Basic Networking Crash Course

2017 RVTEC Meeting

University of Minnesota Duluth, Large Lakes Observatory

October 27th - Duluth, Minnesota

Britton Anderson, Office of Information Technology
University of Alaska

Celebrating a century 1917–2017

Objectives

- Describe basic networking components and operations
- Explain the fundamentals of network communication
- Define common networking terms
- Analyze the OSI Model
- Identify the functions of various network services
- Describe functions and challenges of shipboard networks
- Overview of optimizing TCP throughput

An Overview of Computer Concepts

- Most of the devices you encounter when working with a network involve a computer
- Most obvious devices are workstations and network servers
  - These run operating systems such as Windows, Linux, UNIX, and Mac OS
- Also includes routers and switches
  - These are specialized computers used to move data from computer to computer and network to network

Network Components

- Hardware components
  - *Network interface card*—A NIC is module that's built in to the motherboard, or plugged into the motherboard's expansion slot and provides a connection between the computer and the network.
  - *Network medium*—A cable that plugs into the NIC and makes the connection between a computer and the rest of the network. Network media can also be the air waves, as in wireless networks.
  - *Interconnecting*—Interconnecting devices allow two or more computers to communicate on the network without having to be connected directly to one another.

Fundamentals of Network Communication

- A computer network consists of two or more computers connected by some kind of transmission medium, such as a cable or air waves.
- In order to access the Internet, a computer has to be able to connect to a network.

