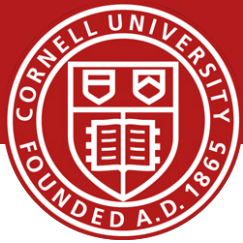

Data Center Virtualization: Open vSwitch

Hakim Weatherspoon

Assistant Professor, Dept of Computer Science

CS 5413: High Performance Systems and Networking

November 10, 2014

Slides from ACM SIGCOMM Workshop on Hot Topics in Networking (HotNets) presentation of “Extending networking into the virtualization layer”

Goals for Today

- Extending networking into the virtualization layer
  - B. Pfaff, J. Pettit, T. Koponen, K. Amidon, M. Casado, S. Shenker. ACM SIGCOMM Workshop on Hot Topics in Networking (HotNets), October 2009.

Outline

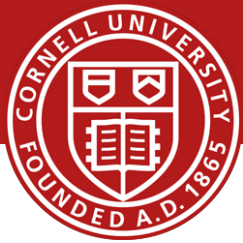

- Motivation
- Design
- Applications
- Implementation
- Evaluation
- Discussion/Future Work

Motivation

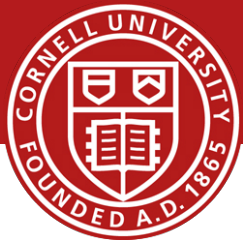

- Virtualization is pervasive
  - At the time of writing in 2009:
    - 12% of workloads were virtual
    - Gartner predicted by 2013, 61% of workloads would be virtual
  - Scott Shenker's talk: Virtual switches is more numerous than physical switches today
  - Intel: All endhosts should be virtualized

Motivation

Networking in virtual environments is important

- Clouds routinely host 40, 60, 120 or more virtual hosts per physical host
  - 128VM's per host is more than 2 racks full of machines

Motivation

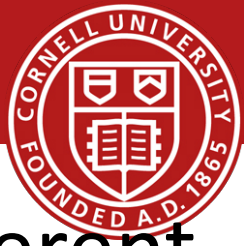

Networking in virtual environments is different

- Challenges
  - Scalability (10^5 VMs or much more)
  - Isolation
  - Mobility

- Conveniences
  - Hypervisor info
  - Introspection
  - Leaf nodes

Motivation

Networking in virtual environments is different

- Challenges
  - Scalability (10^5 VMs or much more)
  - Isolation
  - Mobility
- Conveniences
  - Hypervisor info
  - Introspection
  - Leaf nodes

