Lesson 08-02:
Link Layer 2
VLAN and Data Center Network

CS 356 Computer Networks

Mikyung Han

mhan@cs.utexas.edu

Internet Reference Model

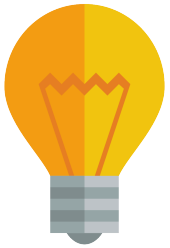

Example Protocols		Responsible for
FTP, HTTP, SMTP	Application	application specific needs
TCP, UDP	Transport	process to process data transfer
IP	Network	host to host data transfer across different network
Ethernet, WiFi	Link	data transfer between physically adjacent nodes
802.3 PHY	Physical	bit-by-bit or symbol-by-symbol delivery

2

Outline

1. VLAN

Why Virtual LANs (VLANs)?

Q: what happens as LAN sizes scale, users change point of attachment?

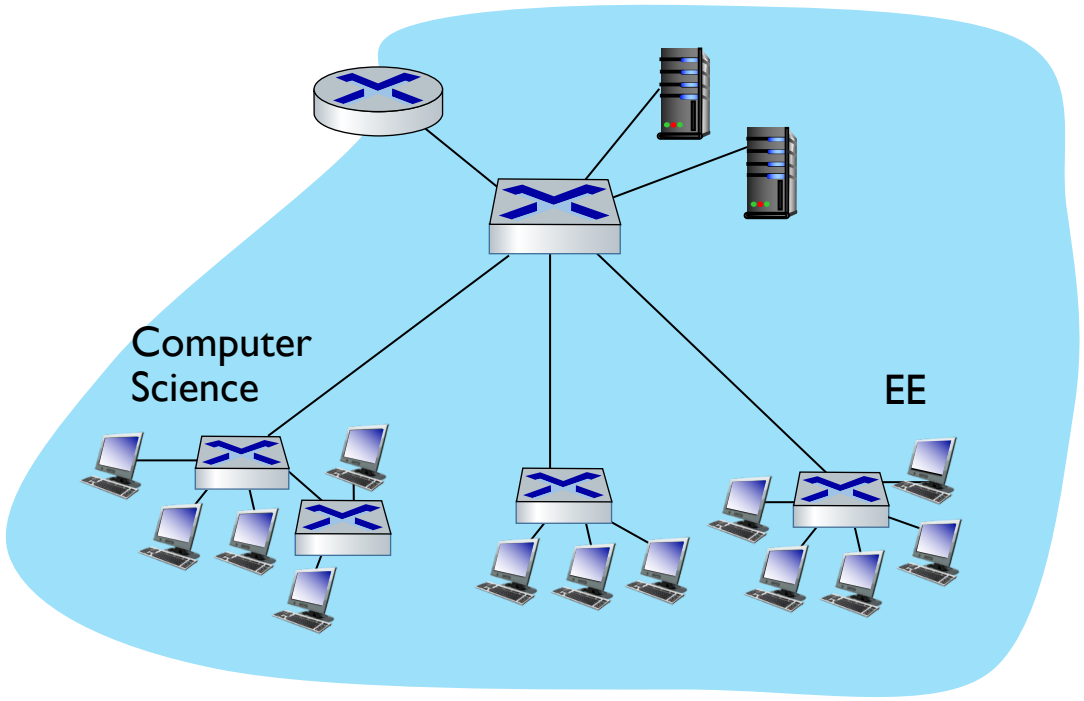

single broadcast domain:

- scaling: all layer-2 broadcast traffic (ARP, DHCP, unknown MAC) must cross entire LAN
- efficiency, security, privacy issues

Why Virtual LANs (VLANs)?

Q: what happens as LAN sizes scale, users change point of attachment?

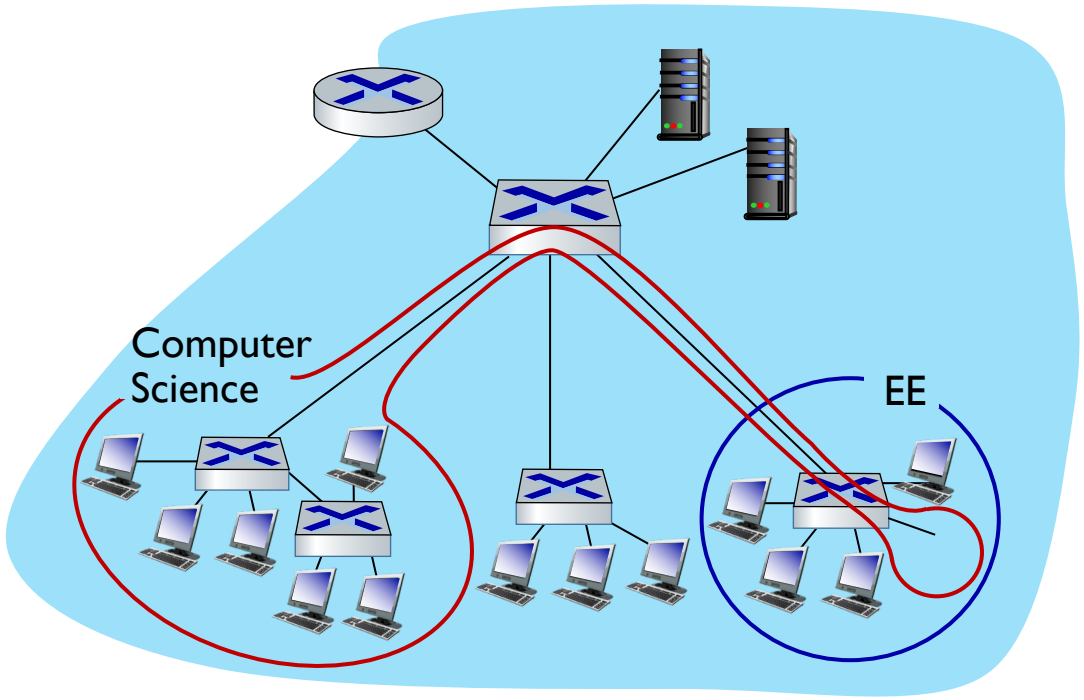

single broadcast domain:

- scaling: all layer-2 broadcast traffic (ARP, DHCP, unknown MAC) must cross entire LAN
- efficiency, security, privacy, efficiency issues

administrative issues:

- CS user moves office to EE - physically attached to EE switch, but wants to remain logically attached to CS switch

Port-based VLANs

Virtual Local Area Network (VLAN)

switch(es) supporting VLAN capabilities can be configured to define multiple virtual LANS over single physical LAN infrastructure.

port-based VLAN: switch ports grouped (by switch management software) so that single physical switch ……