A Typical Home Network

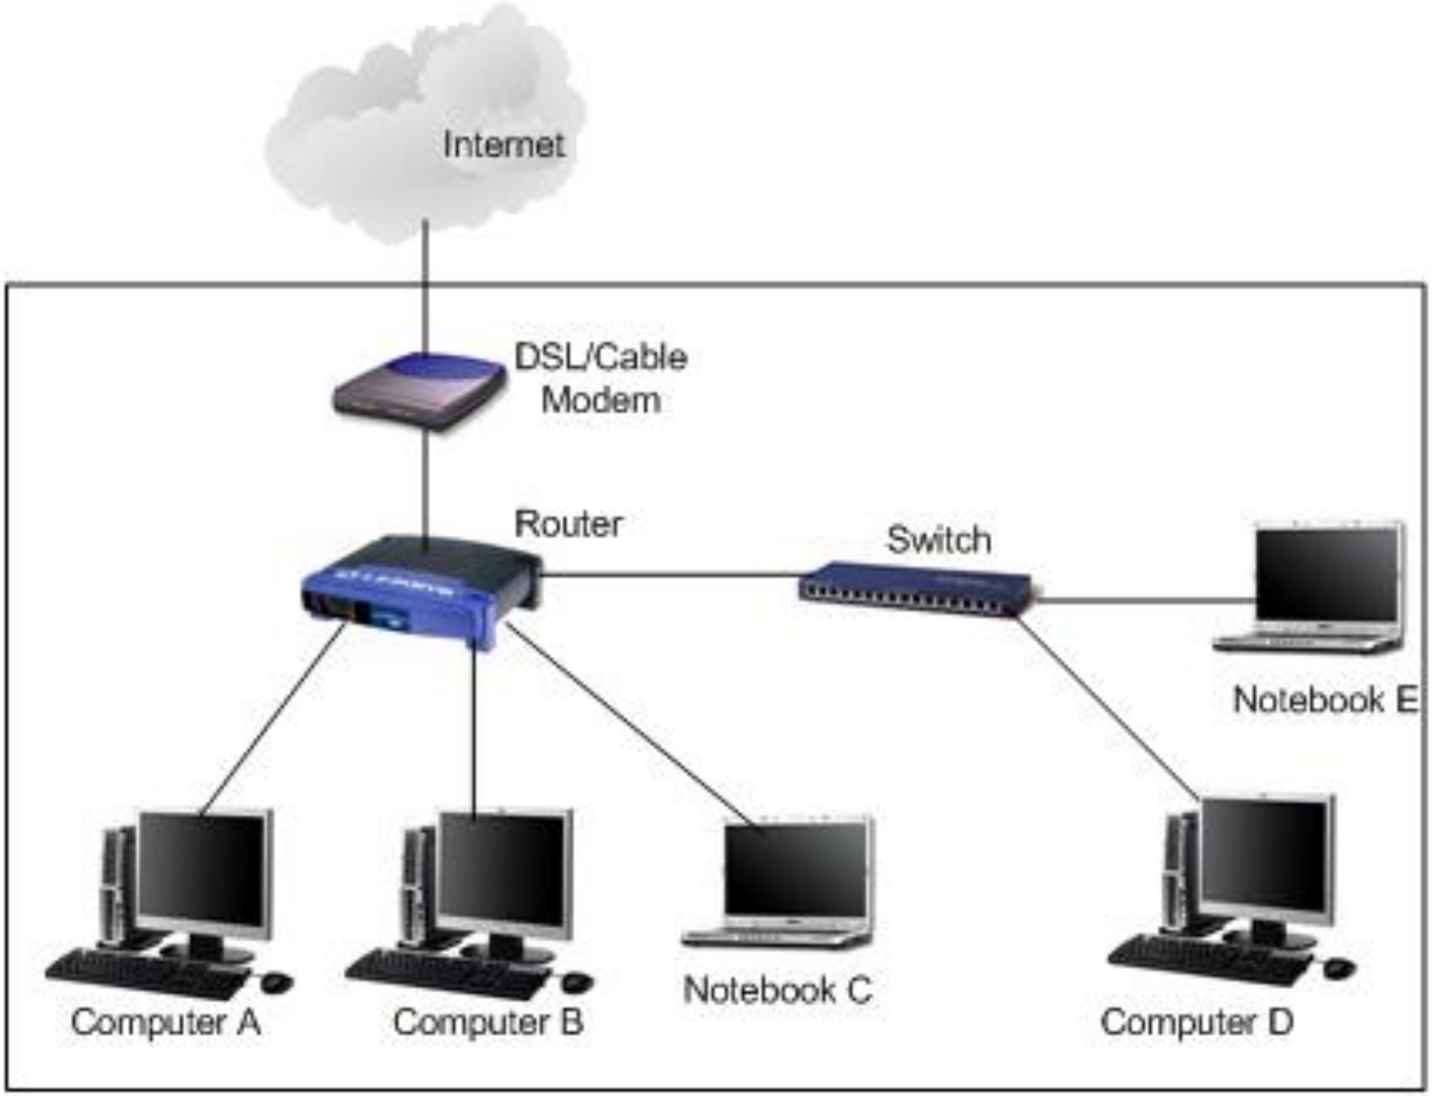

A Research Vessel Network

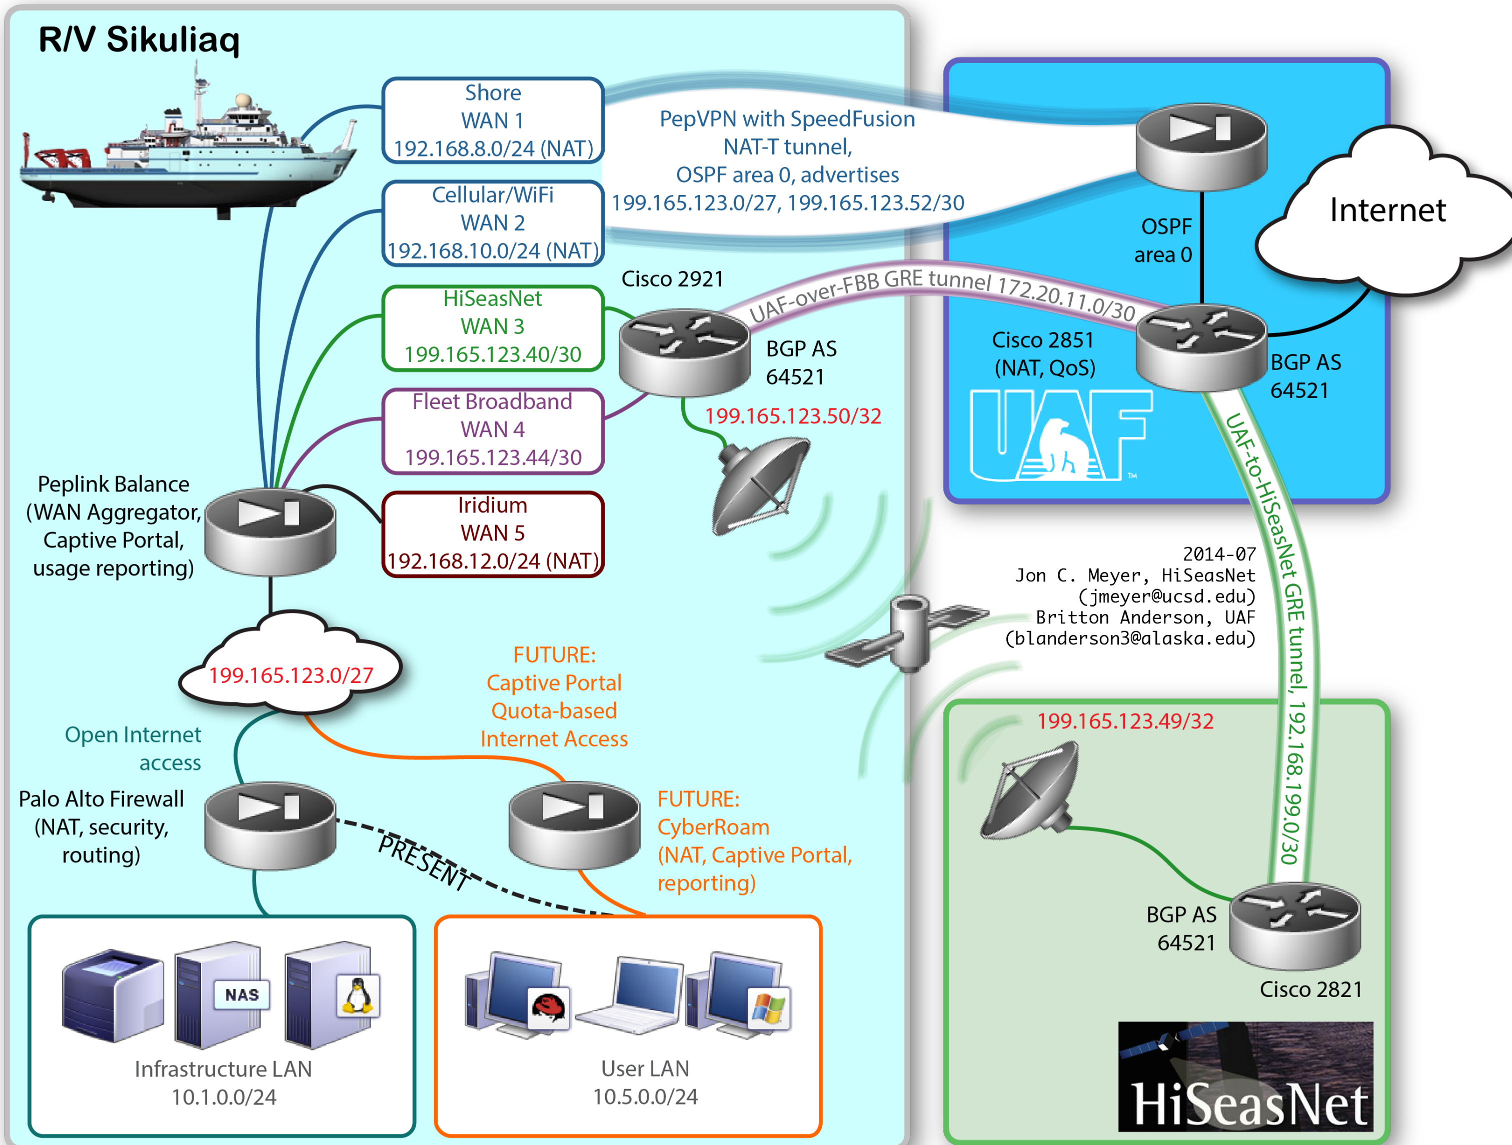

Network Components

- Software Components
  - *Network clients and servers*—**Network client software** requests information that's stored on another network computer or device. **Network server software** allows a computer to share its resources by fielding resource requests generated by network clients.
  - *Protocols*—**Network protocols** define the rules and formats a computer must use when sending information across the network. Think of it as a language that all devices on a network understand.
  - *Network interface*—The network access interface that transmits and receives data from the network medium

Layers of the Network Communication Process

- Each step required for a client to access network resources is referred to as a “layer”
- Each layer has a task and all layers work together

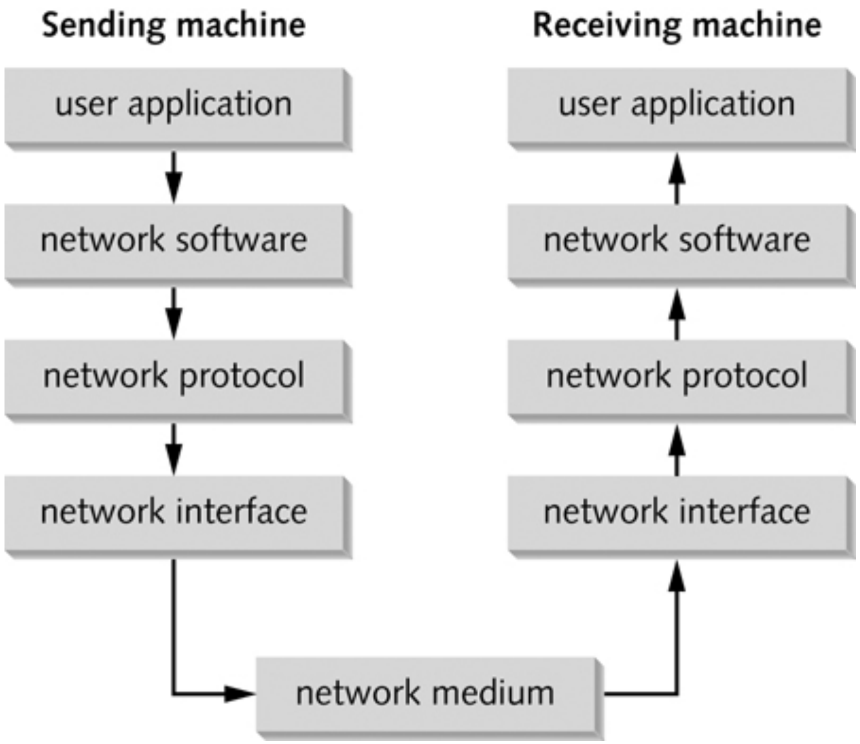

Network Terms

- Every profession has its own language and acronyms
- Need to know the language of networks to be able to properly communicate needs and issues off ship.

LANs, Internetworks, WANs

- Local area network (LAN) – small network, limited to a single collection of machines and connected by one or more interconnecting devices in a small geographic area

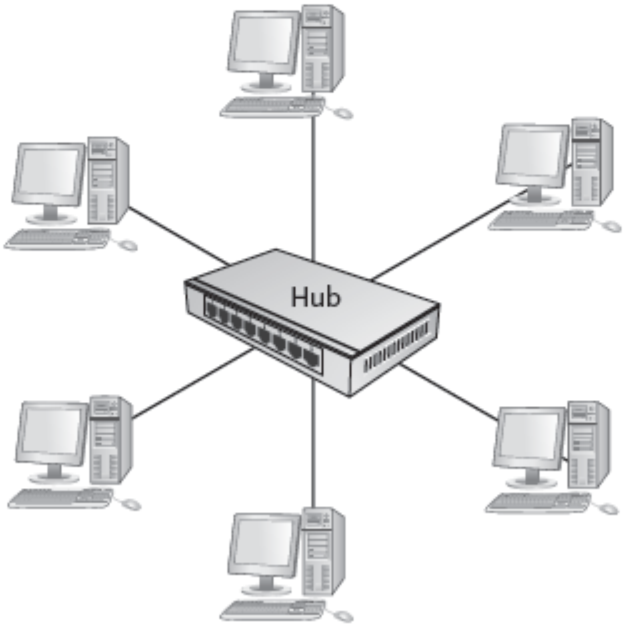

Figure 1-13 A LAN with computers interconnected by a hub

LANs, Internetworks, WANs

- An internetwork is a networked collection of LANs tied together by devices such as routers
- Reasons for being:
  - Two or more groups of users and their computers need to be logically separated but still need to communicate
  - Number of computers in a single LAN has grown and is no longer efficient
  - The distance between two groups of computers exceeds the capabilities of most LAN devices

