TCOM 515

IP Routing

Lecture 1

Class Information

- Instructor: Wei Wu
  - Lectures and Lab session 1 (Monday)
  - Email tcomclass@gmail.com
- Instructor: Cong Tham
  - Lab sessions 2 and 3 (Tuesday and Wednesday)
  - Email tcom.gmu@gmail.com
- TA:
  - Email:
- Class Syllabus

Lecture 1

- OSI model and TCP/IP overview
- IP addresses
- ARP and ICMP
- How routing works
- Directly connected routes
- Static routes

OSI Model

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

TCP/IP

- Developed before OSI
- Actual implementation, OSI model is framework
- Consists of following layers
  - Application
  - Host-to-host or transport
  - Internet
  - Network interface

OSI and TCP/IP

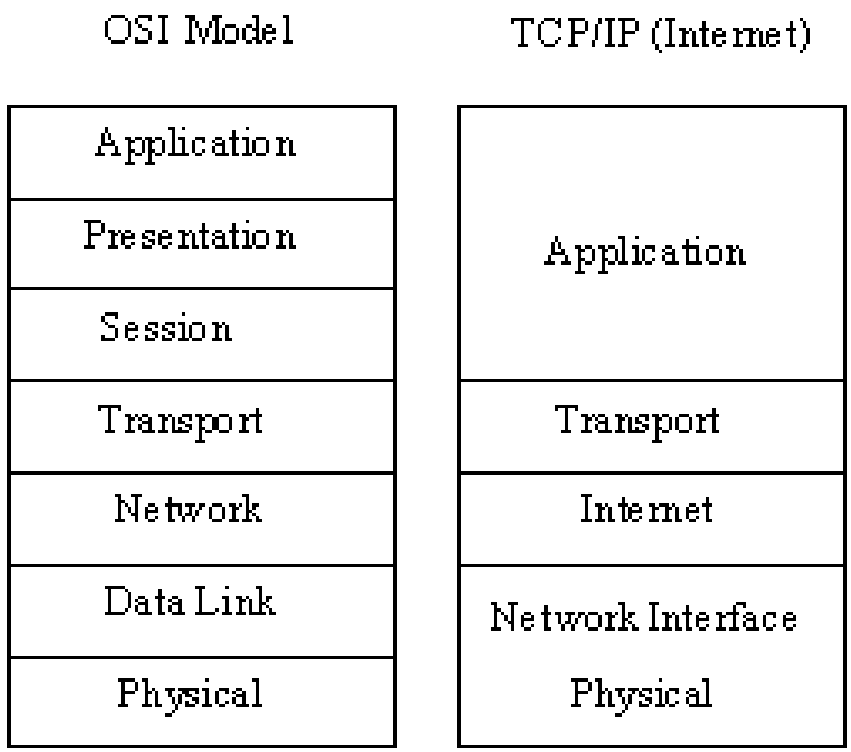

TCP/IP Protocols

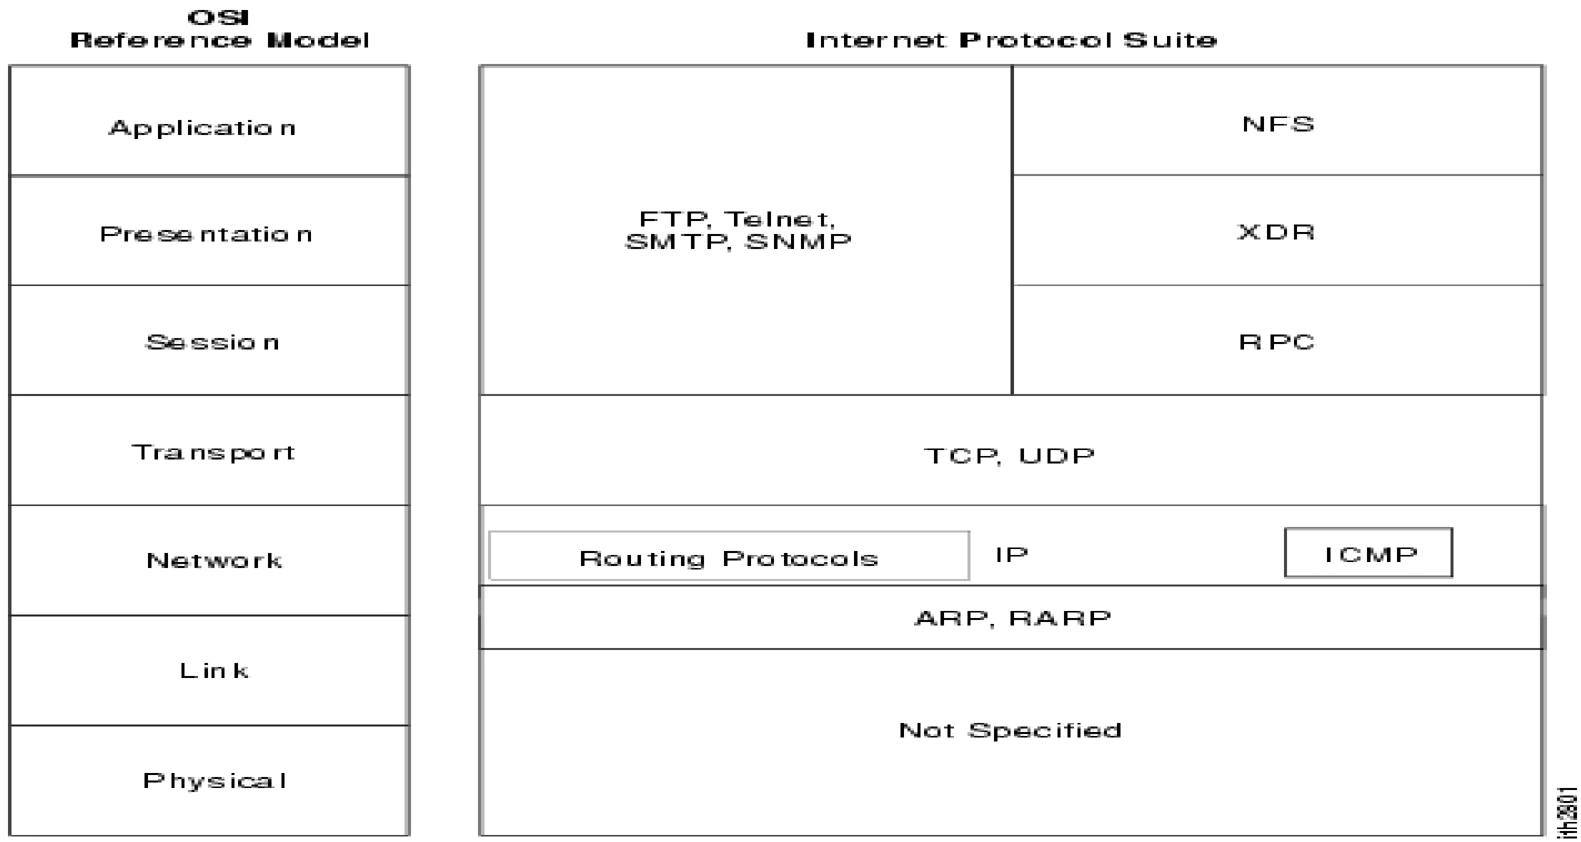

TCP-Transmission Control Protocol

TCP provides applications with a reliable, connection-oriented service. It provides a point-to-point connection between two hosts. TCP fragments are encapsulated within IP packets.

