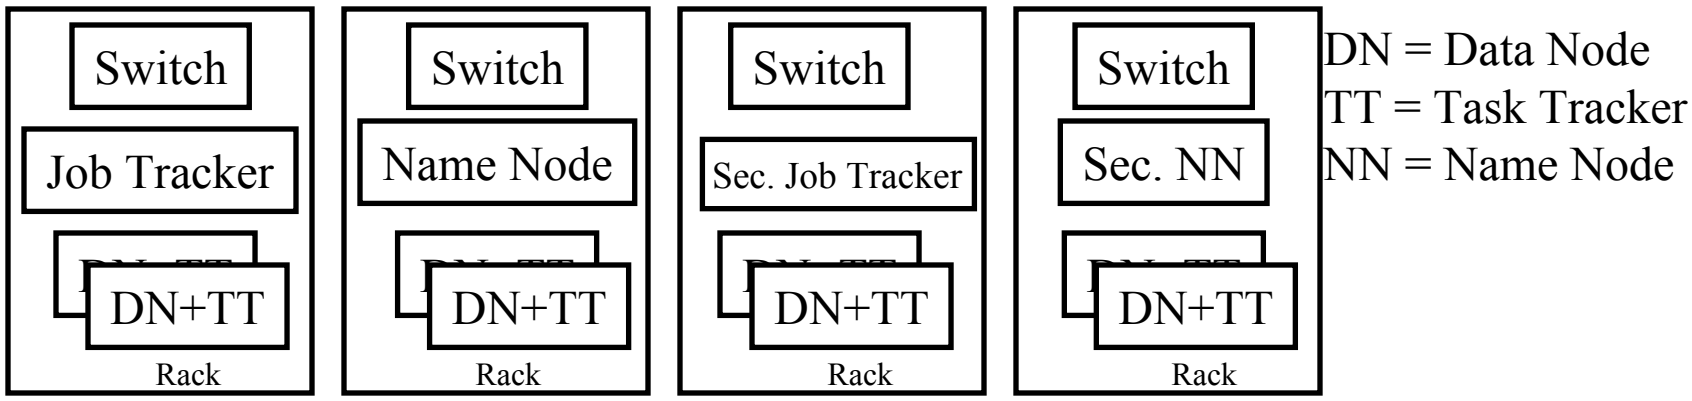

Networking Requirements for Big Data

1. **Code/Data Collocation**: The data for map jobs should be at the processors that are going to map.
2. **Elastic bandwidth**: to match the variability of volume
3. **Fault/Error Handling**: If a processor fails, its task needs to be assigned to another processor.
4. **Security**: Access control (authorized users only), privacy (encryption), threat detection, all in real-time in a highly scalable manner
5. **Synchronization**: The map jobs should be comparables so that they finish together. Similarly reduce jobs should be comparable.

Recent Developments in Networking

1. High-Speed: 100 Gbps Ethernet
   $\Rightarrow$ 400 Gbps $\Rightarrow$ 1000 Gbps
   $\Rightarrow$ Cheap storage access. Easy to move big data.
2. Virtualization
3. Software Defined Networking
4. Network Function Virtualization

Virtualization

"*Virtualization means that Applications can use a resource without any concern for where it resides, what the technical interface is, how it has been implemented, which platform it uses, and how much of it is available.*"

-Rick F. Van der Lans
in Data Virtualization for Business Intelligence Systems

Virtualization (Cont)

- Recent networking technologies and standards allow:
  1. Virtualizing Computation
  2. Virtualizing Storage
  3. Virtualizing Rack Storage Connectivity
  4. Virtualizing Data Center Storage
  5. Virtualizing Metro and Global Storage