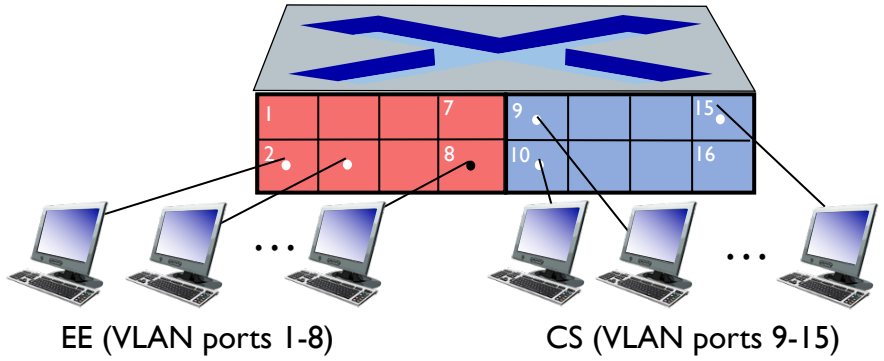

… operates as multiple virtual switches

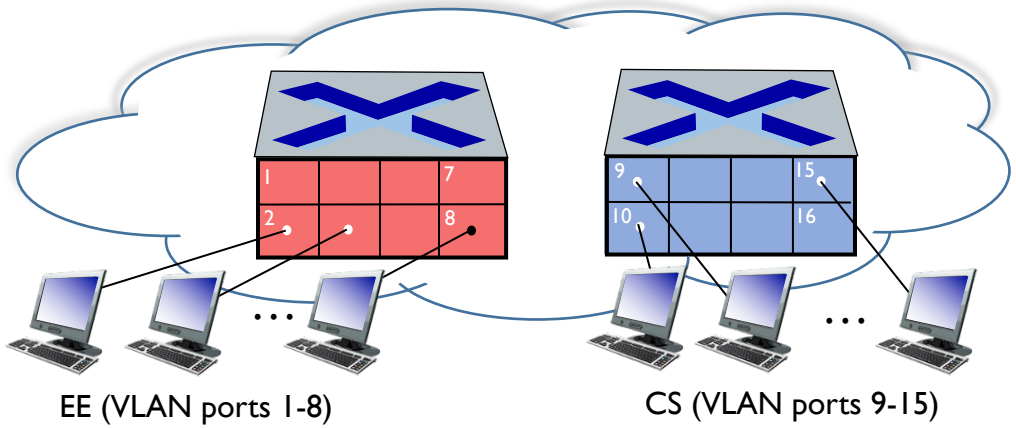

Port-based VLANs

- traffic isolation: frames to/from ports 1-8 can only reach ports 1-8
  - can also define VLAN based on MAC addresses of endpoints, rather than switch port
- dynamic membership: ports can be dynamically assigned among VLANs
- forwarding between VLANS: done via routing (just as with separate switches)
  - in practice vendors sell combined switches plus routers

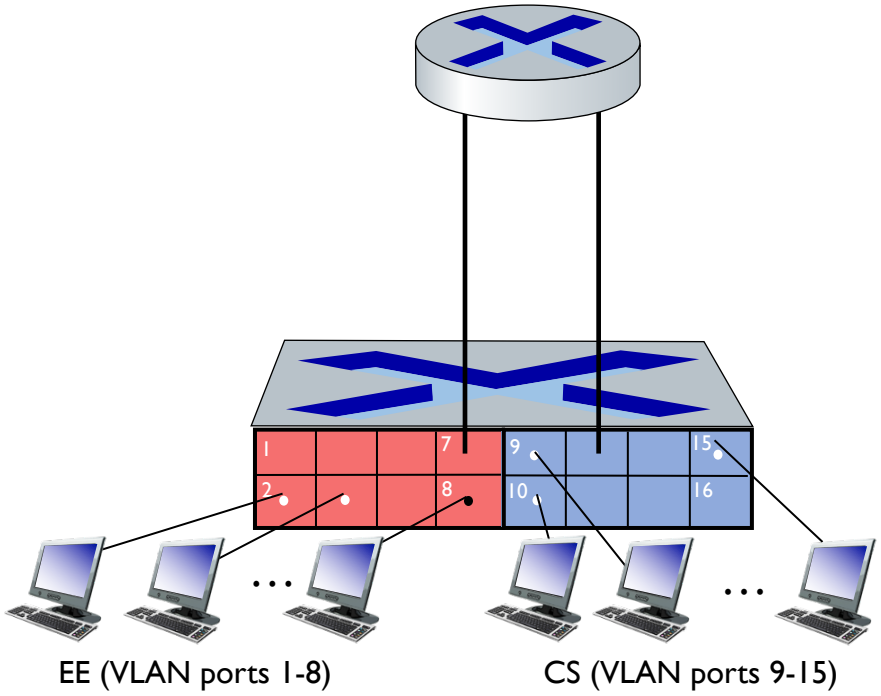

Outline

1. VLAN
2. Data Center Networking

What is Datacenter Networks?

100's of thousands of hosts, often closely coupled, in close proximity

- e-business (e.g. Amazon)
- content-servers (e.g., YouTube, Akamai, Apple, Microsoft)
- search engines, data mining (e.g., Google)

challenges:

- multiple applications, each serving massive numbers of clients
- reliability
- managing/balancing load, avoiding processing, networking, data bottlenecks

