CS144
An Introduction to Computer Networks

Abstractions and Virtualization

Tags, Tunnels and Translation

Nick McKeown
Professor of Electrical Engineering
and Computer Science, Stanford University

The term “Virtual” is (over) used a lot…

- Virtual LANs (VLAN)
- Virtual Private Network (VPN)
- Network Virtualization (used by cloud providers)
- Network Function Virtualization (NFV)

Learning goals of this class

- To learn how tags, tunnels and translation can be used to provide new abstractions in a network.
- To learn about the match + action abstraction
- To learn about three examples:
  Virtual LANs (VLANs), VPNs, and NATs.
- To learn what network virtualization is.
- To learn how overlay network virtualization works.
- To learn what network function virtualization (NFV) is.

What do we mean by an abstraction?

Example: IP datagram delivery

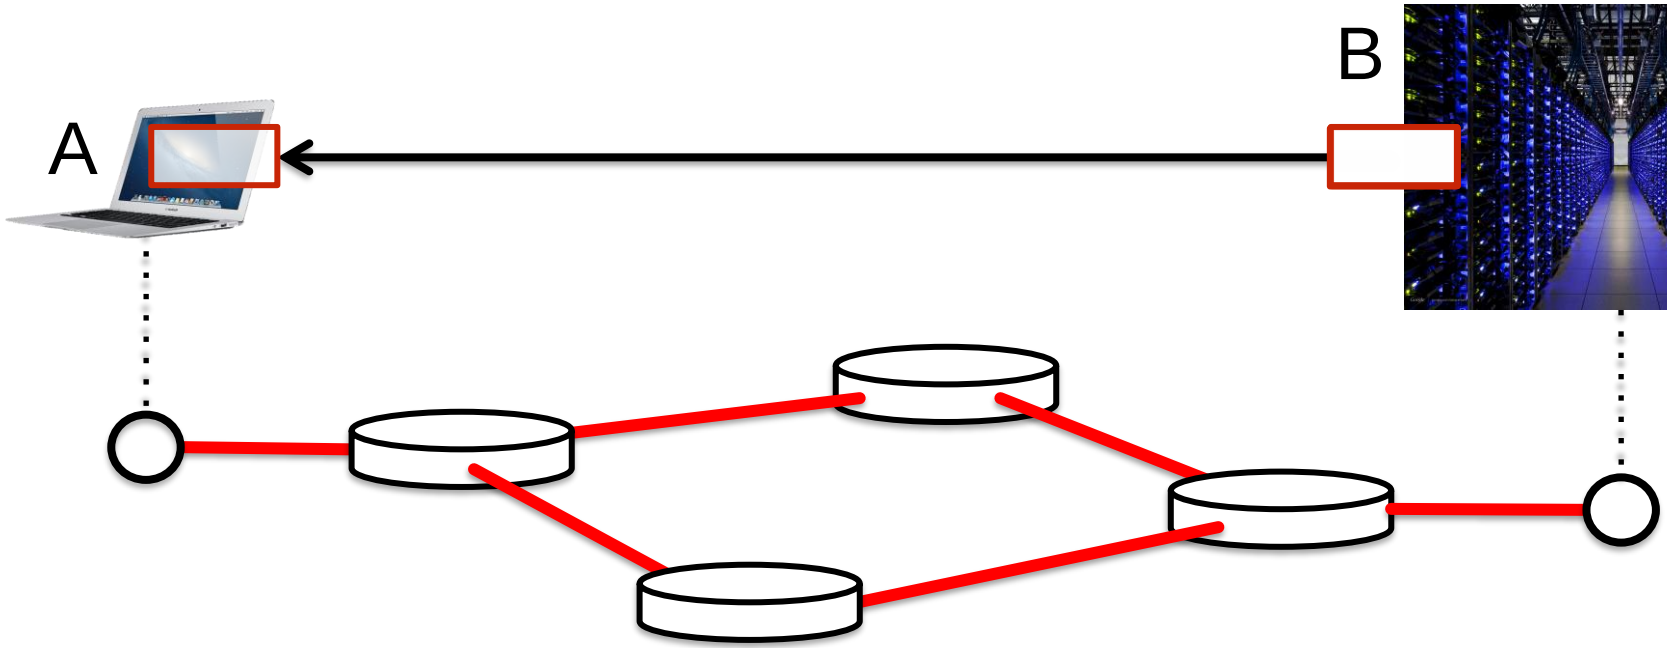

Example: IP datagram delivery

A

B

Abstraction: Packets with IP DA = B are delivered to B (with best effort)

The details of how it is accomplished are hidden from us.

IP Forwarding Abstraction

B

Match	Action
DA==B ?	Deliver to B

Abstraction

Firewall Abstraction

B

Match **Action**

SA==X ?

Dr

Abstraction

In practice, a firewall rule might match on TCP ports as well.
e.g. If SA==X && dst port==80, then Allow