1. Virtualizing Computation

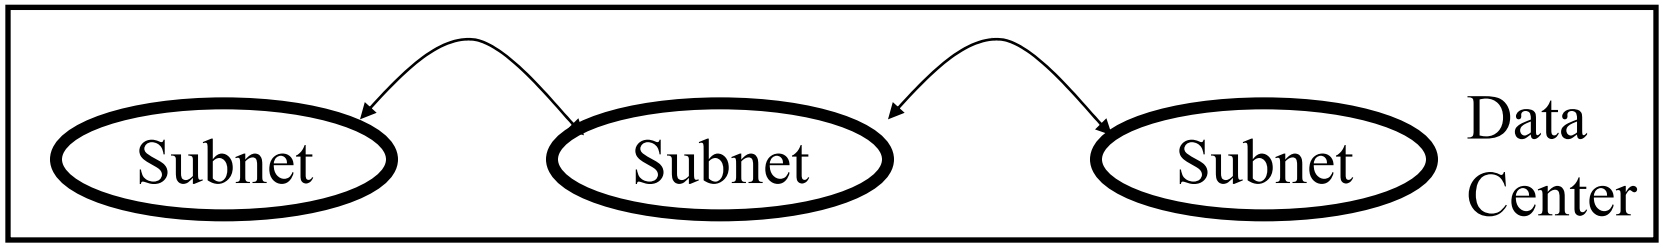

- Initially data centers consisted of multiple IP subnets
  - Each subnet = One Ethernet Network
  - Ethernet addresses are globally unique and do not change
  - IP addresses are locators and change every time you move
  - If a VM moves inside a subnet ⇒ No change to IP address ⇒ Fast
  - If a VM moves from one subnet to another ⇒ Its IP address changes ⇒ All connections break ⇒ Slow ⇒ Limited VM mobility
- IEEE 802.1ad-2005 Ethernet Provider Bridging (PB), IEEE 802.1ah-2008 Provider Backbone Bridging (PBB) allow Ethernets to span long distances ⇒ Global VM mobility

2. Virtualizing Storage

- Initially data centers used Storage Area Networks (Fibre Channel) for server-to-storage communications and Ethernet for server-to-server communication

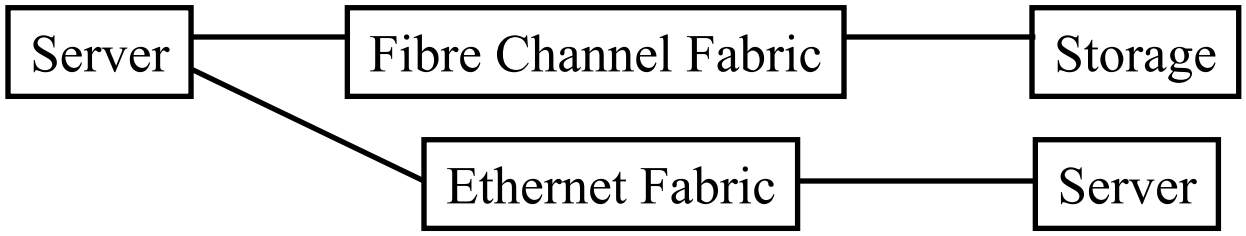

- IEEE added 4 new standards to make Ethernet offer low loss, low latency service like Fibre Channel:
  - Priority-based Flow Control (IEEE 802.1Qbb-2011)
  - Enhanced Transmission Selection (IEEE 802.1Qaz-2011)
  - Congestion Control (IEEE 802.1Qau-2010)
  - Data Center Bridging Exchange (IEEE 802.1Qaz-2011)
- Result: Unified networking ⇒ Significant CapEx/OpEx saving

3. Virtualizing Rack Storage Connectivity

- MapReduce jobs are assigned to the nodes that have the data
- Job tracker assigns jobs to task trackers in the rack where the data is.
- High-speed Ethernet can get the data in the same rack.
- Peripheral Connect Interface (PCI) Special Interest Group (SIG)'s Single Root I/O virtualization (**SR-IOV**) allows a storage to be virtualized and shared among multiple VMs.

