AN OVERVIEW OF NETWORK VIRTUALIZATION

January 14, 2009

Mosharaf Chowdhury

What is Virtualization?

- Transparent abstraction of computing platform and resources
  - Multiple logical interpretations of the physical characteristics
- Additional level of indirection
  - Indirect access to hardware
  - Hides implementation details
  - Controls mappings from abstract view to implementation

> "Any problem in computer science can be solved with another layer of indirection"
> - *David Wheeler*

Example: Virtual Machines

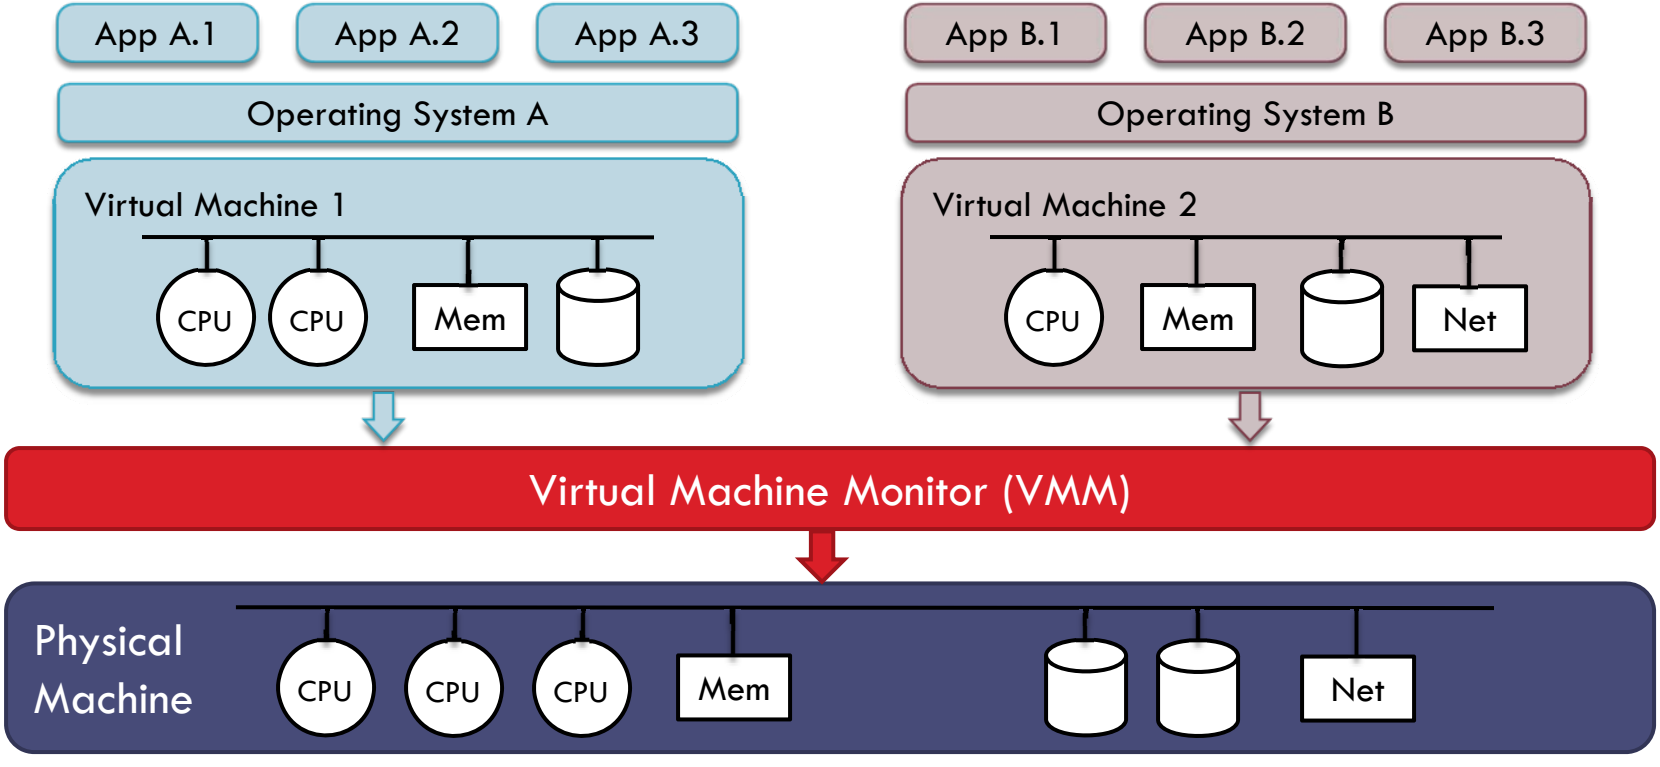

The Good,

- Virtualization adds flexibility, allows heterogeneity, and improves manageability of the computing infrastructure