Inside a 40-ft Microsoft container, Chicago data center

Datacenter networks: network elements

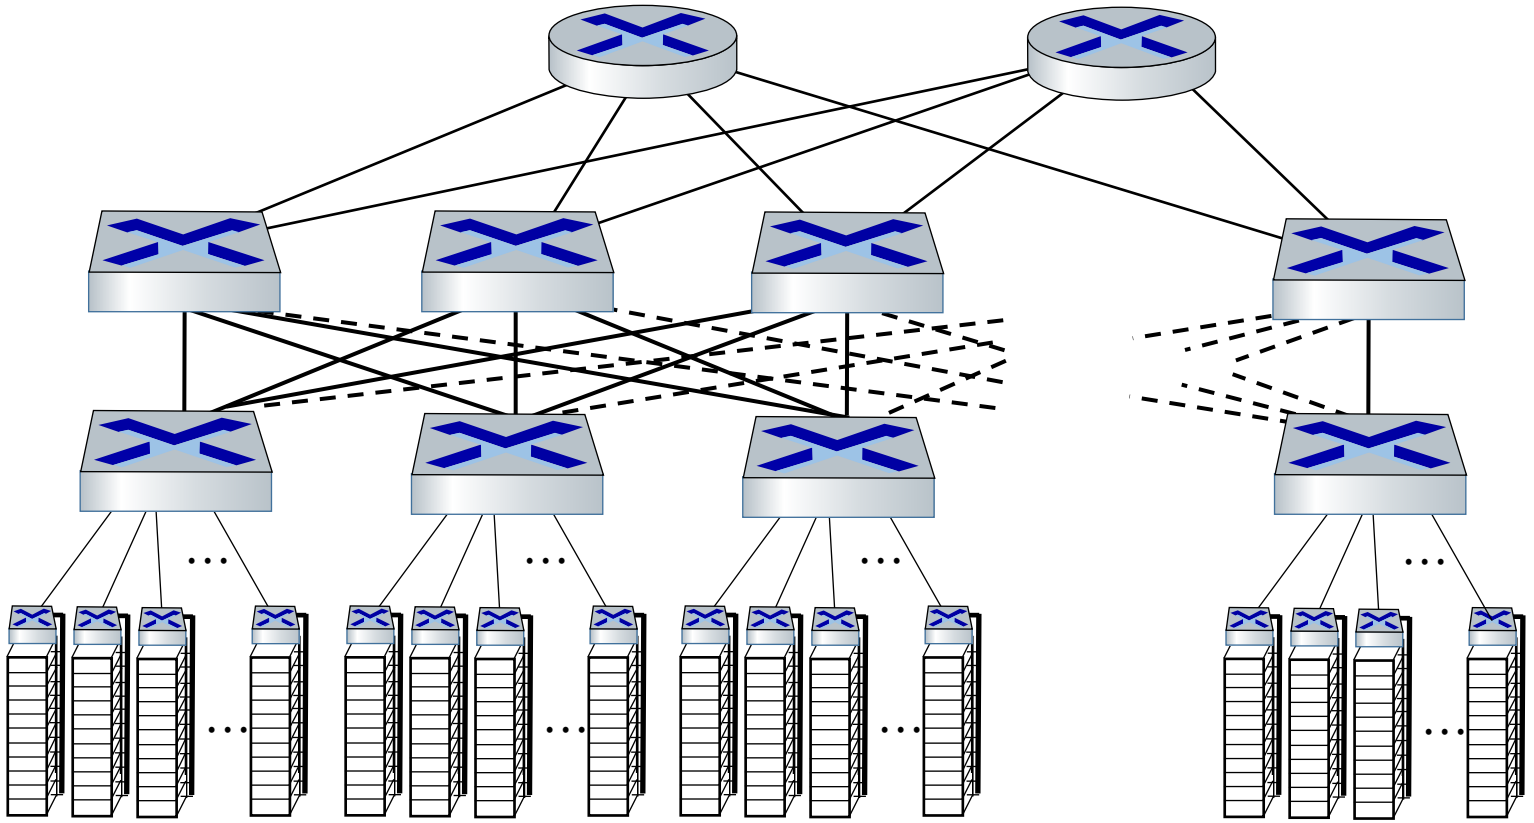

Border routers

- connections outside datacenter

Tier-1 switches

- connecting to ~16 T-2s below

Tier-2 switches

- connecting to ~16 TORs below

Top of Rack (TOR) switch

- one per rack
- 40-100Gbps Ethernet to blades

Server racks

- 20- 40 server blades: hosts

Datacenter networks: network elements

Facebook F16 data center network topology:

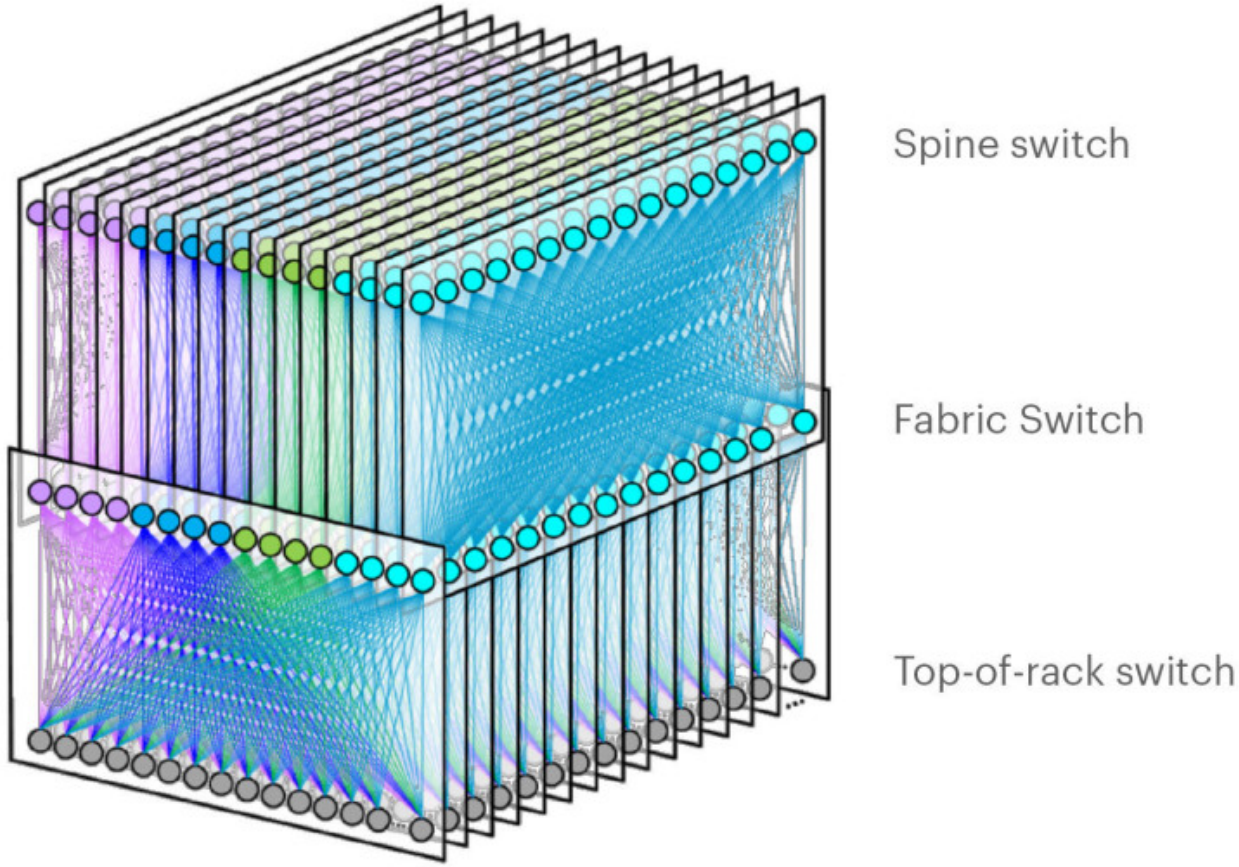

https://engineering.fb.com/data-center-engineering/f16-minipack/ (posted 3/2019)

Datacenter networks: multipath

- rich interconnection among switches, racks:
  - increased throughput between racks (multiple routing paths possible)
  - increased reliability via redundancy