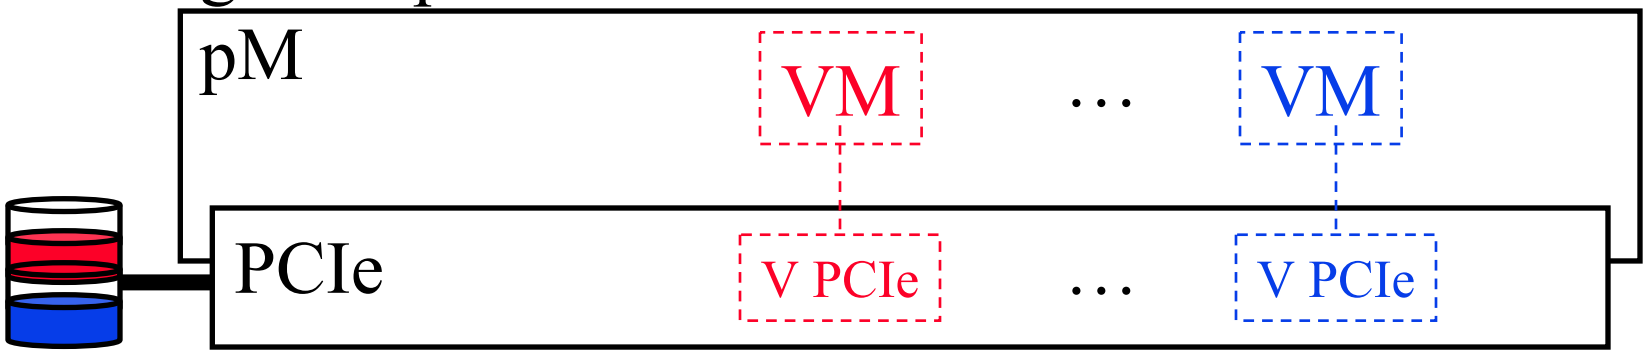

Multi-Root IOV

- PCI-SIG Multi-Root I/O Virtualization (**MR-IOV**) standard allows one or more PCIe cards to serve multiple servers and VMs in the same rack
- Fewer adapters ⇒ Less cooling. No adapters ⇒ Thinner servers

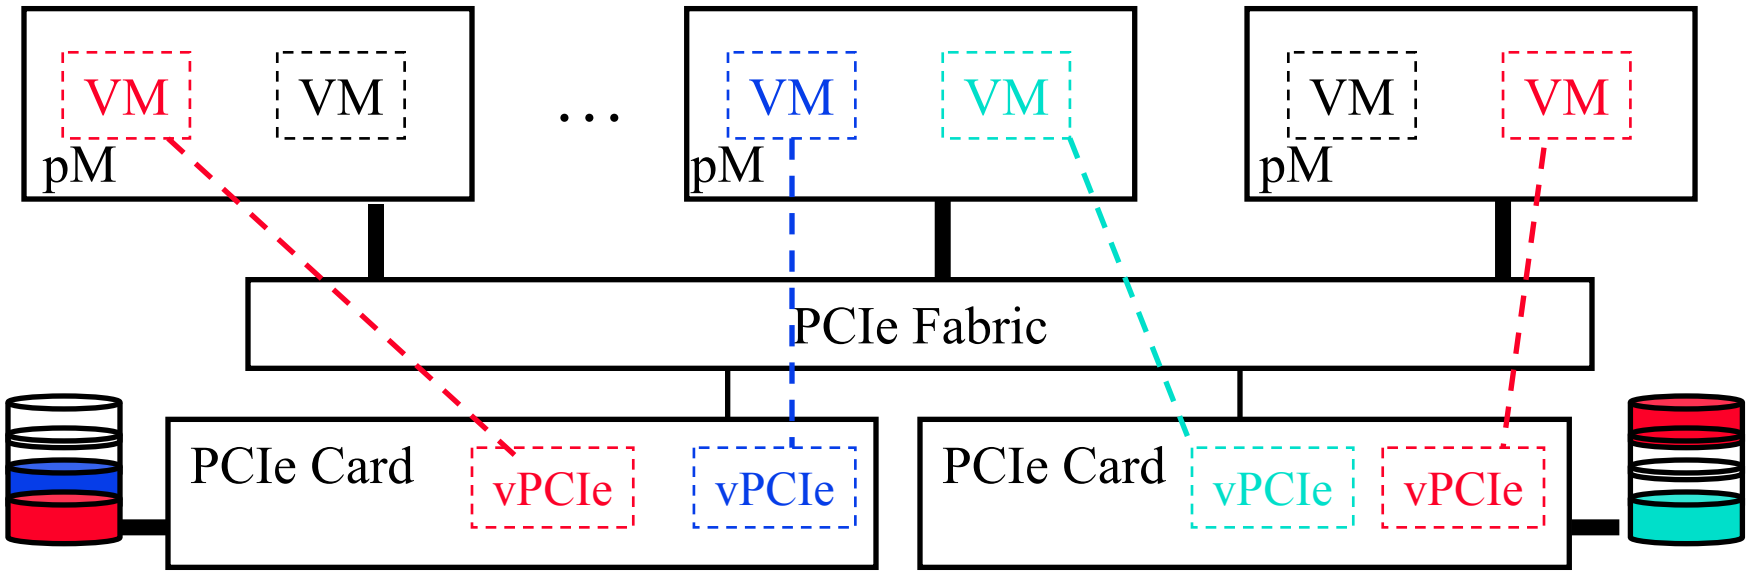

4. Virtualizing Data Center Storage

- IEEE 802.1BR-2012 Virtual Bridgeport Extension (VBE) allows multiple switches to combine in to a very large switch
- Storage and computers located anywhere in the data center appear as if connected to the same switch

