

# Open vSwitch: Extending Networking into the Virtualization Layer

Ben Pfaff  
Justin Pettit  
Teemu Kojonen  
Keith Amidon  
Martin Casado  
Nicira Networks, Inc.

Scott Shenker  
UC Berkeley, Computer Science Division

# Outline

- Virtualization and Networking
- Open vSwitch approach
- Applications
- Implementation

# Virtualization Will Be Pervasive

Gartner:

12% of workloads are virtual today.  
61% by 2013.

Intel:

All end hosts should be virtualized.

# Networking in Virtual Environments is Important

One cloud is planning to run 128 VMs per host.  
That's 2+ full racks in one machine.



# Networking in Virtual Environments is Different

## Challenges

- Scalability ( $10^5$  VMs)
- Isolation
- Mobility
- ...

## Conveniences

- Hypervisor info
- Introspection
- Leaf nodes
- ...

# Networking in Virtual Environments is Different

## Challenges

- Scalability ( $10^5$  VMs)
- Isolation
- Mobility
- ...

## Conveniences

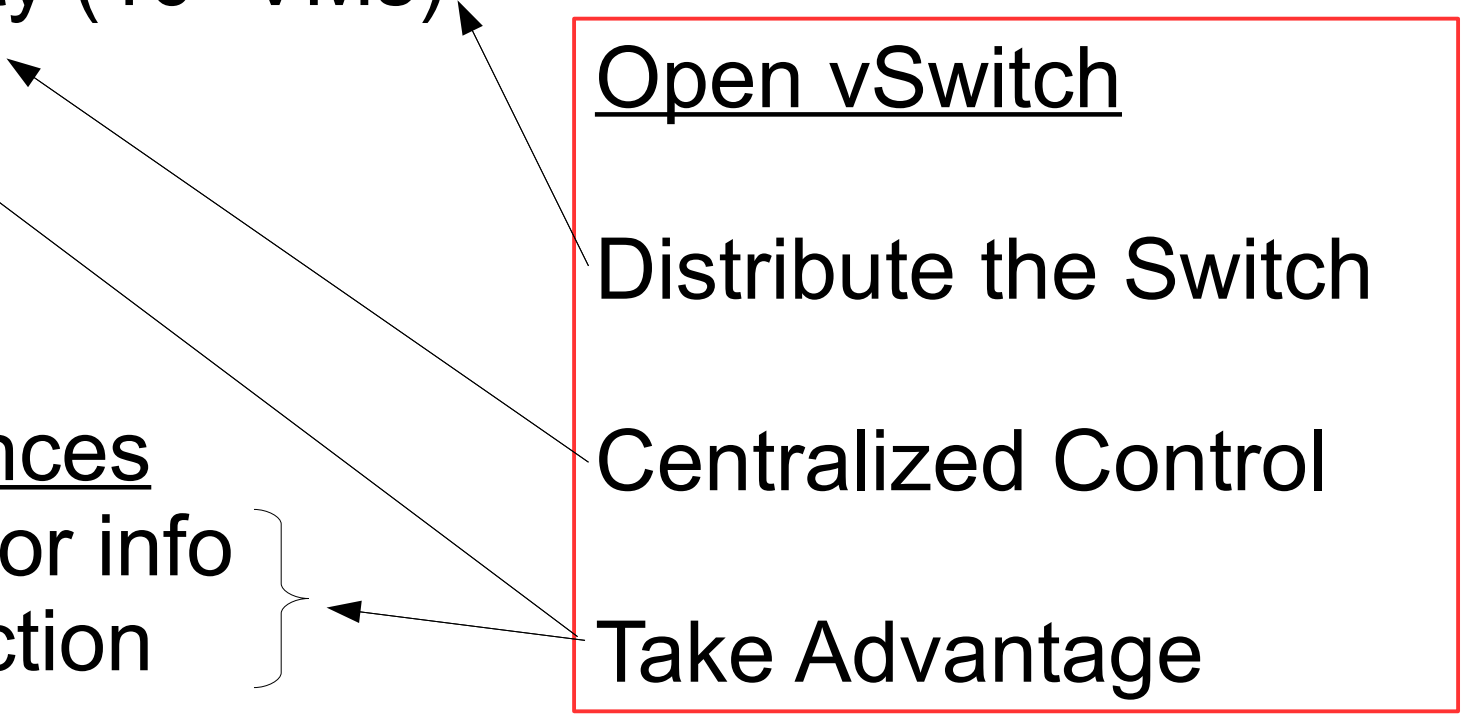
- Hypervisor info
- Introspection
- Leaf nodes
- ...