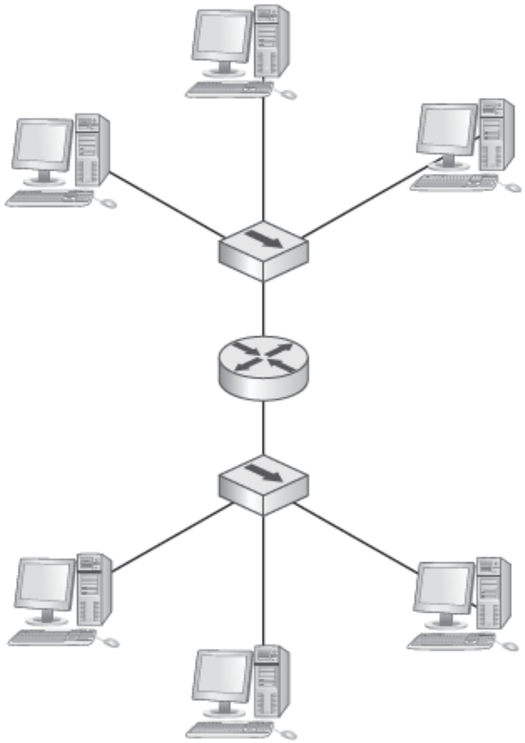

Figure 1-16 An internetwork with two LANS connected by a router

How Two Computers Communicate

- TCP/IP is the most common protocol (language) used on networks
- TCP/IP uses 2 addresses to identify devices on a network
  - Logical address (called IP address)
  - Physical address (called MAC address)
- Just as a mail carrier needs an address to deliver mail, TCP/IP needs an address in order to deliver data to the correct device on a network
- Think of the Logical address as a zip code and the Physical address as a street address

Packets and Frames

- Computers transfer information across networks in shorts bursts of about 1500 bytes of data
- Data is transferred in this way for a number of reasons:
  - The pause between bursts might be necessary to allow other computers to transfer data during pauses
  - The pause allows the receiving computer to process received data, such as writing it to disk
  - The pause allows the receiving computer to receive data from other computers at the same time
  - The pause gives the sending computer an opportunity to receive data from other computers and to perform other processing tasks
  - If an error occurs during transmission of a large file, only the chunks of data involved in the error have to be sent again, not the entire file

Packets

- Chunks of data sent across the network are usually called packets or frames, with packets being the more well-known term
- **Frames are packets** with source and destination MAC addresses, and error checking added to it
- Using the USPS analogy, you can look at a packet as an envelope containing the data that has a street address on it.

Frames

- A **frame** is outside a packet with the source and destination MAC addresses added to it
- The frame is built with the MAC addresses on the beginning and an error-checking code on the end. In between them is the packet
- A frame is like the mail carrier moving your envelope and your letter from place to place
- The process of adding IP addresses and MAC addresses to packets and frames to chunks of data is called **encapsulation**
- Information added to the front of the data is called a **header** and information added to the end is called a **trailer**

Communication Between Two Computers

1. A user at Comp A types ping 10.1.1.2 at a command prompt
2. The network software creates a ping message
3. The network protocol packages the message by adding IP address of sending and destination computers and acquires the destination computer's MAC address
4. The network interface software adds MAC addresses of sending and destination computers and sends the message
5. Comp B receives message, verifies that the addresses are correct and then sends a reply to Comp A using Steps 2 – 4

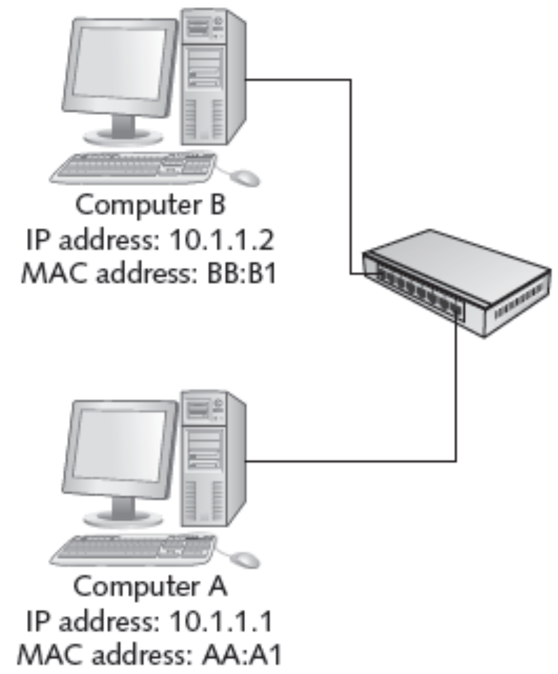

Figure 1-7 Communication between two computers

Clients and Servers

- A **client** can be a workstation running a client OS or it can also refer to the network software on a computer that requests network resources from a server
- The word “client” is usually used in these three contexts:
  - Client operating system: The OS installed on a computer
  - Client computer: Primary role is to run user applications and access network resources
  - Client software: The software that requests network resources from server software running on another computer

Clients and Servers

- A computer becomes a **server** when software is installed on it that provides a network service to client computers
- The term "server" is also used in three contexts:
  - Server operating system: When the OS installed on a computer is designed mainly to share network resources and provide other network services
  - Server computer: When a computer's primary role in the network is to give client computers access to network resources and services
  - Server software: Responds to requests for network resources from client software running on another computer

Network Models

- A **network model** is a framework to conceptually divide network functions progressively in a logical reference.
- Two major models exist
  - **TCP/IP Model** Often referred to as the DOD model since it was originally designed for them
  - **OSI Network Model** developed by the International Standards Organization as a standard called the Open Systems Interconnection (OSI) reference model.