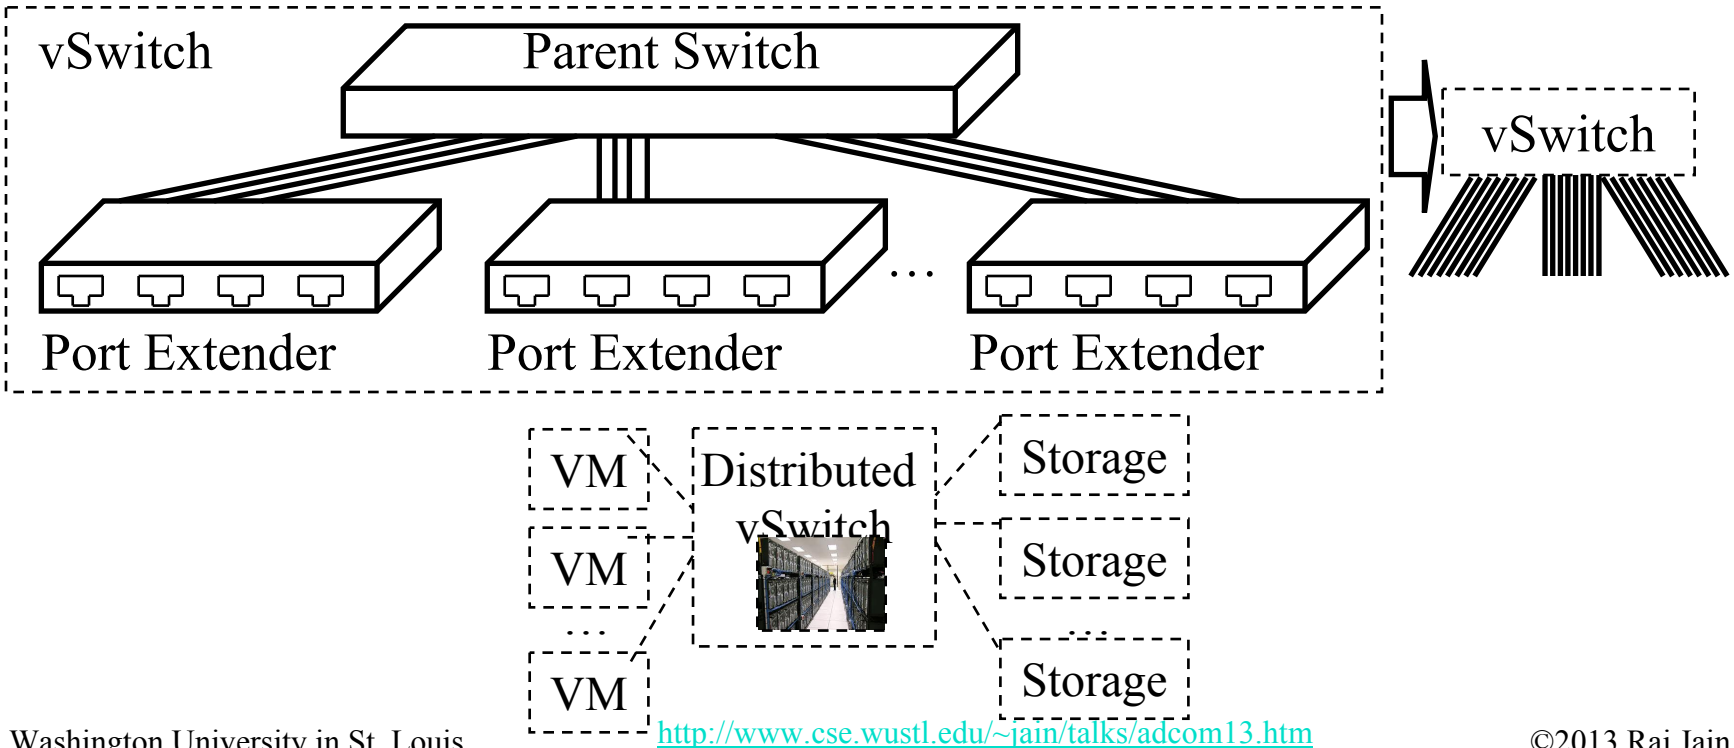

5. Virtualizing Metro Storage

- Data center Interconnection standards:
  - Virtual Extensible LAN (VXLAN),
  - Network Virtualization using GRE (NVGRE), and
  - Transparent Interconnection of Lots of Link (TRILL)