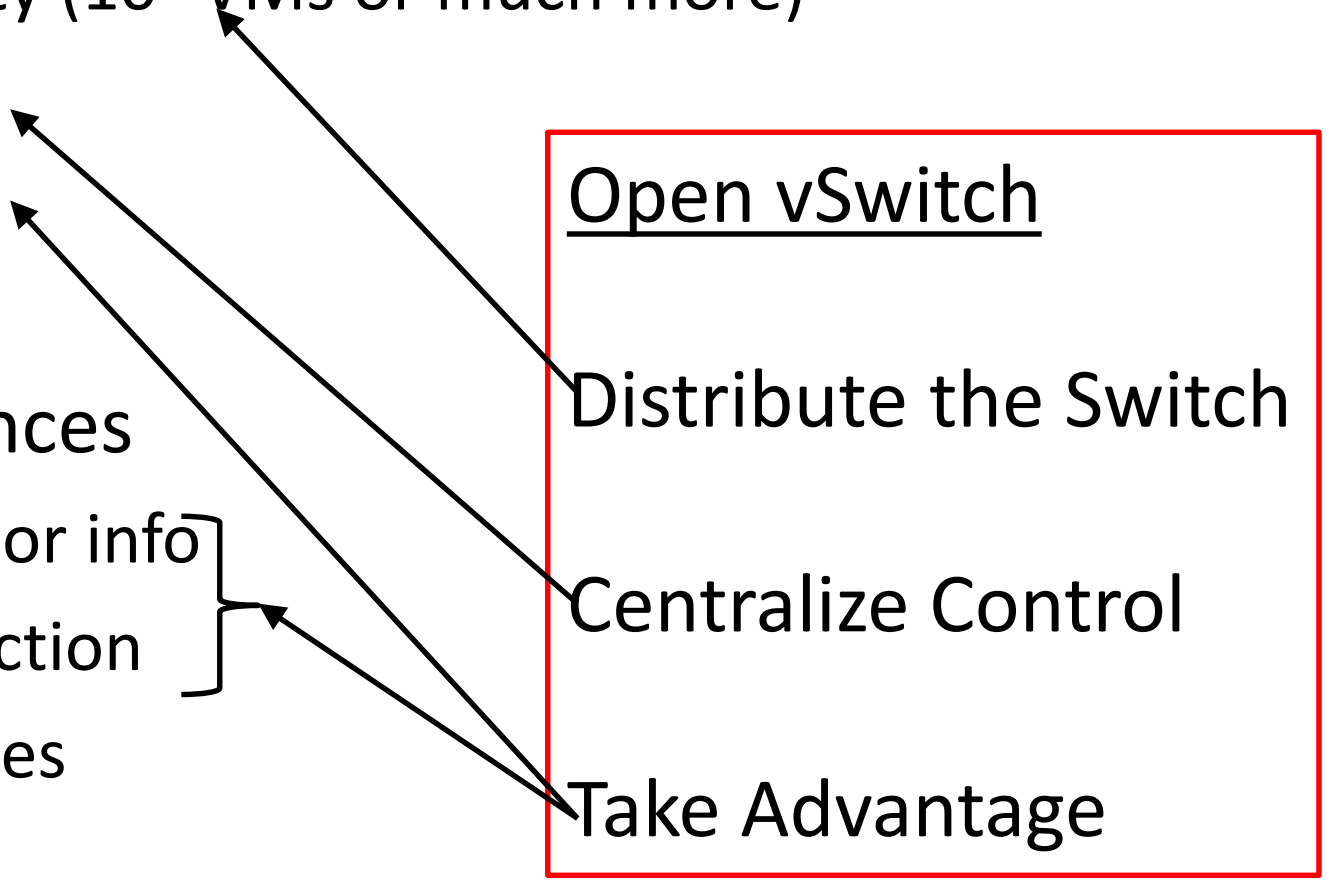

Outline

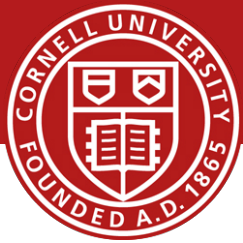

- Motivation
- Design
- Applications
- Implementation
- Evaluation
- Discussion/Future Work