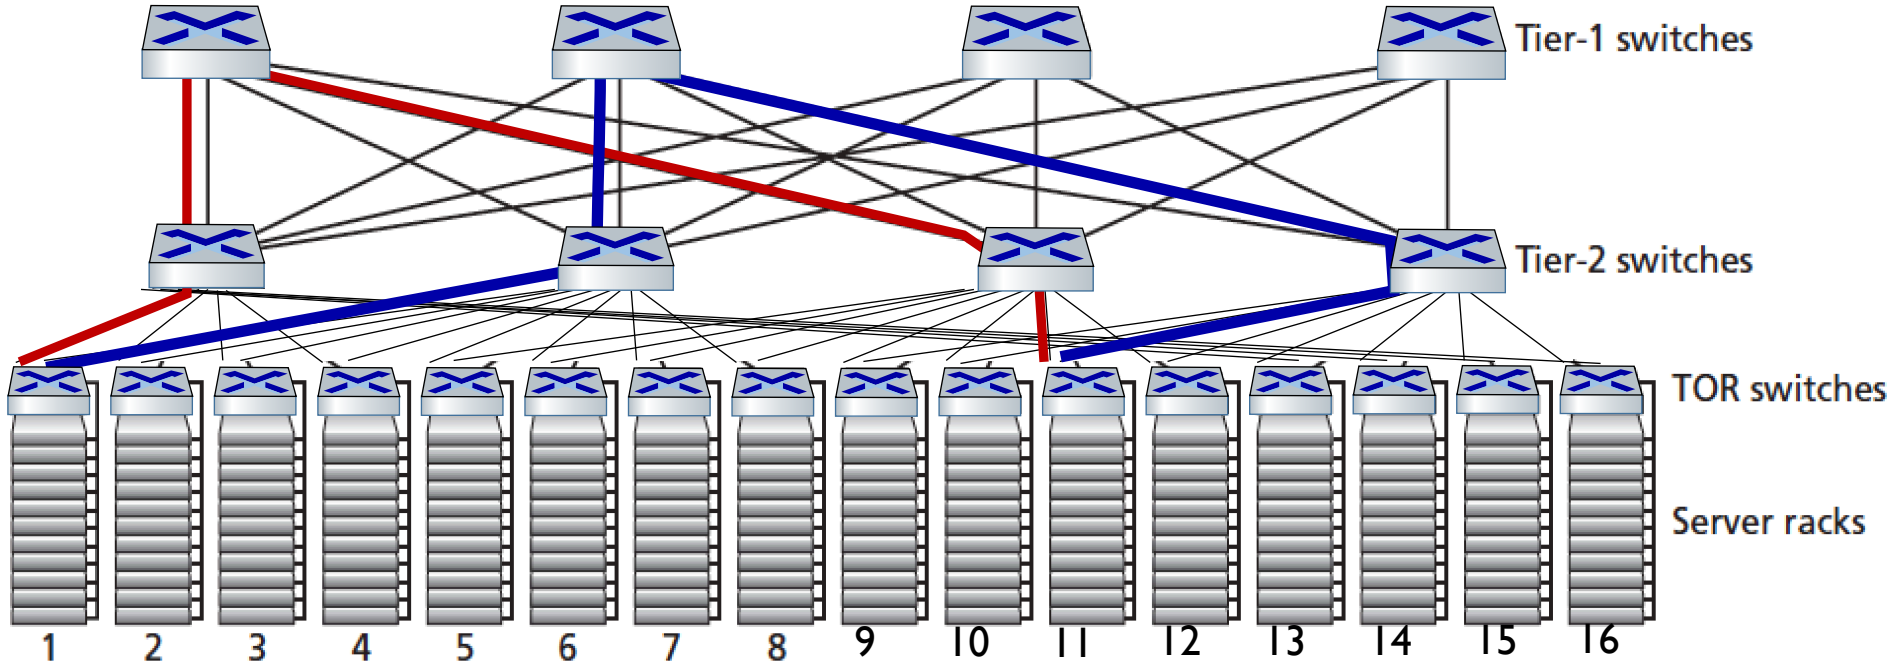

two disjoint paths highlighted between racks 1 and 11

Datacenter networks: protocol innovations

- link layer:
  - RoCE: remote Direct Memory Access (RDMA) over Ethernet
- transport layer:
  - ECN (explicit congestion notification) used for congestion control (DC TCP, DC Quantized Congestion Notification)
  - experimentation with hop-by-hop (backpressure) congestion control
- routing, management:
  - SDN widely used within/among organizations' datacenters
  - place related services, data as close as possible (e.g., in same rack or nearby rack) to minimize tier-2, tier-1 communication

Outline

1. VLAN
2. Data Center Networking
3. Summary

Let's reflect on the course goals

Course goals

1. Understand HOW Internet works
2. Understand WHY behind its design
3. Know the fundamentals

15

First objective:
Understand HOW Internet works

16

BTS Jungkook's post reached 1M people in just 10 min!

Video of Jeon Jungkook singing Lauv's "Never Not" via Twitter (@BTS_twt)

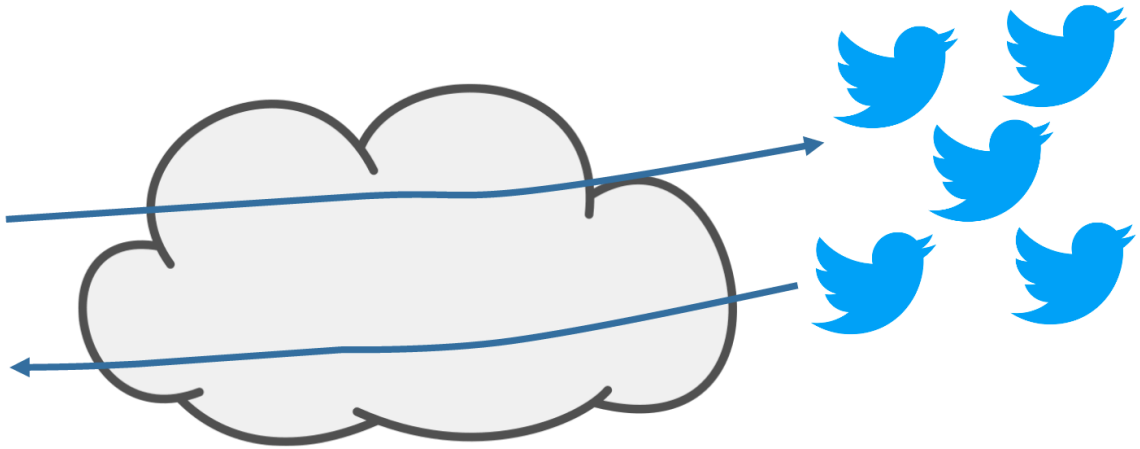

17

One: Learn HOW Internet works

Video of Jeon Jungkook singing Lauv's "Never Not" via Twitter (@BTS_twt)

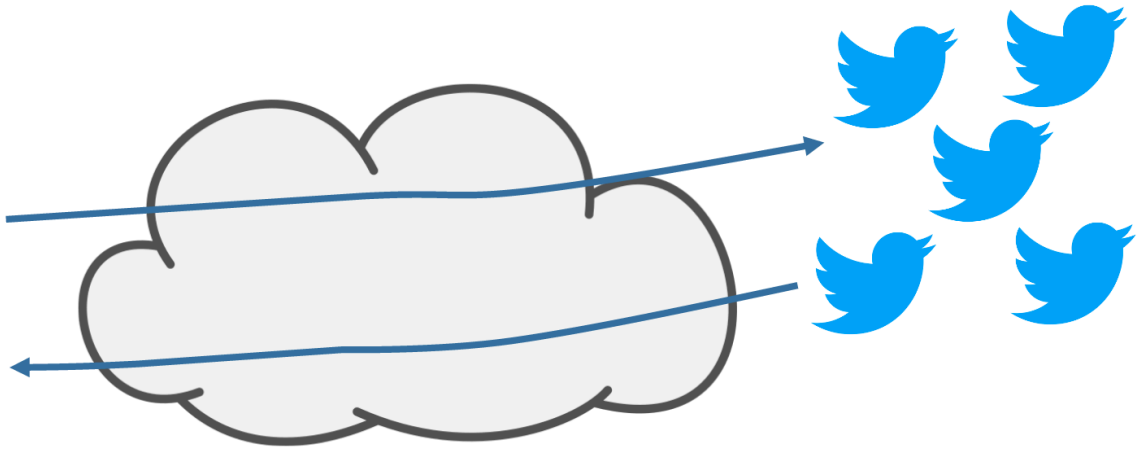

What really happened behind the scene?