- Lower cost of ownership
  - Fewer computing resources
  - More resilient and simpler to manage

The Bad,

- Performance penalty
  - Overhead due to the indirection layer
- Too much abstraction
  - Hidden details

And the Ugly?

Historical Perspective

Network Virtualization for *Dummies*

- Making a physical network appear as multiple logical ones

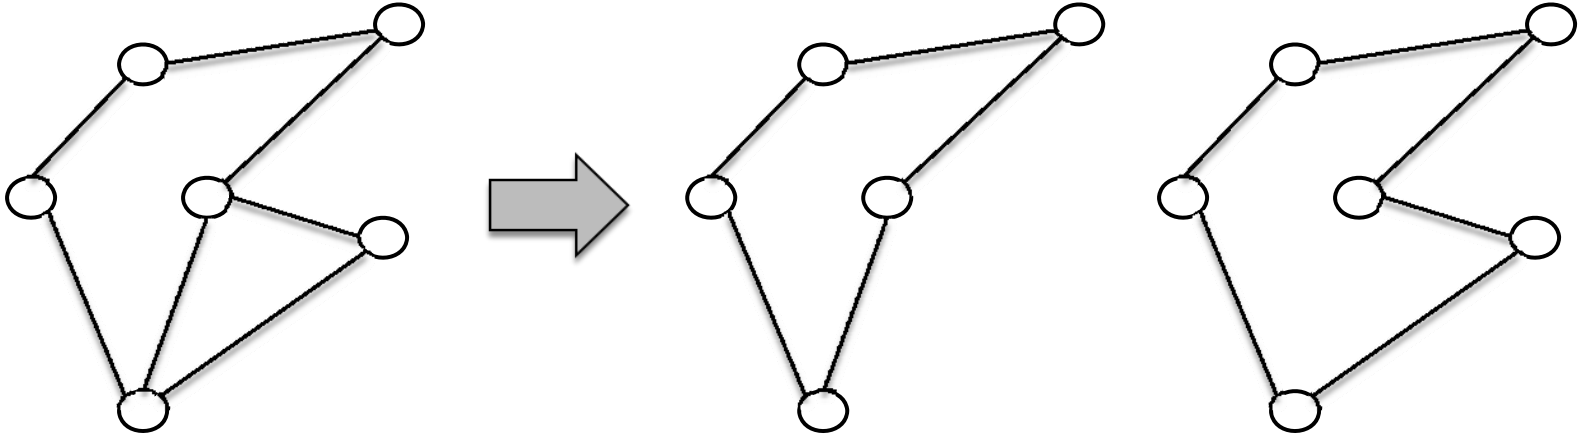

Physical Network

Virtualized Network - 1

Virtualized Network - 2

Related Concepts

1. Virtual Local Area Networks (VLAN)
2. Virtual Private Networks (VPN)
3. Active and Programmable Networks
4. Overlay Networks

Virtual Local Area Networks (VLAN)

- Group of logically networked hosts
  - Single broadcast domain
- Advantages
  - Ease of network administration and management
  - Elevated levels of trust, security, and isolation

Virtual Private Networks (VPN)

- Virtual network connecting distributed sites
  - Works over public communication networks

- VPN classification (based on the protocol used in the VPN data plane)
  1. Layer 3 VPN
  2. Layer 2 VPN
  3. Layer 1 VPN

Major VPN Classification

- L3VPN
  - CE-based VPN using tunneling
    - Network is unaware
  - PE-based VPN
    - States in the network
- L2VPN
  - Agnostic to higher level protocols
  - No control plane
- L1VPN
  - Rise due to advances in optical networking technologies
  - Independent Layer 1 resource view, separate policies, and complete isolation

Active and Programmable Networks

- Customized network functionalities
- Active Networks
  - Customization of network services at packet transport granularity
  - More flexibility with increased security risk
- Programmable Networks
  - Defined programming interfaces
  - More secured than active networks
- Requires changes to existing hardware

Overlay Networks

- Logical network on top of another existing network
  - Internet was an overlay on the telecommunications network
- Application layer virtual networks
- Extravagantly used in the Internet
  - Ensuring performance and availability of Internet routing
  - Enabling Multicasting
  - Providing QoS guarantees
- P2P networks are overlays

Downsides of Overlay Networks

- Largely used as narrow fixes for specific problems
  - No holistic view
- Most overlays are designed in the application layer
  - Cannot support radically different concepts

Anderson et al.

