COMP 204 – Principles of Computer Networks

Week 6

FRANKLIN UNIVERSITY

Agenda

- Review this week's learning outcomes
- Presentation of this week's material
- Introduce homework problems
- Final Exam FAQs
- Q & A session

FRANKLIN UNIVERSITY

This Week's Outcomes

- Categorize Ethernet functionality in layers 1 and 2 of the OSI model.
- Describe the purpose and operation of ARP.
- Simulate collision resolution on Ethernet.
- List cabling types, standards, and ports for common physical connections.

Review – Application and Transport

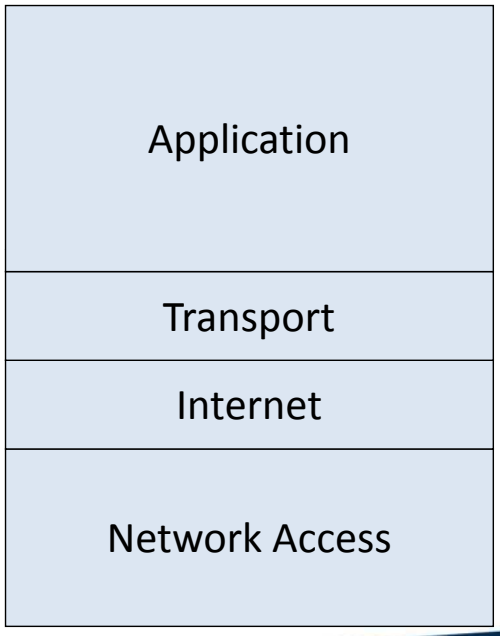

FRANKLIN UNIVERSITY

Review – Application Layer

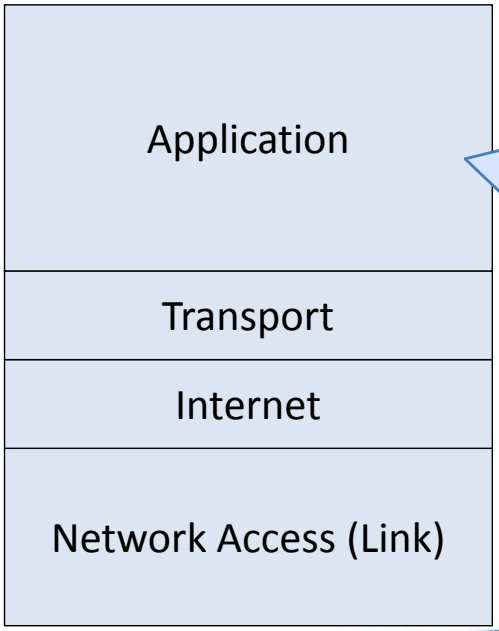

OSI layers 5, 6, 7. User applications, services, and application layer protocols. "messages"

- Ex: browser, httpd, and HTTPD working together.
- App message -> app protocol -> transport layer.

FRANKLIN UNIVERSITY

Review – Transport Layer

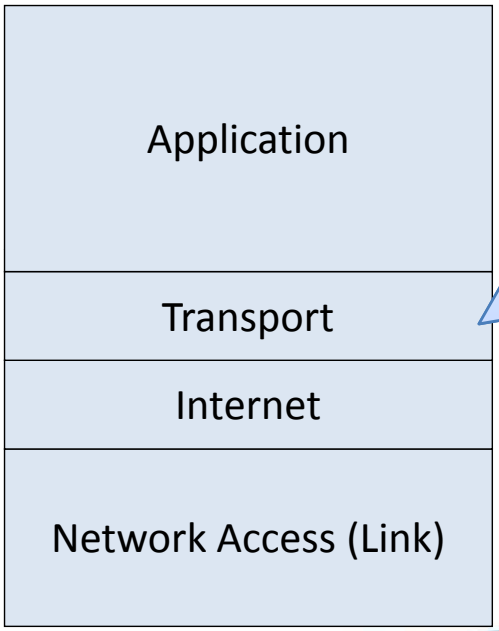

OSI layer 4. Typical examples are UDP and TCP. "segments"

- Port – specifies target service/app listening for messages
- UDP – "connectionless"
- TCP – "connection oriented"
- Header differences
- Handshaking in TCP

FRANKLIN UNIVERSITY

Review – Internet Layer

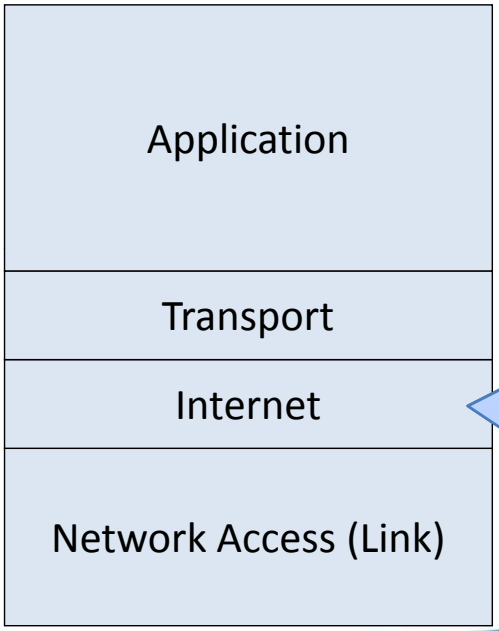

How messages get from host to host. Layer 3. "packets"

- Ex: IP (Internet Protocol) and ICMP (Internet Control Message Protocol)
- Routing
- VLSM/CIDR for subnets

FRANKLIN UNIVERSITY

Review - Link Layer

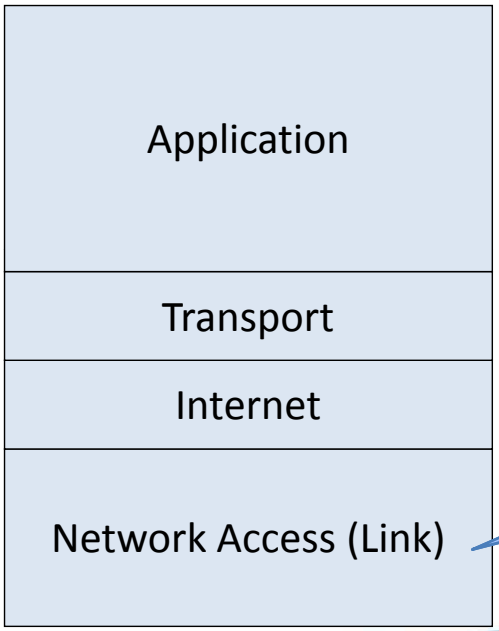

How messages get from device to device. OSI Layers 2 and 1. "Frames."

- Ex: Ethernet, PPP, ATM
- "Bits on a wire"
- LLC/MAC
- ARP/RARP

FRANKLIN UNIVERSITY

Chapter 9 and 10

- Ethernet deep dive
- Cabling and devices

Split Layer 2

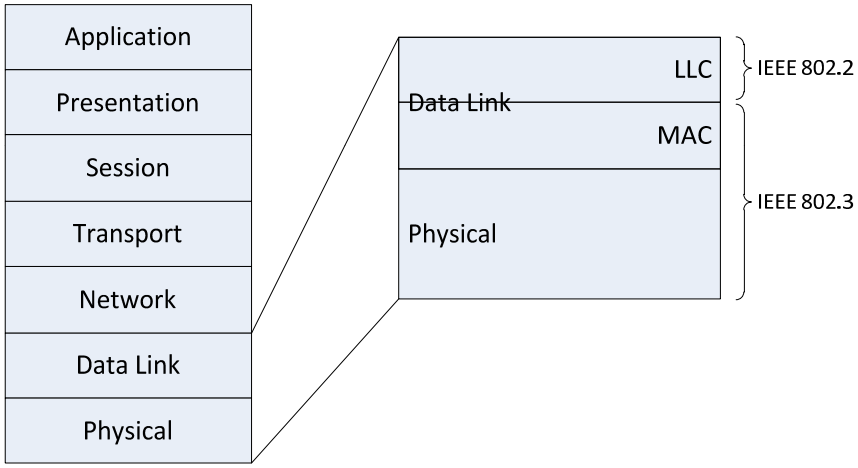

Split Layer 2

Logical Link Control: lets multiple layer 3 protocols exist (IP, IPX, ICMP, etc) on the same network.

- Ethertype field in Ethernet frames
- Physical address (copied into RAM)
- Framing

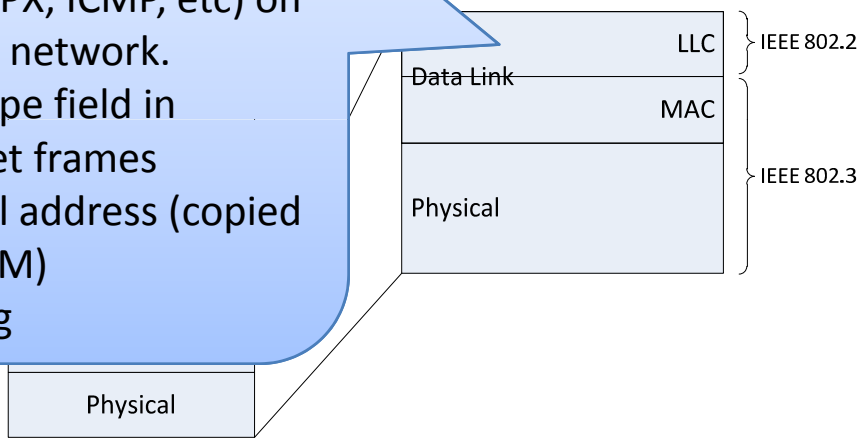

Split Layer 2

Logical Link Control: lets multiple layer 3 protocols exist (IP, IPX, ICMP, etc) on the same network.

Media Access Control: hardware.