Model Comparison

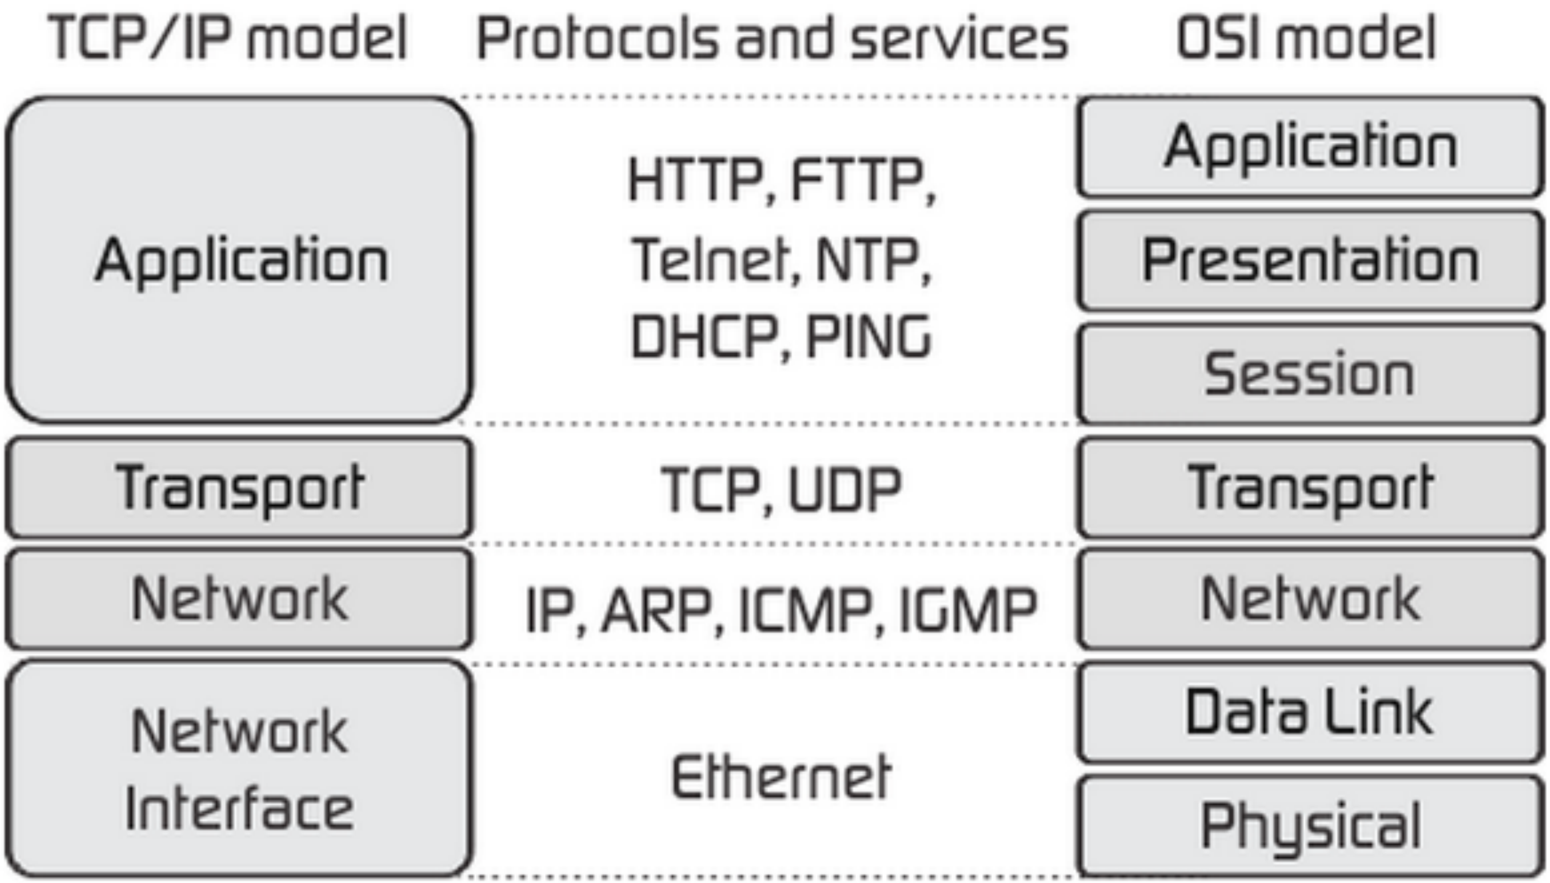

Layer 1 – The Physical Layer

- In networking, data is transmitted in bits
  - A pulse of 5 volts of electricity can represent a 1 bit and a pulse of 0 volts can represent a 0 bit
  - With fiber-optic cable, a 1 bit is represented by the presence of light and a 0 bit by the absence of light
  - WiFi transmits and receives radio wave pulses in either 2.4GHz or 5GHz frequencies.
- A “byte” is a collection of 8 bits

Layer 1 – Media

- Layer 1 is primarily known for the physical network cabling. While copper and fiber are the de facto standards, different types exist for each.
- Fiber optics – where the differences matter
  - Multimode
    - 62.5um-–FDDI/OM1, 50um--OM2, OM3, OM4
  - Singlemode
    - OS1, OS2
- UTP/STP Copper Cabling
  - CAT5/5e/6/6A/7
- Coax

Layer 1 – Devices

- **Layer 1 devices are purely electrical**
- **Repeaters**
  - In line devices that repeat signals to overcome distance limitations.
  - Versions exists for all types of media
- **Hubs**
  - Like a repeater, it repeats signals received from one source, but to all other connected destinations
- **Network Interface Cards (NICs)**
  - Also partially a layer 2 device

Layer 1 - Troubleshooting

- Link testers
  - Fluke Networks
  - NetScout
  - NetTool.io

Layer 2 – Data Link

- Standardized transmission/reception
  - Ethernet
  - MPLS
  - Frame Relay
- Standardizes hardware media access control (MAC) addresses
  - 48 bit addresses, consisting of a 24-bit Organizational Unit Identifier (OUI), and a 24-bit unique address.
  - OUI identifies the originating manufacturer of the NIC.
- Error detection and correction
- Spanning Tree

Layer 2 - Devices

- Switches
  - Maintains an internal table identifying MAC addresses through corresponding ports.
  - Uses the Source/Destination MAC address in the frame to make intelligent decisions to move frames.
  - Faster than routing, not as scalable.
  - Trunks/uplinks will commonly see many MAC addresses
  - Can segment networks into Virtual LANs (VLANs).
- Network Interface Cards (NICs)
  - Converts bits and data into signals for transmission on network media. Converts signals back to bits for reception.

Layer 2 - Troubleshooting

- View the MAC t

blanderson3 — ssh 172.16.47.6 — 73×49

```
anderson-s1#show mac address-table
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
 All    0100.0ccc.cccc    STATIC      CPU
 All    0100.0ccc.cccd    STATIC      CPU
 All    0100.0ccd.cddc    STATIC      CPU
 All    0180.c200.0000    STATIC      CPU
 All    0180.c200.0001    STATIC      CPU
 All    0180.c200.0002    STATIC      CPU
 All    0180.c200.0003    STATIC      CPU
 All    0180.c200.0004    STATIC      CPU
 All    0180.c200.0005    STATIC      CPU
 All    0180.c200.0006    STATIC      CPU
 All    0180.c200.0007    STATIC      CPU
 All    0180.c200.0008    STATIC      CPU
 All    0180.c200.0009    STATIC      CPU
 All    0180.c200.000a    STATIC      CPU
 All    0180.c200.000b    STATIC      CPU
 All    0180.c200.000c    STATIC      CPU
 All    0180.c200.000d    STATIC      CPU
 All    0180.c200.000e    STATIC      CPU
 All    0180.c200.000f    STATIC      CPU
 All    0180.c200.0010    STATIC      CPU
 All    ffff.ffff.ffff    STATIC      CPU
  20    0001.5c62.6446    DYNAMIC     Gi0/1
  20    24a4.3c05.de10    DYNAMIC     Gi0/14
  10    0017.8812.bc47    DYNAMIC     Gi0/11
  10    001b.218e.cb8d    DYNAMIC     Gi0/3
  10    0cee.e685.0ea6    DYNAMIC     Gi0/8
  10    1077.b19e.b37f    DYNAMIC     Gi0/7
  10    1077.b19e.b380    DYNAMIC     Gi0/7
  10    24a4.3c05.de0f    DYNAMIC     Gi0/13
  10    2832.c5ed.5408    DYNAMIC     Gi0/7
  10    3c52.82a0.4f94    DYNAMIC     Gi0/8
  10    40cb.c0b2.b448    DYNAMIC     Gi0/7
  10    7073.cbdc.e8ff    DYNAMIC     Gi0/8
  10    709e.29bb.1046    DYNAMIC     Gi0/7
  10    802a.a856.6d56    DYNAMIC     Gi0/8
  10    8489.ad64.a259    DYNAMIC     Gi0/8
  10    902b.345e.e4dd    DYNAMIC     Po2
  10    902b.345e.e4df    DYNAMIC     Po2
  10    c869.cd52.71e6    DYNAMIC     Gi0/8
  10    d003.4b57.5cac    DYNAMIC     Gi0/8
  10    f431.c373.3e27    DYNAMIC     Gi0/8
Total Mac Addresses for this criterion: 41
anderson-s1#
```

Layer 2 - Troubleshooting

- View Spanning Tree

blanderson3 — -bash — 77×10

```
uaf-lchr-1#show spanning-tree root


                                        Root Hello Max Fwd
Vlan                 Root ID            Cost  Time Age Dly  Root Port
---------------- -------------------- ------ ----- --- ---  ----------------
                                        Root Hello Max Fwd
Vlan                 Root ID            Cost  Time Age Dly  Root Port
---------------- -------------------- ------ ----- --- ---  ----------------
VLAN0006         32774 2cab.ebbf.0100         6     2   20  15   Gi2/19
```

blanderson3 — telnet uaf-lchr-1 — 88×23

```
uaf-lchr-1#show spanning-tree vlan 6

VLAN0006
  Spanning tree enabled protocol rstp
  Root ID    Priority    32774
             Address     2cab.ebbf.0100
             Cost        6
             Port        147 (GigabitEthernet2/19)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32774  (priority 32768 sys-id-ext 6)
             Address     b414.8961.1f00
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 480

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi2/10              Desg FWD 4         128.138  P2p
Gi2/19              Root FWD 4         128.147  P2p
Gi3/3               Desg FWD 4         128.259  P2p
Te7/4               Desg FWD 2         128.772  P2p

uaf-lchr-1#
```