Open vSwitch

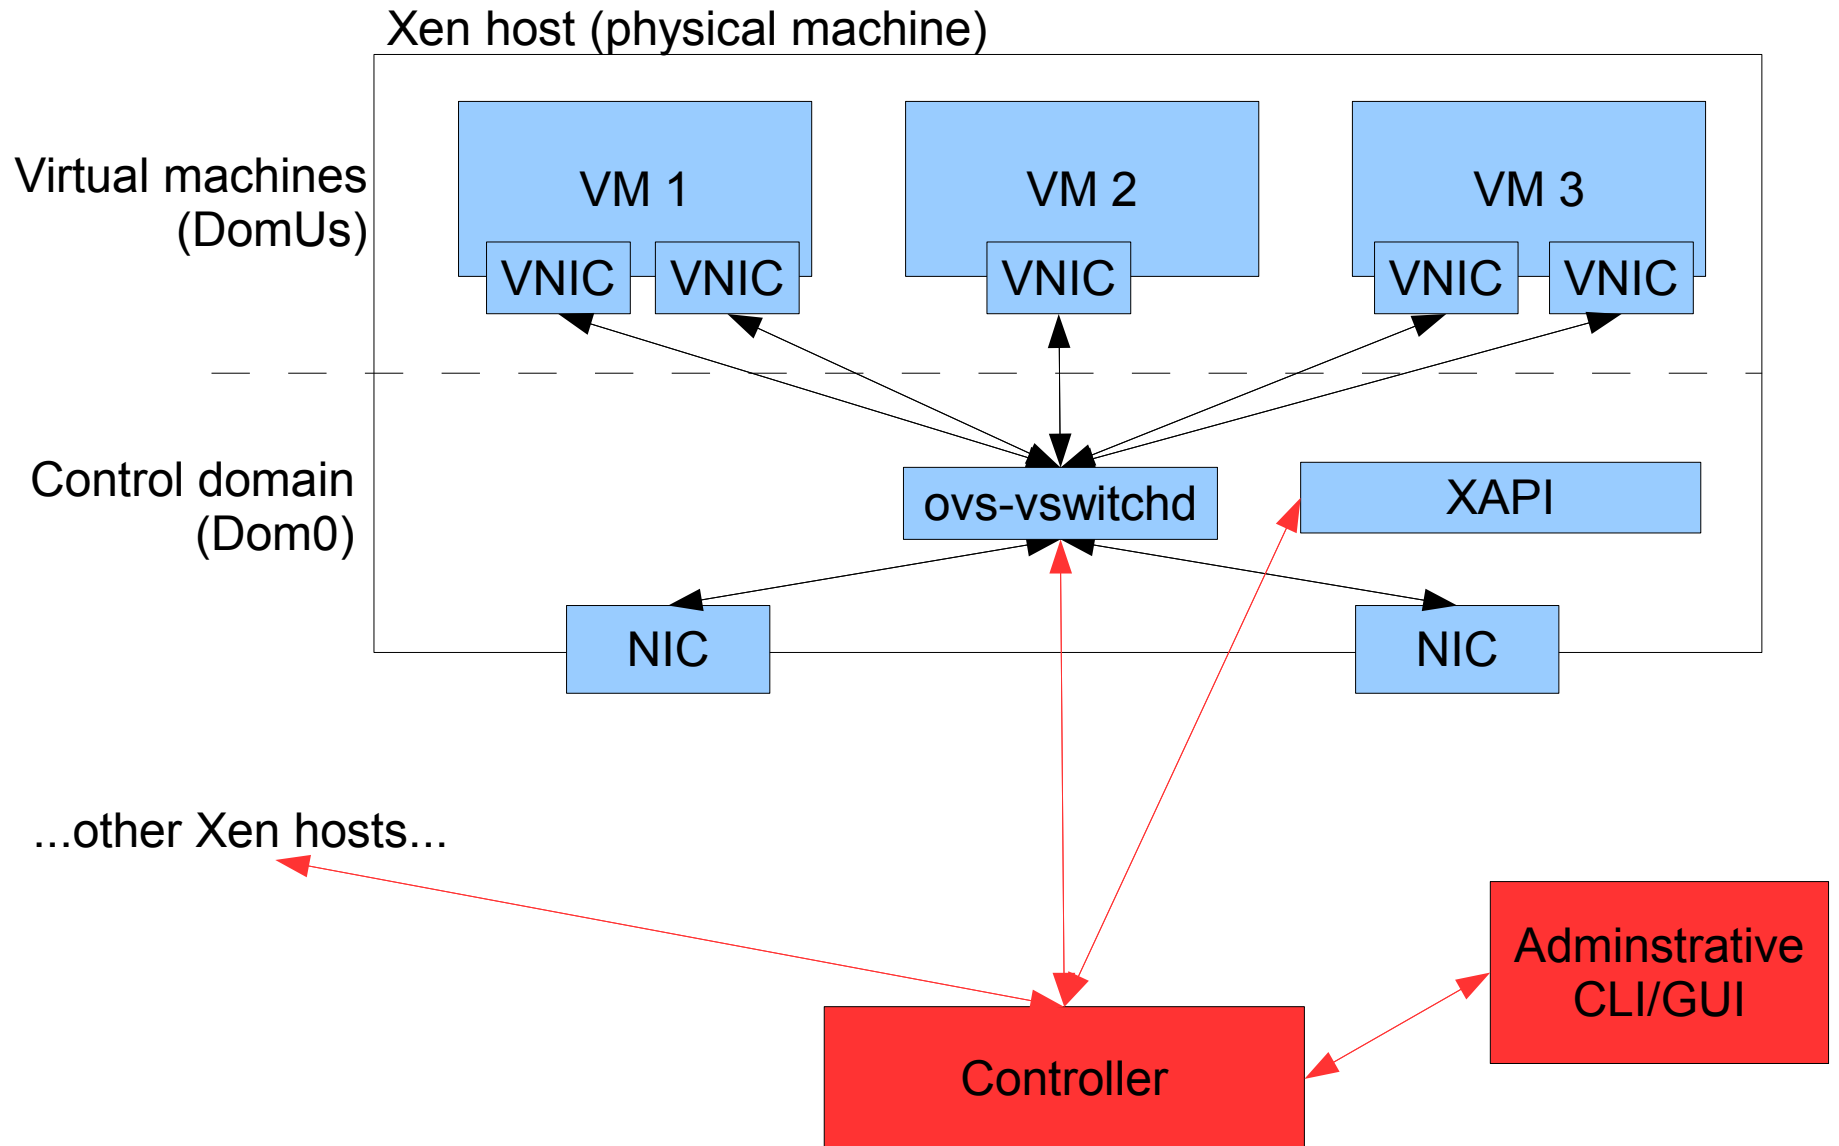
Distribute the Switch

Centralized Control

Take Advantage



# Basic Design (Xen)

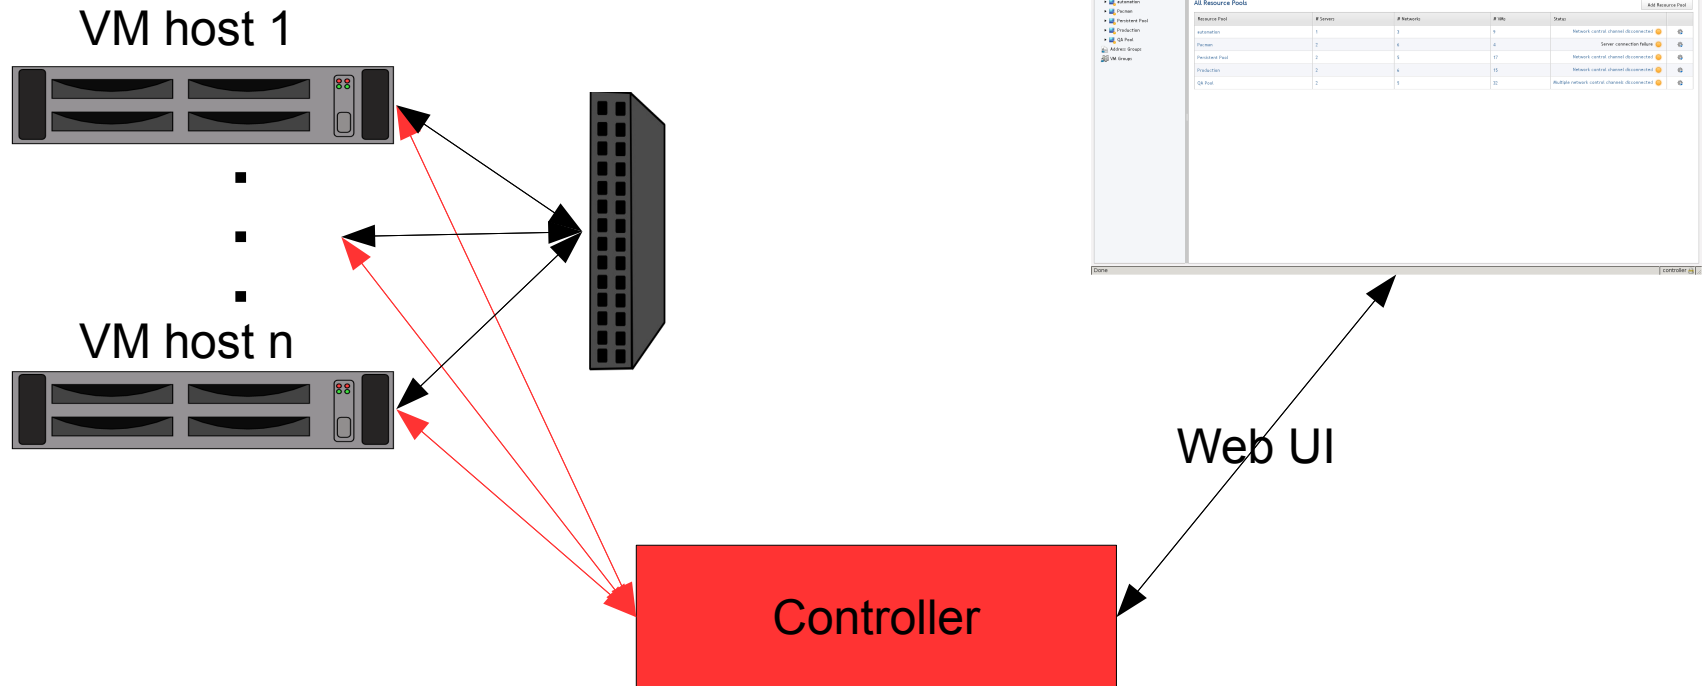


# Open vSwitch

- Controller:
  - Configuration
  - OpenFlow
- Features:
  - VLAN
  - Port mirroring
  - ACLs
  - NetFlow
  - Bonding
  - QoS
  - Anything\*