- Frame delimiting
- Addressing
- Error detection (collisions & CRC)

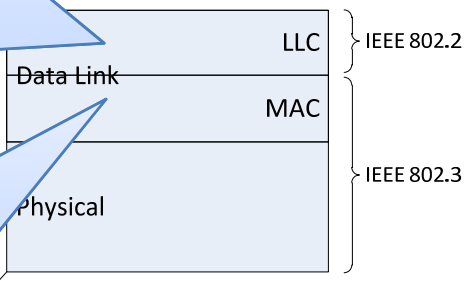

FRANKLIN UNIVERSITY

Split Layer 2

Logical Link Control: lets multiple layer 3 protocols exist (IP, IPX, ICMP, etc) on the same network.

- Ethertype field in

Media Access Control: hardware.

- Frame delimiting
- Addressing
- Error detection (collisions & CRC)

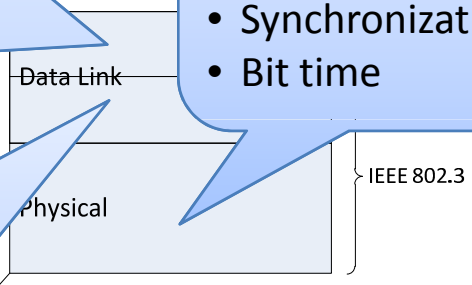

The "wire"

- Encoding
- Signaling
- Synchronization
- Bit time

Legacy Ethernet

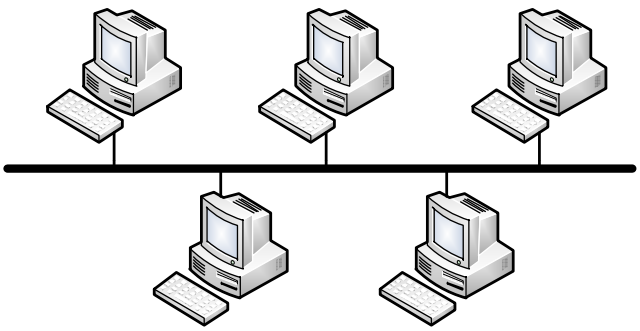

Physical bus, logical bus. Shared media. Half duplex. Collisions are an issue.

Legacy Ethernet

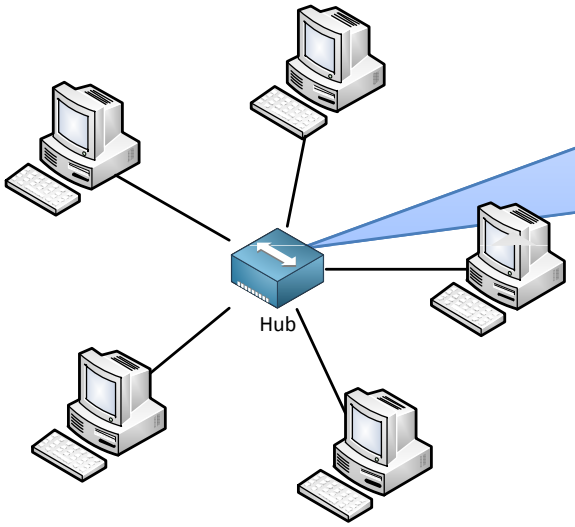

Physical star, logical bus. Shared media. Half duplex. Collisions are an issue.

FRANKLIN UNIVERSITY

Legacy Ethernet

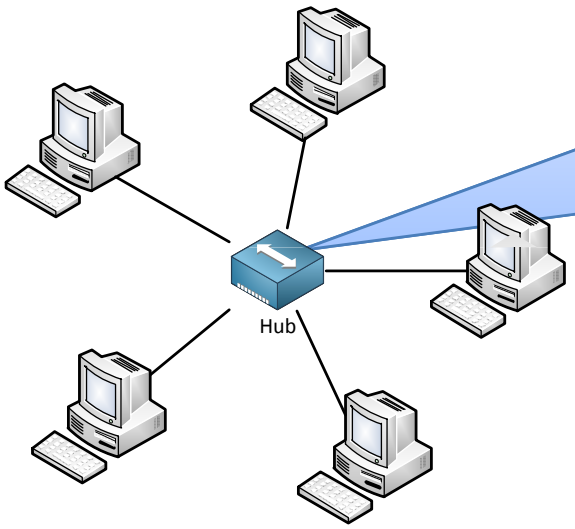

Physical star, logical bus. Shared media. Half duplex. Collisions are an issue.

Each device connected to a hub enters the same collision domain. What happens if hubs are connected to hubs?

FRANKLIN UNIVERSITY

Legacy Ethernet

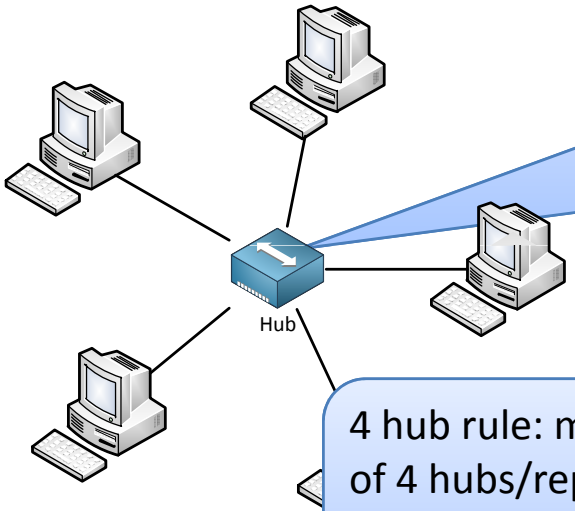

Physical star, logical bus. Shared media. Half duplex. Collisions are an issue.

Each device connected to a hub enters the same domain. What if hubs are ed to hubs?

4 hub rule: maximum of 4 hubs/repeaters between any two devices.

FRANKLIN UNIVERSITY

Switched Ethernet

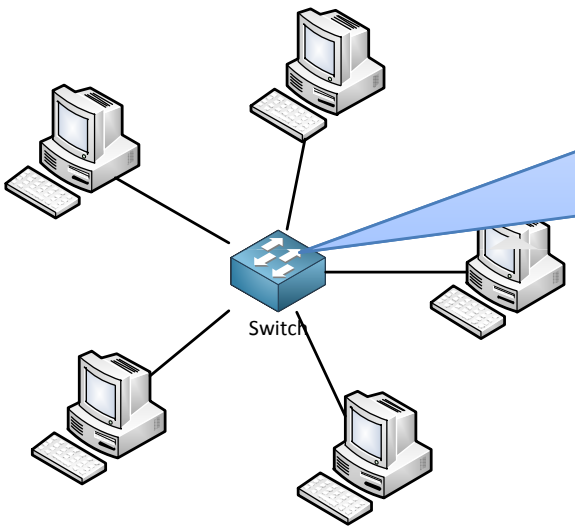

Physical star, logical point-to-point. No shared media. Full duplex. Collisions don't happen.

FRANKLIN UNIVERSITY

Switched Ethernet

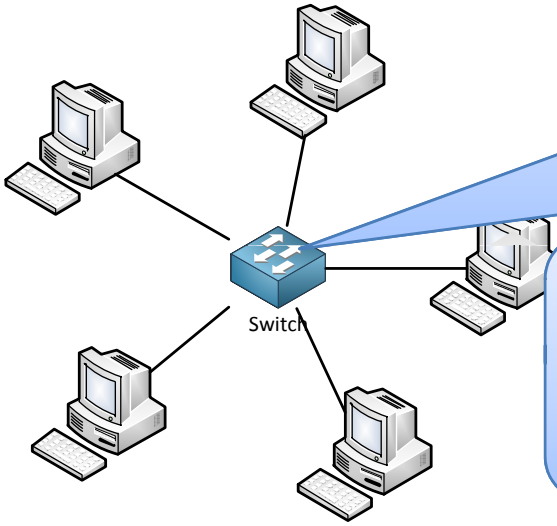

Physical star, logical point-to-point. No shared media. Full duplex. Collisions don't happen.

Switches keep a table of MAC addresses and ports so it knows which to send out on.

Switched Ethernet

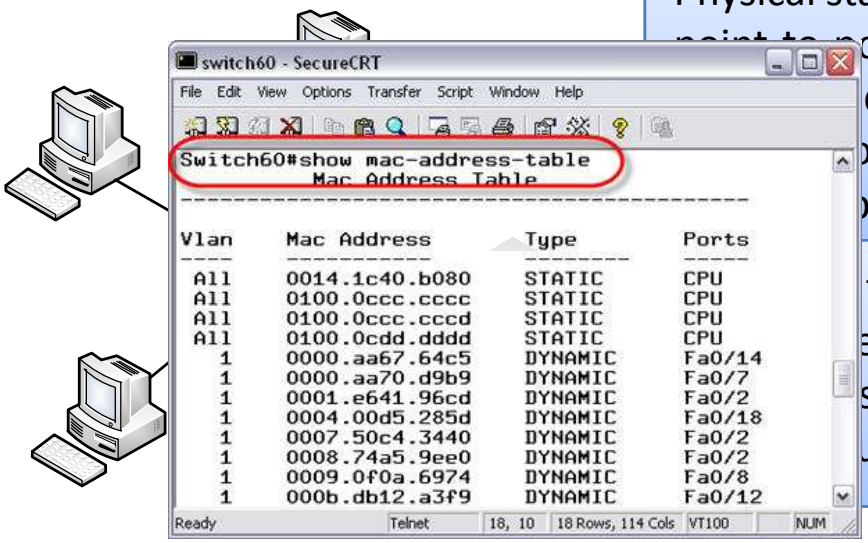

Physical star, logical
oint. No
dia. Full
llisions
en.

table
es and
s
ut on.