Layer 3 - Network

- The most complex layer in the OSI model.
  - Also one that presents the most problem areas.
- TCP – The most common protocol suite used in networking. UDP – Very prevalent in streaming data.
  - IPv4 – Still most common addressing suite in use, however exhausted. 32-bit based addresses
    - 4.3 billion addresses globally
  - IPv6 – Standardized for nearly two decades, not seeing wide adoption, but rollout gaining stream. 128-bit
    - 3.4e38 addresses globally

Layer 3 – IP Addressing

- IPv4
  - 32-bit addresses, dotted decimal octets. Most common.
  - Subnet mask delimiter segments IP networks.
  - Ex. Subnet mask of 255.255.255.0 and an IP address of 10.11.12.13 segments the first three octets for the network ID, and the last octet for hosts in the network.
  - Private reserved IP ranges to preserve exhausted public ranges
- IPv6
  - 128-bit addresses in 16-bit hexadecimal segments
    - Subnet mask represented with the address.
    - Trailing zeros summarized with ::
    - Ex 2607:f318::/32 == 2607:f318:0000:0000:0000:0000:0000:0000/32

Layer 3 - IP Addressing

- DHCP – Dynamic Host Control Protocol
  - Allows for automated IPv4 configuration to hosts on your network.
  - Provisions IP address, subnet mask, default gateway, DNS servers at a minimum.
  - Can also allow DNS registration, NTP configuration, limited automated configuration parameters.
  - DHCPv6 exists for IPv6 control
- SLAAC – Stateless Automated Address Configuration
  - Automated IPv6

Layer 3 - ARP

- Address resolution protocol binds IP addresses to MAC addresses.
- As a packet reaches a subnet, a broadcast message is sent to all connected hosts to discover what MAC address has the destination IP address.
- If the IP address does not match the network as designated by the subnet mask, an ARP request is sent for the address of the default gateway.
- ARP entries are stored in a cache table so broadcasts don't have to continually be sent out for each frame.

Layer 3 – ARP Table

- Router – show ip

```
blanderson3 — root@triton:~ — ssh root@triton.sw.alaska.edu — 72...
~ — -bash | ...on:~ — -bash | ....alaska.edu | ~ — -bash
rv-sikuliaq-core>show ip arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.2.0.1                0   001b.1700.0112  ARPA   Vlan4
Internet  10.2.0.3              185   001b.218c.b780  ARPA   Vlan4
Internet  10.2.0.4              259   0025.90c4.e302  ARPA   Vlan4
Internet  10.2.0.5              116   0025.90c4.e59e  ARPA   Vlan4
Internet  10.2.0.7              203   0025.90a2.bdbc  ARPA   Vlan4
Internet  10.2.0.8              247   0025.90a2.bd8c  ARPA   Vlan4
Internet  10.2.0.10              42   0006.f661.deaf  ARPA   Vlan4
Internet  10.2.0.11              30   6c41.6a24.3541  ARPA   Vlan4
Internet  10.2.0.12             246   6c41.6a24.9d41  ARPA   Vlan4
Internet  10.2.0.13              33   6c41.6a24.9a41  ARPA   Vlan4
Internet  10.2.0.14             101   6c41.6a24.1841  ARPA   Vlan4
Internet  10.2.0.15              43   6c41.6a27.c941  ARPA   Vlan4
Internet  10.2.0.16              54   6c41.6a1f.e1c1  ARPA   Vlan4
Internet  10.2.0.17              50   6c41.6a24.15c1  ARPA   Vlan4
 --More--
```

- Computer – arp -

```
blanderson3 — root@triton:~ — ssh root@triton.sw.alaska.edu — 7...
~ — -bash | ...:~ — -bash | ....alaska.edu | ~ — -bash
[root@triton ~]# arp -an
? (137.229.0.194) at 00:3e:e1:c6:9c:35 [ether] on eth0
? (137.229.0.137) at 00:1d:a2:f4:c9:b2 [ether] on eth0
? (137.229.0.130) at c8:f9:f9:96:a0:00 [ether] on eth0
? (137.229.0.241) at 00:22:55:2a:a2:48 [ether] on eth0
? (137.229.0.162) at 40:6c:8f:b9:c2:97 [ether] on eth0
? (137.229.4.99) at c8:f9:f9:96:9c:00 [ether] on eth1
? (137.229.0.249) at ac:87:a3:12:98:16 [ether] on eth0
? (137.229.0.139) at 00:50:56:8e:38:bb [ether] on eth0
? (137.229.0.243) at 00:22:55:2a:a2:48 [ether] on eth0
? (137.229.0.235) at b0:83:fe:d0:56:e9 [ether] on eth0
? (137.229.0.177) at 98:5a:eb:df:d2:47 [ether] on eth0
? (137.229.0.244) at 00:22:55:2a:a2:48 [ether] on eth0
? (137.229.0.157) at 58:97:1e:0e:96:41 [ether] on eth0
? (137.229.0.138) at e0:db:55:1c:79:b0 [ether] on eth0
? (137.229.0.129) at 00:00:5e:00:01:69 [ether] on eth0
```

Layer 3 – Network Address Translation

- Private IPv4 ranges to preserve exhausted public IP space—RFC 1918
  - 10.0.0.0/8 = 16.78 Million IP addresses
  - 172.16.0.0/12 = 1.04 Million IP addresses
  - 192.168.0.0/16 = 65,536 IP addresses
- Allows firewalls to associate a public IP to a private IP as needed – 1:1
  - Host (Private IP) <> Firewall <> Public IP <> Internet
  - As more traffic becomes internet dependant, NAT becomes less useful as 1:1 relationship uses similar resources.

Layer 3 – Domain Name Service

- Domain Name Service (DNS) is a basic fundamental necessity of every day life.
- Brings accessibility by allowing internet navigation using text-based names (domains)
- Larger trusted structure worldwide indexes all names.
- DNS servers are responsible for translating domain names into IP addresses
  - First thing to occur when navigating to any website

Layer 3 - Routing

- Makes up the internet – responsible for ensuring data moves through effective paths to its destination.
- Several standard routing protocols exist to automate the provisioning of network routes.
  - Interior Gateway Protocol (IGP)
    - Open Shortest Path First (OSPF)
    - Enhanced Interior Gateway Routing Protocol (EIGRP)
    - Routing Information Protocol (RIP/RIPv2)
    - Primarily what we ship-going folks are concerned with
  - Exterior Gateway Protocol (EGP)
    - Used to advertise routes to the public internet. Can not advertise private IP addresses externally.
    - Border Gateway Protocol (BGP)

Layer 3 - Hardware

- Routers
  - Specialized hardware with few ports
  - Designed to table large route tables and direct traffic efficiently
  - Peplink Balance appliance is a special type of router with proprietary functions.
- Multilayer (Layer 3) Switching
  - Switches with beefier memory and cpu to both switch and route traffic.
  - Designed for specific functions in a small environment, like a ship!

Layer 3 - Routing

- Routing Table

blanderson3 — root@triton:~ — ssh root@triton.sw.alaska.edu — 88×20

~ — -bash | ...@triton:~ — ssh root@triton.sw.alaska.edu

```
Gateway of last resort is 137.229.2.17 to network 0.0.0.0

     137.229.0.0/28 is subnetted, 1 subnets
C       137.229.2.16 is directly connected, GigabitEthernet0/0.22
     172.20.0.0/30 is subnetted, 1 subnets
C       172.20.11.0 is directly connected, Tunnel11
     199.165.123.0/24 is variably subnetted, 6 subnets, 3 masks
O E2    199.165.123.0/27
           [110/10000] via 199.165.123.60, 09:18:17, GigabitEthernet0/0.501
C       199.165.123.56/29 is directly connected, GigabitEthernet0/0.501
S       199.165.123.48/30 is directly connected, Tunnel10
B       199.165.123.40/30 [200/0] via 192.168.199.2, 13:27:50
B       199.165.123.44/30 [200/0] via 172.20.11.2, 03:18:54
S       199.165.123.32/29 [1/0] via 199.165.123.59
     192.168.199.0/30 is subnetted, 1 subnets
C       192.168.199.0 is directly connected, Tunnel10
     192.168.96.0/32 is subnetted, 1 subnets
S       192.168.96.198 is directly connected, Tunnel10
S*   0.0.0.0/0 [1/0] via 137.229.2.17
swf-sikuliaq-2851#
```