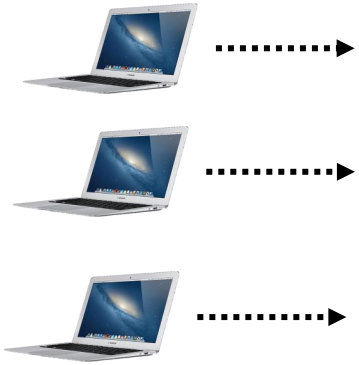

Virtual LAN Abstraction

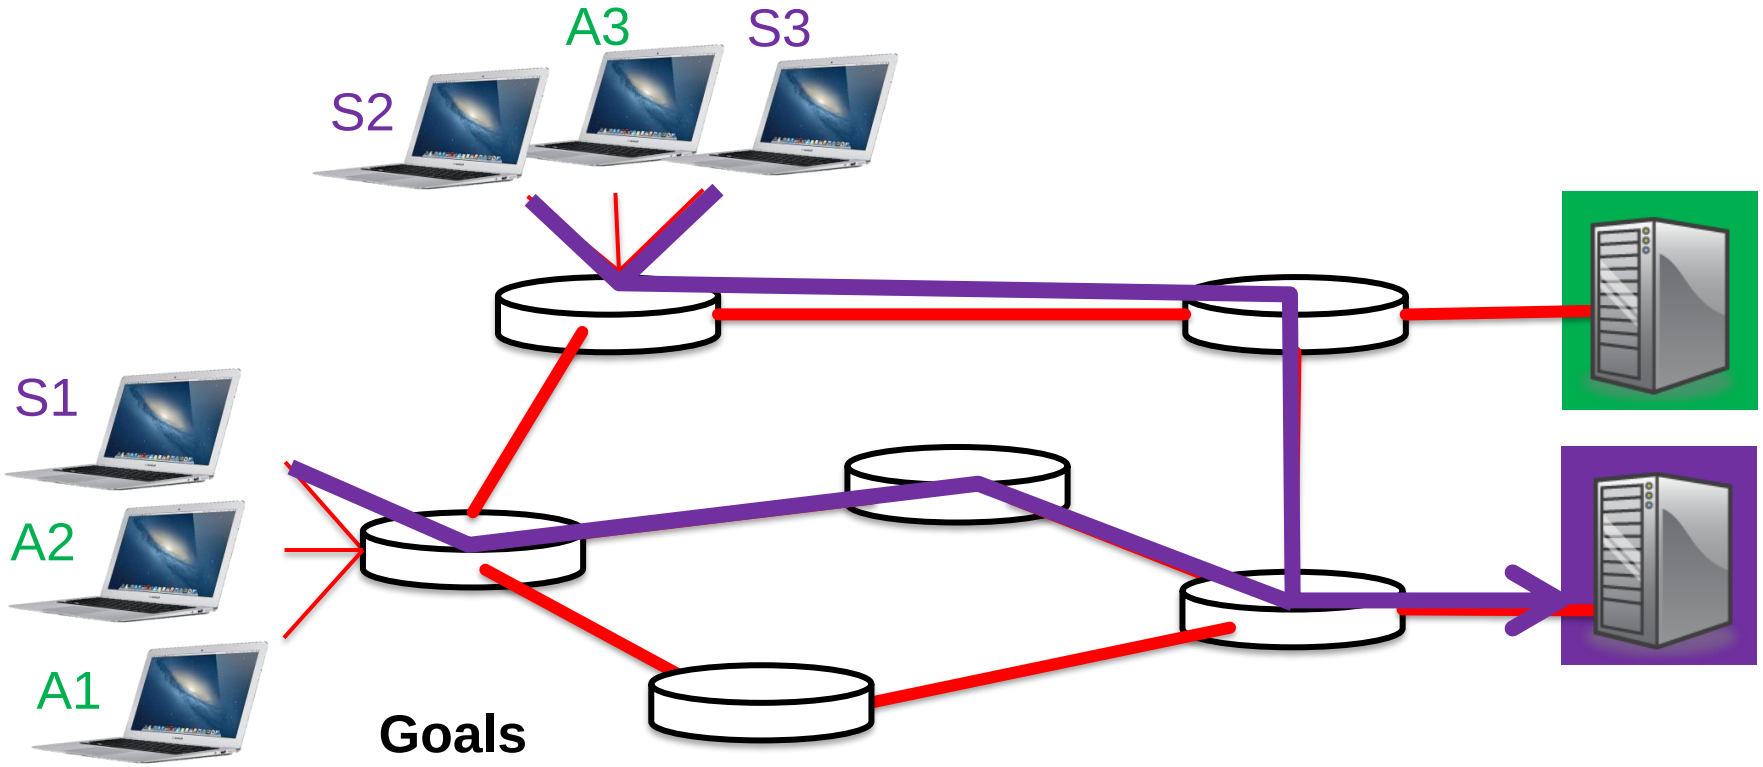

Goals

Packets on VLAN A never delivered to hosts on VLAN S

Packets in each VLAN follow their own spanning tree

Virtual LAN Abstraction

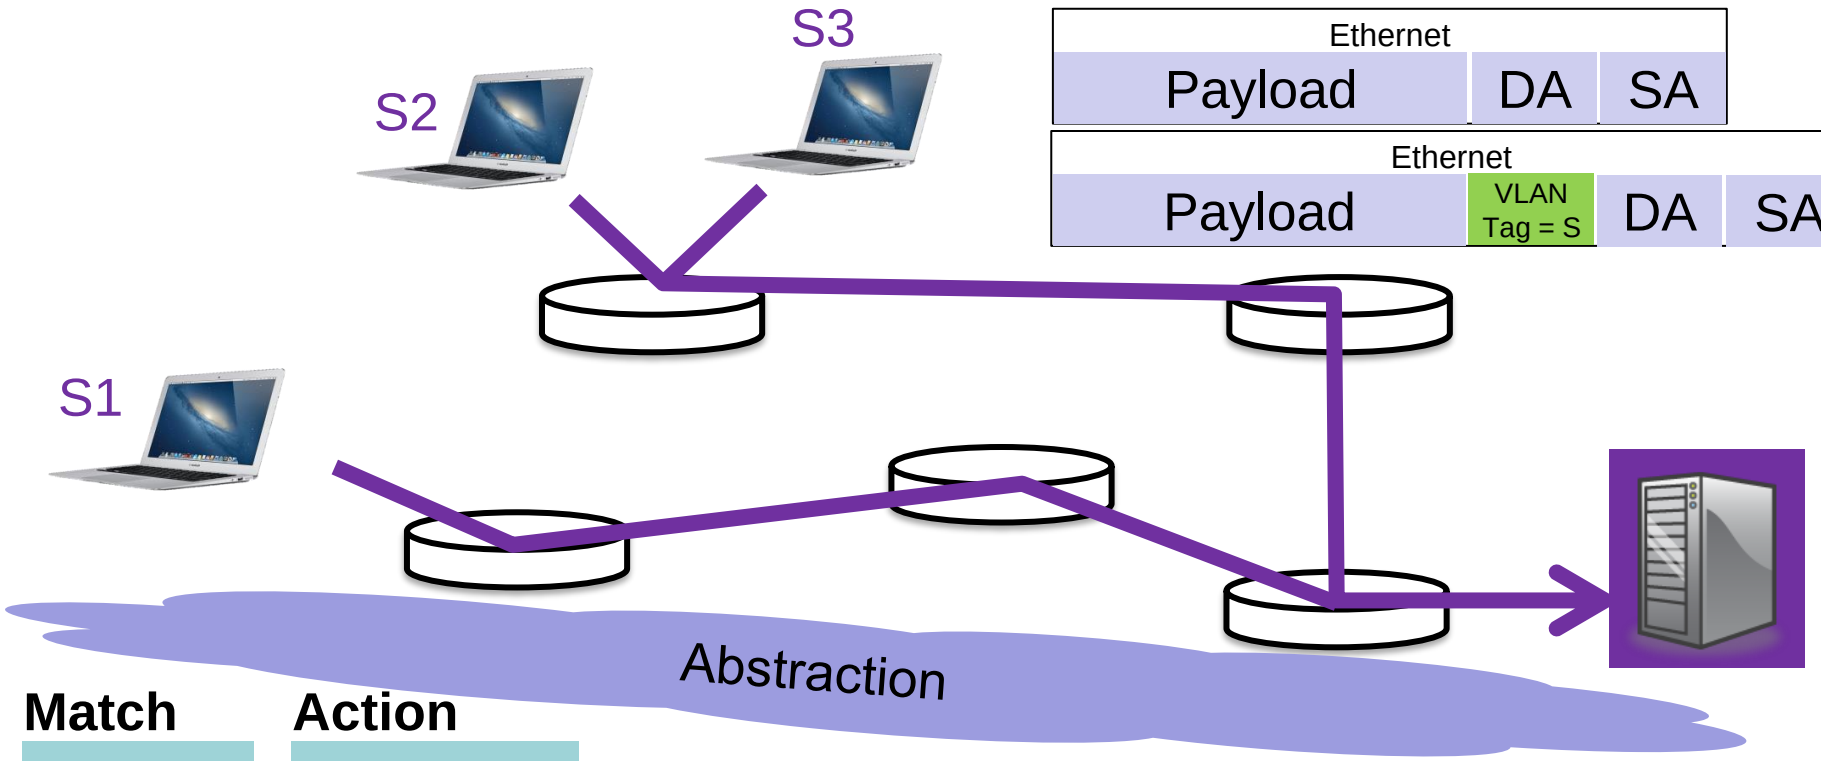

Match	Action
VLAN Tag==S ?	Forward on VLAN S

Example: Virtual Private Network (VPN)