Source: http://www.petri.co.il/images/arp_cache_03.jpg

Switched Ethernet

- 5 operations on switches
  - Learning
  - Aging
  - Flooding
  - Selective forwarding
  - Filtering

Switched Ethernet

- 5 operations on switches
  - Learning
  - Aging
  - Flooding
  - Selective forwarding
  - Filtering

As a frame arrives on port X, examine the source MAC address A. Update an internal table so it knows that A is off of port X.

Switched Ethernet

- 5 operations on switches
  - Learning
  - Aging
  - Flooding
  - Selective forwarding
  - Filtering

Table entries more than a certain number of seconds old are discarded. Why?

Switched Ethernet

- 5 operations on switches
  - Learning
  - Aging
  - Flooding
  - Selective forwarding
  - Filtering

If a frame arrives for a destination address that doesn't exist in the table, send it out *all* ports *except* the one it arrived on.

Switched Ethernet

- 5 operations on switches
  - Learning
  - Aging
  - Flooding
  - Selective forwarding
  - Filtering

If a frame arrives for a destination that does exist in the table, then send it out only that specific port.

Switched Ethernet

- 5 operations on switches
  - Learning
  - Aging
  - Flooding
  - Selective forwar
  - Filtering

Don't forward packets back to the same port. Drop corrupt frames. Drop according to MAC address filtering (configured).

Ethernet Addresses

- 48 bit address in 2 24-bit fields
  - First 24 bits: Organizational Unique Identifier
    - Ex: Intel has 009027, 00A083, 00A0C9, 00AA00…
  - Last 24 bits: Vendor assigned (unique within vendor). A kind of "serial number."

Ethernet Unicast

- Sending from one device to another on the same network.

Preamble	Destination	Source	Type	Data	CRC
8 bytes	6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes

Put destination address here

Put source address here

Ethernet Broadcast

- Sending from one device to all other devices on the same network.

Preamble	Destination	Source	Type	Data	CRC
8 bytes	6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes

Put FF:FF:FF:FF:FF:FF here

Put source address here

Ethernet Broadcast

- Sending from one device to all other devices on the same network.

Preamble	Destination	Source	Type	Data	CRC
8 bytes	6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes

Put FF:FF:FF:FF:FF:FF here

Put source address here

Switches must send this out every port.

FRANKLIN UNIVERSITY

Ethernet Multicast

- Sending from one device to all subscribed devices on the same network.

Preamble	Destination	Source	Type	Data	CRC
8 bytes	6 bytes	6 bytes	2 bytes	46-1500 bytes	4 bytes

Put 01:00:5E:IP:IP:IP here

Put source address here

The 23-bit IP multicast address goes here.

FRANKLIN UNIVERSITY

Ethernet Timing

- "How long is a nanosecond?"
  - Time it takes light to travel 30 cm.
  - Time it takes an electrical signal to travel 20 cm.

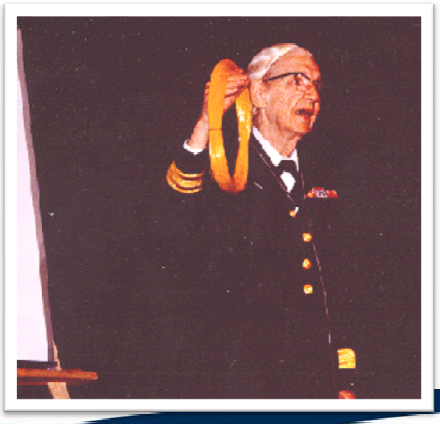

FRANKLIN UNIVERSITY

Ethernet Timing

- Bit time
  - How much time does it take to put a bit on the media? Or sense a bit on the media?

Speed	Bit time (ns)
10 Mbps	100
100 Mbps	10
1 Gbps	1
10 Gbps	0.1

Ethernet Timing

- Bit time
  - How much time does it take to put a bit on the media? Or sense a bit on the media?

$$\frac{10^7 \text{ bits}}{1 \text{sec}} * \frac{1 \text{sec}}{10^9 \text{ ns}} = 1 \text{ bit/100ns}$$

Speed	Bit time (ns)
10 Mbps	100
100 Mbps	10
1 Gbps	1
10 Gbps	0.1

FRANKLIN UNIVERSITY

Ethernet Timing

- Bit time
  - How much time does it take to put a bit on the media? Or sense a bit on the media?

Speed	Distance (m) for 1 bit
10 Mbps	20 m
100 Mbps	2 m
1 Gbps	20 cm
10 Gbps	2 cm

Ethernet Timing

- Bit time
  - How much time does
    media? Or sense a bit on

$$\frac{1\,\text{bit}}{100\,\text{ns}} * \frac{1\,\text{ns}}{20\,\text{cm}} * \frac{100\,\text{cm}}{1\,\text{m}} = 1\,\text{bit/20 m}$$

Speed	Distan (m) for 1 bit
10 Mbps	20 m
100 Mbps	2 m
1 Gbps	20 cm
10 Gbps	2 cm

FRANKLIN UNIVERSITY

Ethernet Timing

- Bit time
  - How much time does it take to put a bit on the media? Or sense a bit on the media?

Speed	Distance for smallest frame (m)
10 Mbps	10,240
100 Mbps	1,024
1 Gbps	102
10 Gbps	10

Ethernet Timing

- Bit time
  - How much time does it take to put a bit on the media? Or sense a bit on the media?

$$\frac{512\text{ bit}}{1\text{ bit}} * \frac{20\text{ m}}{} = 10{,}240\text{ m}$$

Speed	Distance (m) smallest frame
10 Mbps	10,240
100 Mbps	1,024
1 Gbps	102
10 Gbps	10

On UTP wire, max lengths are 100 m.

MAC vs. IP addresses

- Layer 3 – host to host communication
  - Want to get data from one host to the other host on the same network. (Q: how do you know if you're on the same network?)
    - No intervening routers, right?
    - So this is direct delivery.
    - Therefore we need to know the destination device's MAC address to properly format the frame.
    - How do we get it?

MAC vs. IP addresses

- Address Resolution Protocol (ARP)
  - Gets packets to hosts on the same network by finding out out the MAC address for direct delivery.
  - How:
    - Sender sends Ethernet broadcast "who has this IP address?"
    - All devices listen, and target host replies
    - Sender stores the MAC address in an ARP cache.

MAC vs. IP addresses

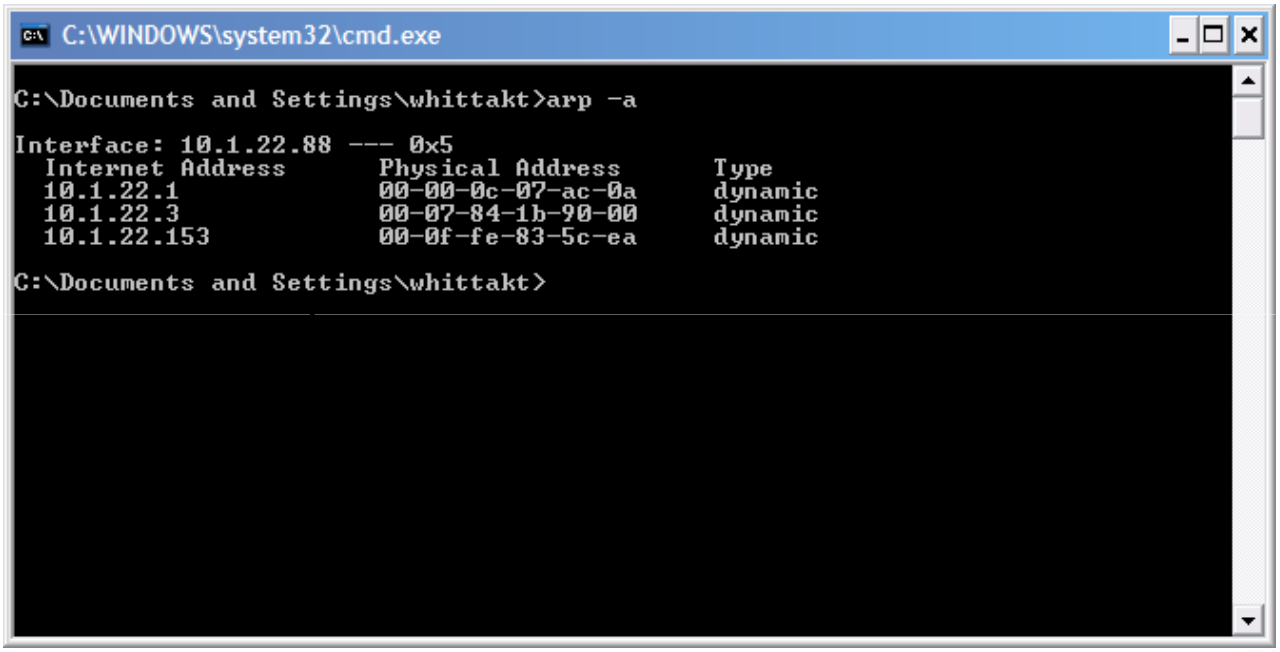

FRANKLIN UNIVERSITY

MAC vs. IP addresses

- ARP Poisoning
  - Attacker fakes ARP replies on the network and then receives all frames destined for someone else.

  OR
  - Attacker spoofs MAC addresses and sends out fake ARP requests (thus poisoning the receiver's ARP cache).