Open vSwitch: Design

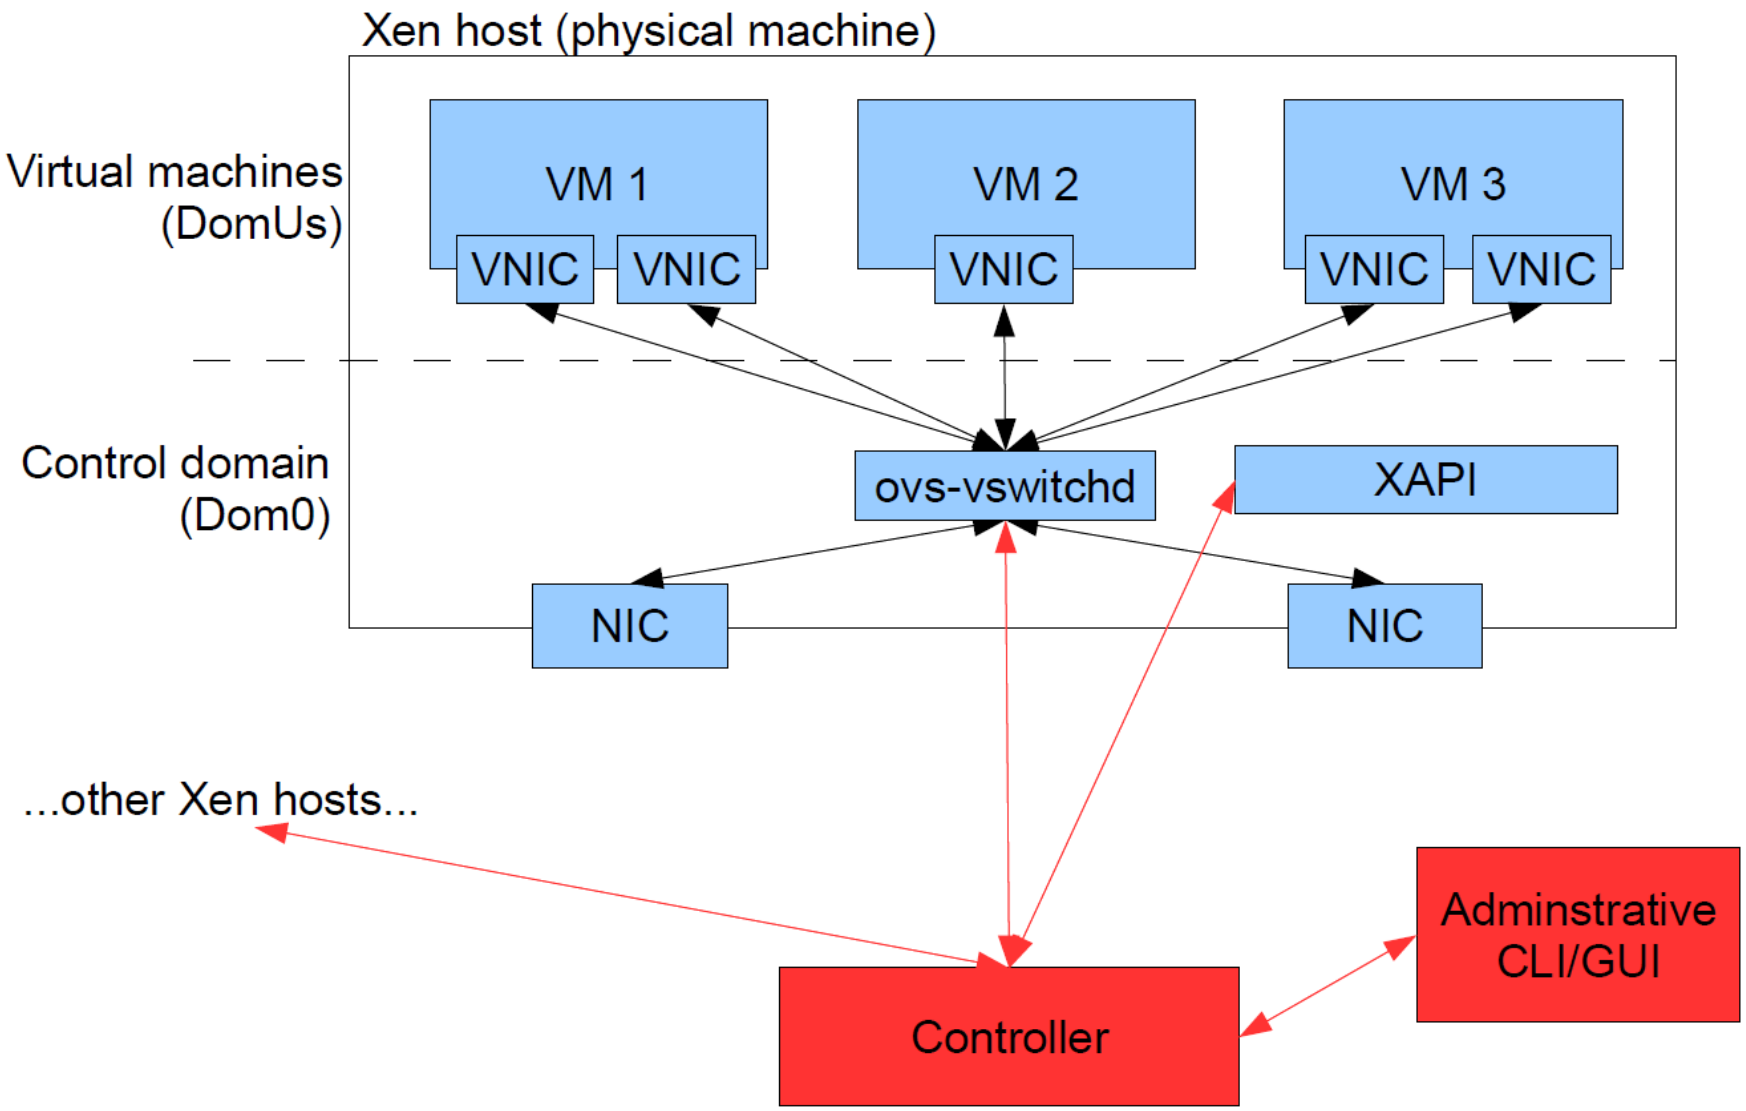

Open vSwitch

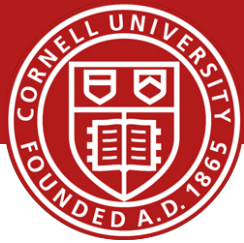

- Control Plane/Controller/OpenFlow
  - Configuration
  - Connectivity Management
    - creating switches, managing virtual interface (VIF) connectivity, and managing physical interface (PIF) connectivity
      - for each connected VIF, a logical port is added to the switch
  - Features
    - VLAN, Port Mirroring, ACLs, NetFlow, Bonding, QoS, Anything*
- Data Plane/Forwarding/OpenFlow
  - OpenFlow controller remotely controls forwarding table
  - Defines how packets handled based on L2,L3,L4 headers

Outline

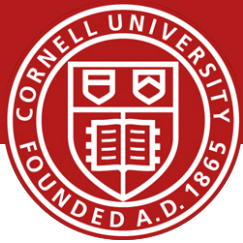

- Motivation
- Design
- Applications
- Implementation
- Evaluation
- Discussion/Future Work