Remote client “appears to be” on corporate network

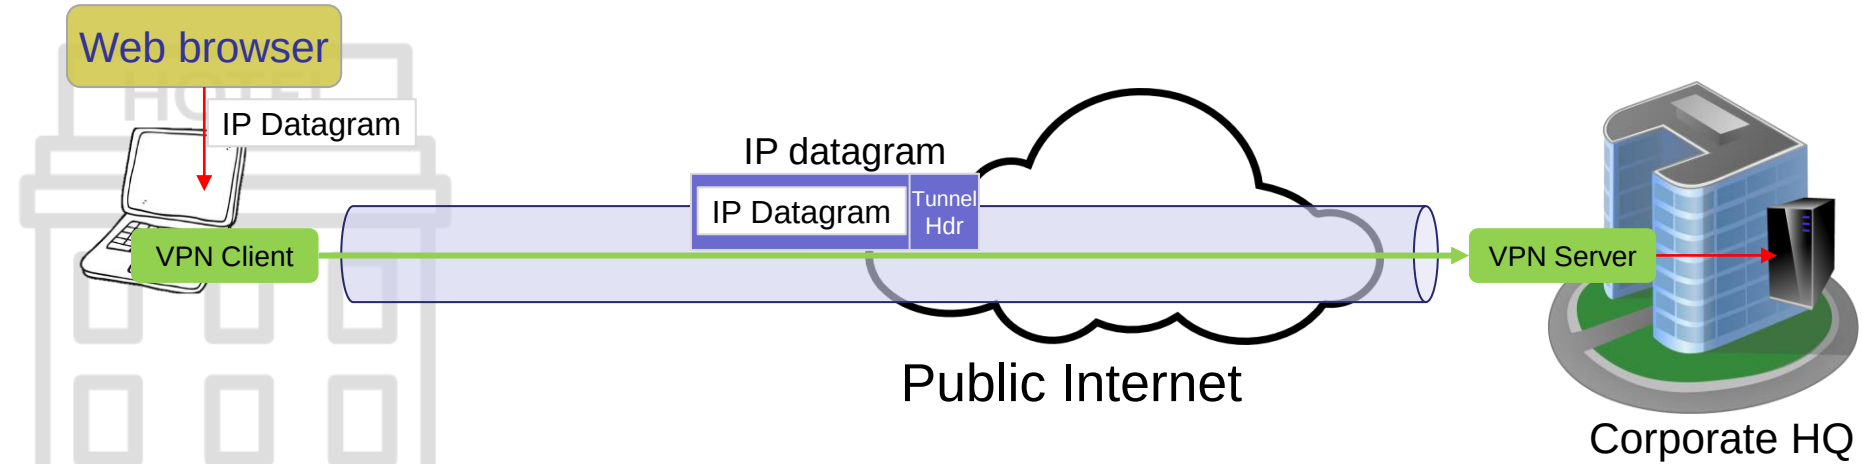

Example: Virtual Private Network (VPN)