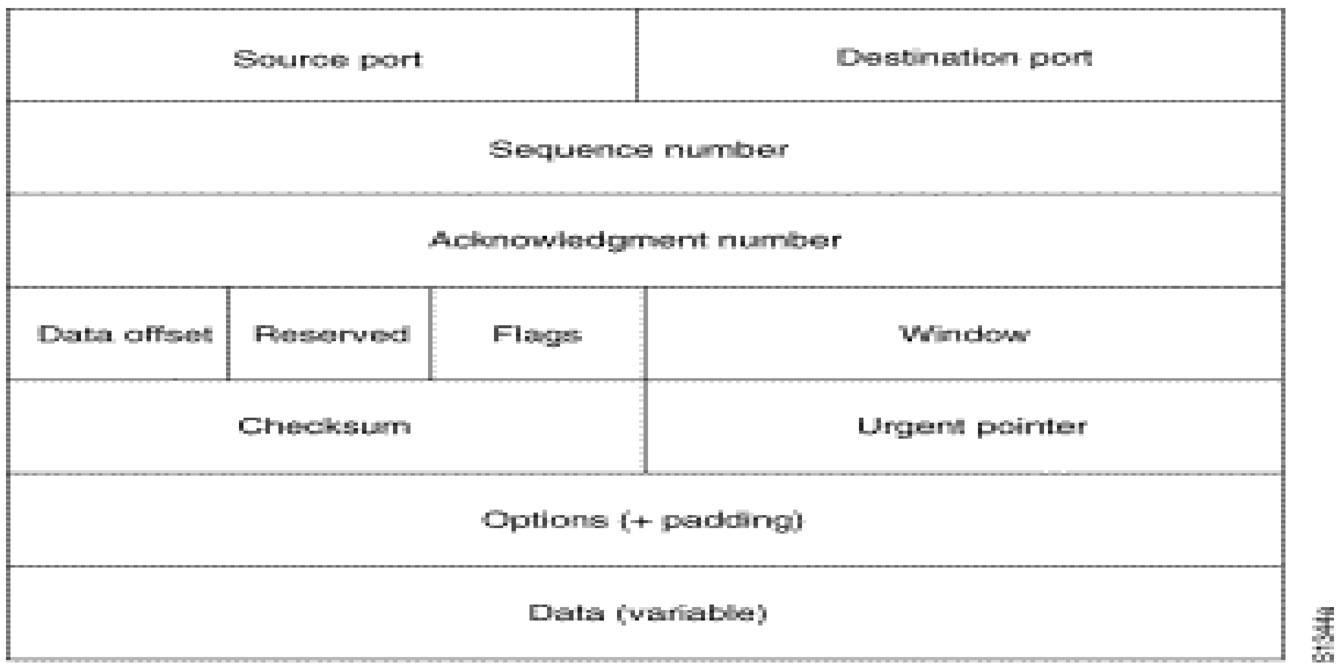

UDP - User Datagram Protocol

- UDP is a unreliable, best effort protocol. There is no error checking or acknowledgement.
- UDP datagrams are also encapsulated in IP packets.

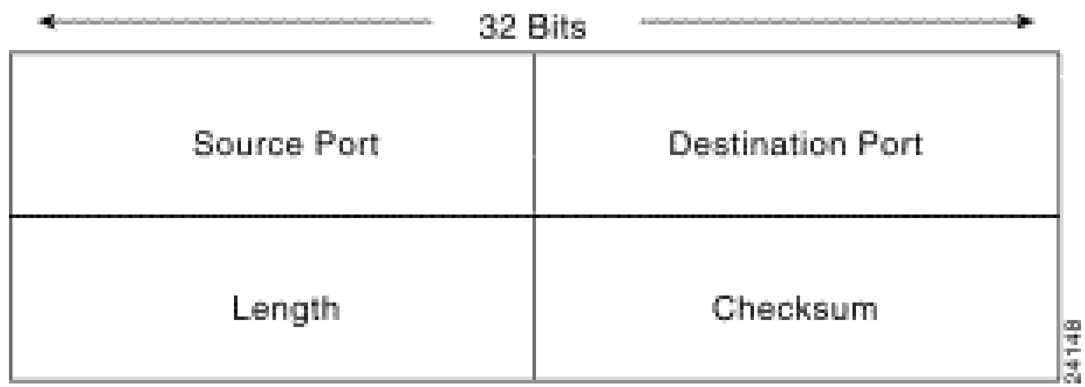

IP Packet Header

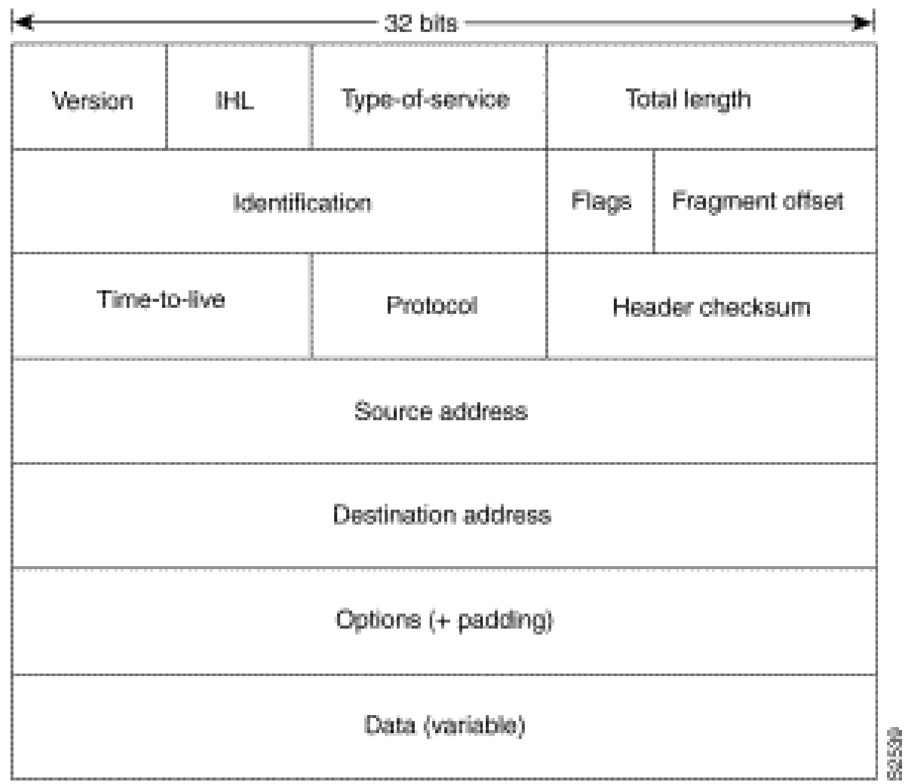

IP Packet Header Fields

Version - 4 bit - IP version IPv4 (0100) or IPv6 (0110)

IP Header Length (IHL) - 4 bit - indicates the datagram header length in 32-bit words, including options.