Layer 3 - Troubleshooting

- Ping

blanderson3 — -bash — 88×15

~ — -bash | ...root@triton.sw.alaska.edu | ~ — -bash

```
MacBook-Pro:~ blanderson3$ ping 199.165.123.1
PING 199.165.123.1 (199.165.123.1): 56 data bytes
64 bytes from 199.165.123.1: icmp_seq=0 ttl=57 time=295.424 ms
64 bytes from 199.165.123.1: icmp_seq=1 ttl=57 time=278.082 ms
64 bytes from 199.165.123.1: icmp_seq=2 ttl=57 time=450.256 ms
64 bytes from 199.165.123.1: icmp_seq=3 ttl=57 time=369.510 ms
64 bytes from 199.165.123.1: icmp_seq=4 ttl=57 time=295.263 ms
64 bytes from 199.165.123.1: icmp_seq=5 ttl=57 time=510.341 ms
64 bytes from 199.165.123.1: icmp_seq=6 ttl=57 time=742.278 ms
64 bytes from 199.165.123.1: icmp_seq=7 ttl=57 time=653.353 ms
^C
--- 199.165.123.1 ping statistics ---
8 packets transmitted, 8 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 278.082/449.313/742.278/163.664 ms
MacBook-Pro:~ blanderson3$
```

Layer 3 - Troubleshooting

- Traceroute

blanderson3 — -bash — 89×28

~ — -bash | ...root@triton.sw.alaska.edu | ~ — -bash

```
MacBook-Pro:~ blanderson3$ traceroute google.com
traceroute to google.com (216.58.193.78), 64 hops max, 52 byte packets
 1  172.20.10.1 (172.20.10.1)  4.701 ms  3.420 ms  3.425 ms
 2  172.26.96.161 (172.26.96.161)  66.346 ms  67.322 ms  80.038 ms
 3  172.16.232.228 (172.16.232.228)  63.935 ms
    172.16.232.252 (172.16.232.252)  83.378 ms  83.249 ms
 4  12.83.186.161 (12.83.186.161)  79.961 ms  87.537 ms  88.103 ms
 5  12.83.186.145 (12.83.186.145)  71.743 ms  73.212 ms  55.936 ms
 6  12.123.159.49 (12.123.159.49)  73.467 ms  59.769 ms  79.932 ms
 7  12.247.252.14 (12.247.252.14)  87.995 ms  93.235 ms
    12.247.252.10 (12.247.252.10)  69.997 ms
 8  108.170.244.2 (108.170.244.2)  47.217 ms
    108.170.243.197 (108.170.243.197)  55.553 ms
    108.170.243.175 (108.170.243.175)  75.375 ms
 9  209.85.251.241 (209.85.251.241)  64.060 ms
    209.85.241.124 (209.85.241.124)  56.003 ms
    209.85.249.136 (209.85.249.136)  83.128 ms
10  72.14.233.183 (72.14.233.183)  113.206 ms
    72.14.239.209 (72.14.239.209)  130.204 ms
    72.14.233.111 (72.14.233.111)  278.623 ms
11  216.239.50.38 (216.239.50.38)  111.270 ms
    209.85.248.92 (209.85.248.92)  97.856 ms
    216.239.62.18 (216.239.62.18)  135.754 ms
12  108.170.245.113 (108.170.245.113)  111.591 ms  127.704 ms *
13  209.85.242.39 (209.85.242.39)  112.214 ms  99.464 ms
    209.85.242.37 (209.85.242.37)  104.368 ms
14  sea15s07-in-f14.1e100.net (216.58.193.78)  101.987 ms  127.030 ms  112.135 ms
MacBook-Pro:~ blanderson3$
```

Layer 3 - Troubleshooting

- nslookup

blanderson3 — -bash — 72×17

~ — -bash | ...triton.sw.alaska.edu | ~ — -bash

```
MacBook-Pro:~ blanderson3$ nslookup www.google.com
Server:         137.229.5.193
Address:        137.229.5.193#53

Non-authoritative answer:
Name:   www.google.com
Address: 74.125.28.147
Name:   www.google.com
Address: 74.125.28.104
Name:   www.google.com
Address: 74.125.28.106
Name:   www.google.com
Address: 74.125.28.99
Name:   www.google.com
Address: 74.125.28.105
Name:   www.google.com
Address: 74.125.28.103
```

Layer 3 - Troubleshooting

- Host

blanderson3 — -bash — 72×17

~ — -bash | ...triton.sw.alaska.edu | ~ — -bash

```
MacBook-Pro:~ blanderson3$ host google.com
google.com has address 216.58.216.142
google.com has IPv6 address 2607:f8b0:400a:800::200e
google.com mail is handled by 20 alt1.aspmx.l.google.com.
google.com mail is handled by 10 aspmx.l.google.com.
google.com mail is handled by 50 alt4.aspmx.l.google.com.
google.com mail is handled by 30 alt2.aspmx.l.google.com.
google.com mail is handled by 40 alt3.aspmx.l.google.com.
MacBook-Pro:~ blanderson3$
```

OSI Model Recap

Network (3)	Reads the IP address form the packet.	Routers, Layer 3 Switches
Data Link (2)	Reads the MAC address from the data packet	Switches
Physical (1)	Send data on to the physical wire.	Hubs, NICS, Cable

Layer 4 - Transport

- Where applications become identified – based on port numbers
- Standard set of port numbers for well-known applications (0-1024 reserved as standards)
  - TCP/22 – SSH
  - TCP/80 – HTTP
  - TCP/443 – HTTPS
  - UDP/53 – DNS
  - Many many many more (and many more after that)
- 65,535 ports per IP address
- IP address and port together is a socket

Layer 4 - Transport

- Firewalls - application identification
  - Basis for securing networks to allow specific applications in/out specific networks.
  - Allows for application specific rules to deny certain applications but not others while allowing others.
  - Next-gen firewalls (NGFW) use packet inspection to identify applications’ traffic pattern signatures and can identify those using non-standard ports.

Layer 4 – Port Address Translation

- Supplants the Network Address Translation function at Layer 3 to use ports to translate many IP addresses to one.
  - Common in home networking.
  - Only allows one inside server to be reachable on a given port due to port forwarding.
  - Best at conserving public IP addresses when many hosts access internet resources - most common on ships.
- Host (rhp)<>Router<>FW<>internet host(dst p)
- Firewall translates the rhp to another rhp
  - Firewall tracks the connection state to forward outside port to inside port.