FRANKLIN UNIVERSITY

MAC vs. IP addresses

- Reverse ARP (RARP)
  - Gets an IP address for hosts that are booting up and don't have one yet.
    - Obsolete, replaced by DHCP and BOOTP

Tracing the Full Path

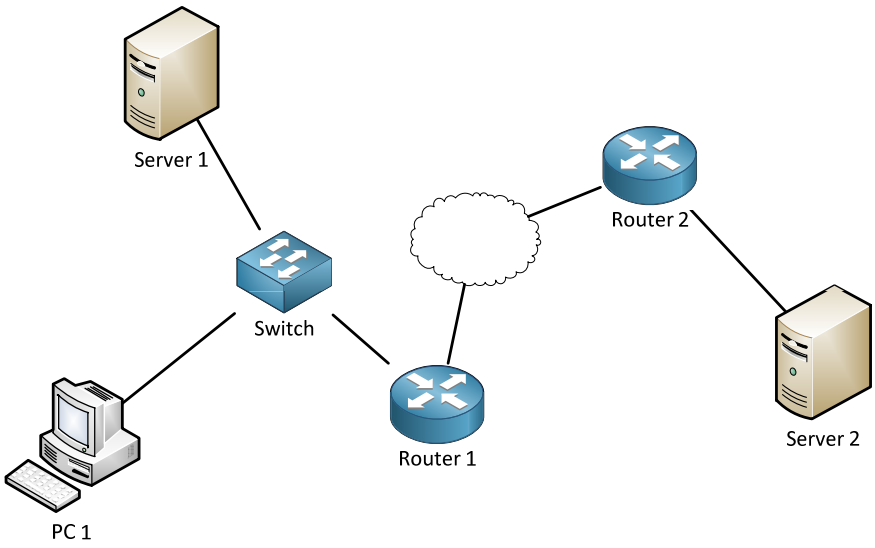

- How does PC 1 talk to Server 1?

Tracing the Full Path

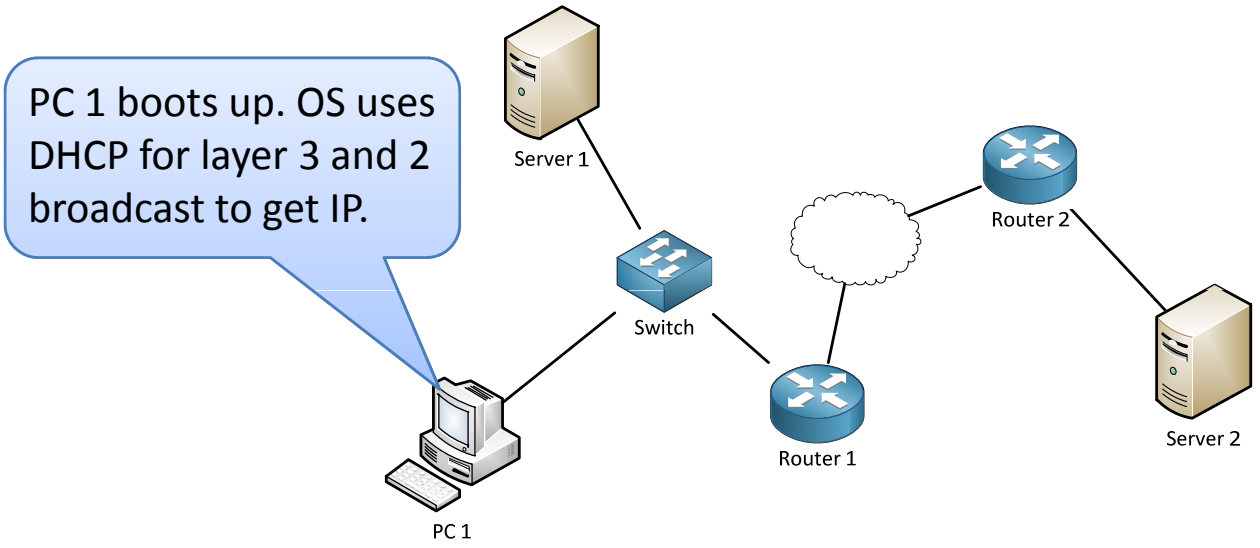

- How does PC 1 talk to Server 1?

Tracing the Full Path

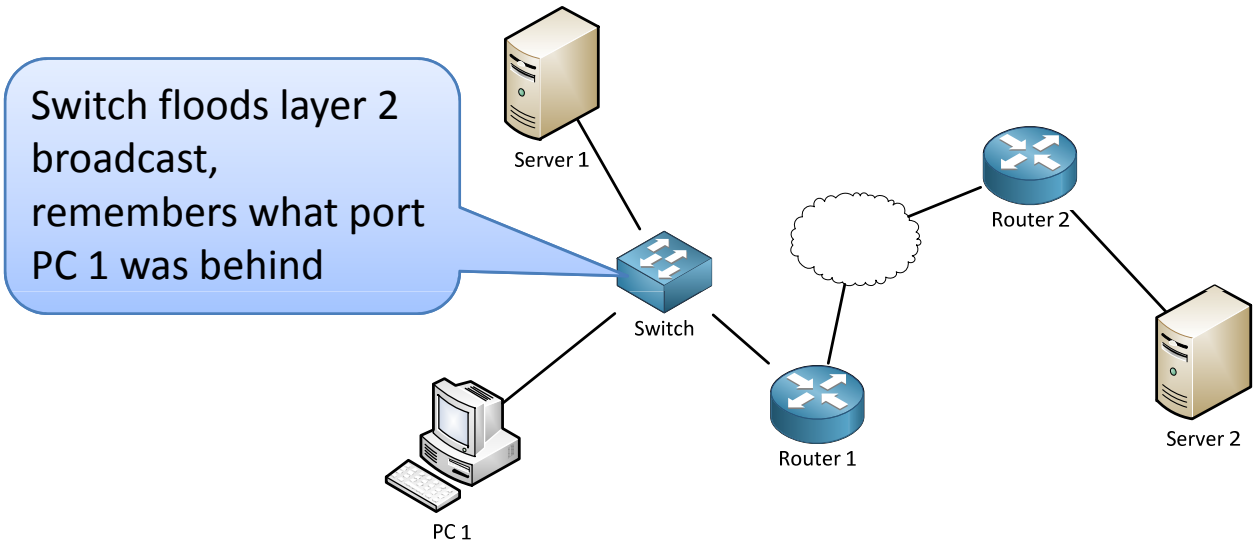

- How does PC 1 talk to Server 1?

Tracing the Full Path

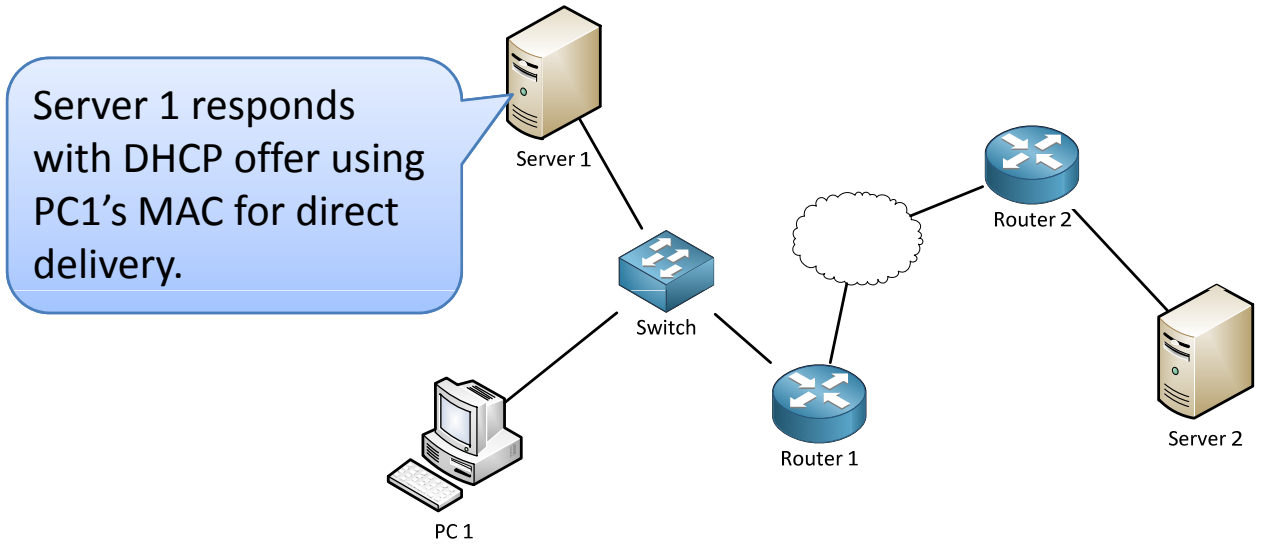

- How does PC 1 talk to Server 1?

Tracing the Full Path

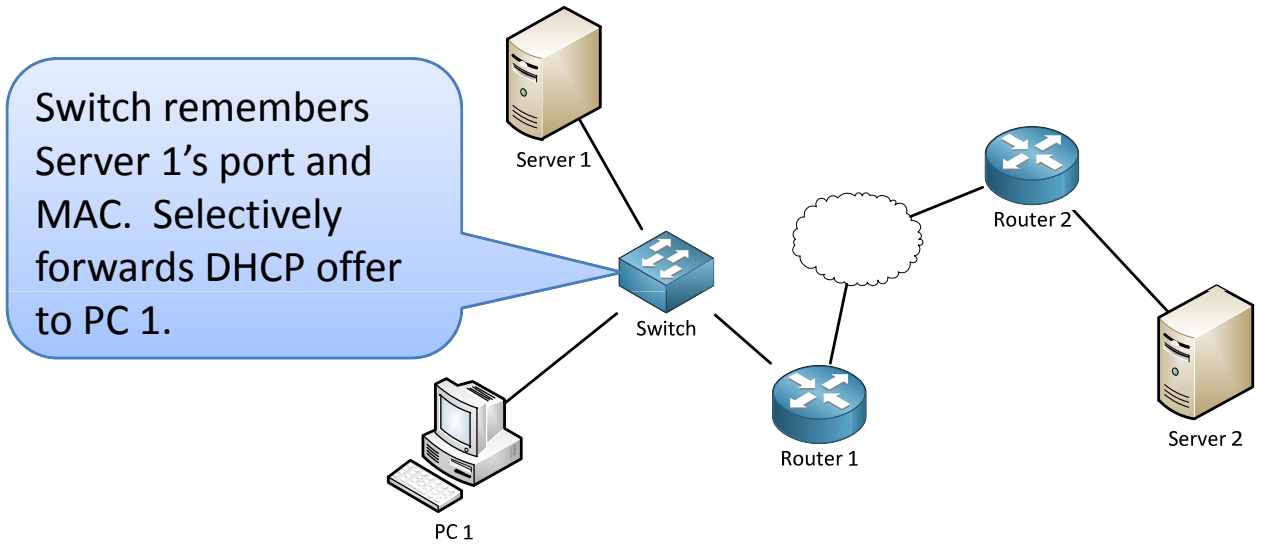

- How does PC 1 talk to Server 1?