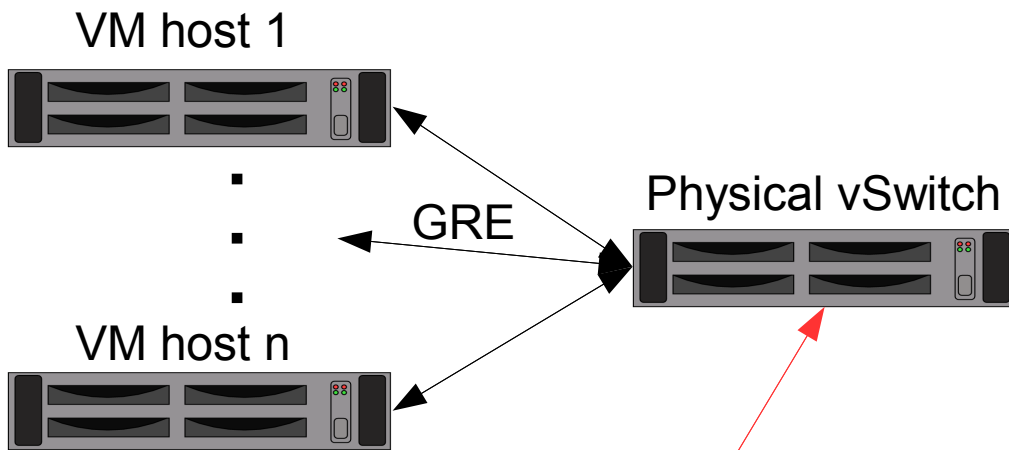


# Open vSwitch Application: Single Distributed Switch

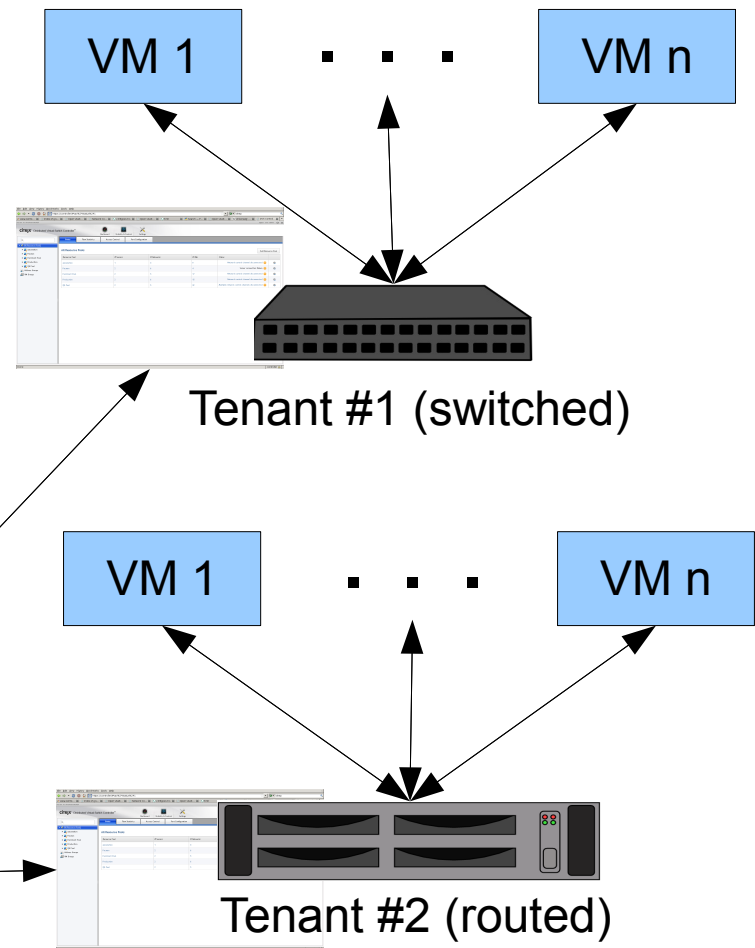


# Open vSwitch Application: Multiple