Type-of-Service - 8 bit -information for higher layer protocols on handling and priority levels, broken into precedence and TOS

Total Length - 16 bit - length of IP packet, in terms of bytes, including both data and header.

Identification - 16 bit - Contains an integer that identifies the current datagram. This field is used to help piece together datagram fragments.

Flags - 3 bit - First bit is unused. The second bit is the DF, Do Not Fragment bit, which generates an error and drops the packet if fragmentation is needed. The last bit is the MF, More Fragments bit, which tells the router whether or not more fragments are coming. MF = 0 means that it is the last fragment.

Fragment Offset - 13 bit - Specifies fragmentation offset in terms of eight octet units from the beginning of the IP header. Used to order fragments once all are received.

IP Packet Header Fields

Time-to-Live - 8 bit - Hop count that is decremented by each router that touches the packet. TTL is used to prevent endless loops and routing problems. Once TTL = 1 or 0, the packet is dropped and ICMP error message is sent to source.

Protocol - 8 bit - Specifies the next higher level protocol of host-to-host or transport protocol. TCP (00000110) UDP (00010001)

Header Checksum - 16 bit - One's complement checksum of original packet, compared to calculation of receiver's packet. Recalculated at each router because of TTL change.

Source Address - 32 bit - IP address of sender.

Destination Address - 32 bit - IP address of intended receiver

Options - Variable Length - allows for further information including, routing info types: loose or strict, record route and timestamps. Uses padding to finish 32 bit words.

Data - higher layer (TCP or UDP) datagram

IP Addressing Review

IP Addresses have 32 bits broken into 4 octets. They are expressed in dotted decimal notation. Addresses may range from 0.0.0.0 to 255.255.255.255. Address ranges are assigned by Internet Assigned Numbers Authority IANA, who in turns assigns IP ranges to Regional Internet registry.

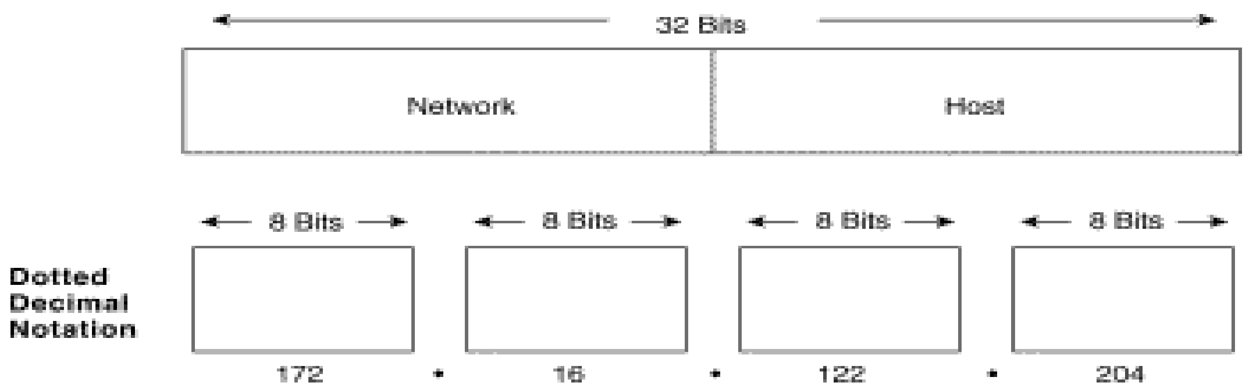

IP Address Classes

Class	Range	Purpose	Number of Hosts
A	1.0.0.0 -> 126.0.0.0	Large Organizations	16,777,214
B	128.1.0.0 -> 191.254.0.0	Medium Organizations	65,543
C	192.0.1.0 -> 223.255.254.0	Small Organizations	254
D	224.0.0.0 -> 239.255.255.255	Multicast	N/A
E	240.0.0.0 -> 254.255.255.255	Experimental	N/A

Classful vs Classless Addressing

Classful addressing divides the IP space into A, B, C classes, network masks are 255.0.0.0 (/8), 255.255.0.0 (/16), and 255.255.255.0 (/24).