Tracing the Full Path

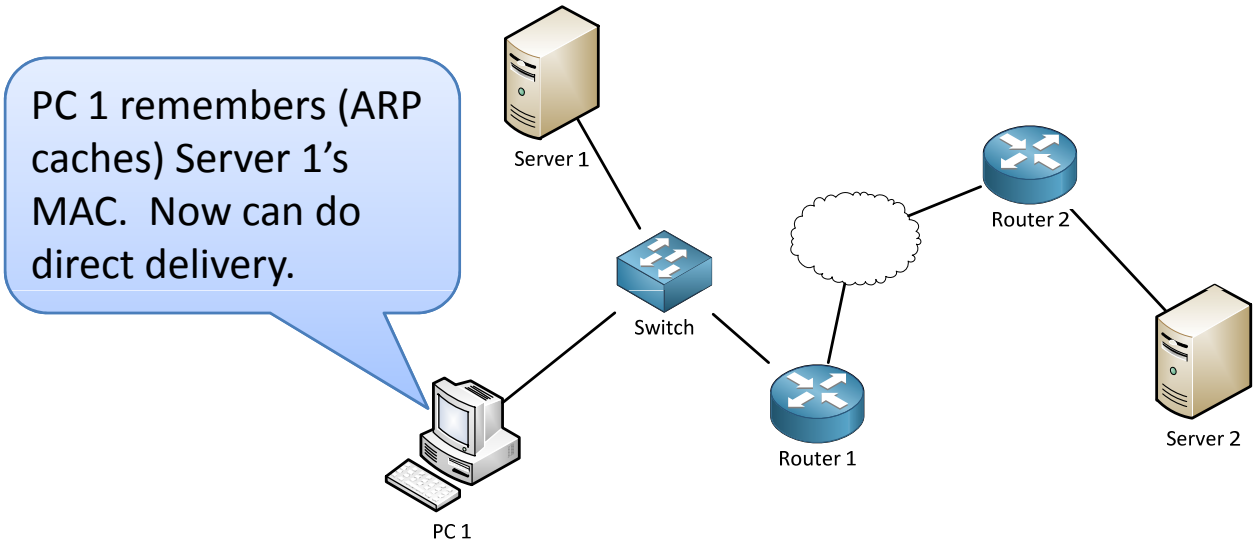

- How does PC 1 talk to Server 1?

Tracing the Full Path

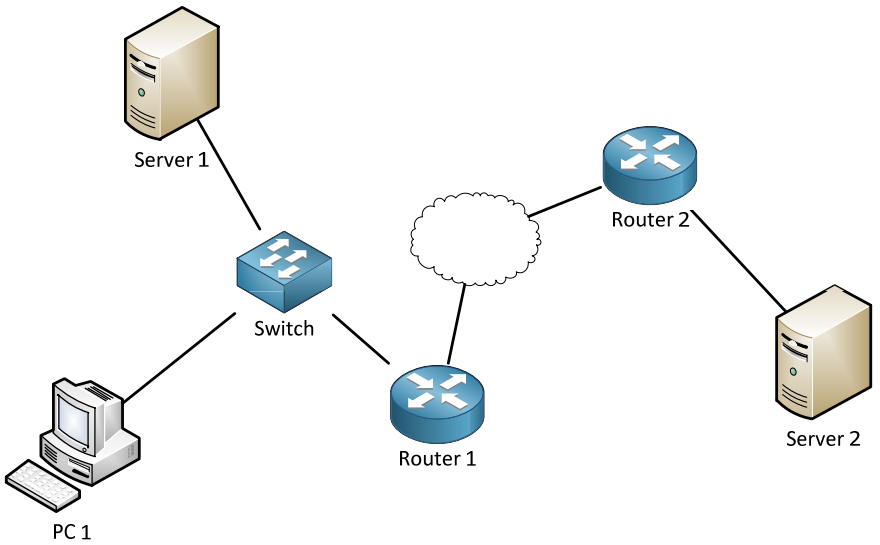

- How does PC 1 talk to Server 2?

Tracing the Full Path

PC 1 uses DNS to look up server 2 (lots hidden here).

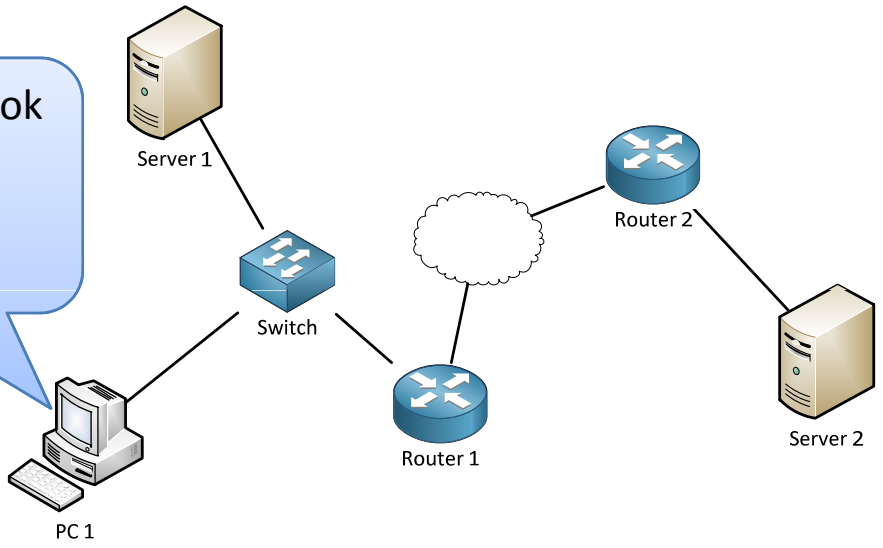

- How does PC 1 talk to Server 2?

Tracing the Full Path

PC 1 compares its network to Server 2 IP. Not on same network. Constructs L3 packet destined for Server 2, but L2 frame destined for Router 1.

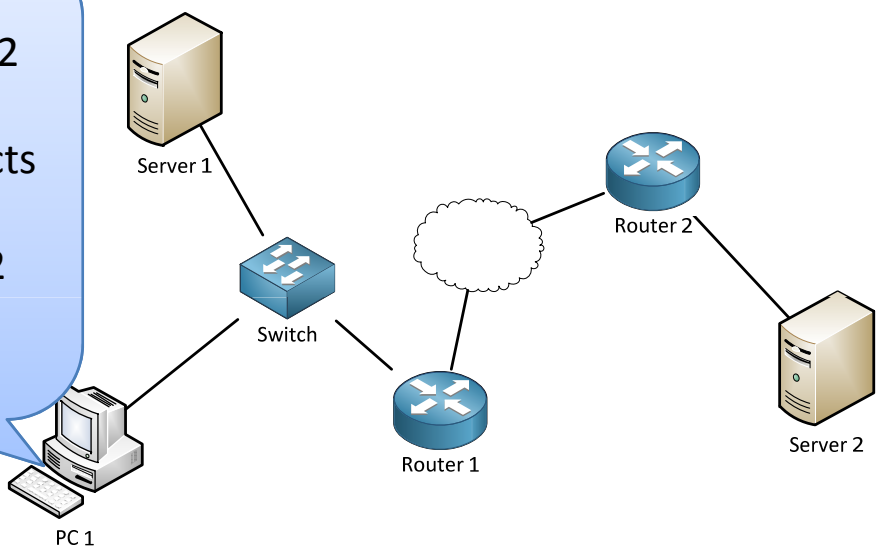

- How does PC 1 talk to Server 2?

Tracing the Full Path

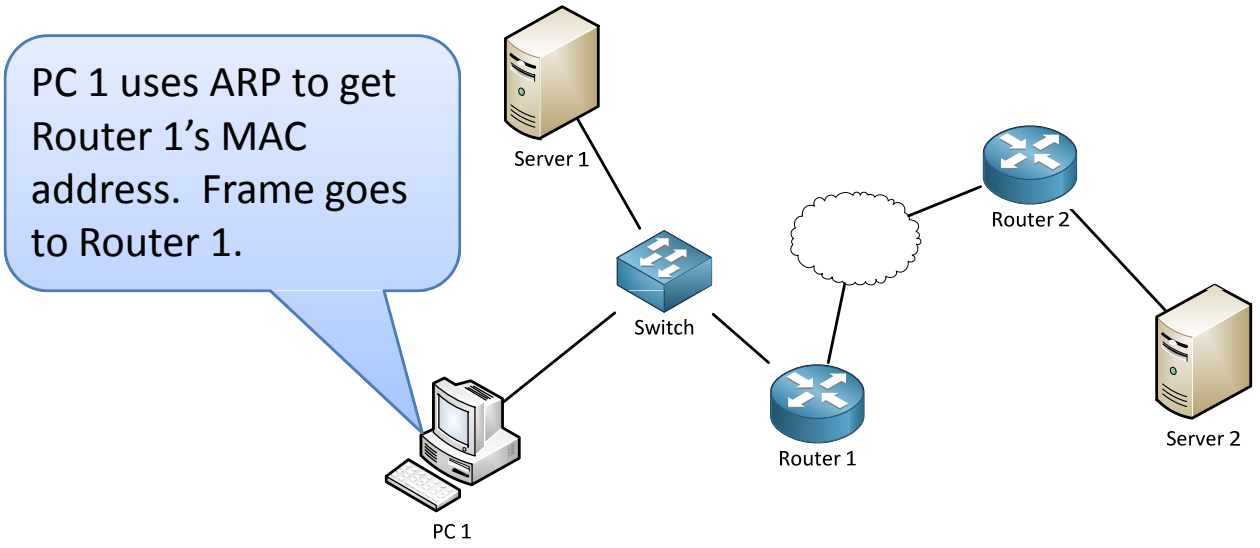

- How does PC 1 talk to Server 2?