Network Virtualization Environment

What is Network Virtualization?

- Transparent abstraction of networking platform and resources
  - Multiple logical interpretations of the physical characteristics
- Additional level of indirection
  - Indirect access to network resources
- Resource partitioning and isolation
  - Physical and logical
  - Dynamic provisioning and configuration

Why Virtualize the Network?

- Internet is *almost* ossified
  - Lots of band-aids and makeshift solutions (e.g., overlays)
  - A new architecture (aka clean-slate) is needed
- Hard to come up with a *one-size-fits-all* architecture
  - Almost impossible to predict what future might unleash
- Why not create an *all-sizes-fit-into-one* instead!
  - Open and expandable architecture
  - Coexistence of heterogeneous architectures

Network Virtualization Environment (NVE)

- Virtual Network
- Business Model
- Principles
- Architecture
- Design Goals

What is a Virtual Network (VN)?

- A collection of virtual nodes and virtual links forming a virtual topology
  - Subset of physical topology
  - Basic entity of the NVE
- A virtual node is hosted on a particular physical node
  - Multiple virtual nodes can coexist
- A virtual link spans over a physical path
  - Includes a portion of the underlying physical resources

Business Model

Players

- Infrastructure Providers (*InP*)
  - Manage underlying physical networks
- Service Providers (*SP*)
  - Create and manage virtual networks
  - Deploy customized end-to-end services
- End Users
  - Buy and use services from different service providers
- Brokers
  - Mediators/Arbiters

Relationships

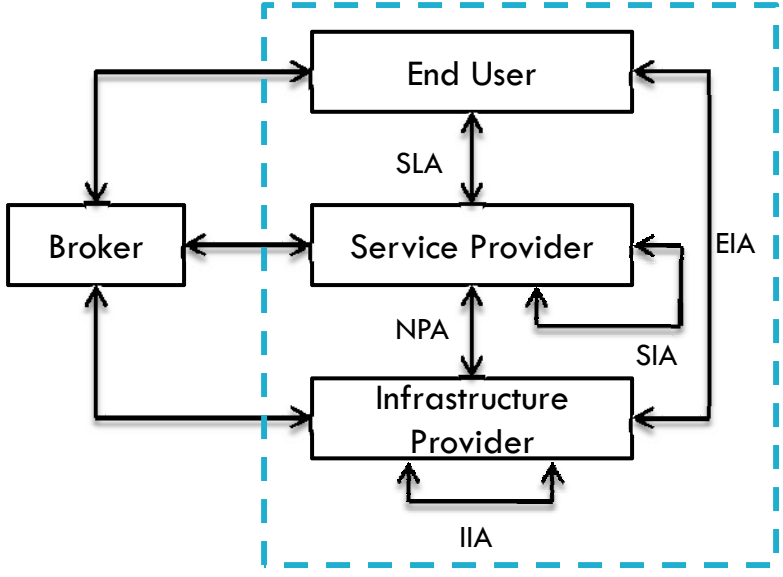

Principles

- *Coexistence* of multiple heterogeneous virtual networks
  - Introduces diversity
- *Recursion* of virtual networks
  - Opens the door for network virtualization economics
- *Inheritance* of architectural attributes
  - Promotes value-addition
- *Revisitation* of virtual nodes
  - Simplifies network operation and management