  ⇒ data centers located far away to appear to be on the same Ethernet

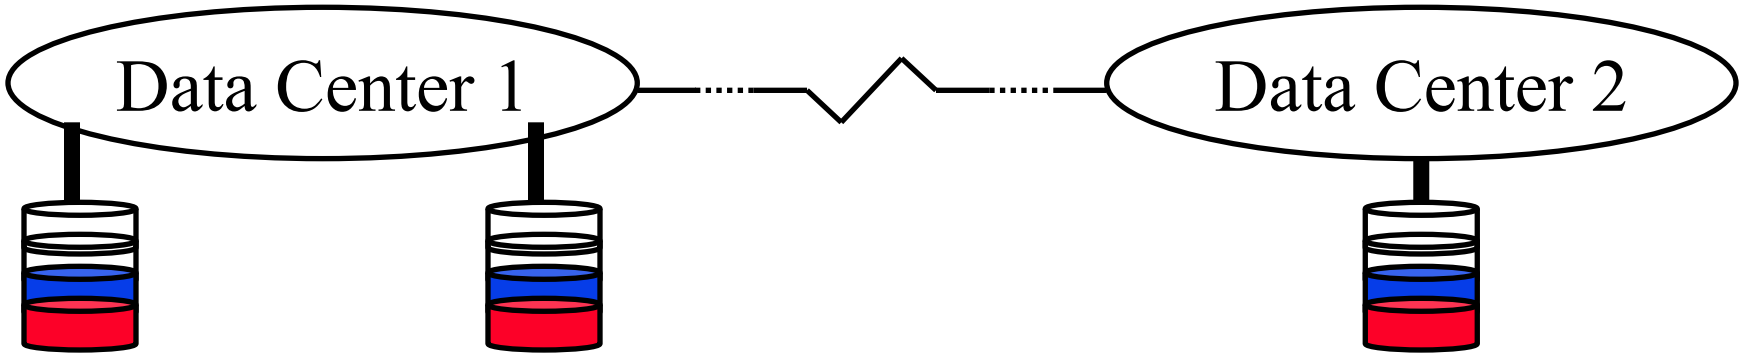

Ref: http://tools.ietf.org/html/draft-mahalingam-dutt-dcops-vxlan-04, http://tools.ietf.org/html/draft-sridharan-virtualization-nvgre-03, RFC 5556

Virtualizing the Global Storage

- Energy Science Network (ESNet) uses virtual switch to connect members located all over the world
- Virtualization ⇒ Fluid networks ⇒ The world is flat ⇒ You draw your network ⇒ Every thing is virtually local