Applications

Distributed Switch

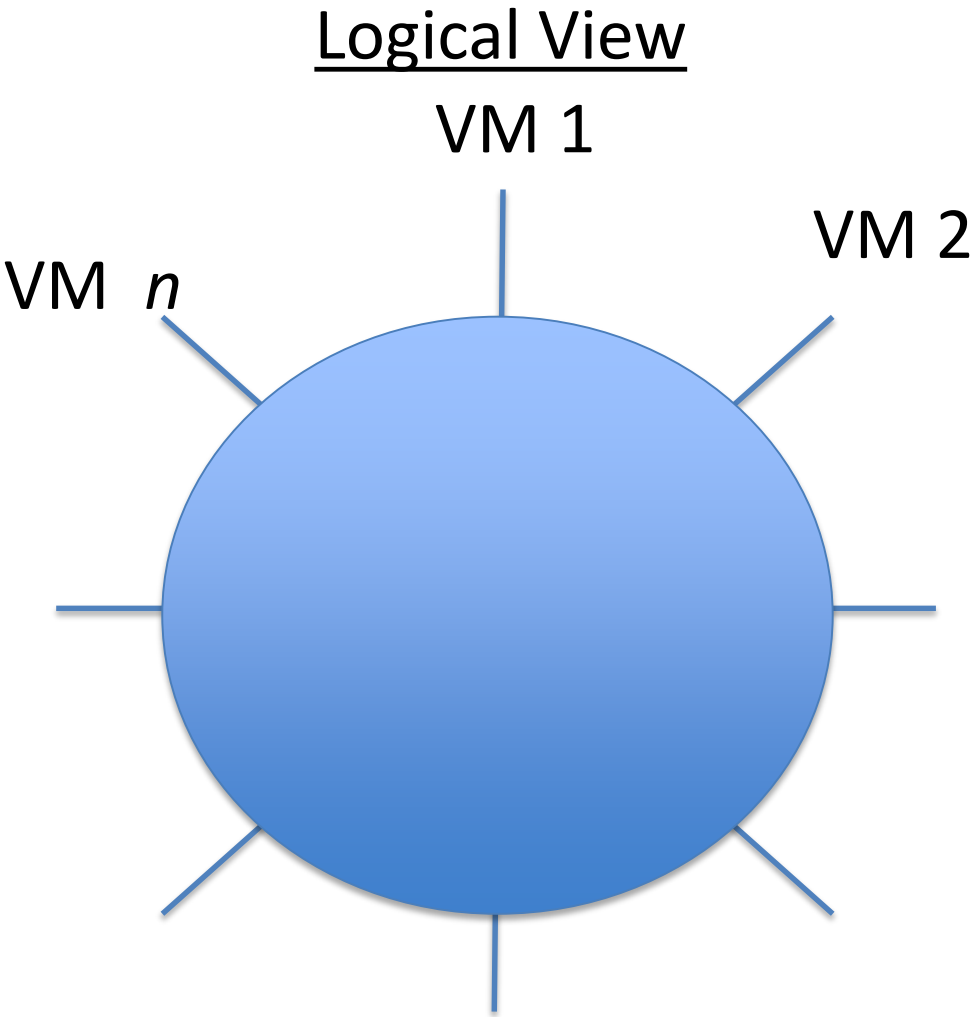

Applications

Distributed Switch

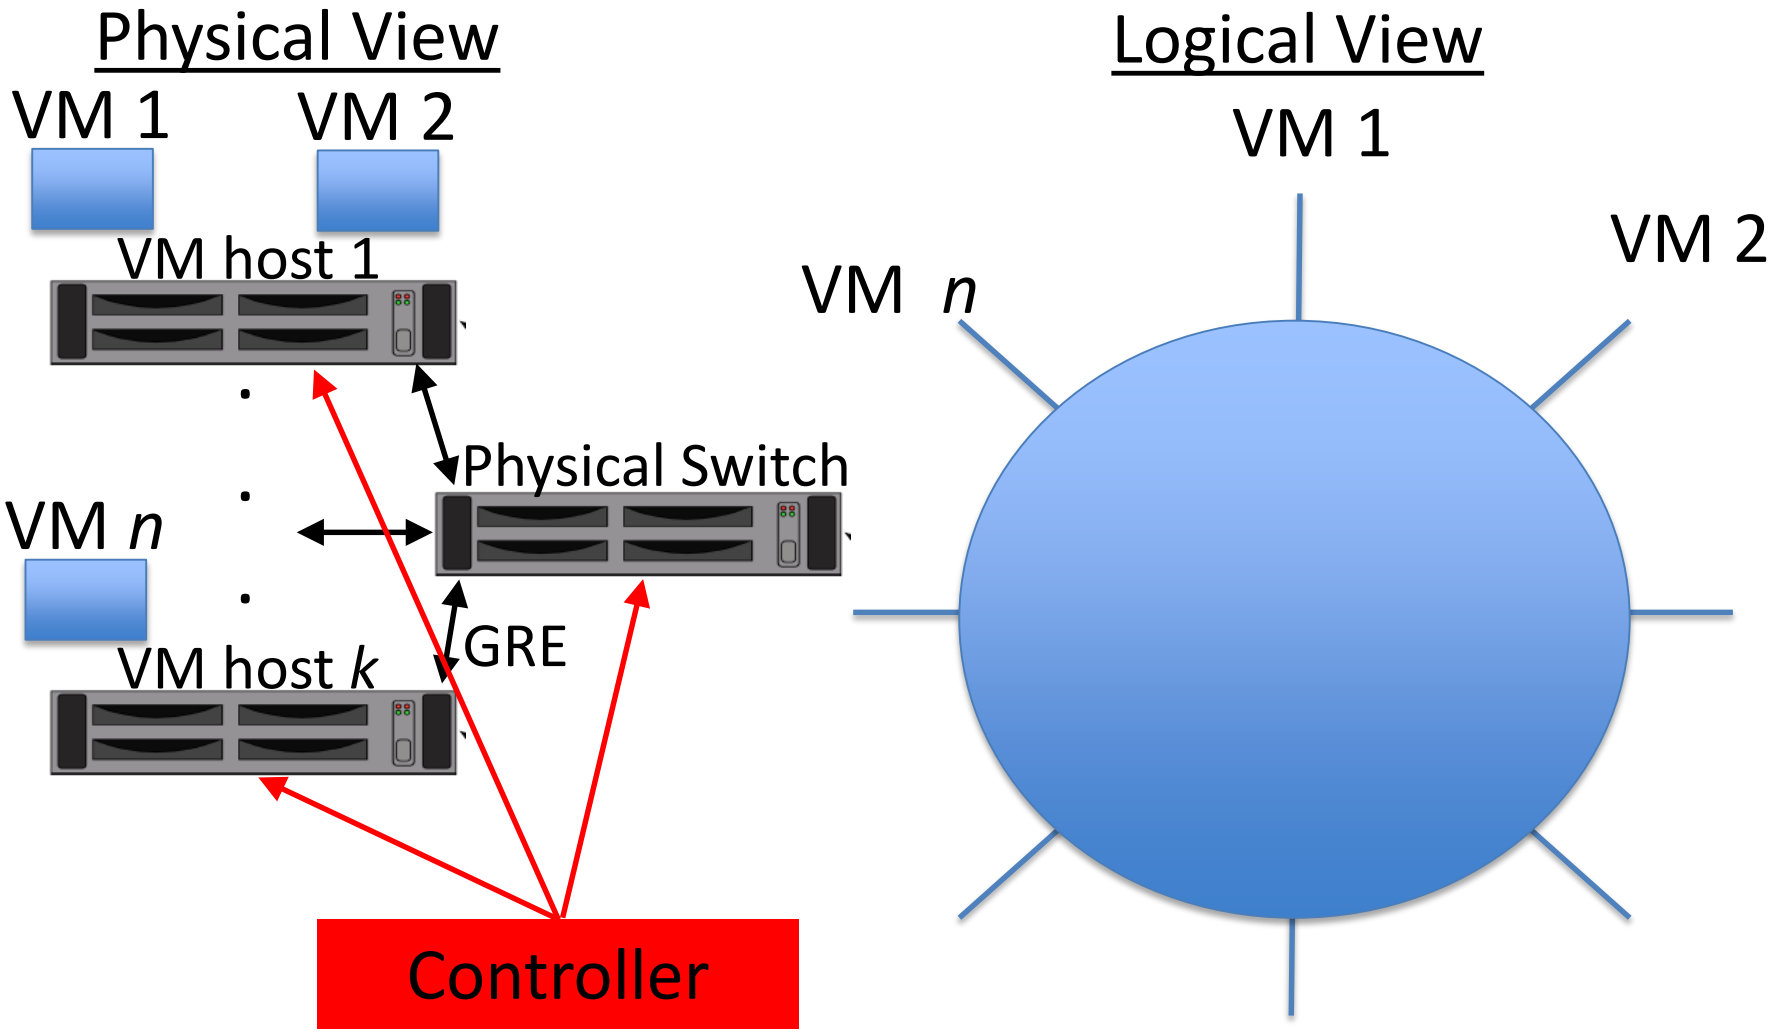