Architecture

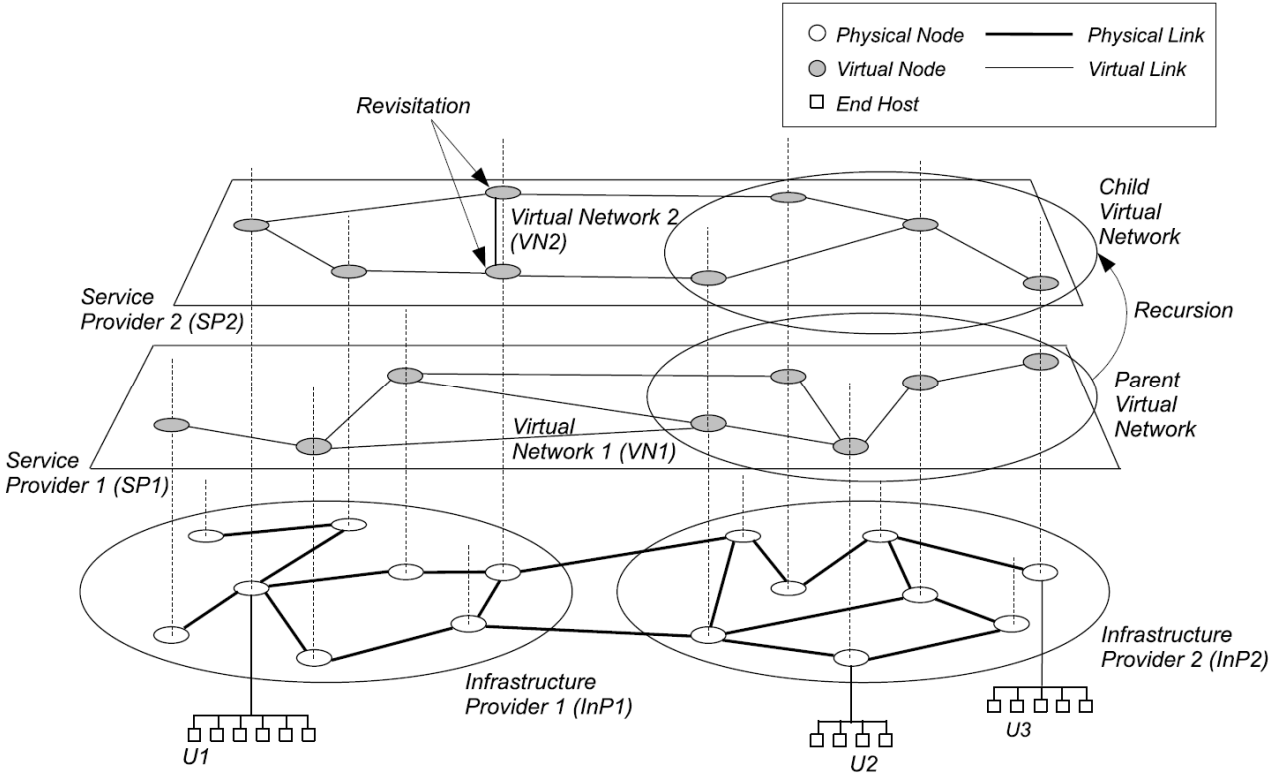

Hierarchy of Roles

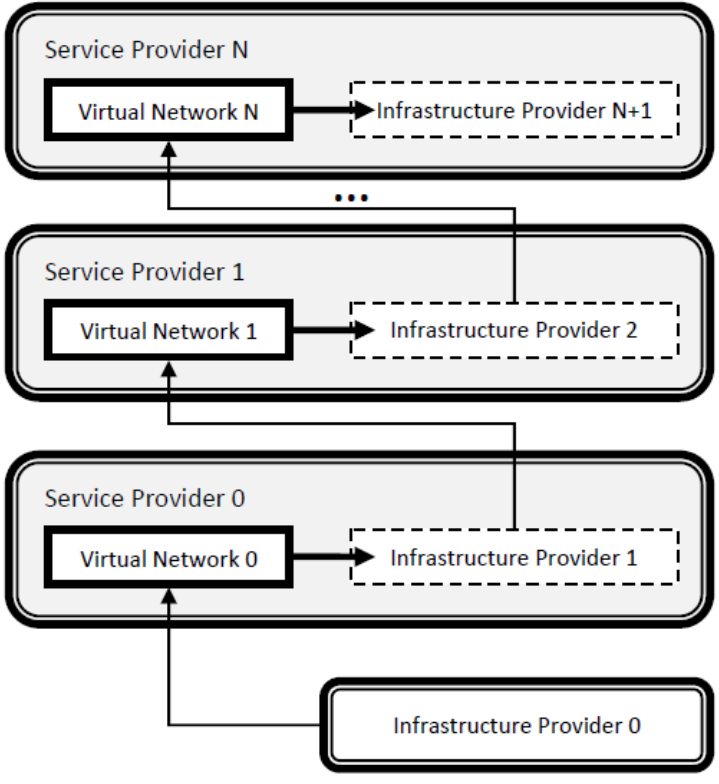

Design Goals

- Flexibility
  - Service providers can choose
    - arbitrary network topology,
    - routing and forwarding functionalities,
    - customized control and data planes
  - No need for co-ordination with others
    - IPv6 fiasco should never happen again

Design Goals (Cont.)