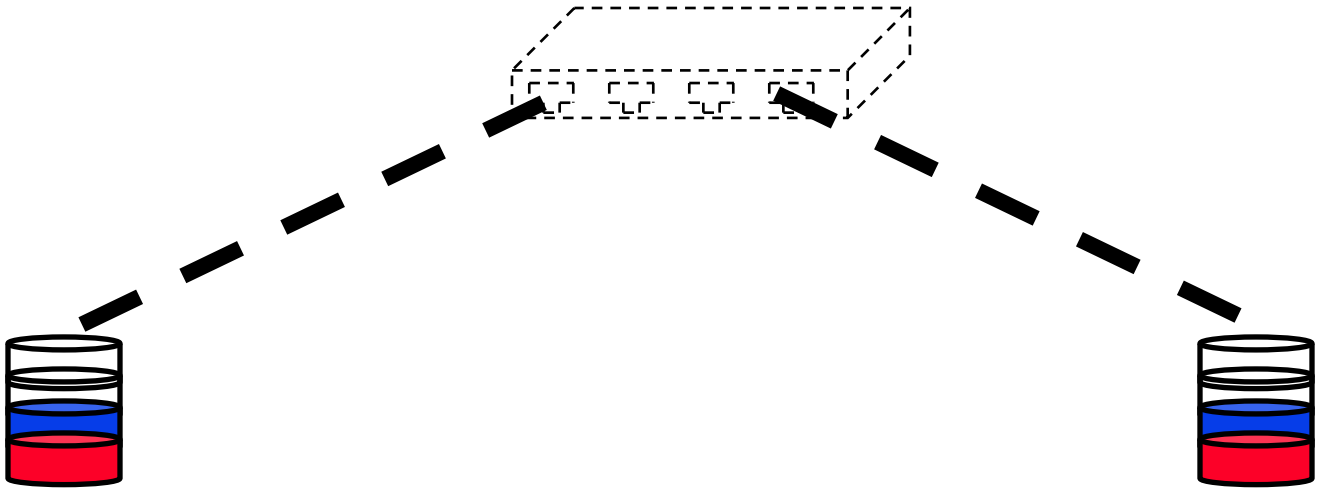

Ref: I. Monga, "Software Defined Networking for Big-data Science," http://www.es.net/assets/pubs_presos/Monga-WAN-Switch-SC12SRS.pdf

Software Defined Networking

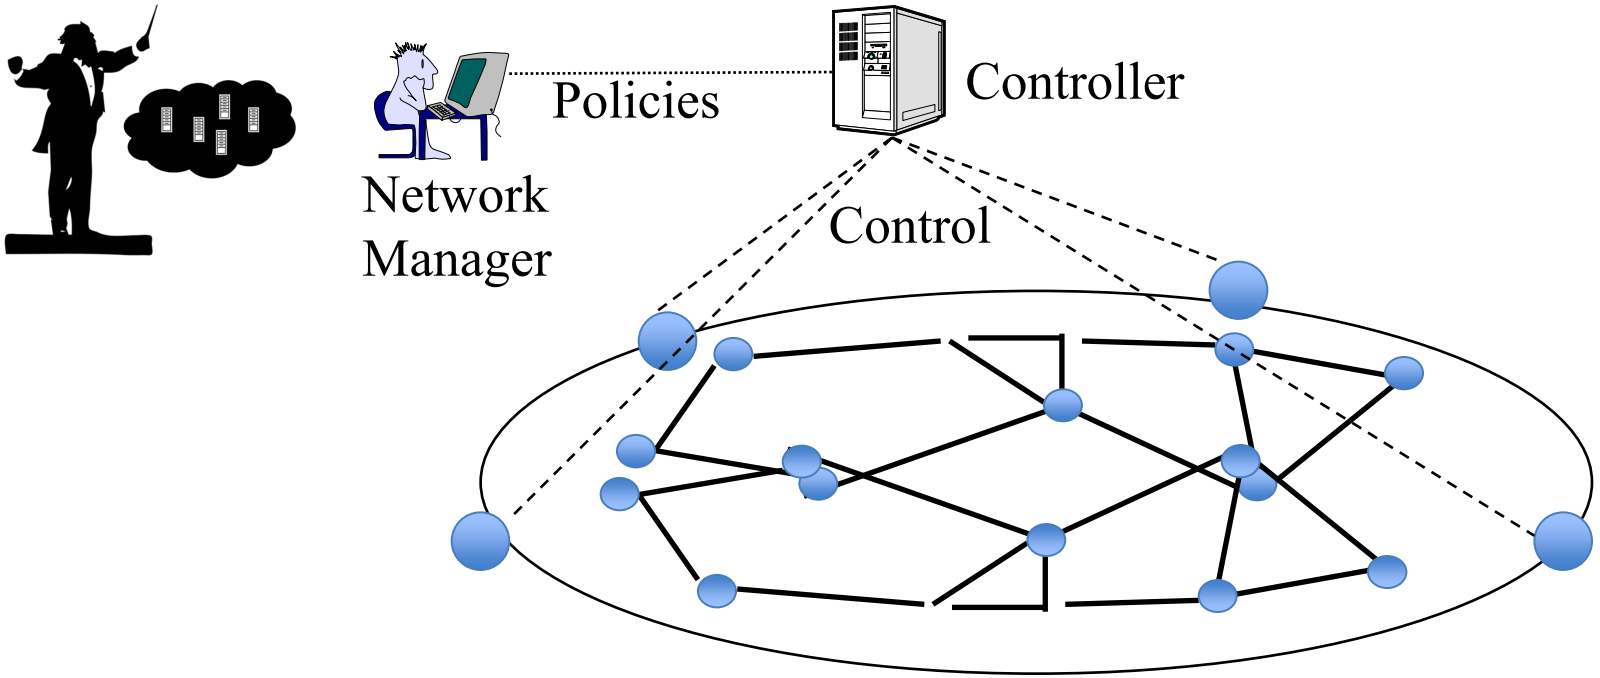

- Centralized Programmable Control Plane
- Allows automated orchestration (provisioning) of a large number of virtual resources (machines, networks, storage)
- Large Hadoop topologies can be created on demand