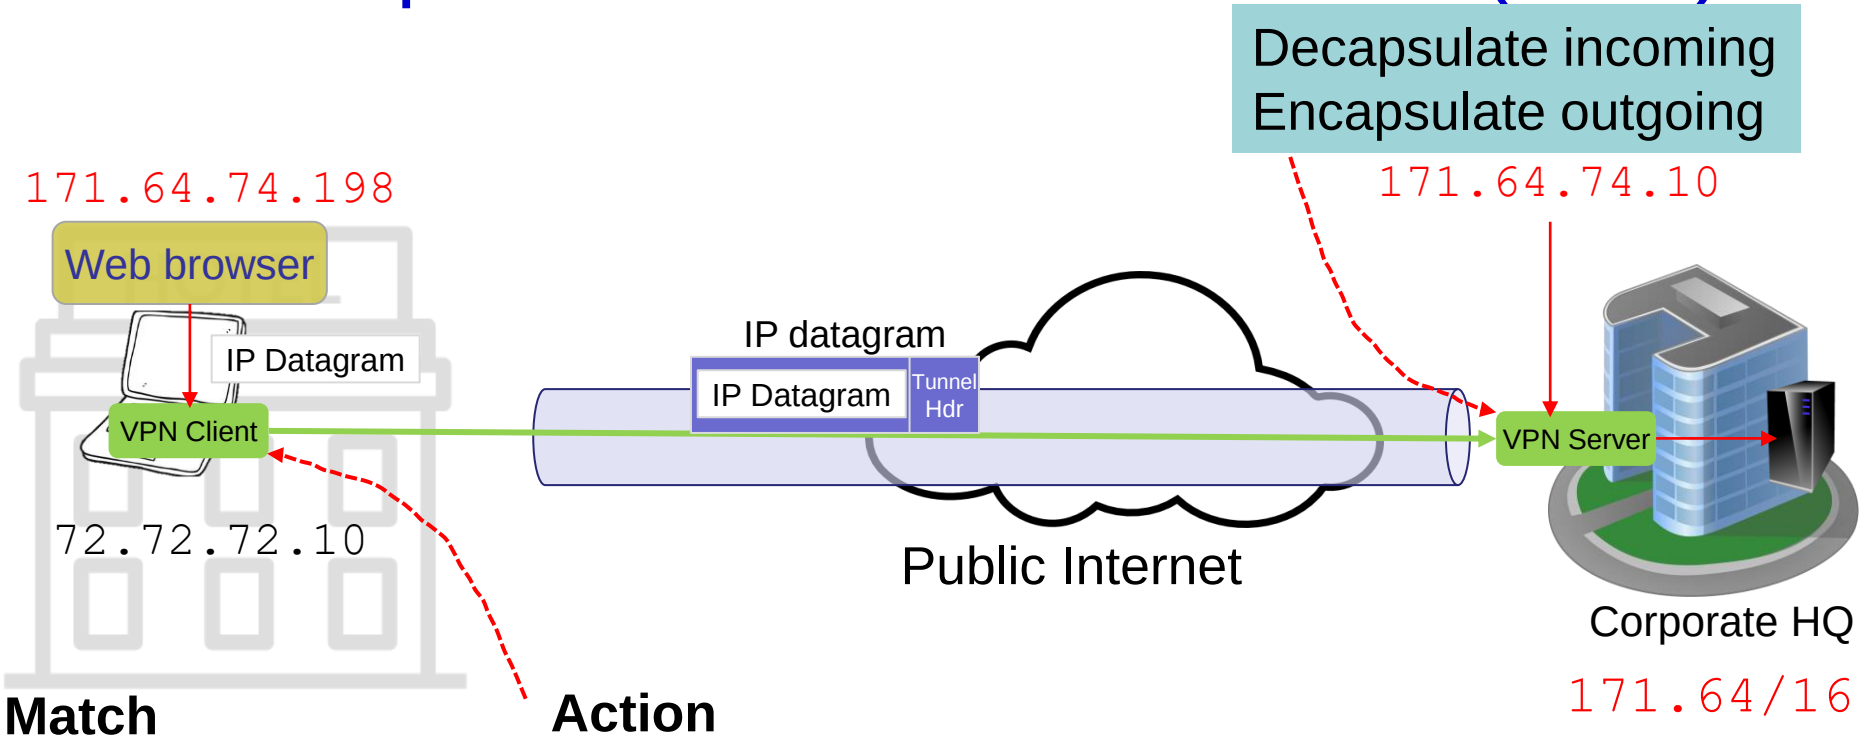

Match

IP SA == 171.64.74.198

Action

Encapsulate in IP DA=171.64.74.10, SA=72.72.72.10
Forward to 171.64.74.10

Example: Network Address Translation (NAT)

Multiple clients share a common IP address

Q: Why does NAT use translation instead of tags or tunnels?

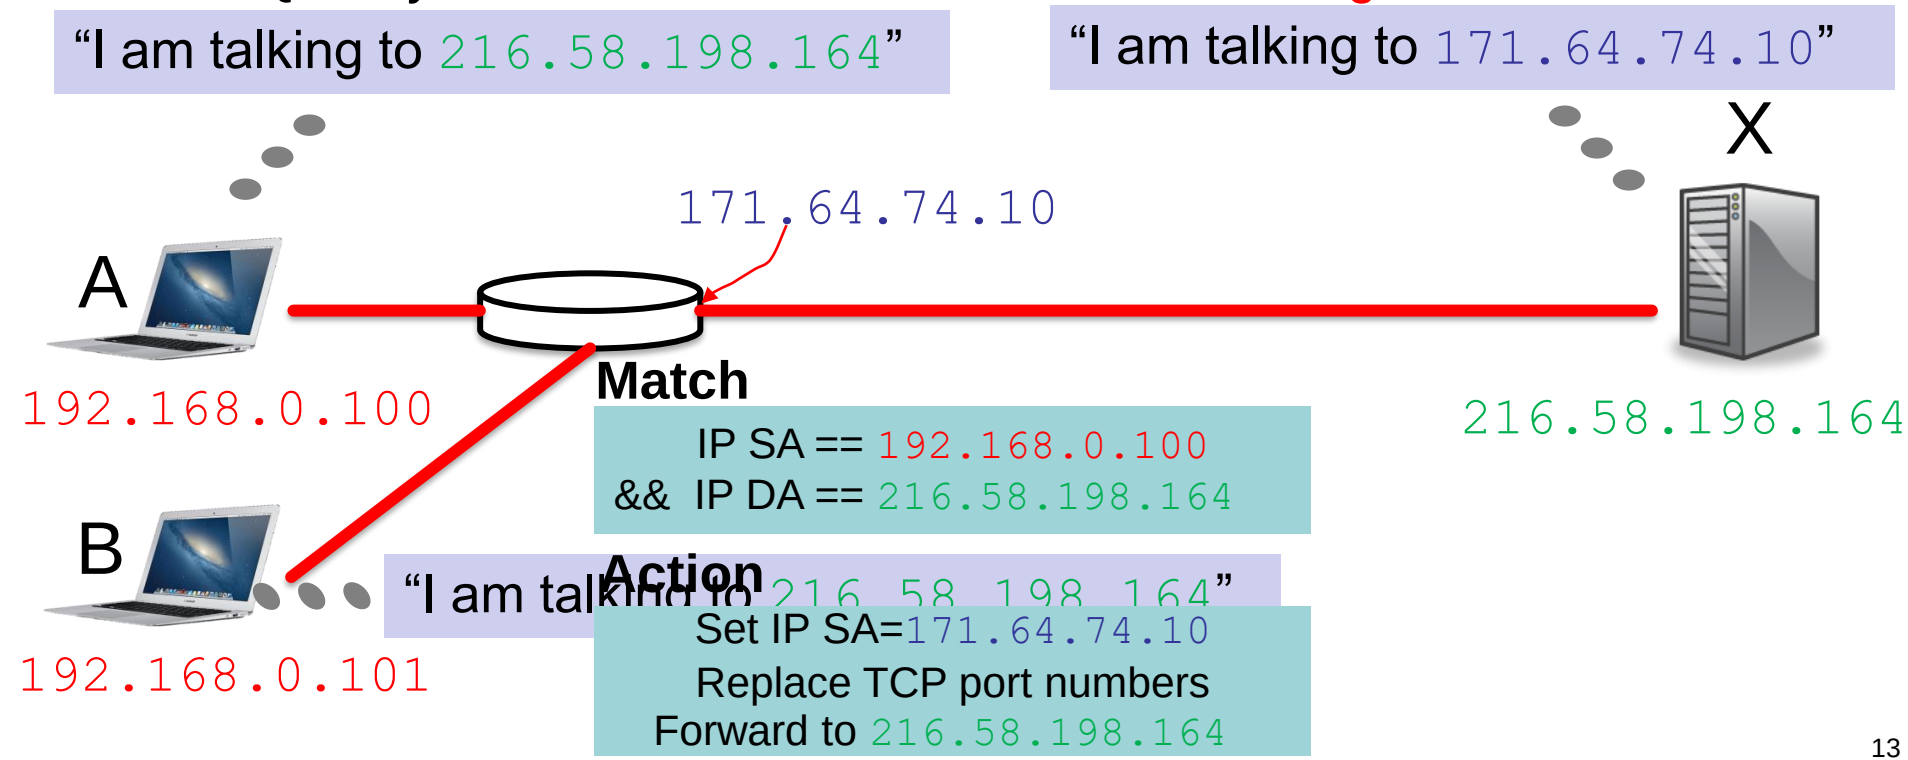

“Modularity based on abstraction is the way things are done!”

Barbara Liskov (MIT)
Turing Award Lecture 2009

Learning goals of this class


To learn how tags, tunnels and translation can be used to provide new abstractions in a network.


To learn about the match + action abstraction


To learn about three examples:
Virtual LANs (VLANs), VPNs, and NATs.

- To learn what network virtualization is.
- To learn how overlay network virtualization works.
- To learn what network function virtualization (NFV) is.