Layer 4 – PAT Example

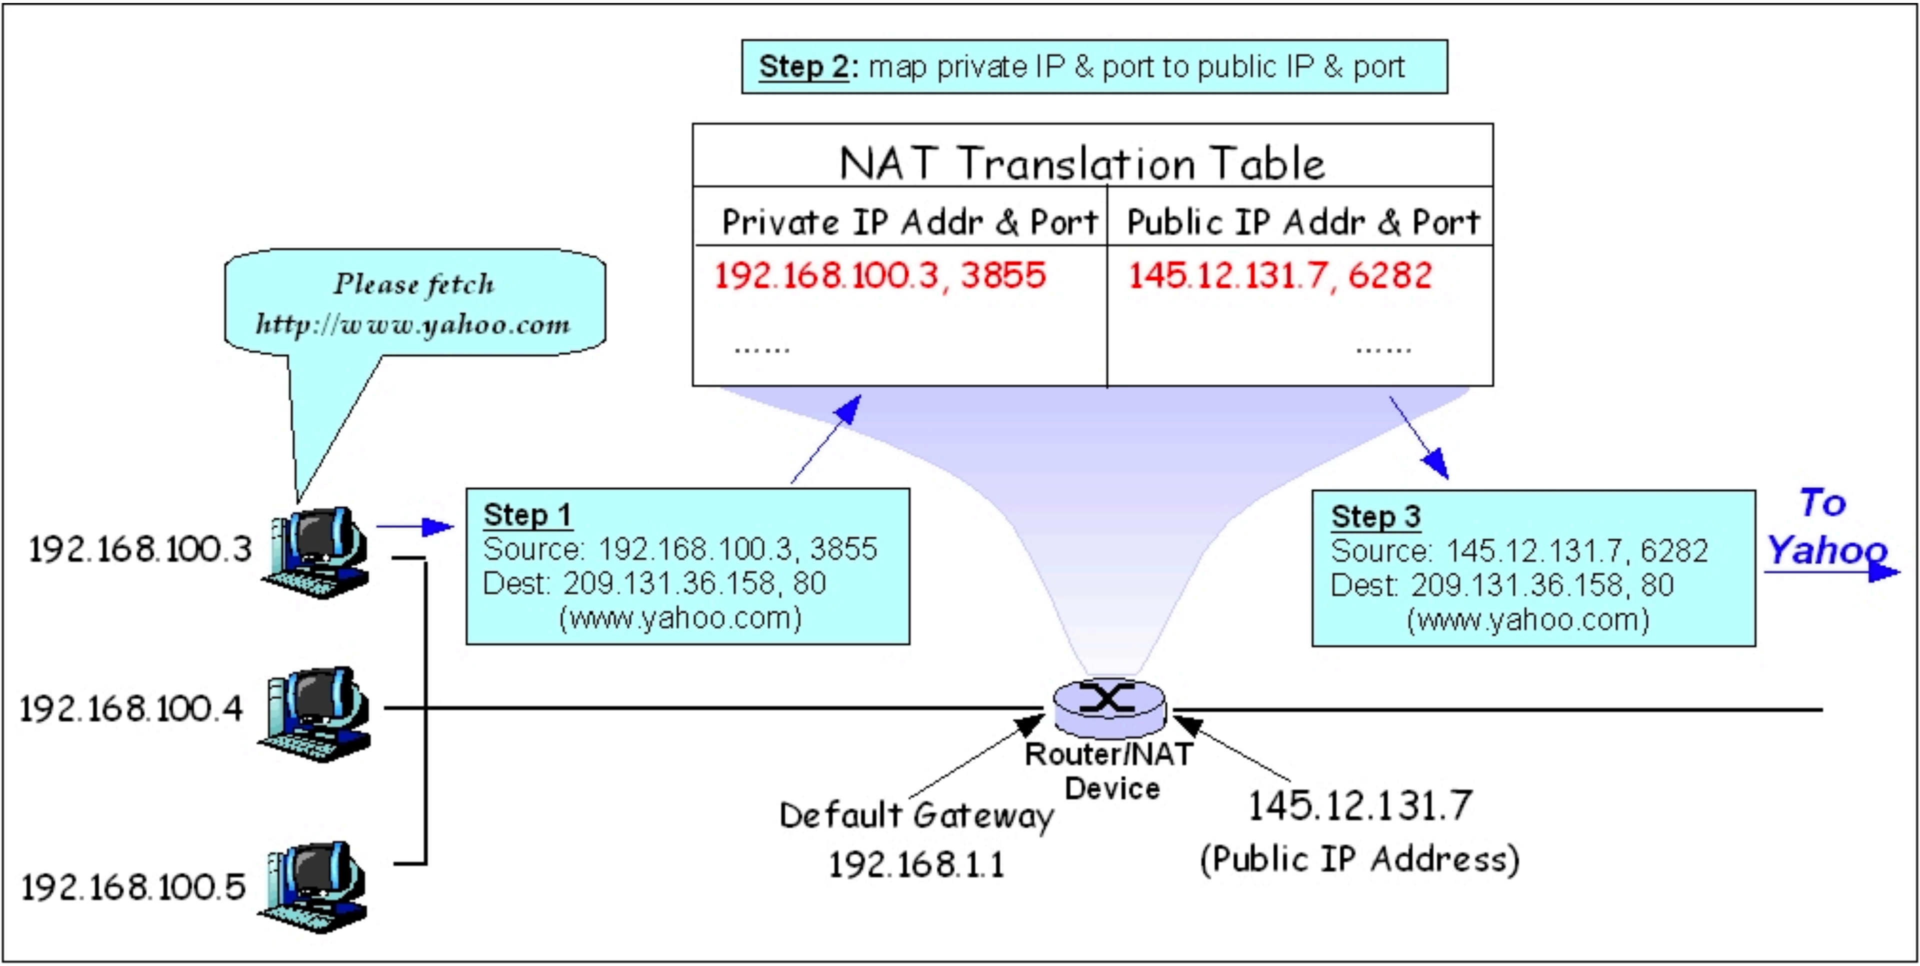

Source: Wikibooks

Layers 5, 6 & 7

- Layer 5 – Session
  - Ensures both TCP sessions and any system or user network sessions (ex logging into your bank) are timed out appropriately.
  - Where port numbers are originated and synchronized for source and destination to Transport layer.
- Layer 6 – Presentation
  - Where data is collected and prepared for the application. Also where encryption/decryption happens
- Layer 7 – Application
  - This pptx presentation file displayed in front of you.

OSI Model Recap

7. Application
6. Presentation
5. Session
4. Transport
3. Network
2. Data Link
1. Physical

7. Acumen / At
6. Palin's / Patient
5. Sarah / Same
4. Teach / Teachers
3. Not / Needs
2. Do
1. Please / Phil

Bandwidth Delay Product

- TCP is the original protocol of the internet as built in the late 80s-early 90s.
  – Not particularly efficient with today's workloads or today's bandwidth.
  – TCP receive window (RWIN) scales via Slow Start
    - Scaling occurs slowly, and latency fluctuations (jitter) often cause it to restart.
    - Results in single flows crawling over highly latent and fluctuating links.

Bandwidth Delay Product

- The BDP is a formula that can both determine maximum possible throughput given latency and loss, as well as unscaled RWIN values to reach desired throughput.
  - Bandwidth (Kbps) * Latency (ms) = RWIN (b) / 8 = RWIN (B)
  - For example: 2000Kbps * 500ms = 1,000,000 / 8 = 125,000 bytes = 122.07KB RWIN -> 128KB RWIN

Congestion Window

- Sender controlled
- Window managed by congestion algorithm
- Input is varied by system and algorithm