Distributed Switches

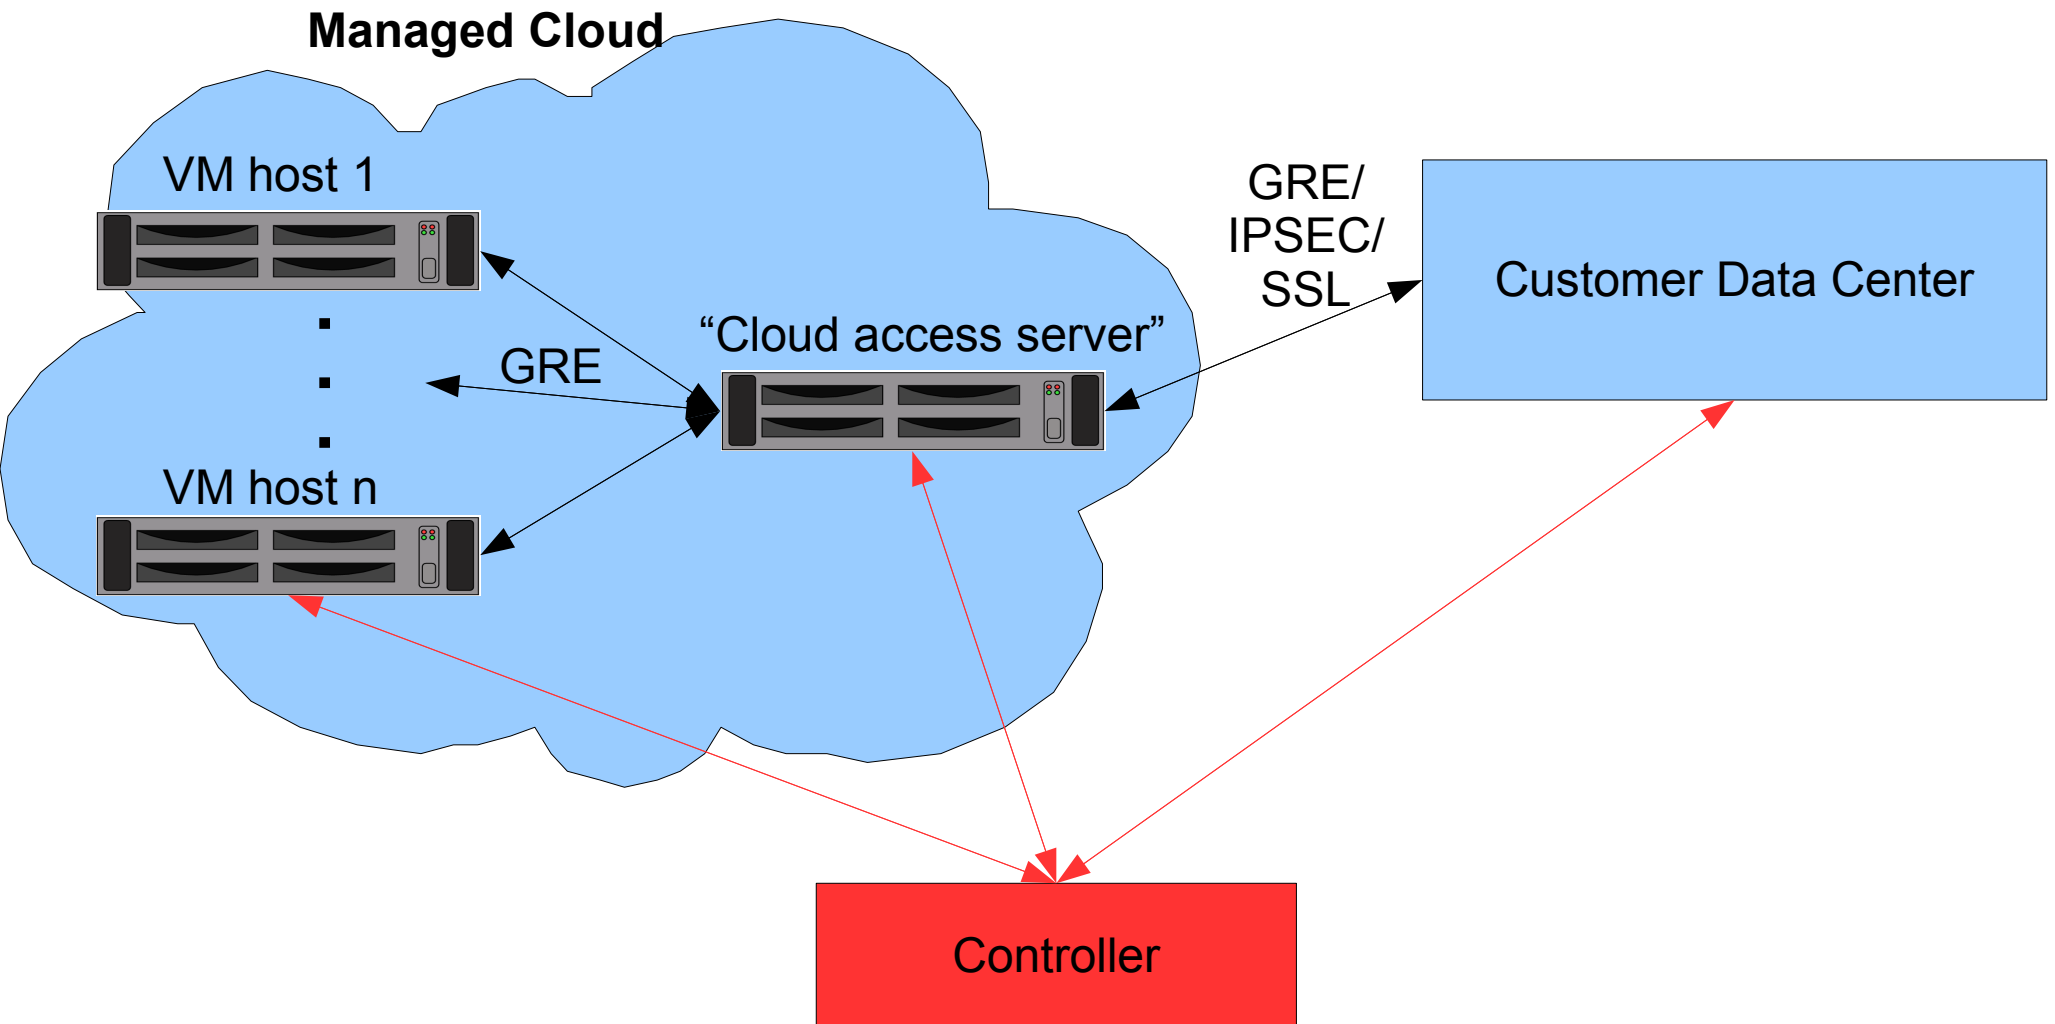
Physical



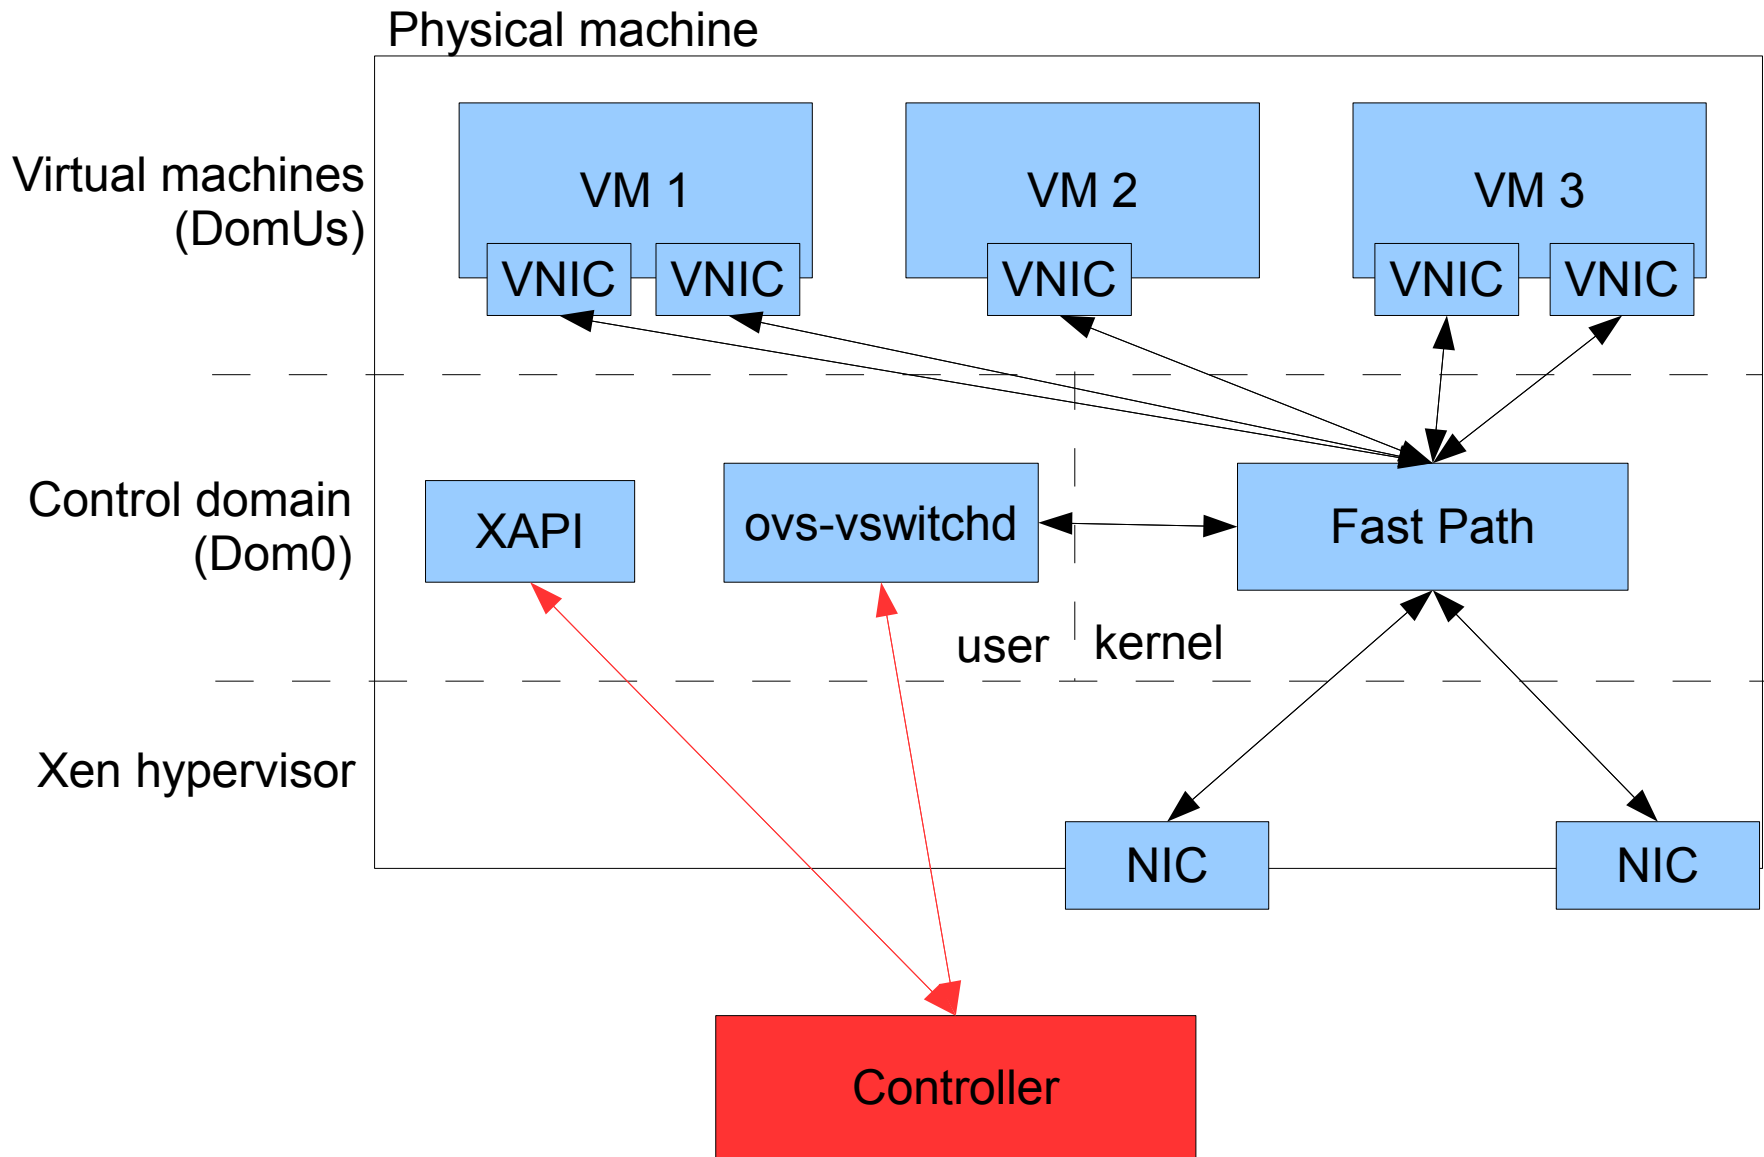
Logical



# Open vSwitch Application: Extending Data Center into