Network Function Virtualization (NFV)

- Fast standard hardware ⇒ Software based Devices
  Virtual networking modules (DHCP, Firewall, DNS, …)
  running on standard processors
- Modules can be combined to create any combination of function for data privacy, access control, …
- Virtual Machine implementation ⇒ Quick provisioning
- Standard Application Programming Interfaces (APIs)
  ⇒ Networking App Market
  ⇒ Privacy and Security for Big data in the multi-tenant clouds

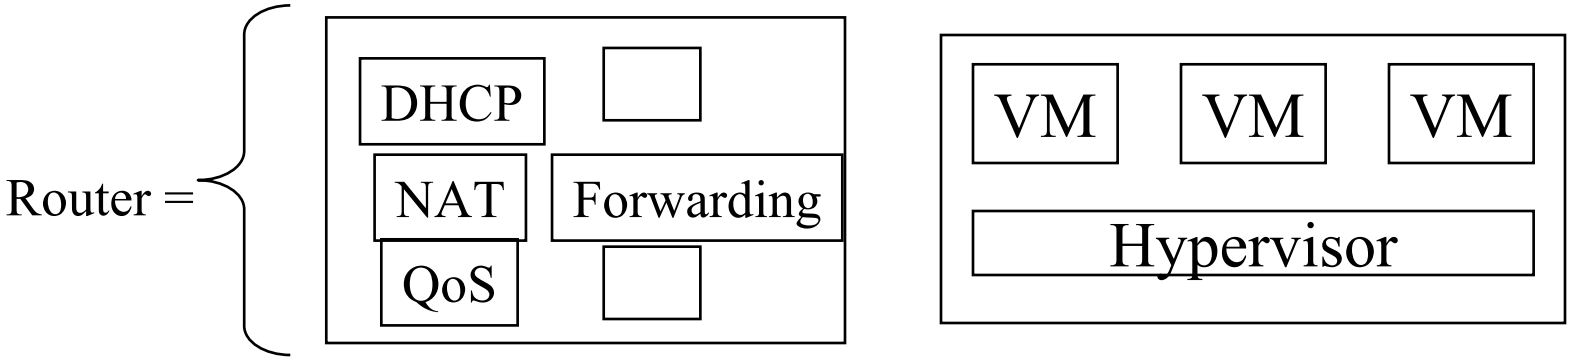

Big Data for Networking

- Today's data center:
  - Tens of tenants
  - Hundreds of switches and routers
  - Thousands of servers
  - Hundreds of administrators
- Tomorrow:
  - 1k of clients
  - 10k of pSwitches
    $\Rightarrow$ 100k of vSwitches
  - 1M of VMs
  - Tens of Administrators
- Need to monitor traffic patterns and rearrange virtual networks connecting millions of VMs in real-time
  $\Rightarrow$ Managing clouds is a real-time big data problem.
- Internet of things $\Rightarrow$ Big Data generation and analytics

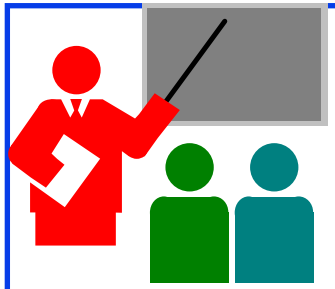

Summary

1. Virtualization has made networking, computing, and storage to be liquid. You can shape it (virtually) in any way you like.
2. I/O virtualization allows all storage in the rack to appear local to any VM in that rack ⇒ Solves the co-location problem of MapReduce
3. Network virtualization allows storage anywhere in the data center or even other data centers to appear local
4. Software defined networking allows orchestration of a large number of resources ⇒ Dynamic creation of Hadoop clusters
5. Network function virtualization will allow these clusters to have special functions and security in multi-tenant clouds.

References

- Raj Jain, "CSE570S: Recent Advances in Networking - Data Center Networking, Virtualization, Software Defined Networking, Big Data, Cloud Computing, Internet of Things (Fall 2013)," http://www.cse.wustl.edu/~jain/cse570-13/index.html