Applications

Extending the Data Center into the Cloud

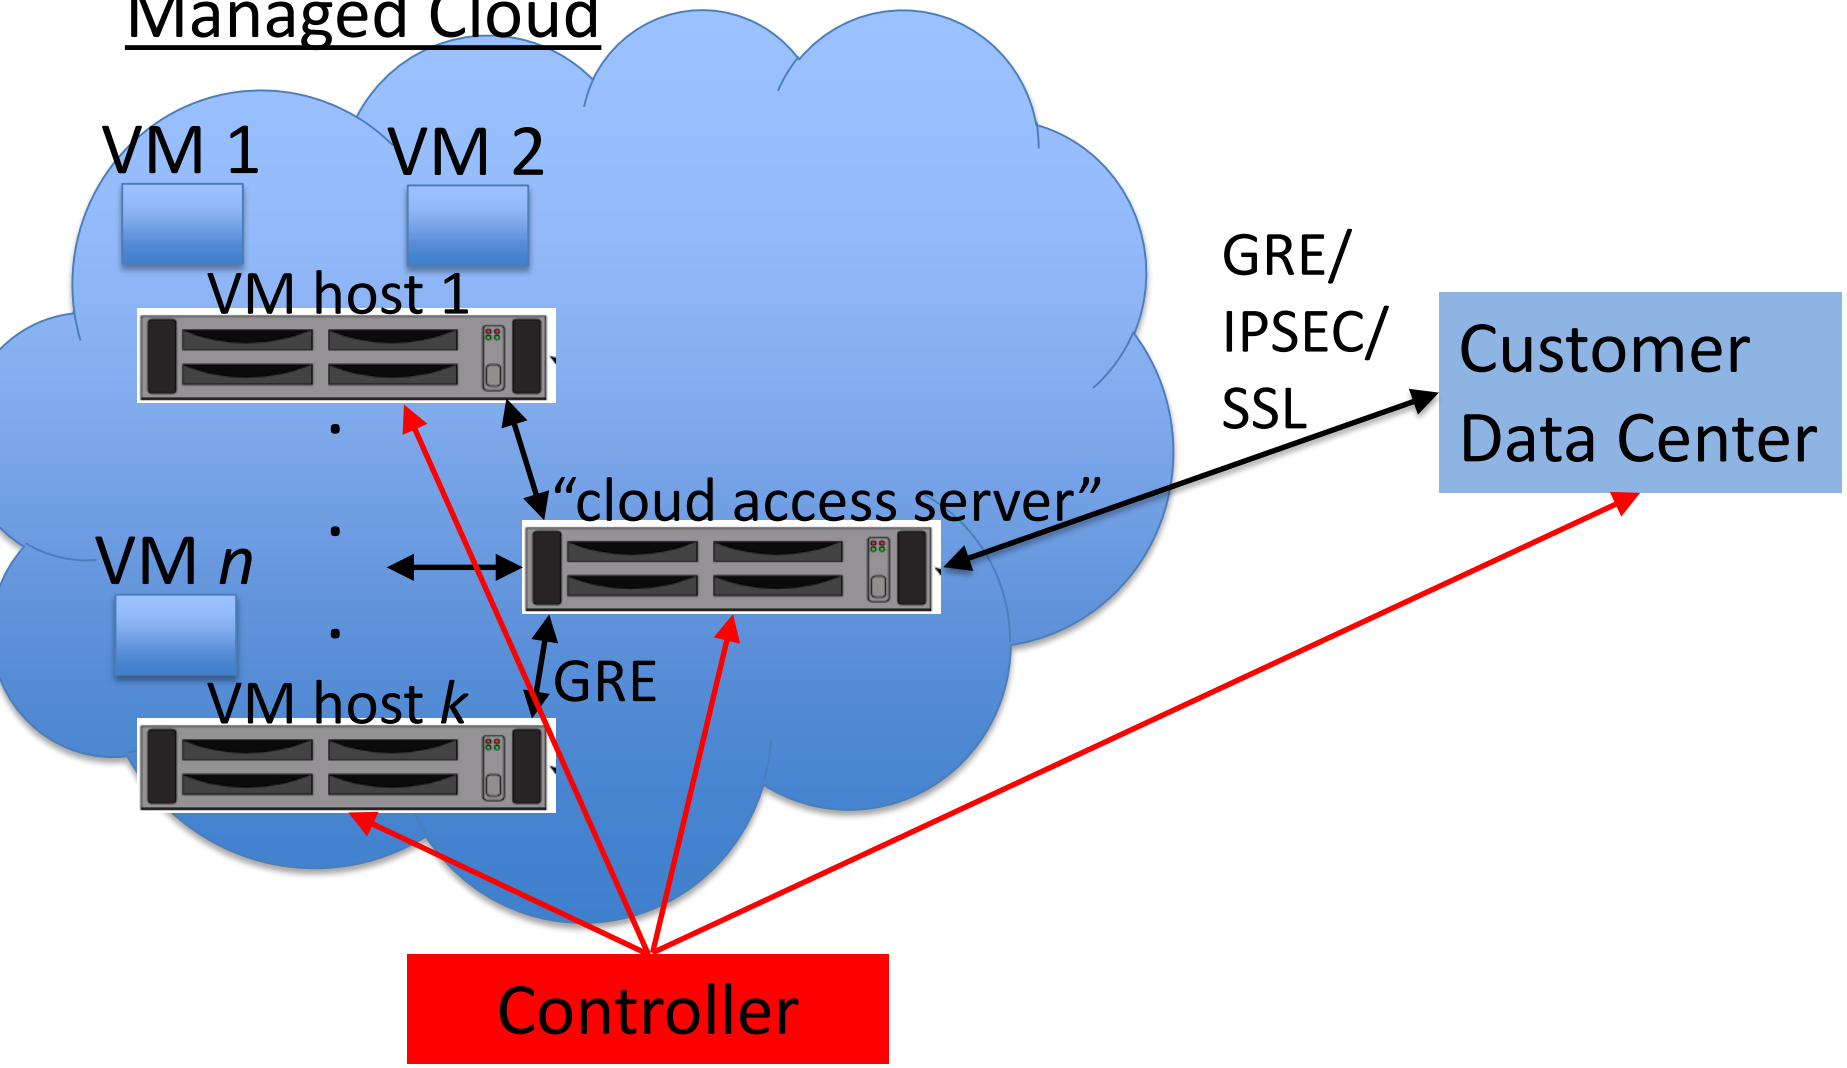

Outline

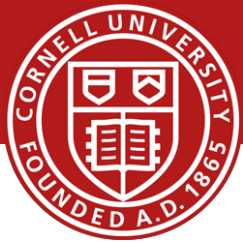

- Motivation
- Design
- Applications
- Implementation
- Evaluation
- Discussion/Future Work