18

A day in the life: scenario

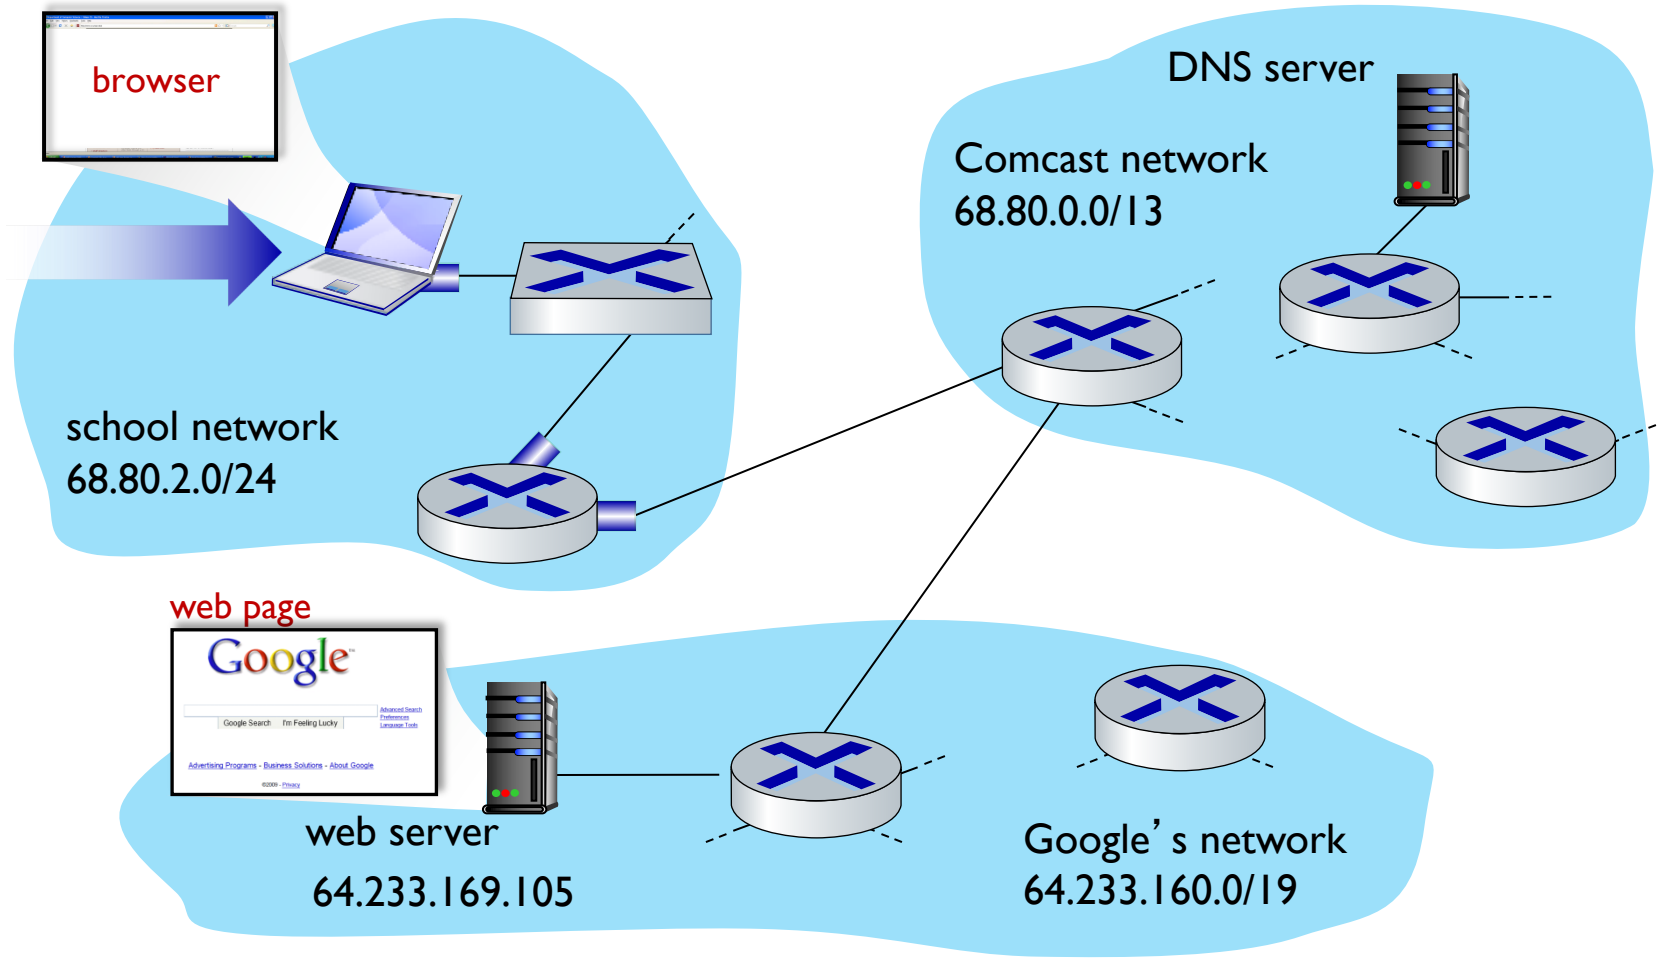

scenario:

- arriving mobile client attaches to network …
- requests web page: www.google.com

Sounds simple!

A day in the life: connecting to the Internet

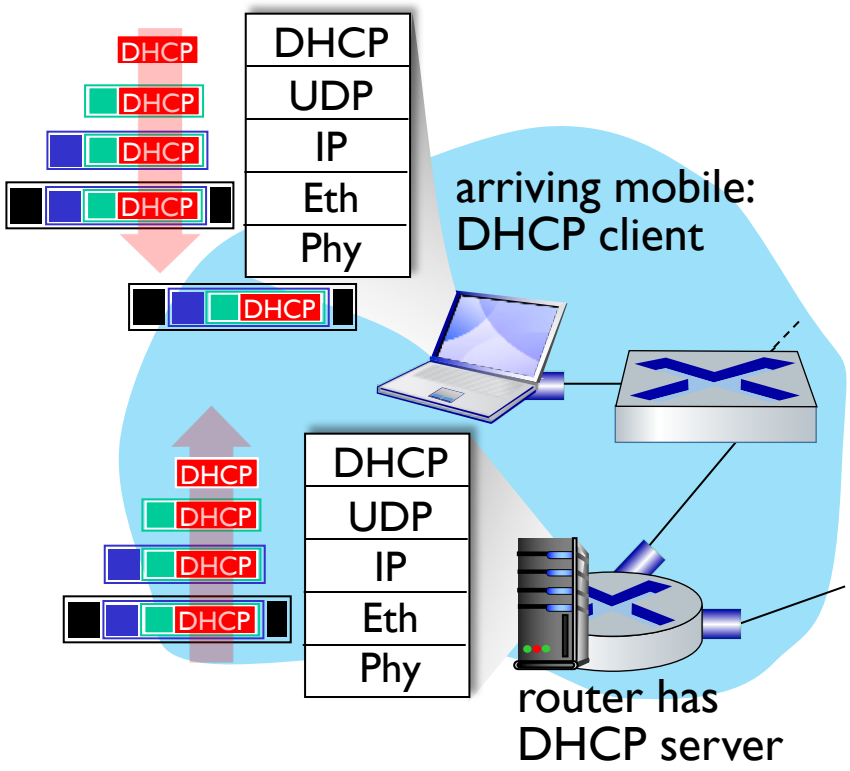

- connecting laptop needs to get its own IP address, addr of first-hop router, addr of DNS server: use DHCP
- DHCP request encapsulated in UDP, encapsulated in IP, encapsulated in 802.3 Ethernet
- Ethernet frame broadcast (dest: FFFFFFFFFFFF) on LAN, received at router running DHCP server
- Ethernet demuxed to IP demuxed, UDP demuxed to DHCP