- Manageability
  - Clear separation of policy from mechanism
  - Defined *accountability* of infrastructure and service providers
  - Modular management
- Scalability
  - Maximize the number of co-existing virtual networks
  - Increase resource utilization and amortize CAPEX and OPEX

Design Goals (Cont.)

- Isolation
  - Complete isolation between virtual networks
    - *Logical* and *resource*
  - Isolate faults and misconfigurations

- Stability and Convergence
  - Instability due to
    - Errors and misconfigurations
    - Instability in InP algorithms
  - Quick convergence to stable state

Design Goals (Cont.)

- Programmability
  - Of network elements (e.g., routers)
  - Answer *"How much"* and *"how"*
  - Easy and effective without being vulnerable to threats

- Heterogeneity
  - Networking technologies
    - Optical, sensor, wireless etc.
  - Virtual networks
  - End user devices

Design Goals (Cont.)

- Experimental and Deployment Facility
  - PlanetLab, GENI, VINI etc.
  - Directly deploy services in real world from the testing phase
- Legacy Support
  - Consider the existing Internet as a member of the collection of multiple virtual Internets
  - Very *important* to keep all concerned parties satisfied

What is Network Virtualization? (Revisited)

Network virtualization is a networking environment that allows multiple service providers to dynamically compose multiple heterogeneous virtual networks that coexist together in isolation from each other, and to deploy customized end-to-end services on-the-fly as well as manage them on those virtual networks for the end-users by effectively sharing and utilizing underlying network resources leased from multiple infrastructure providers.

Basic Concepts

Principles

- Concurrence
- Recursion
- Inheritance
- Revisitation

Design Goals

- Flexibility
- Manageability
- Scalability
- Isolation
- Stability and Convergence
- Programmability
- Heterogeneity
- Experimental and Deployment Facility
- Legacy Support

Existing Projects

Classification

- *Networking technology*
  - Targeted technology for virtualization
- *Layer of virtualization*
  - Particular layer in the network stack where virtualization is introduced
- *Architectural domain*
  - Specific problem domain that virtualization addresses
- *Level of virtualization*
  - Granularity at which virtualization is realized

Existing Projects

Project	Architectural Domain	Networking Technology	Layer of Virtualization	Level of Virtualization
VNRMS	Virtual network management	ATM/IP		Node/Link
Tempest	Enabling alternate control architectures	ATM	Link	
NetScript	Dynamic composition of services	IP	Network	Node
Genesis	Spawning virtual network architectures		Network	Node/Link

Existing Projects (Cont.)

Project	Architectural Domain	Networking Technology	Layer of Virtualization	Level of Virtualization
VNET	Virtual machine Grid computing		Link	Node
VIOLIN	Deploying on-demand value-added services on IP overlays	IP	Application	Node
X-Bone	Automating deployment of IP overlays	IP	Application	Node/Link
PlanetLab	Deploy and manage overlay-based testbeds	IP	Application	Node
UCLP	Dynamic provisioning and reconfiguration of lightpaths	SONET	Physical	Link

Existing Projects (Cont.)

Project	Architectural Domain	Networking Technology	Layer of Virtualization	Level of Virtualization
AGAVE	End-to-end QoS-aware service provisioning	IP	Network	
GENI	Creating customized virtual network testbeds	Heterogeneous		
VINI	Evaluating protocols and services in a realistic environment		Link	
CABO	Deploying value-added end-to-end services on shared infrastructure	Heterogeneous		Full

Insights

- Shift toward a holistic and generalized network virtualization environment that is
  - Completely virtualized
    - Virtualization of all network elements
  - Highly customizable
    - Virtualization at lower layers of the network stack
  - Technology agnostic
    - Support for heterogeneity

Future Directions

Future Directions

- *Instantiation*
  - Concerned with issues related to successful creation of virtual networks
- *Logistics*
  - Deals with operations of virtual networks and virtual components
- *Management*
  - Manages co-existing virtual networks
- *Interactions*
  - Handles interactions between players in the *network virtualization environment*

Instantiation

- Interfacing
  - Request format for a virtual network
  - Make programmability of the network elements available

- Signaling and Bootstrapping
  - Request for a virtual network
  - Bootstrap the customized network onto the physical network elements
  - Use a *separate* network (e.g. Genesis) or *out-of-band* communication mechanism

Instantiation (Cont.)