Implementation

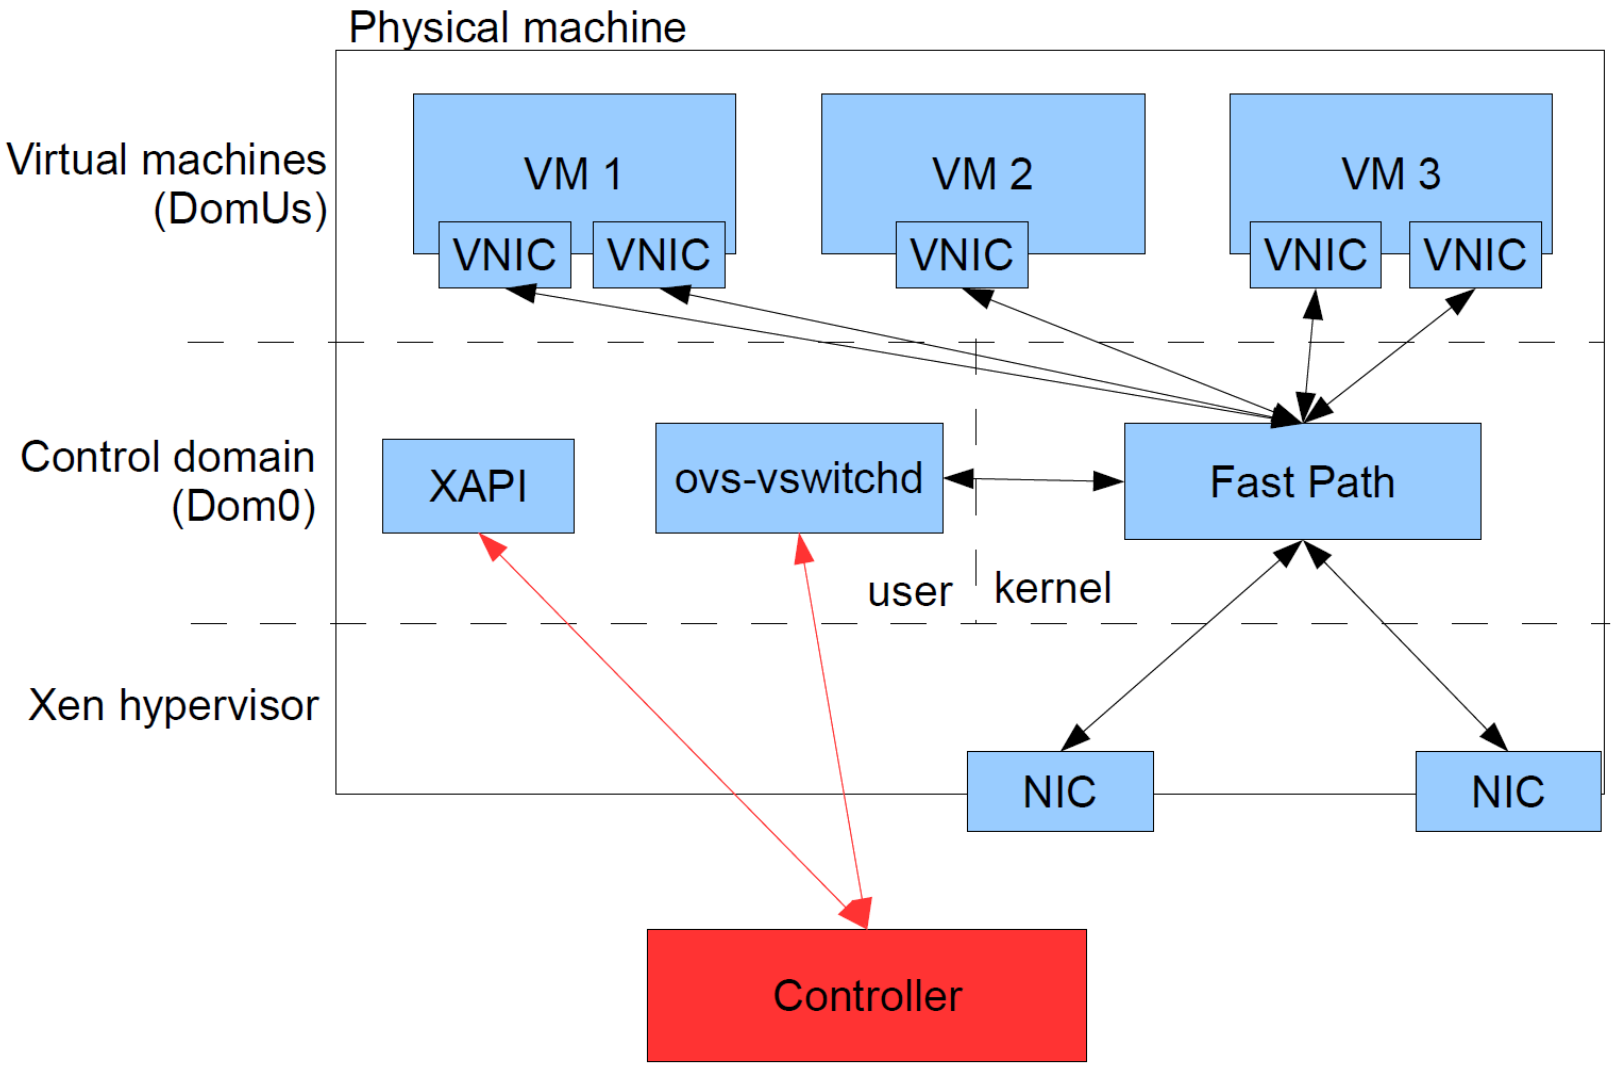

Outline

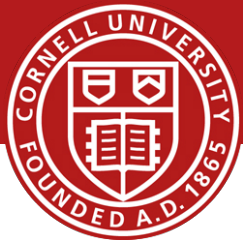

- Motivation
- Design
- Applications
- Implementation
- Evaluation
- Discussion/Future Work

Evaluation

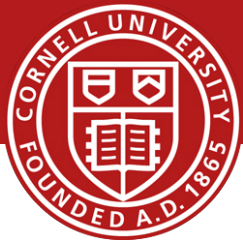

- Compare Open vSwitch to Linux Bridge

- Bandwidth
  - Fast Path: > 1Gbps
  - ovs-vswitchd: 100Mbps
  - Controller: 10Mbps
- Latency
  - Fast Path: < 1 us
  - ovs-vswitchd: < 1 ms
  - Controller: ms+