A day in the life: connecting to the Internet

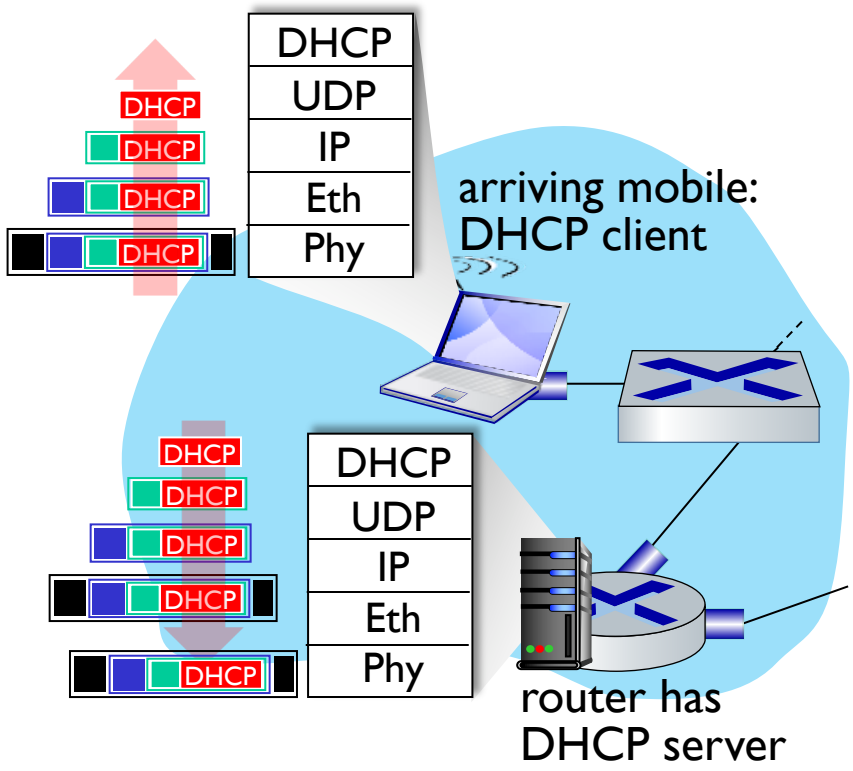

- DHCP server formulates DHCP ACK containing client's IP address, IP address of first-hop router for client, name & IP address of DNS server
- encapsulation at DHCP server, frame forwarded (switch learning) through LAN, demultiplexing at client
- DHCP client receives DHCP ACK reply

Client now has IP address, knows name & addr of DNS server, IP address of its first-hop router

A day in the life… ARP (before DNS, before HTTP)

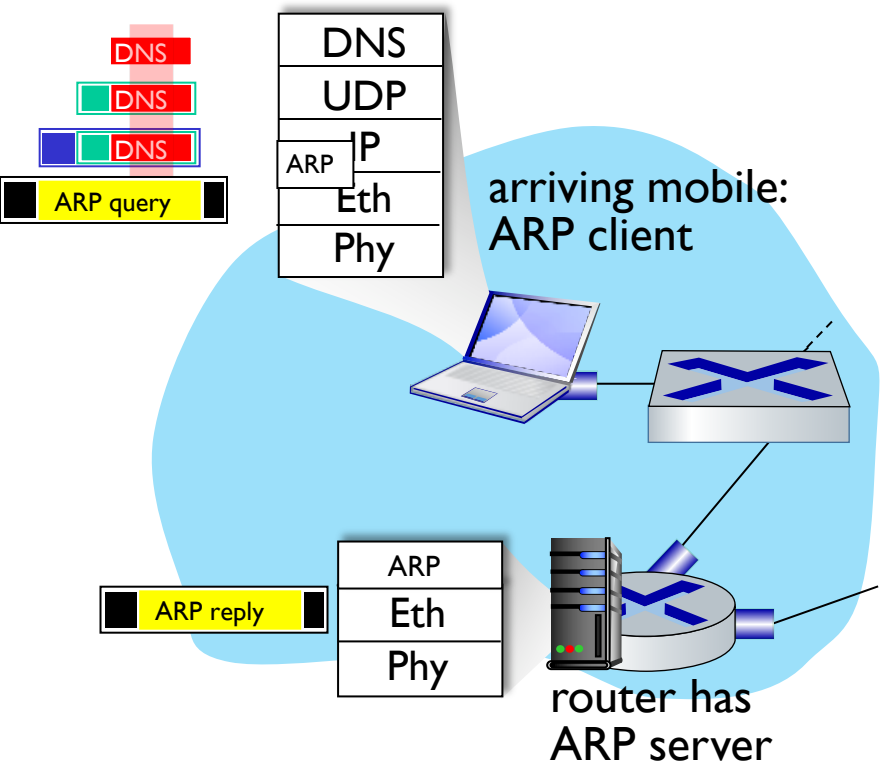

- before sending HTTP request, need IP address of www.google.com: DNS
- DNS query created, encapsulated in UDP, encapsulated in IP, encapsulated in Eth. To send frame to router, need MAC address of router interface: ARP
- ARP query broadcast, received by router, which replies with ARP reply giving MAC address of router interface
- client now knows MAC address of first hop router, so can now send frame containing DNS query