Network Virtualization

Abstractions in computer systems

Virtual memory
Abstract illusion of infinite, private physical memory

File system
Uniform illusion of read/write data store.

Virtual Machine User application cannot tell if it is running on a physical or virtual machine.

…

Virtual Network: The abstraction

The abstraction (or illusion) of a physical network: The user, application (and possibly the network administrator too) cannot tell if the network is physical or virtual.

Virtual Network: The abstraction

A set of VMs operating as if connected to the same physical network.

1. Typically belonging to the same tenant.
2. VMs communicate with each other using their own address space.
3. Virtual networks are isolated from each other: They cannot communicate, except through a gateway.
4. VMs can migrate to a different server without changing IP address.
5. A virtual network has a SLO expressed as a desired quality of service (e.g. data rate, reliability, latency)
6. A VM can operate as if on the tenant's home network.
7. Used for containers too

Virtualized Data Center

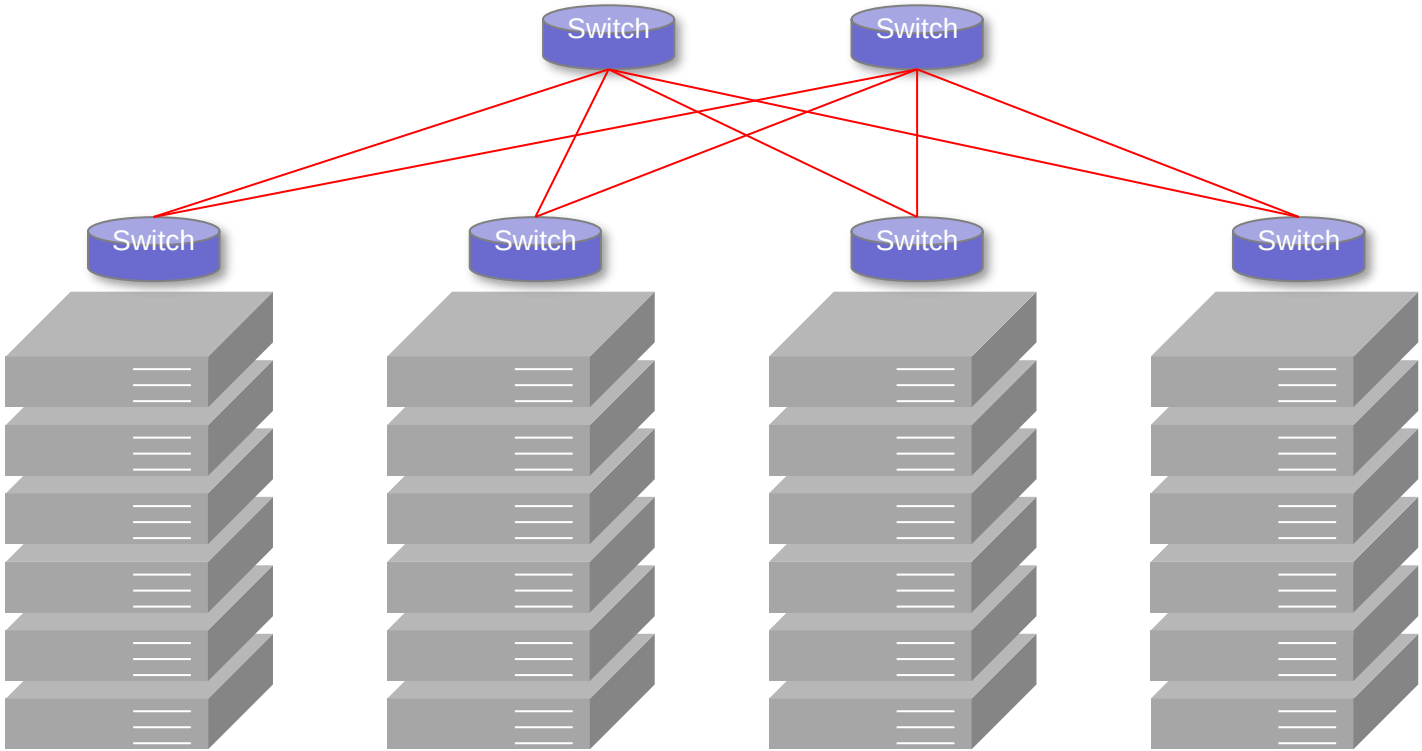

Abstraction for tenant VMs

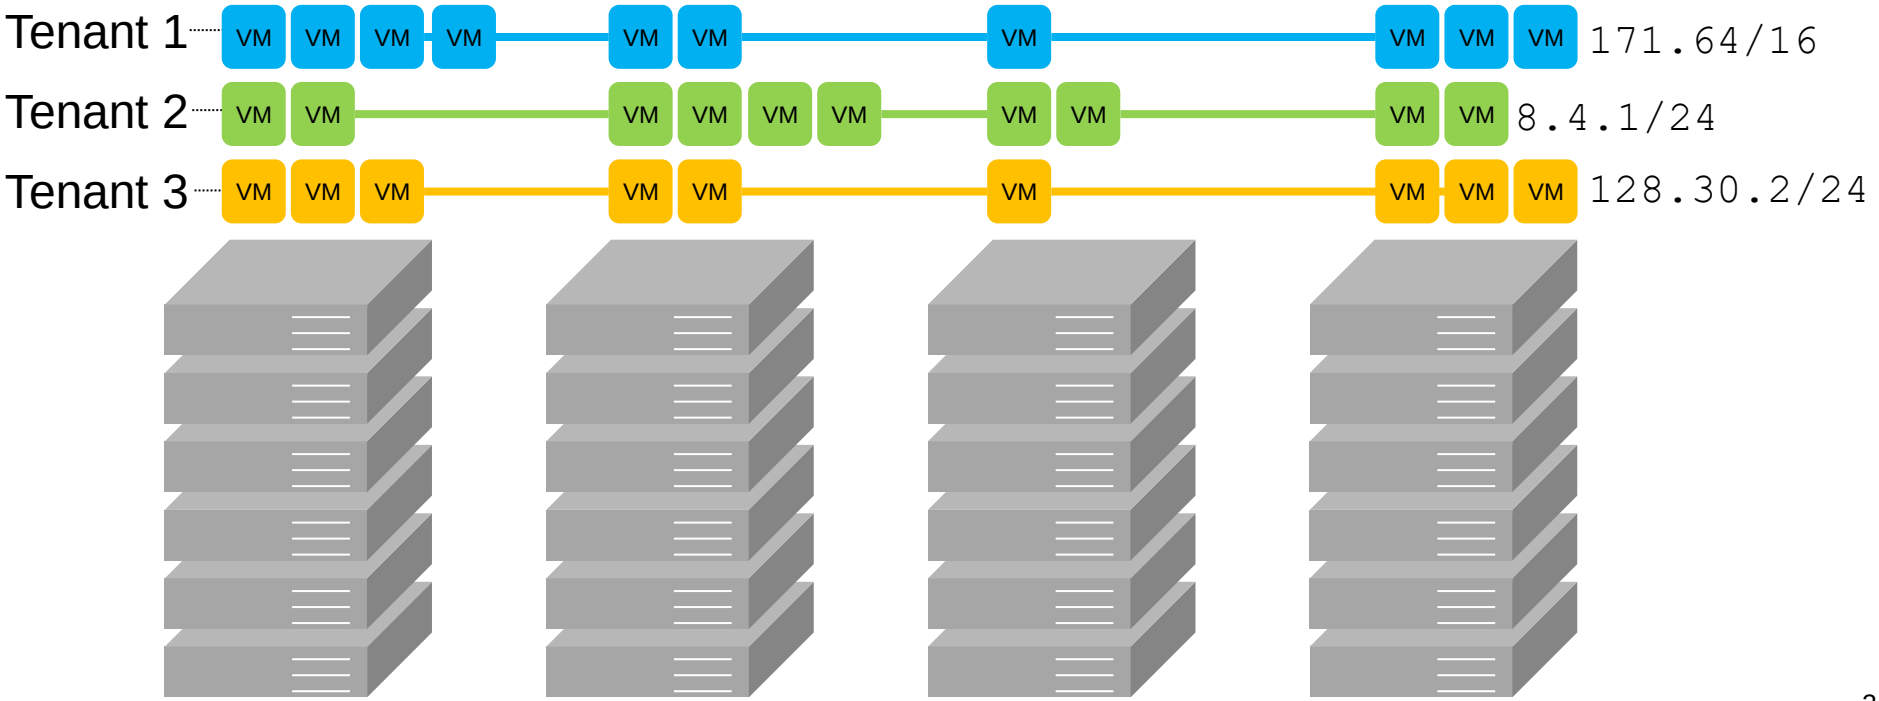

VMs using their own IP addresses

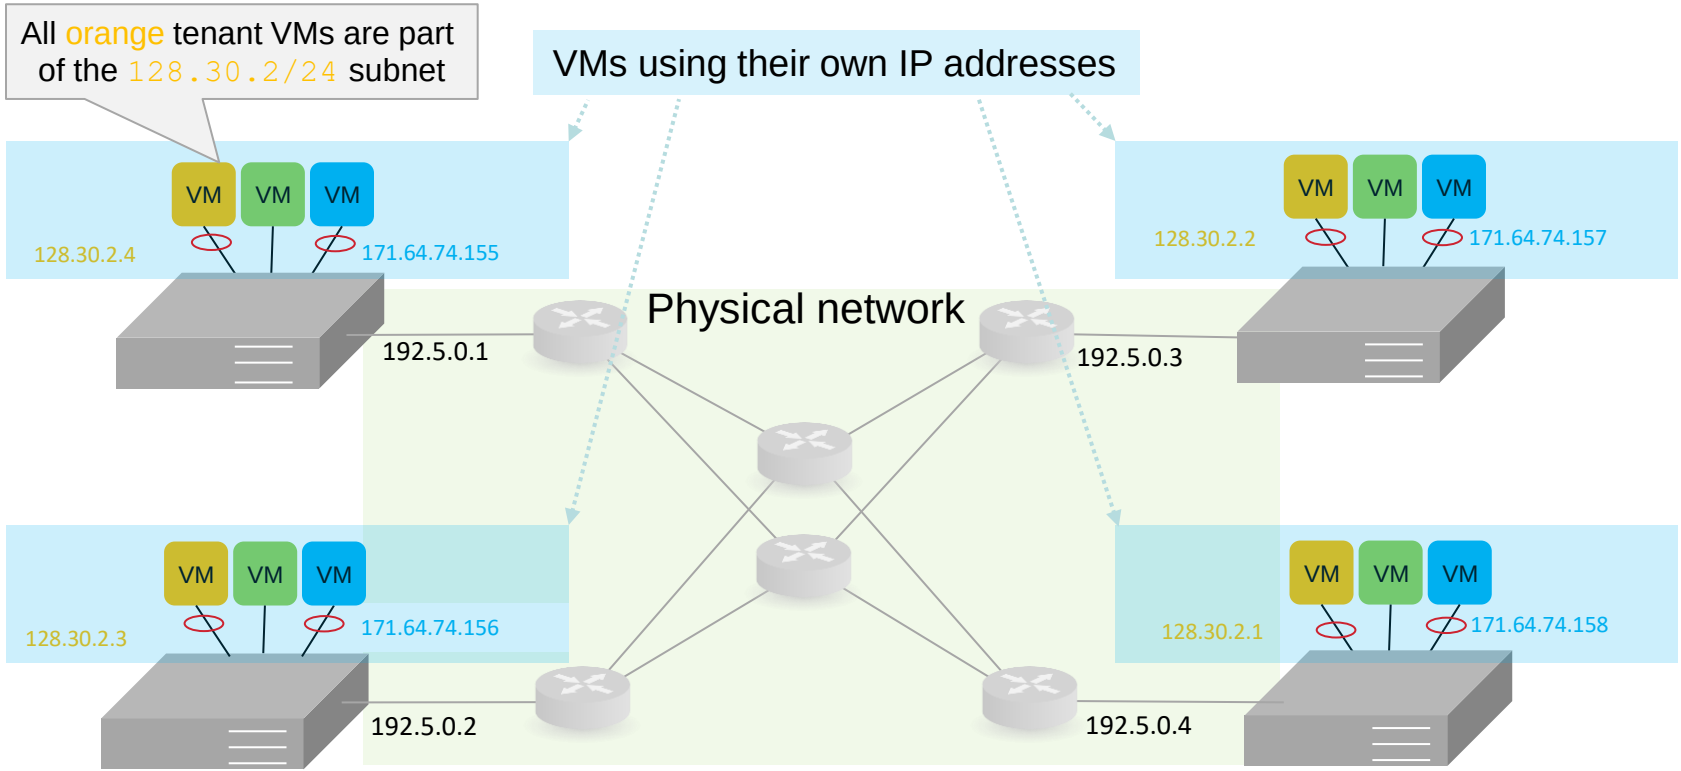

Q: Which mechanism Tag, tunnel or translation?

Mechanism: Tags, Tunnels or Translation?

Any mechanism could be made to work.

Tags: Switches contain a forwarding table per tenant.

- Tag in every packet indicates the tenant and therefore the forwarding table to use.
- But: We need to change the switches to recognize the tag and forward based on it.

Translation: Use NAT, with port numbers identifying VMs.

- But: Both ends behind NATs, therefore need NAT traversal everywhere – complicated.
- But: With thousands of VMs per server, quickly run out of port numbers for mapping.

Tunnel: Create tunnel between every pair of servers.
Forward traffic between VMs through the tunnel.