- Admission Control and Usage Policing
  - Prohibit overbooking of network resources through *admission control*
  - *Distributed rate limiting*
  - Applied on *complete* virtual networks

- Virtual Network Embedding
  - Within single InP domain and across InP boundaries
  - Known to be a *NP-Hard* problem
  - *Heuristic*-based solutions
  - Two versions of the problem
    - *Offline,* where all the requests are known in advance
    - *Online,* where requests arrive dynamically

Operation

- Virtual Nodes
  - Multiple logical routers inside one physical router
  - Issues of interest
    - Performance
    - Scalability
    - Migration (e.g. VROOM)

- Virtual Links
  - Similar to tunnels in VPNs
  - Cross-InP virtual links
  - *Link scheduling* (e.g. DaVinci)

Operation (Cont.)

- Naming and Addressing
  - Generic naming and addressing for all the virtual networks
  - *Überhoming*
    - *Allows end users in a network virtualization environment to simultaneously connect to multiple VNs through multiple InPs using heterogeneous technologies to access different services.*
  - *Identity-based* routing

Operation (Cont.)

- Resource Scheduling
  - Maximize *degree of co-existence*
  - Schedule CPU, Disk and Link b/w
- Topology Discovery
  - Within an InP administrative domain and across InP boundaries
  - *Event-based* and *periodic* topology discovery (e.g., UCLP)
  - Separate discovery plane (e.g., CABO)

Management

- VN Configuration and Monitoring
  - Enable virtualization from the level of NOCs to lower level network elements
    - Concept of MIBlets (e.g., VNRMS)
- Management Frameworks
  - Generic management framework for the service providers
  - Interface between multiple management paradigms
  - Draw clear line between the management responsibilities of the InPs and the SPs

Management (Cont.)

- Mobility Management
  - Geographic mobility of the end user devices
  - Mobility of the virtual routers through migration techniques
  - Logical mobility of the end users in different virtual networks
- Failure Handling
  - Isolate failures
  - Prevent *cascading* failures

Management (Cont.)

- Self-*/Autonomic Properties
  - *Self-configuration* and *self-optimization* for maximizing virtual resource utilization
  - *Self-protection* and *self-healing* to survive malicious attacks

Interactions

- Networking Technology Agnostic Virtualization
  - Virtualization *on* and *across* optical, wireless, and sensor technology among other technologies
  - Transparently create end-to-end virtual networks across heterogeneous technologies
- Inter-VN Communication
  - Sharing of resources and information between multiple virtual networks
  - Creating *compound* virtual networks

Interactions (Cont.)

- Tussles in the NVE
  - Between multiple InPs
  - Between InPs and SPs
- Network Virtualization Economics
  - Trade node resources (e.g. processing power, memory) in addition to bandwidth
  - Centralized, decentralized and hybrid markets

Major Ongoing Projects

Project	Originated In	Link
4WARD	Europe	http://www.4ward-project.eu/
AKARI	Japan	http://akari-project.nict.go.jp/
CABO	USA	http://www.cs.princeton.edu/~jrex/virtual.html
Clean Slate	USA	http://cleanslate.stanford.edu/
GENI	USA	http://www.geni.net/
NouVeau	Canada	http://netlab.cs.uwaterloo.ca/virtual/
PlanetLab	USA	http://www.planet-lab.org/
Trilogy	Europe	http://www.trilogy-project.org/
UCLP	Canada	http://www.uclp.ca/
VINI	USA	http://www.vini-veritas.net/

Reference

- N.M. Mosharaf Kabir Chowdhury, Raouf Boutaba, "A Survey of Network Virtualization", *University of Waterloo Technical Report CS-2008-25, Oct. 2008.*

Questions ?

Mosharaf Chowdhury
http:// www.mosharaf.com/