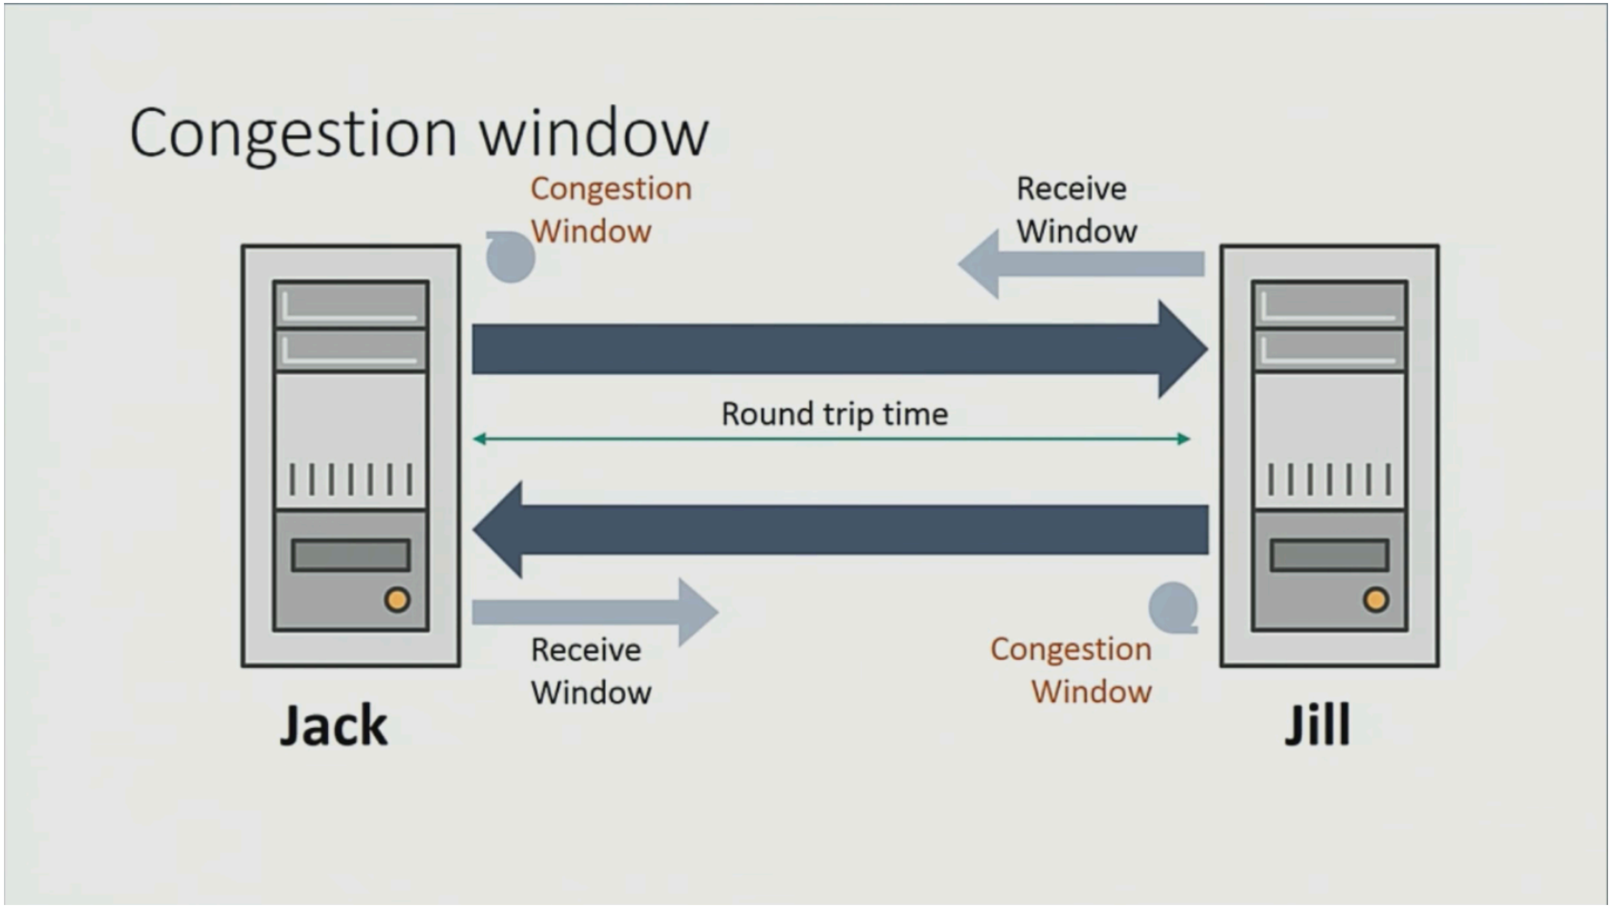

Initial Congestion Window

- How much data to send before expecting to see acknowledgements.
  - Basis of the bandwidth delay product
  - Coordinated values with the TCP RWIN on the receiver end
  - RWIN and CWIN values should be set on both sides for optimum performance.

Impact of Loss

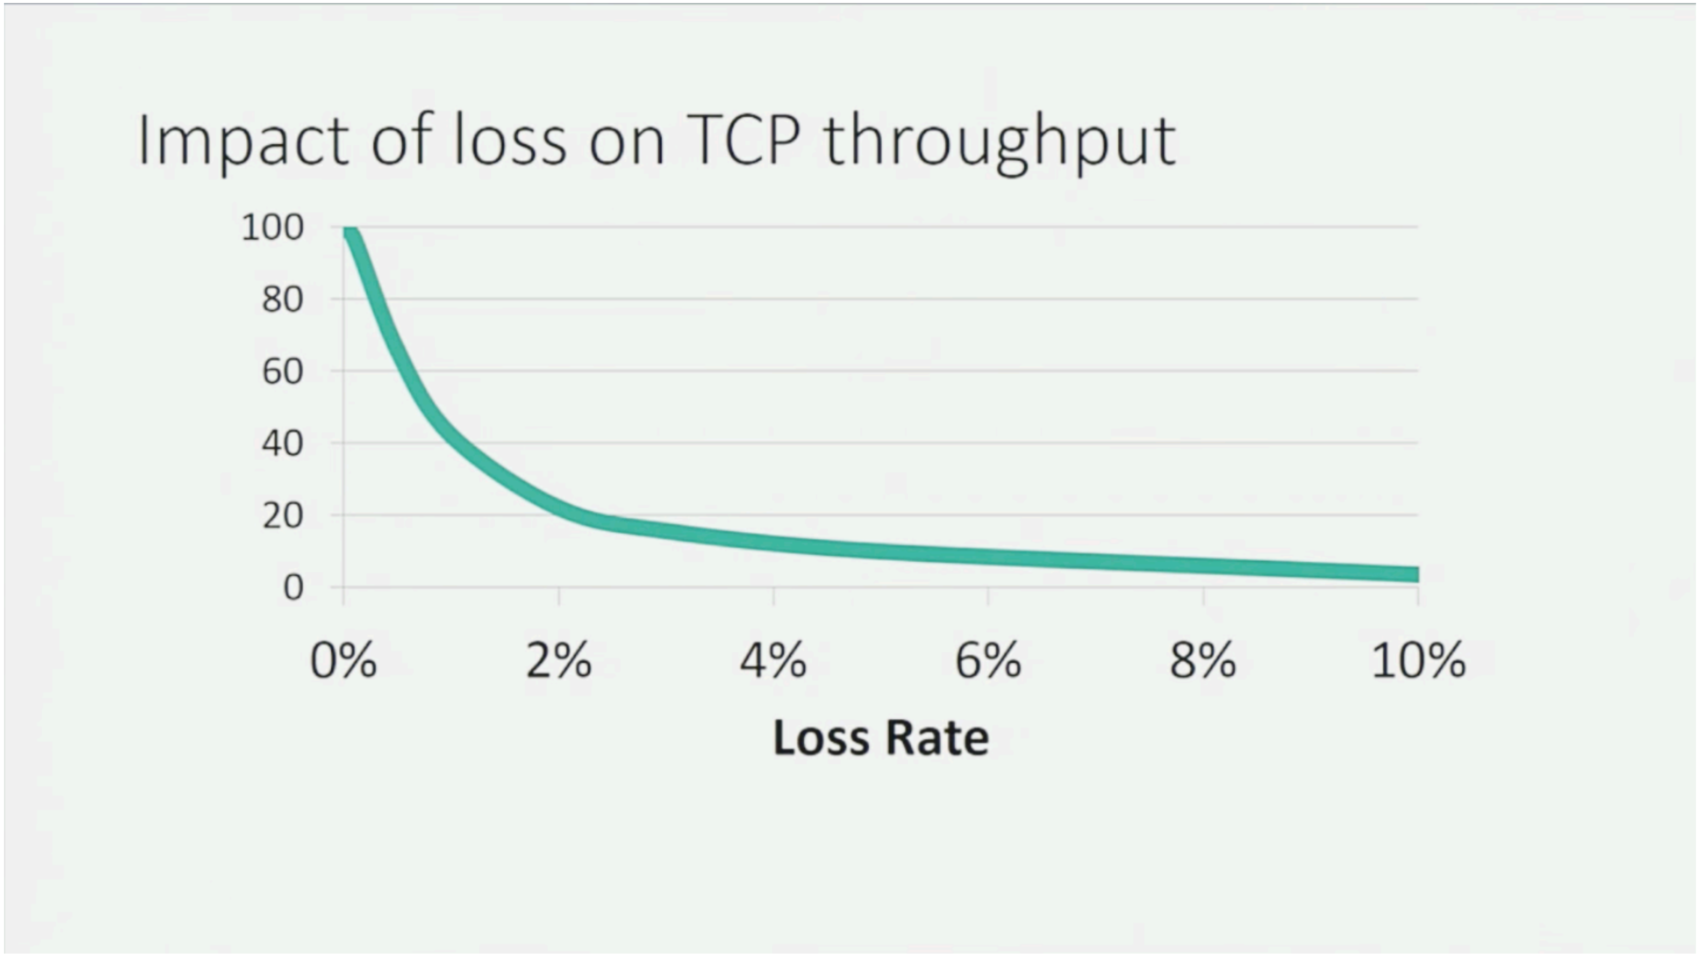

Retransmission Timers

- Input as to when congestion control considers a packet lost.
  - Too low: Retransmit lots of things possibly for no reason
  - Too high: Connections sit for a while timers expire for data to come back

Considerations

- CWIN/RWIN are critical to tune over high latency links like satellites for best performance.
  - CWIN values should be slightly less than BDP
  - RWIN values should be slightly higher
    - Consider maximum average latency to maintain speeds.
- Optimize retransmission timers if necessary to eliminate fake loss.
  - Loss should not be expected, but can be prepared for.

Conclusion

Questions?

UNIVERSITY OF ALASKA FAIRBANKS

Thank You!