- But: We need to change switches to create tunnels.
- But: Server will receive packets for all addresses used by its VMs.

How it is done in virtualized data centers

1: Use the software “vSwitch” in every server

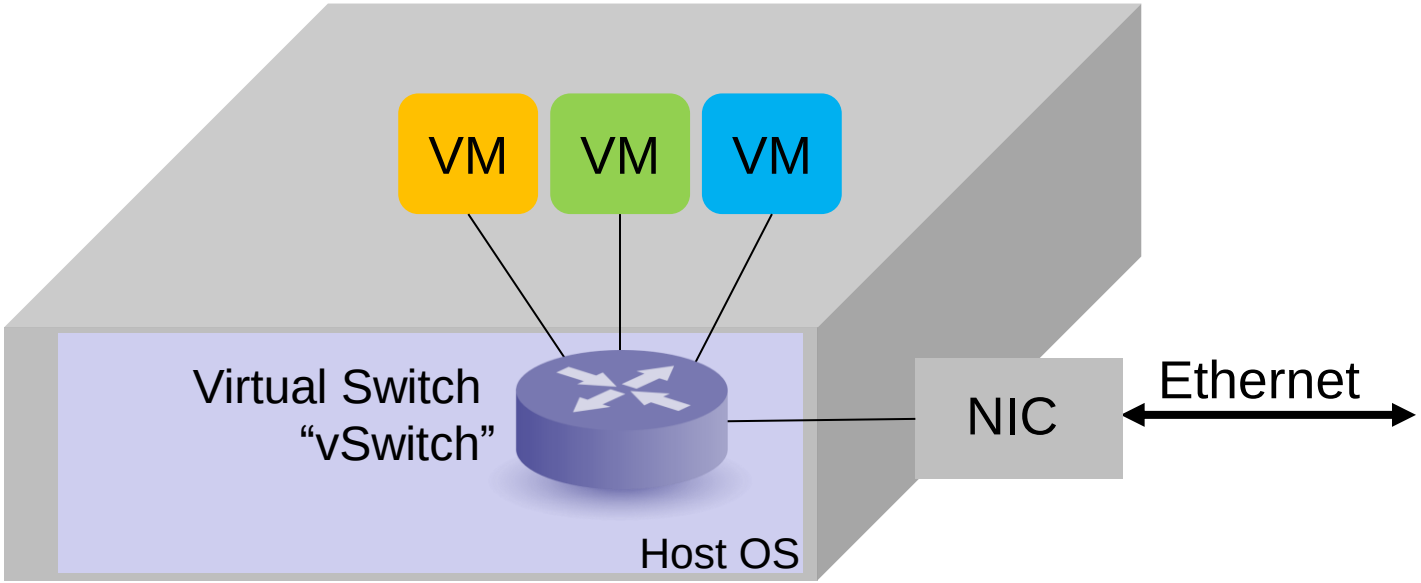

- Maintains tunnel to every other server’s vSwitch
- Tags packets with tenant ID
- Forwards packets into tunnel

2: Forward packets in tunnels between vSwitches

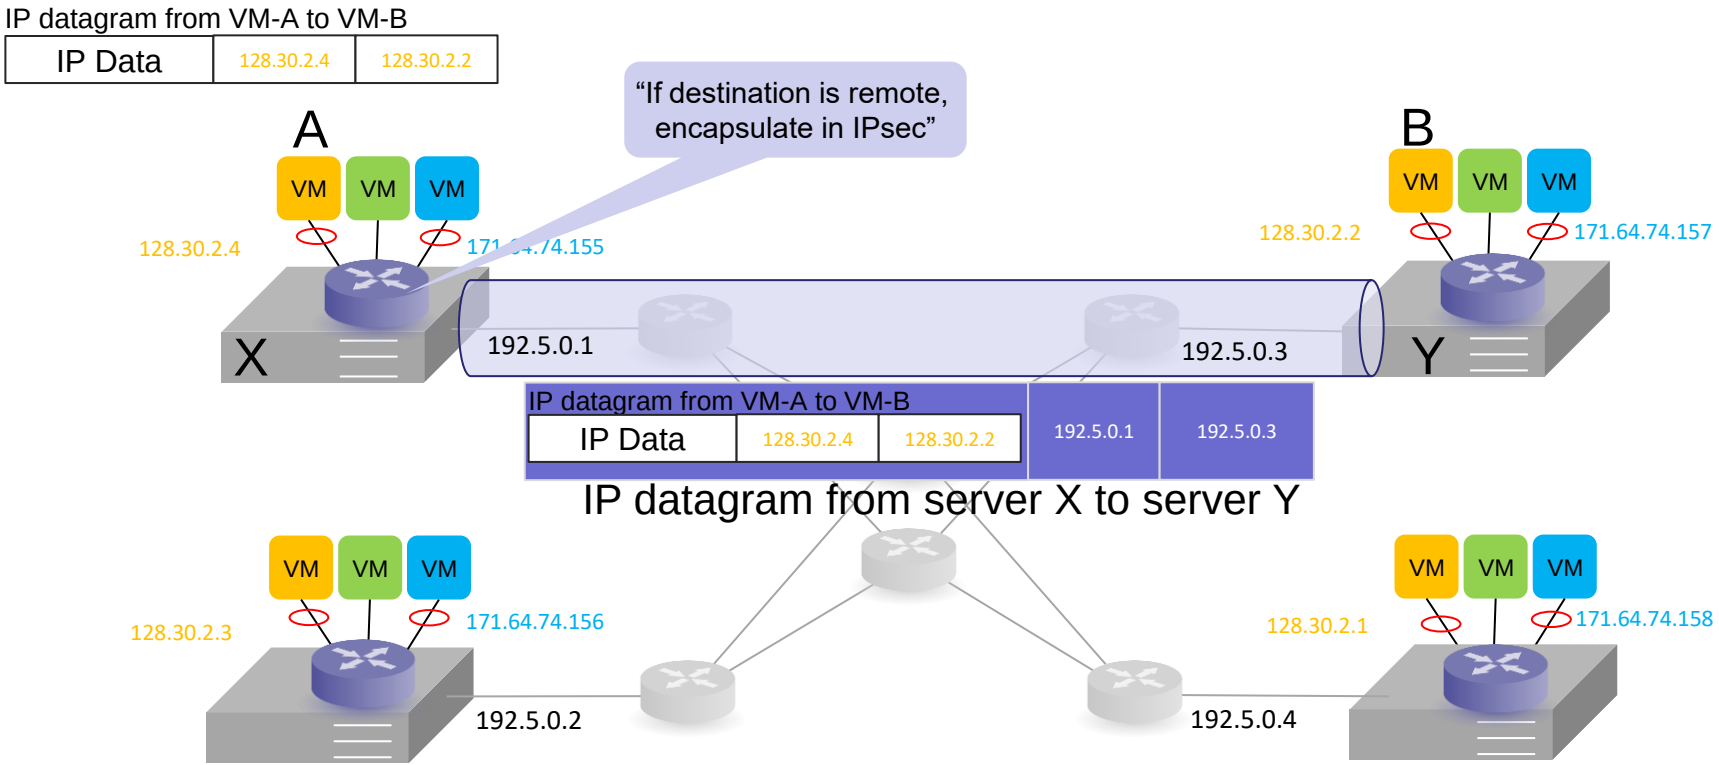

Learning goals of this class


To learn how tags, tunnels and translation can be used to provide new abstractions in a network.