Tracing the Full Path

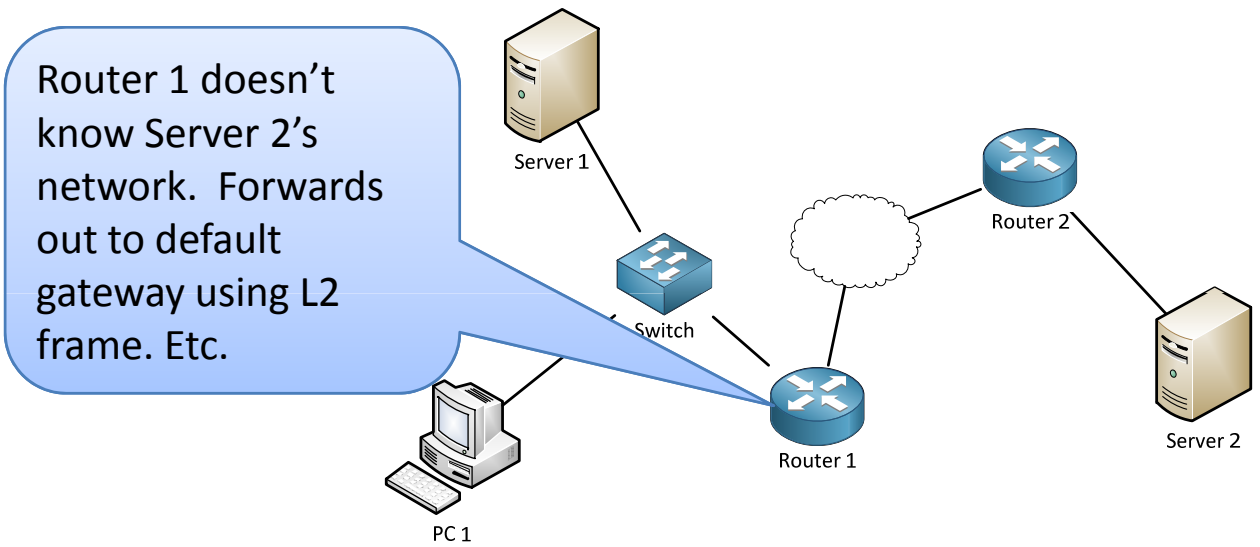

- How does PC 1 talk to Server 2?

LAN and WAN devices

- Hubs
  - Don't use them
- Switches
  - Do use them
- Routers / multilayer switches
  - Connect LANs

Redundancy

- As much as possible, have redundant connections between core routers/switches.

Design – Layered Approach

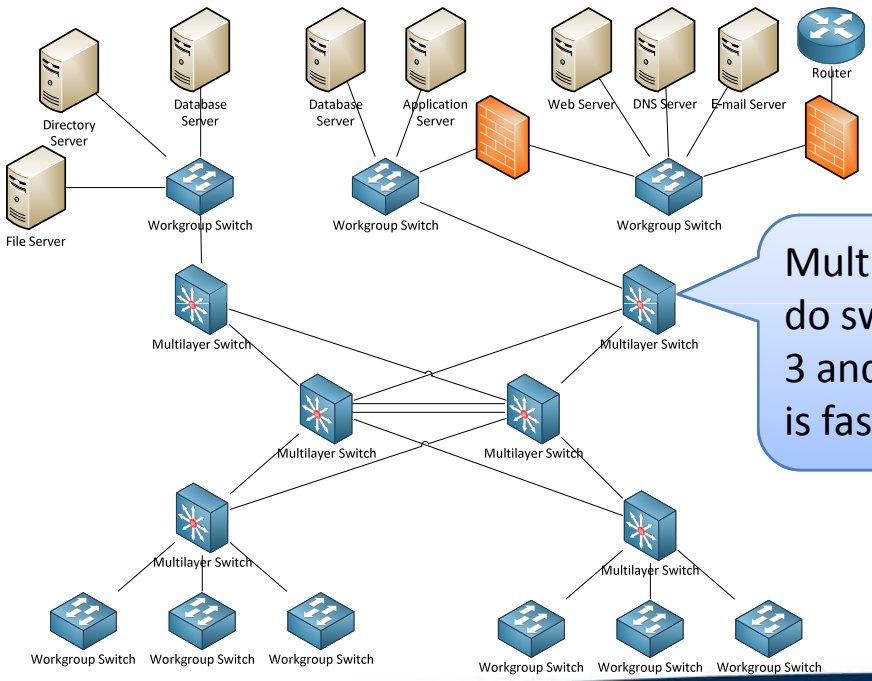

Multilayer switches: do switching at layer 3 and up. Switching is faster than routing.

FRANKLIN UNIVERSITY

Design – Layered Approach

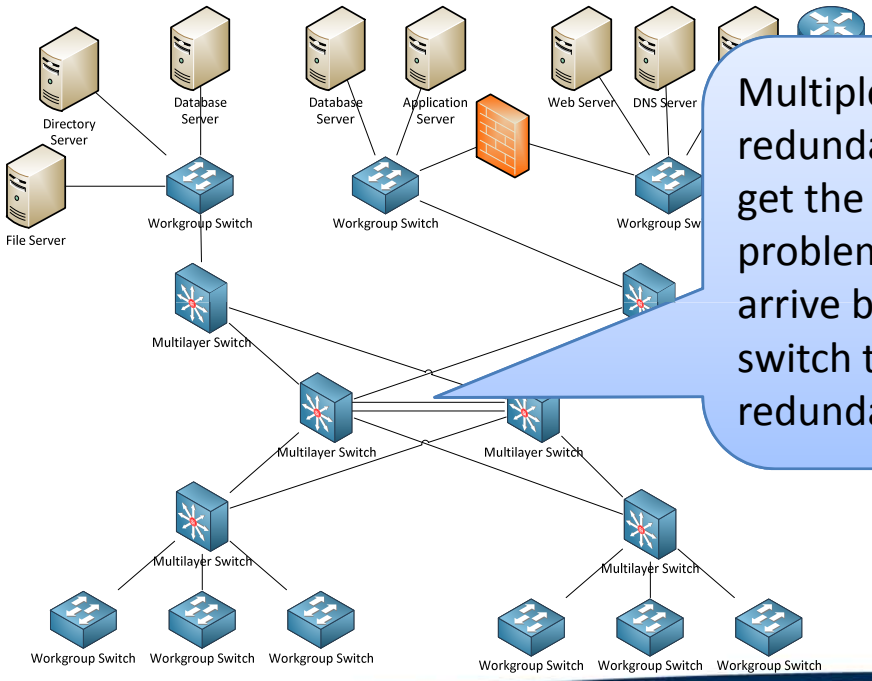

Multiple connections for redundancy, but then we get the flooding problem. L2 broadcasts arrive back at the same switch through the redundant links.

FRANKLIN UNIVERSITY

Design – Layered Approach

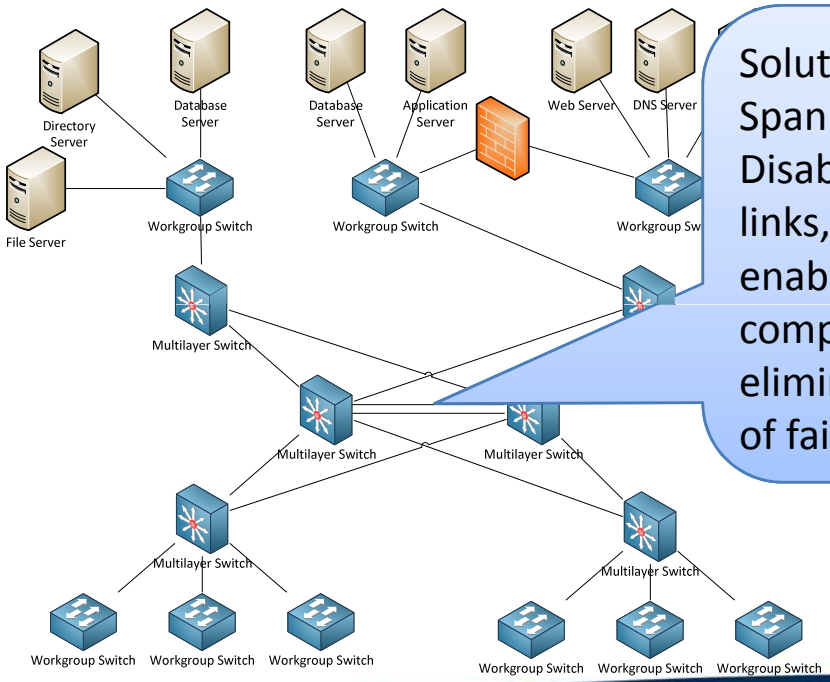

Solution: RSTP – Rapid Spanning Tree Protocol. Disables redundant links, monitors, re-enables as needed. Very complex in practice, but eliminates single points of failure.