A day in the life… using DNS

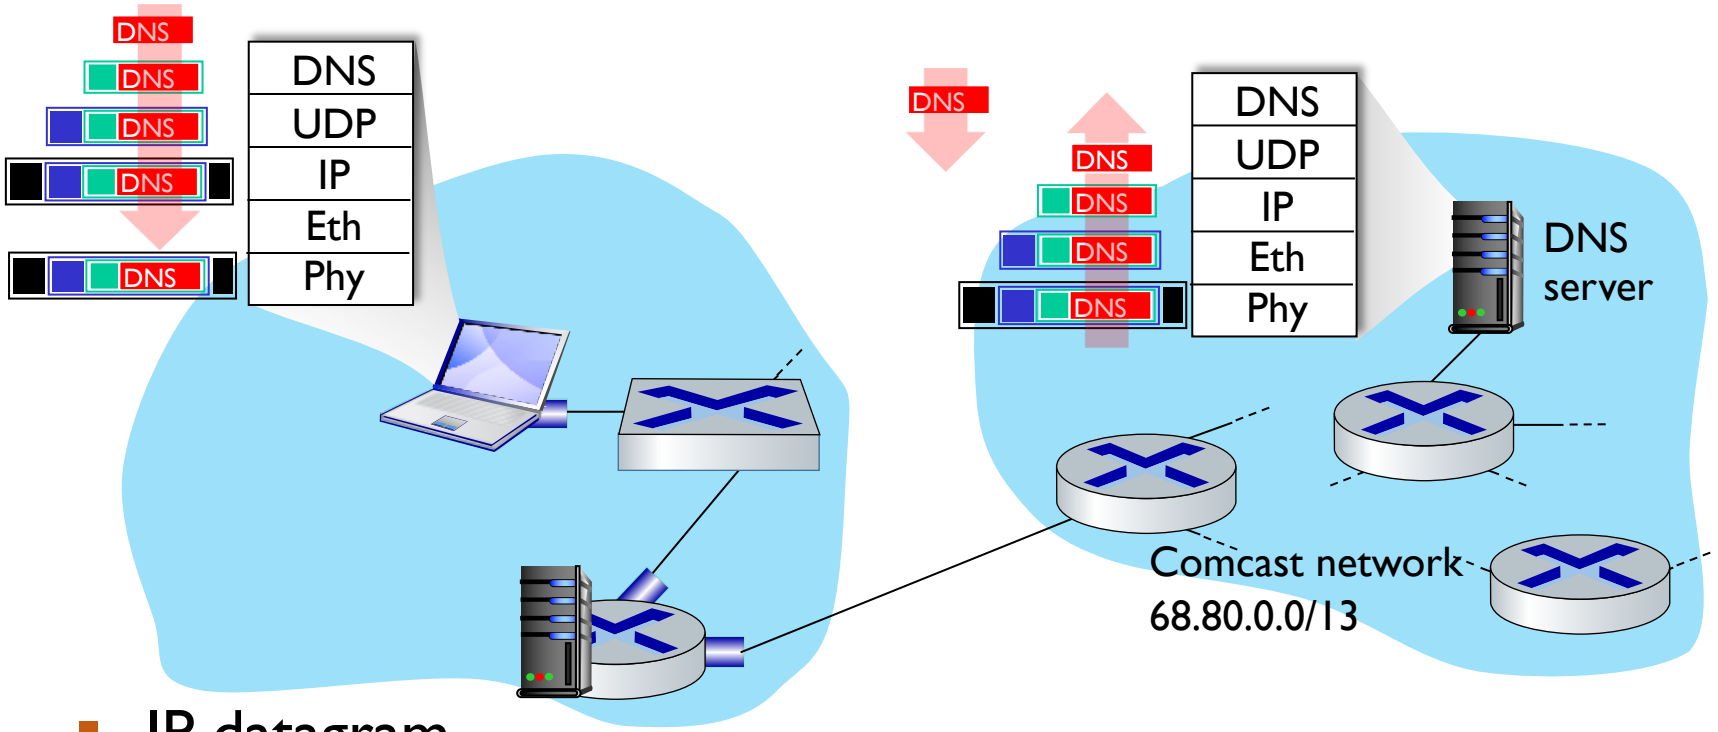

- demuxed to DNS
- DNS replies to client with IP address of www.google.com

- IP datagram containing DNS query forwarded via LAN switch from client to 1st hop router

- IP datagram forwarded from campus network into Comcast network, routed (tables created by RIP, OSPF, IS-IS and/or BGP routing protocols) to DNS server

A day in the life…TCP connection carrying HTTP

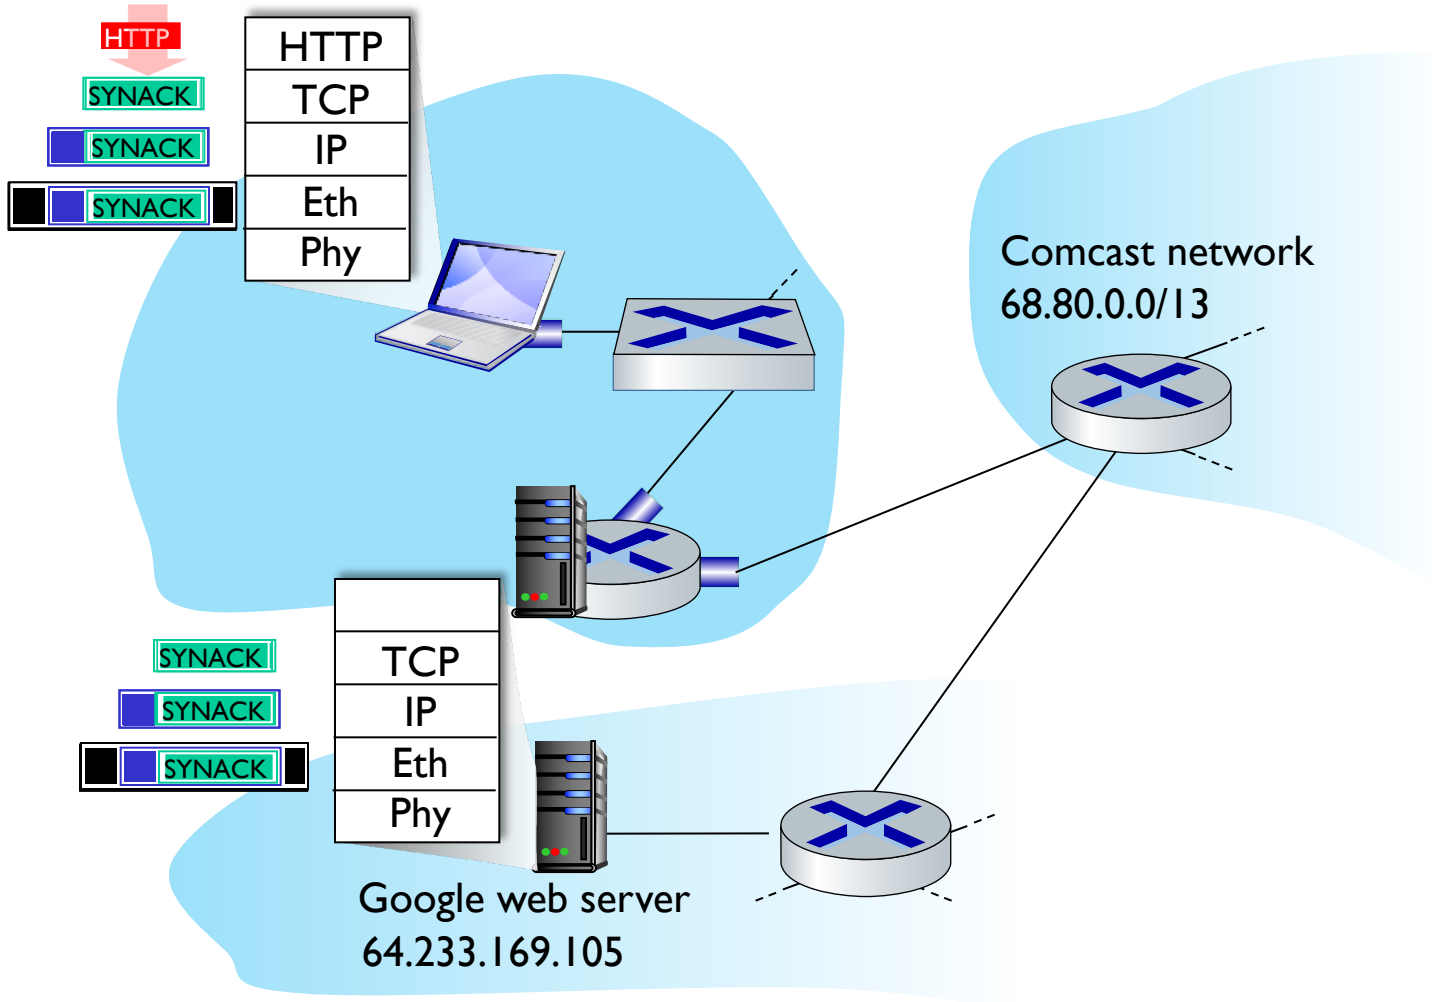

- to send HTTP request, client first opens TCP socket to web server
- TCP SYN segment (step 1 in TCP 3-way handshake) inter-domain routed to web server
- web server responds with TCP SYNACK (step 2 in TCP 3-way handshake)
- TCP connection established!