To learn about the match + action abstraction


To learn about three examples:
Virtual LANs (VLANs), VPNs, and NATs.


To learn what network virtualization is.

To learn how overlay network virtualization works.

- To learn what network function virtualization (NFV) is.

Network Function Virtualization (NFV)

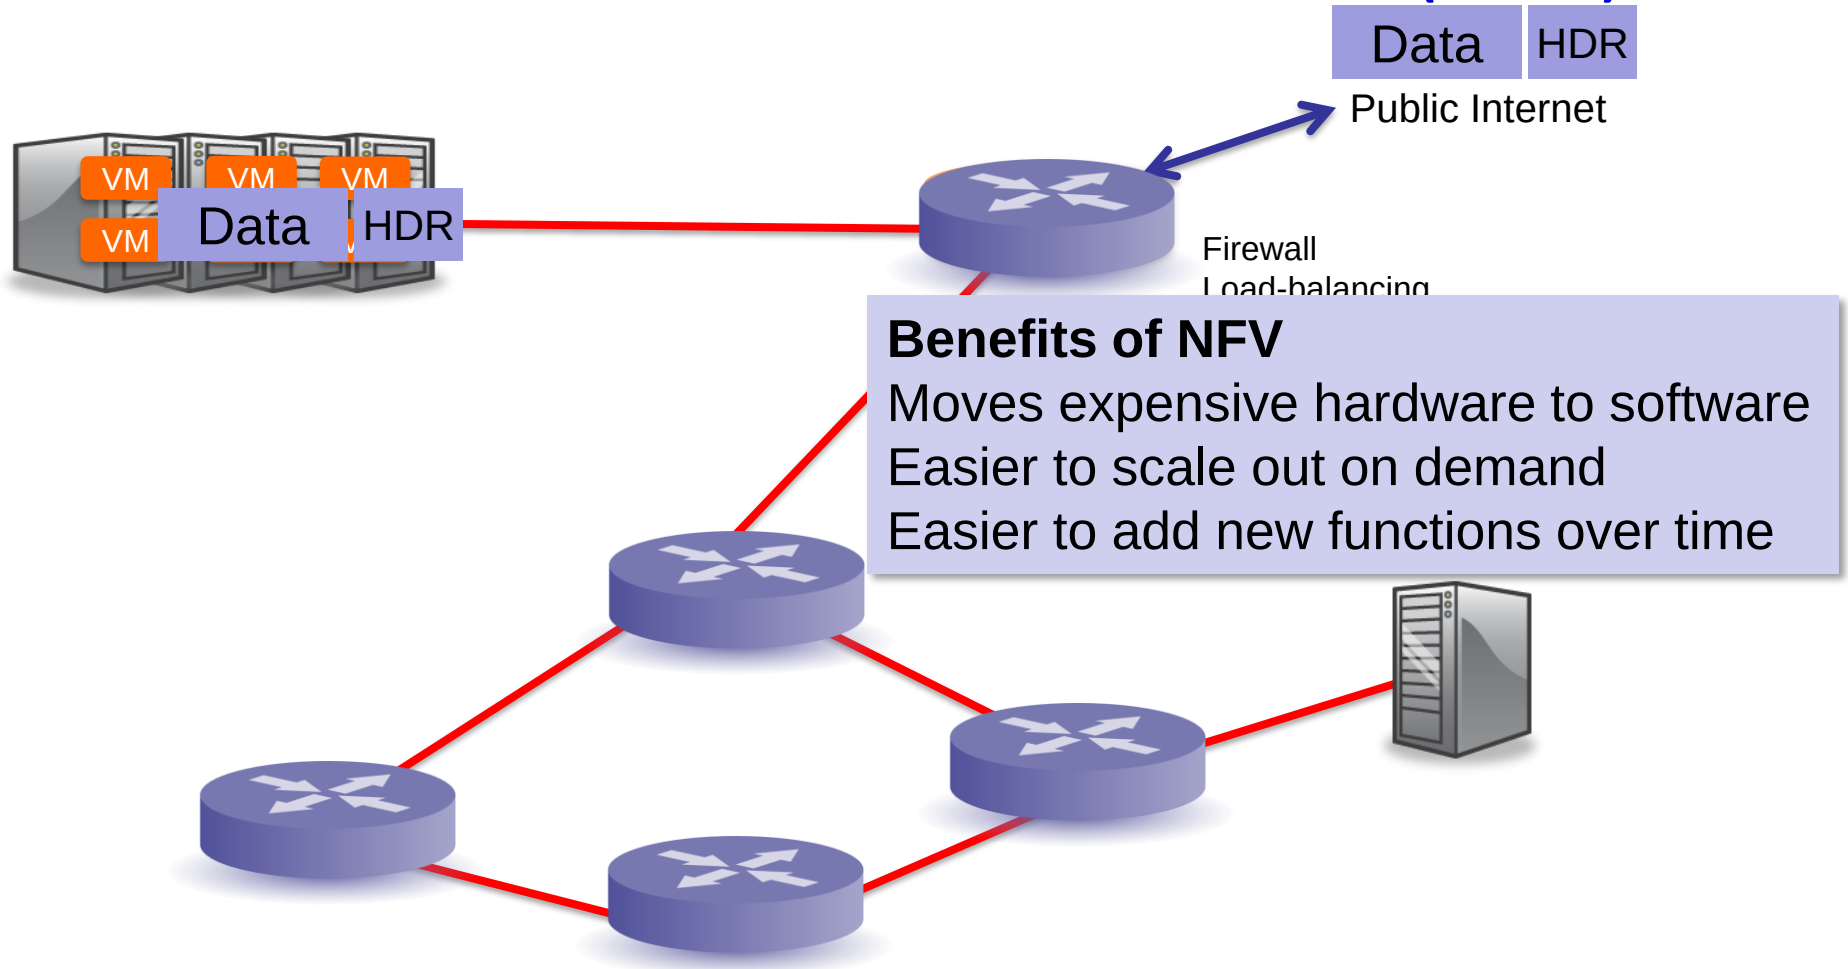

Learning goals of this class

- To learn how tags, tunnels and translation can be used to provide new abstractions in a network.
- To learn about the match + action abstraction
- To learn about three examples:
  Virtual LANs (VLANs), VPNs, and NATs.
- To learn what network virtualization is.
- To learn how overlay network virtualization works.
- To learn what network function virtualization (NFV) is.

Thank you!