FRANKLIN UNIVERSITY

Design – Layered Approach

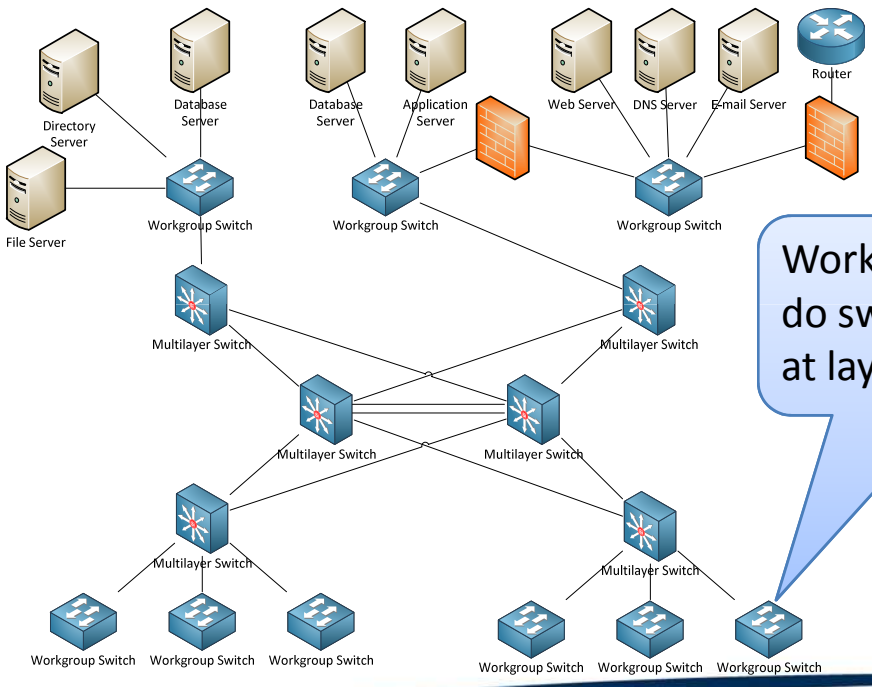

Workgroup switches: do switching strictly at layer 2.

FRANKLIN UNIVERSITY

Design – Layered Approach

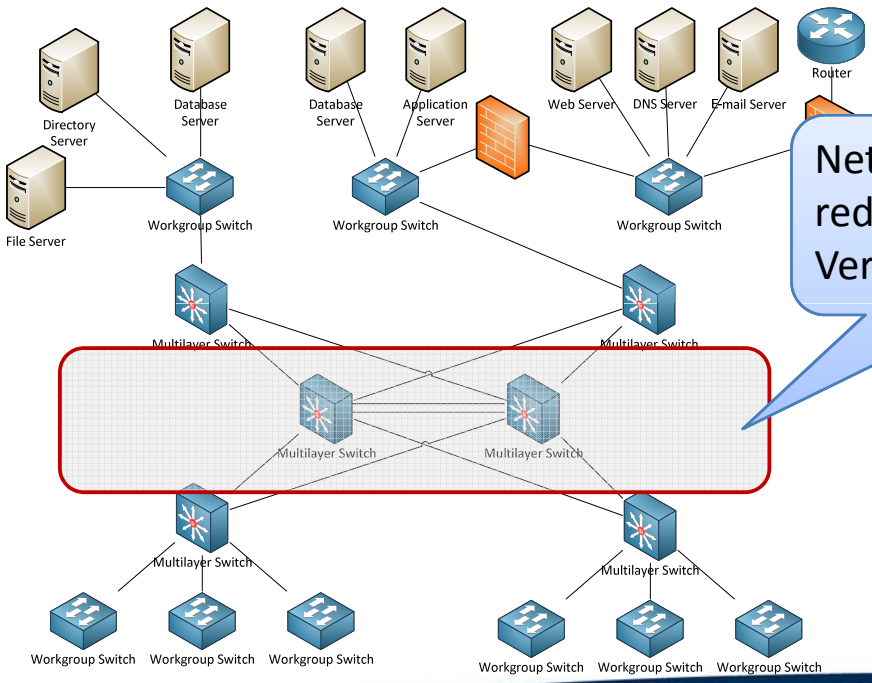

Network core. Highly redundant paths. Very fast/expensive.

FRANKLIN UNIVERSITY

Design – Layered Approach

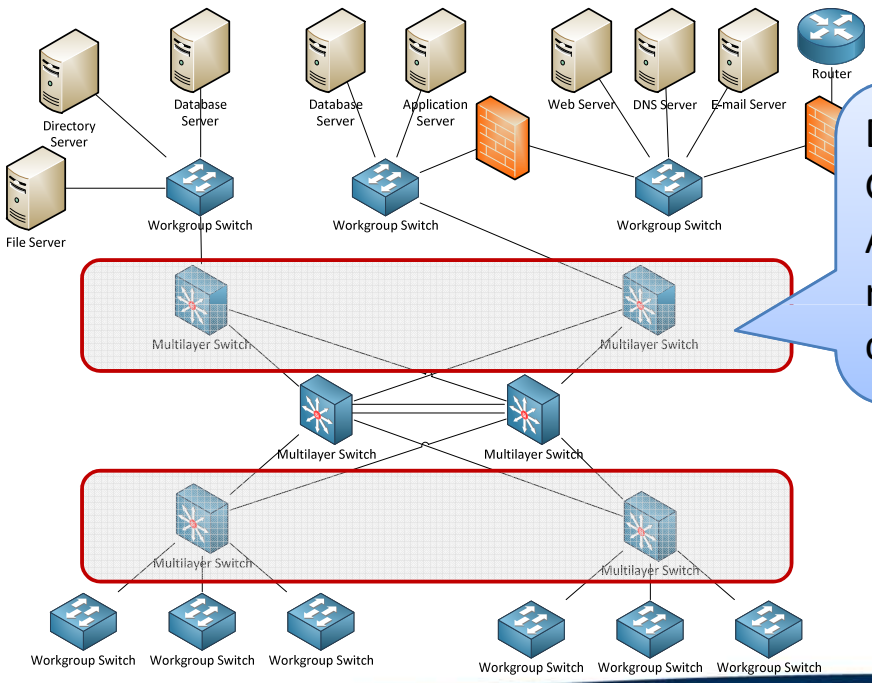

Distribution layer. Gets to buildings. Also highly redundant connections.

FRANKLIN UNIVERSITY

Design – Layered Approach

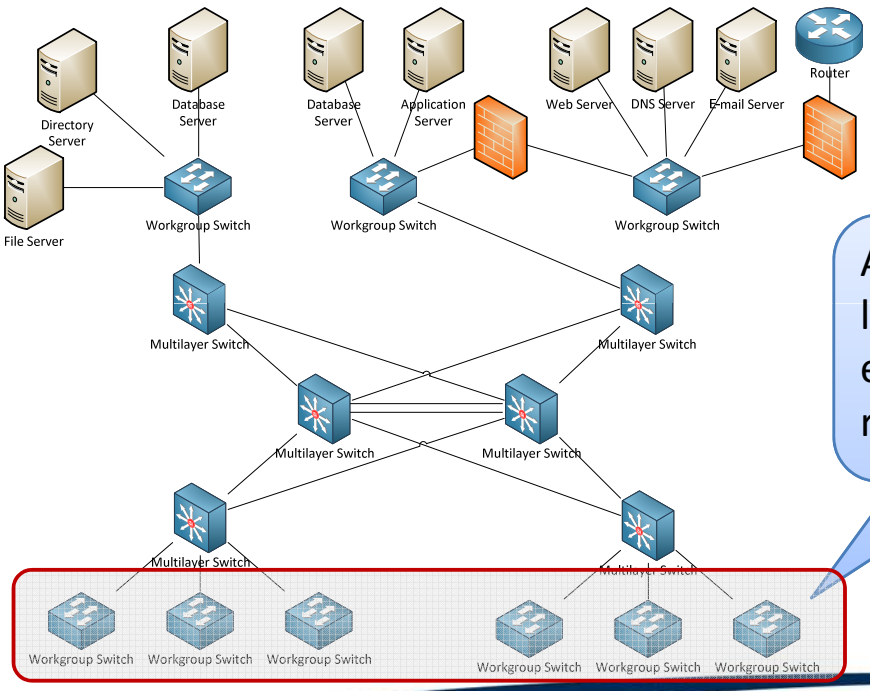

Access (edge) layer. Less expensive. Not redundant.

Design – Layered Approach

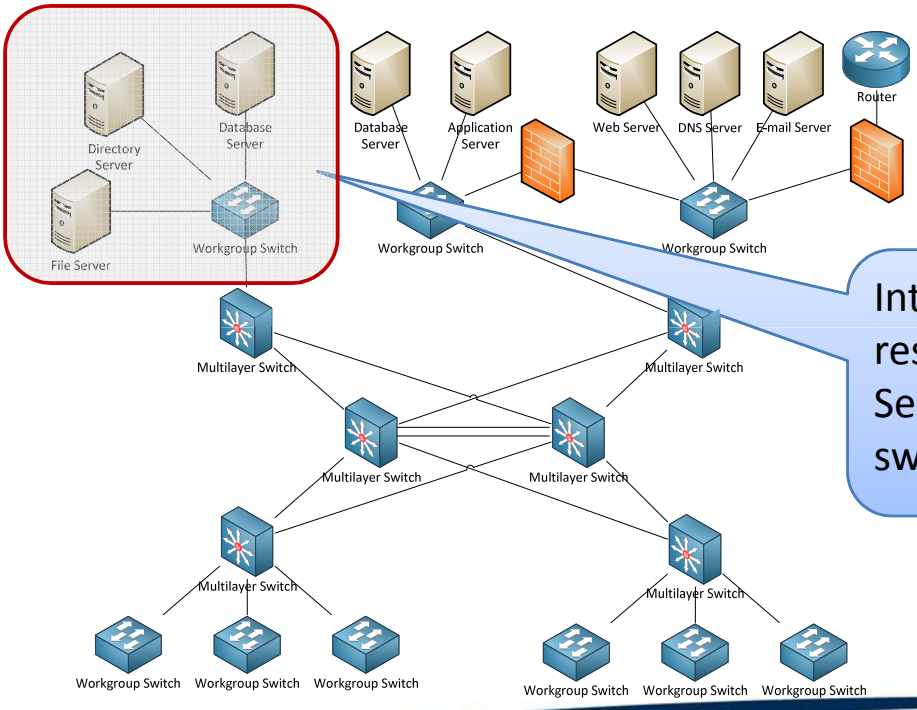

Internal only resources. Security in switches/routers.