A day in the life… HTTP request/reply

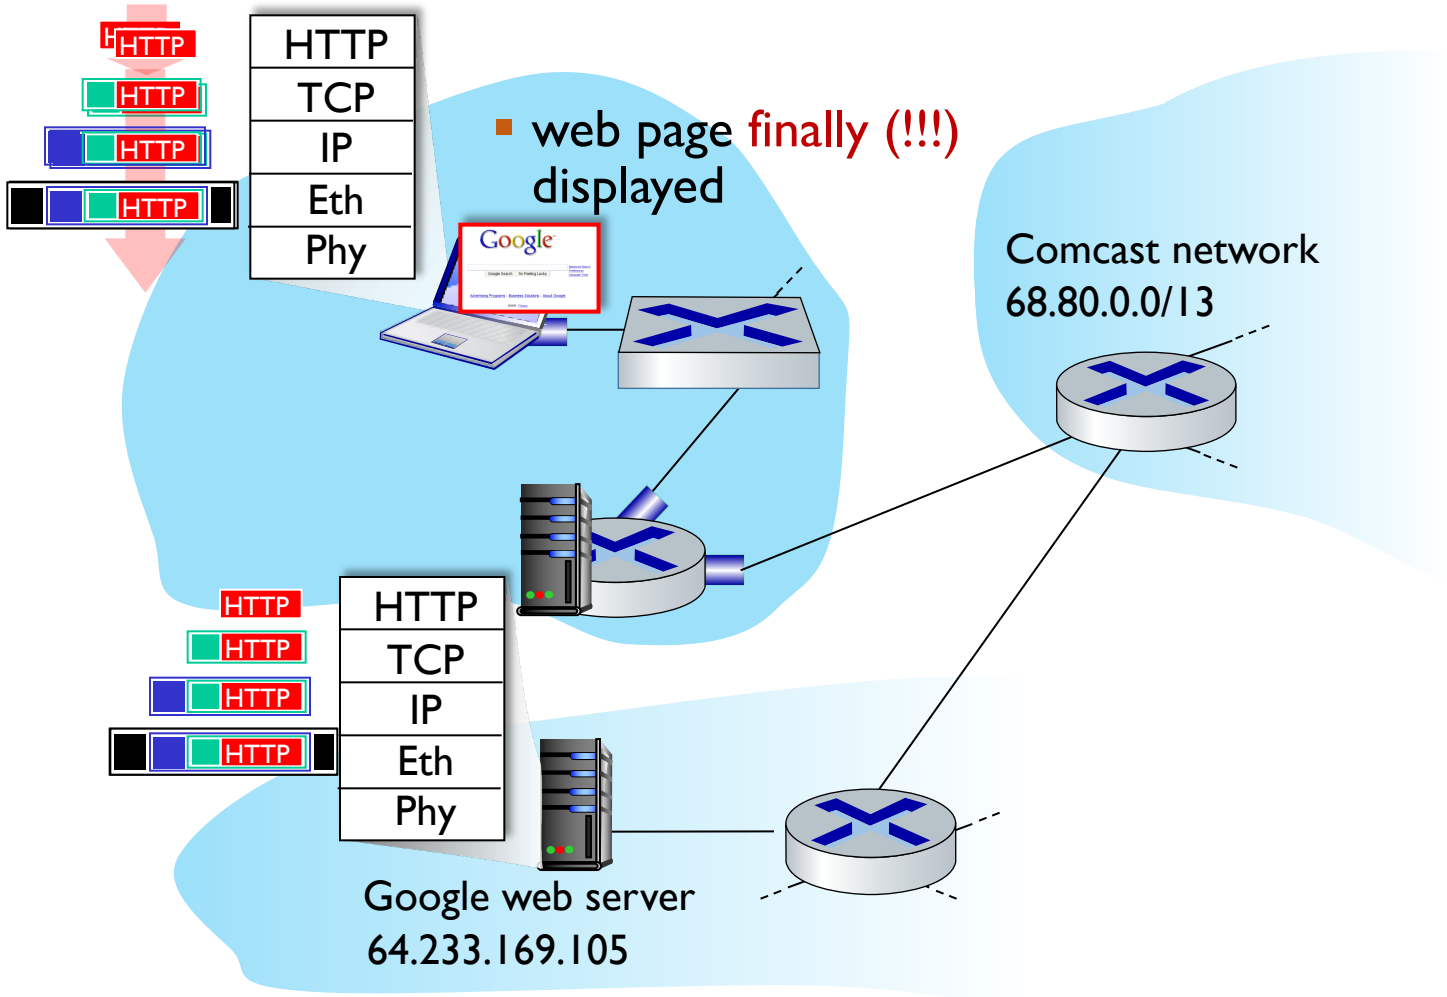

- HTTP request sent into TCP socket
- IP datagram containing HTTP request routed to www.google.com
- web server responds with HTTP reply (containing web page)
- IP datagram containing HTTP reply routed back to client

In addition…

Video of Jeon Jungkook singing Lauv's "Never Not" via Twitter (@BTS_twt)

In addition ….since it's multimedia streaming

- Video encoding
- Web cache
- Dynamic Adaptive Streaming over HTTP (DASH) of Content Distribution Network (CDN)
- Playout buffering
- TLS for security: encryption, message integrity, authentication
- Tor could have been used if users are in a country where Twitter is censored or just proxy
- Bloom filter and distributed hash table could have been used

Second Objective:
Understand WHY behind the Internet design

28

Second objective:
Understand WHY behind what

29

Internet Reference Model

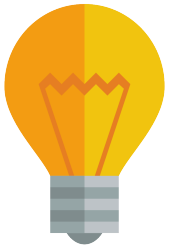

Example Protocols		Responsible for
FTP, HTTP, SMTP	Application	application specific needs
TCP, UDP	Transport	process to process data transfer
IP	Network	host to host data transfer across different network
Ethernet, WiFi	Link	data transfer between physically adjacent nodes
802.3 PHY	Physical	bit-by-bit or symbol-by-symbol delivery

30

Motivation

- Why layers?
- Why CDN?
- Why TCP/UDP?
- Why SDN?
- Why overlay?
- Why Tor?
- Why VLAN?
- ..

Third Objective:
Know the fundamentals of computer networks

- Back of the envelope calculations
- Reliable data transfer
- Stateless vs stateful
- Connectionless vs connection oriented
- Flow control
- Congestion control
- Error detection
- Routing vs switching
- Addressing

32

Summary

- Covered all layers (except PHY)
- Covered major protocols and fundamentals + Tor
- 2 multi-threaded projects
  - UDP client-server blocking and non-blocking IO
  - HTTP Proxy
- 3 labs
  - TCP Buffer Bloat/Network Measurements/Port Scan
- 3 hands on
  - DNS Dig/TCP Wireshark/Traceroute

Thanks for your hard work! ☺

Acknowledgements

Slides are adopted from Kurose' Computer Networking Slides