Cloud



# Implementation (Xen)



# Open vSwitch is Fast

As fast as Linux bridge  
with same CPU usage

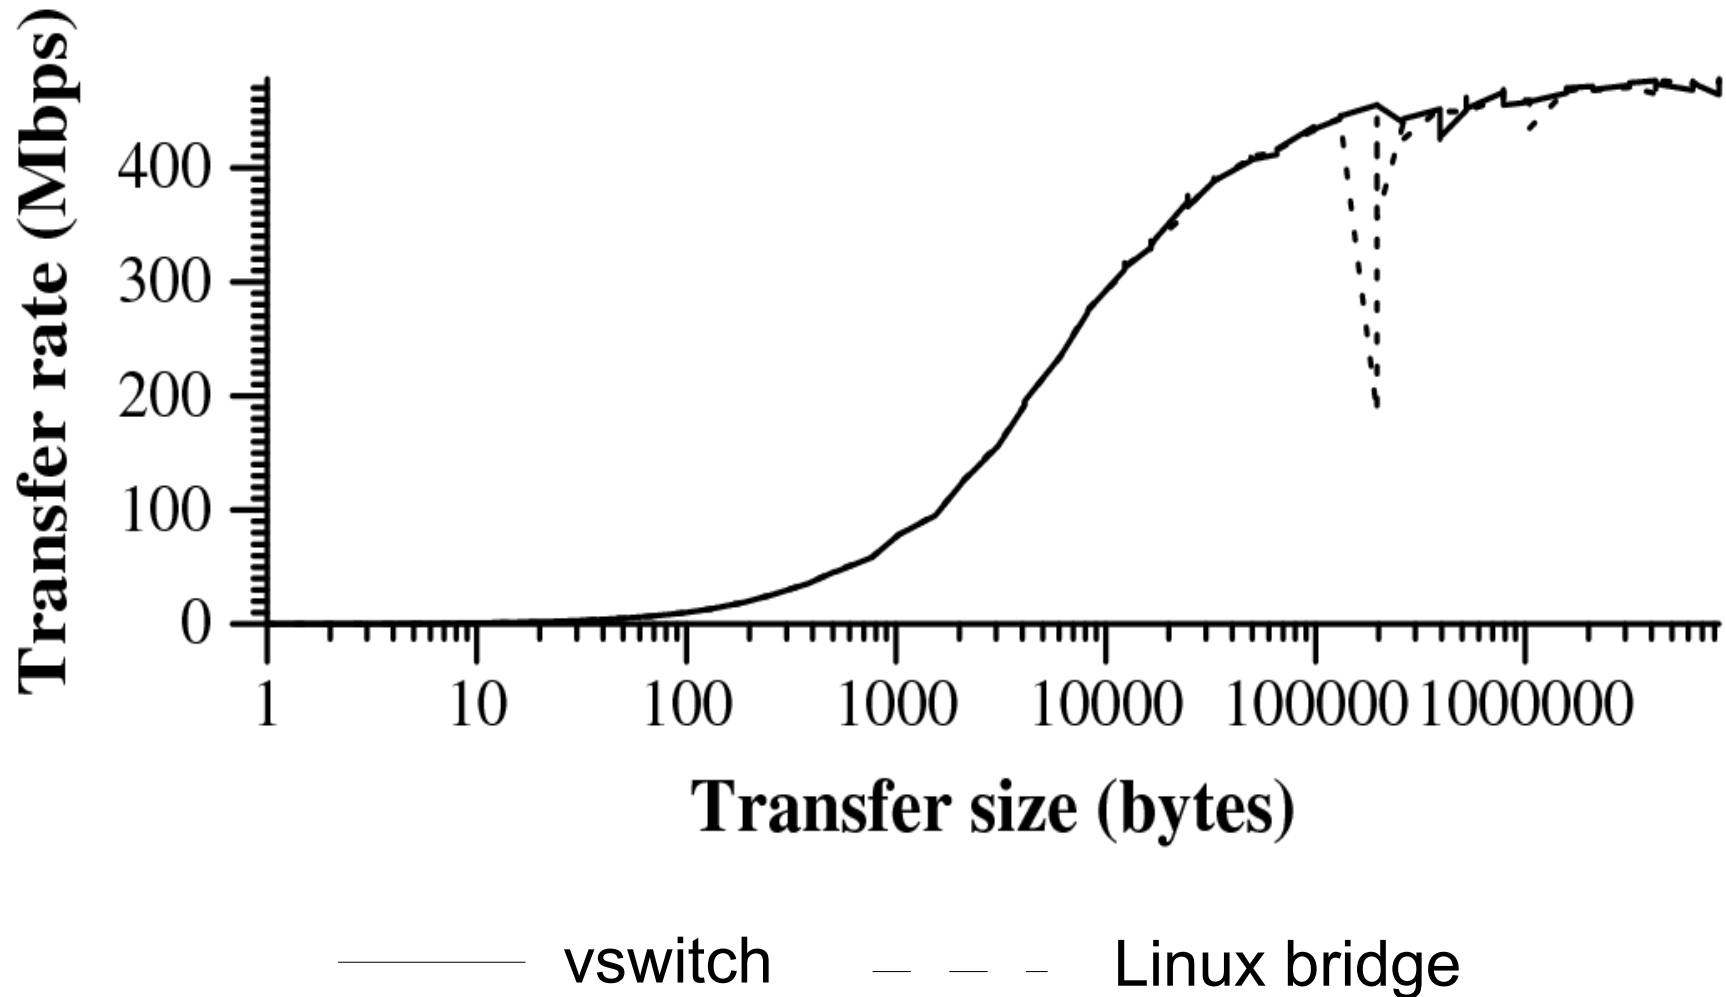
## Bandwidth

Fast Path: > 1 Gbps  
ovs-vswitchd: 100 Mbps  
Controller: 10 Mbps

## Latency

Fast Path: < 1  $\mu$ s  
ovs-vswitchd: < 1 ms  
Controller: ms

# Open vSwitch is Fast



# Hardware Acceleration

- Inevitable
- Netronome: right approach
- VN-Tag: wrong approach
- VEPA: powerless

# Future Directions

- Physical switches
- Upstream kernel integration
- Anything\*



Questions?