Design – Layered Approach

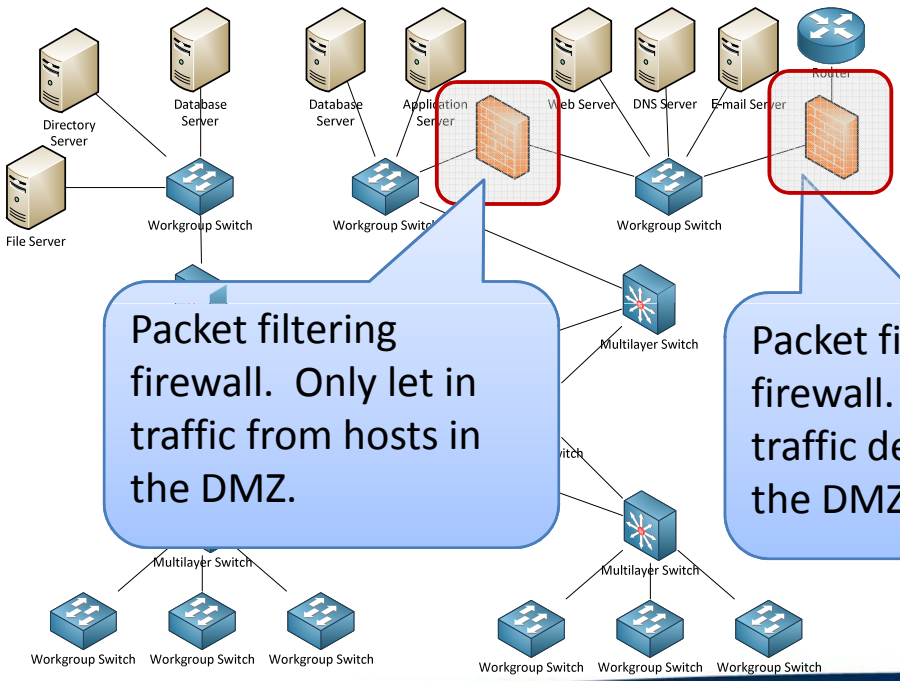

Packet filtering firewall. Only let in traffic from hosts in the DMZ.

Packet filtering firewall. Only let in traffic destined for the DMZ.

FRANKLIN UNIVERSITY

Design – Layered Approach

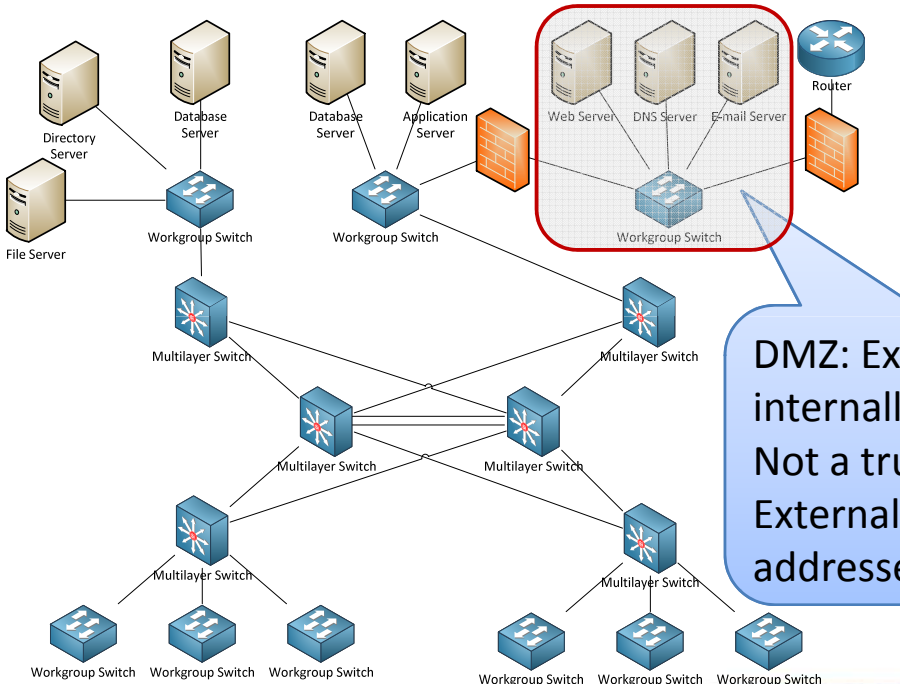

DMZ: Externally and internally accessible. Not a trusted zone. Externally routable IP addresses.

FRANKLIN UNIVERSITY

Design – Layered Approach

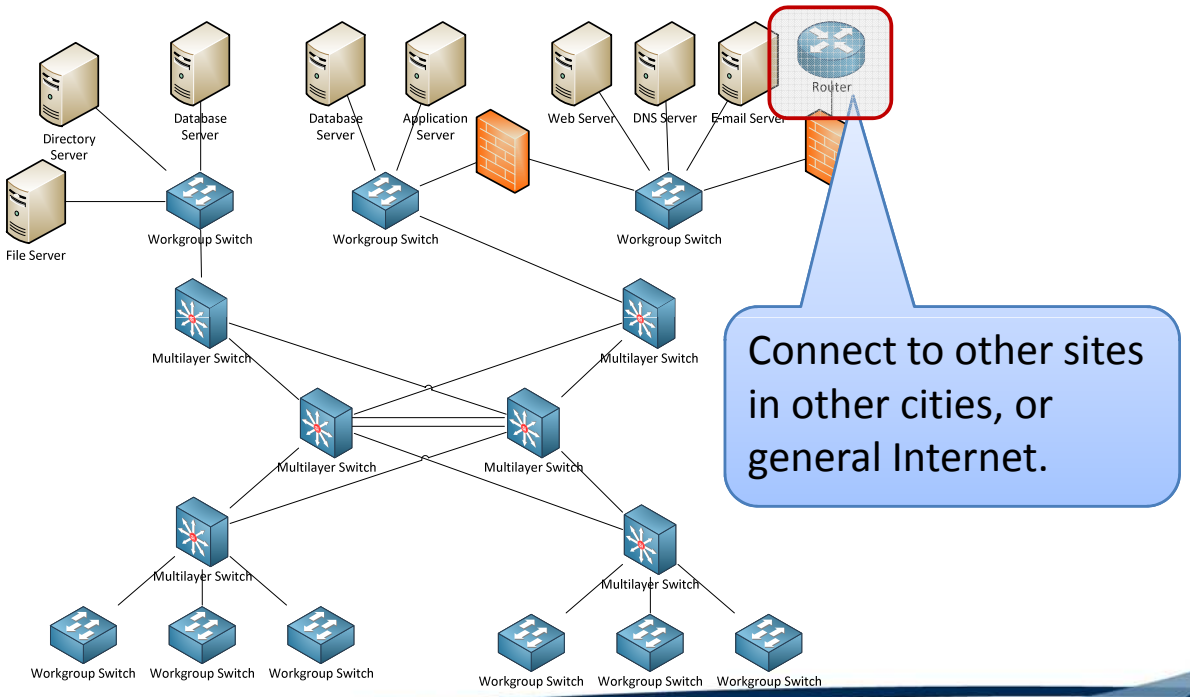

IP Addressing Schemas

- VLSM/CIDR is your friend (as are RFC 1918 addresses)
  - Calculating Subnets (p. 391 – 396)
  - Minimal wasted IP space, but also room to expand.

ITEC 275

- Picks up with extensive design principles
  - Layered networks
  - IP addressing
  - Business goals/constraints
  - Security and management
  - Enterprise networks

This Week's Outcomes

- Categorize Ethernet functionality in layers 1 and 2 of the OSI model ✓
- Describe the purpose and operation of ARP. ✓
- Simulate collision resolution on Ethernet. ✓
- List cabling types, standards, and ports for common physical connections. ✓

This Week's Outcomes

- Categorize Ethernet functionality in layers 1 and 2 of the OSI model
- Describe the purpose and operation of ARP
- Simulate collision resolution on Ethernet.
- List cabling types, standards, and ports for common physical connections.

Actually done last week.

FRANKLIN UNIVERSITY

Self Quiz

- Name and define the five operations that switches perform.
- What were (are) the problems with hubs that switches solved? How do switches accomplish that?
- What is ARP? Why is it used?
- How do switches know where to forward frames?

FRANKLIN UNIVERSITY

Self Quiz

- What are the challenges in ever increasing network speeds?
- What is the advantage of a layered approach to physical network design (core, distribution, access)?
- What does redundancy in networking buy for you?

Due this week

- Homework 5
- Participation 6

FRANKLIN UNIVERSITY

Next week

- Final Exam!
- Lab 3 due
- FranklinLive open office hours (answer whatever questions you have).

Final Exam FAQs

- "Is X on the final?"
  - If we talked about it, did an assignment over it, or it was in the book, yes, it is fair game.
- "Can you narrow what I need to study?"
  - Yes. Look at the following:
    - Outcomes for the course and each module in it.
    - Self Quiz questions on the slides.
    - The "blue pages" in your book (chapter introductions, summaries, and questions).

Final Exam FAQs

- "Can I bring a calculator for the base conversions?"
  - Yes.
- "What will the format of the exam be like?"
  - Short answer / essay (similar to the homework)
  - Problems (subnetting, routing, switching, etc.)
- "How much time will I have?"
  - 2 hours (designed to take 90 minutes).

Q & A

- Questions, comments, concerns?

FRANKLIN UNIVERSITY