Evaluation

- Compare Open vSwitch to Linux Bridge

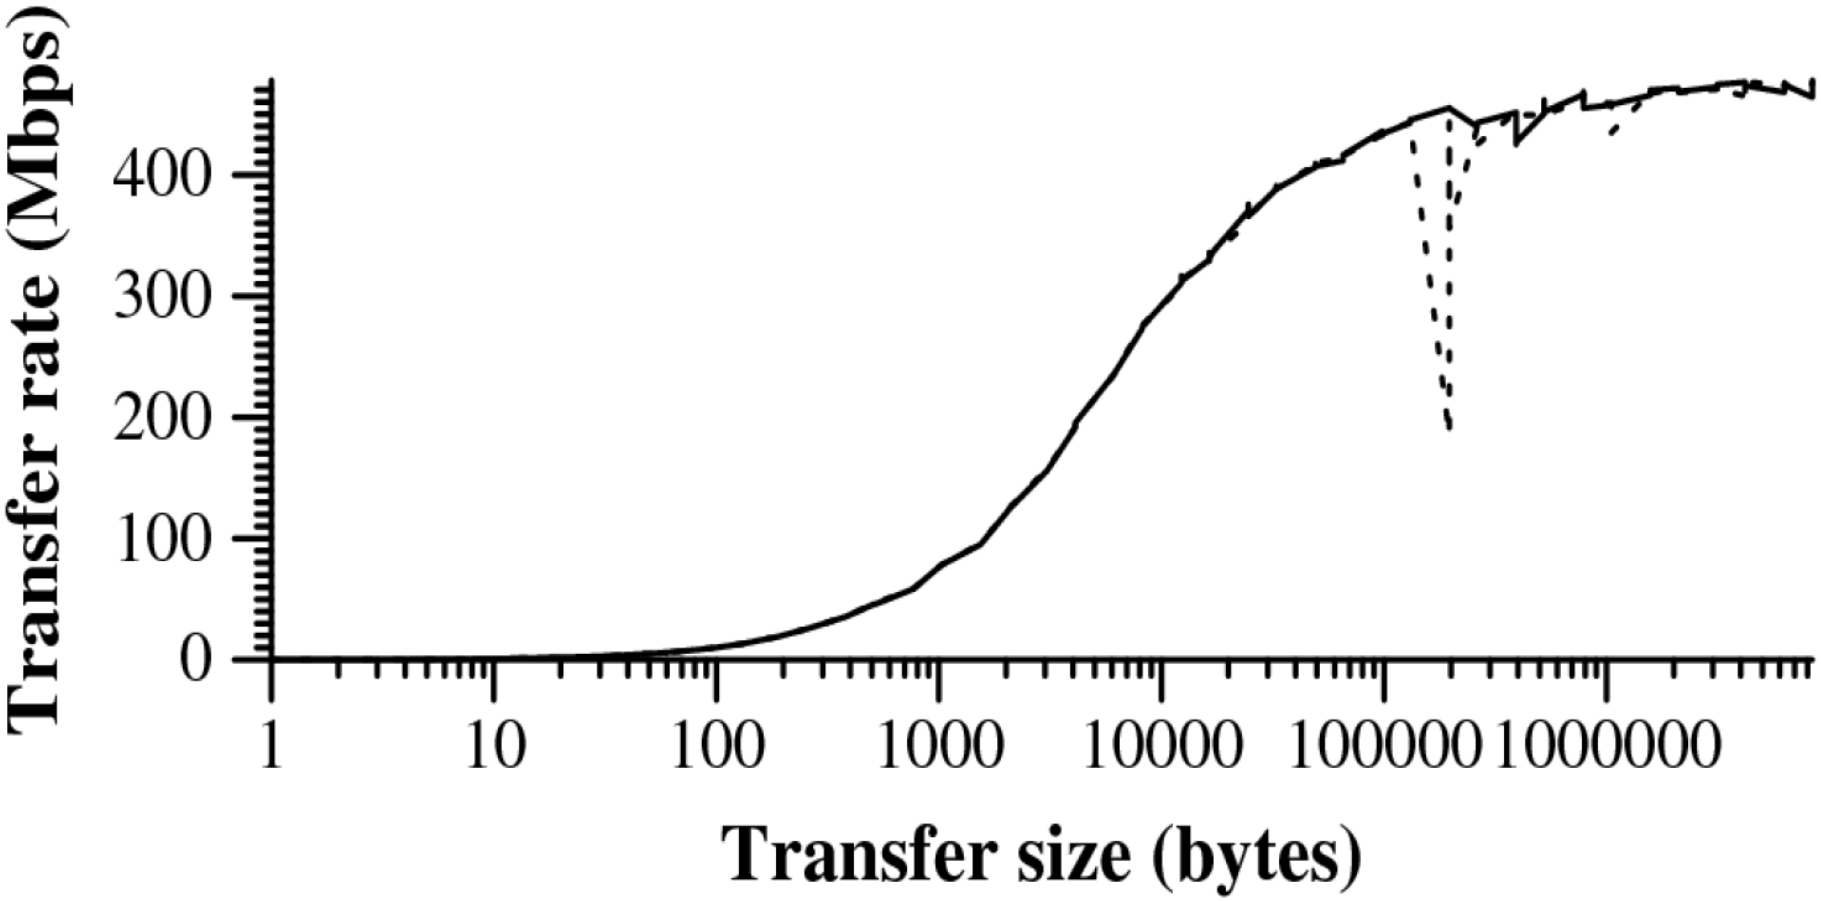

- Same performance as Linux bridge with same CPU

Outline

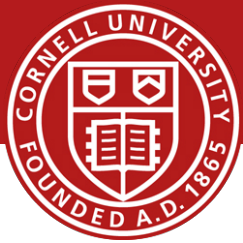

- Motivation
- Design
- Applications
- Implementation
- Evaluation
- Discussion/Future Work

Discussion/Future work

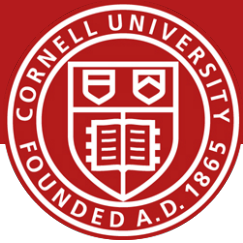

- Hardware acceleration in a virtual world?
  - Netronome, VN-Tag, VEPA

- Performance, performance, performance
- Physical switches integration
- Upstream kernel integration
- Anything*

Before Next time

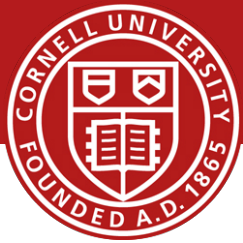

- Project Interim report
  - **Due Monday, November 24.**
  - And meet with groups, TA, and professor
- Fractus Upgrade: Should be back online

- ***Required review and reading for Friday, November 21***
  - VirtualWires for Live Migrating Virtual Networks across Clouds, D. Williams, H. Jamjoom, Z. Jiang, and H. Weatherspoon. IBM Tech. Rep. RC25378, April 2013.
  - http://domino.research.ibm.com/library/cyberdig.nsf/papers/FD9A14E59B138E7E85257B6000572CC3/$File/rc25378.pdf

- Check piazza: http://piazza.com/cornell/fall2014/cs5413
- Check website for updated schedule