Classless (Variable Length Subnet Masking) uses variable network masks.

IP Subnetting

IP addresses need a subnet mask in order to put it into the reference of a network. The subnet mask determines the network address that a specific host IP belongs.

For example:

10.0.0.1 255.0.0.0 which translates to:

00001010.00000000.00000000.00000001 - 10.0.0.1
11111111.00000000.00000000.00000000 – 255.0.0.0

The network address is the host bits that align with the 1 bits of the subnet mask: 10.0.0.0

The number of IP addresses in each subnet can be calculate by 2^N. The first IP address is the **subnet address**, the last IP address is the **broadcast address**. All IP addresses in between are valid IP host addresses. 2^N-2 is the number of valid IP host addresses in the subnet.

The subnet mask can also be represented in bit count format, 255.0.0.0 = /8

Let’s try an example of subnet masking:

A host has an IP address of:

10.2.5.105 255.255.192.0

What is the IP address put into binary form?

00001010.00000010.000000101.01101001 = 10.2.5.105

11111111.11111111.11000000.000000 = 255.255.192.0

Now let's continue the example of: 10.2.5.105 255.255.192.0

00001010.00000010.000000101.01101001 = 10.2.5.105

11111111.11111111.11000000.000000 = 255.255.192.0

What is the network and broadcast address for this host?

Network address is

00001010.00000010.00000000.00000000 = 10.2.0.0

Broadcast address is

00001010.00000010.00.111111.11111111 = 10.2.63.255

Now let's continue the example of: 10.2.5.105 255.255.192.0

00001010.00000010.000000101.01101001 = 10.2.5.105

11111111.11111111.11000000.000000 = 255.255.192.0

00001010.00000010.00000000.00000000 = 10.2.0.0

00001010.00000010.00.111111.11111111 = 10.2.63.255

How many usable IP host addresses are in this network?

Host range is 2^14 -2 (network and broadcast addresses)

16384-2 = 16,382 host IP addresses

One more example: 192.168.1.15 255.255.255.224

What are the network address, broadcast address, and host range?

Network address is

11000000.10101000.00000001.00000000 = 192.168.1.0

Broadcast address is

11000000.10101000.00000001.00011111 = 192.168.1.31

Host range is 2^5 -2 (network and broadcast addresses)

32-2 = 30 host IP addresses

Below is a table to help you understand translation of bits into subnet masks. Each octet of a subnet mask can only have values of the numbers in the blocks on the right-hand side in addition to the possible value of 1.

128	64	32	16	8	4	2	1		
1	0	0	0	0	0	0	0	=	128
1	1	0	0	0	0	0	0	=	192
1	1	1	0	0	0	0	0	=	224
1	1	1	1	0	0	0	0	=	240
1	1	1	1	1	0	0	0	=	248
1	1	1	1	1	1	0	0	=	252
1	1	1	1	1	1	1	0	=	254
1	1	1	1	1	1	1	1	=	255

24148

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ip.htm

ARP- Address Resolution Protocol

ARP is used to find the MAC address of destination IP address on the local subnet. ARP has three variants: Reverse, Proxy, Gratuitous.

RARP- Reverse Address Resolution Protocol - is used find the IP address of the given MAC address.

Proxy ARP - allows another device to answer ARP requests for IP addresses it doesn’t actually own.

Gratuitous ARP – A host ARPing for its own IP address to check for duplicate IP addresses and advertise its own IP.

IP addressed hosts use their subnet mask to determine whether to ARP for the destination or send the packet to their gateway. If the destination IP is in its network, it ARPs. If it is not in its network, it sends the packet to its default gateway.

ICMP Internet Control Message Protocol

type	code	checksum
unused		
IP header + 1st 64 bits of datagram		

ICMP messages are used by network devices to manage communication such as error reporting, queries and responses. A combination of types and codes identify the packet type which is defined in RFC 792 and RFC 1700. The IP header and part of the datagram help the source determine which packet generated the ICMP message.

ICMP Types and Codes

0,0 - Echo reply & 8,0 - Echo request - Ping functionality

3,0-3 - Unreachable messages

3,4 - Fragmentation needed but not allowed

3,6-7 - Destination unknown

3,9-12 - Destination prohibited

11,0-1 Time exceeded - TTL field 1 or 0***Used for traceroute

What Is Routing?

Routing operates at the 3rd layer of the OSI model, the network layer. IP routing occurs at the Internet layer of TCP/IP stack.

Routing is needed to forward a packet from a source across one or more networks to the identified destination.

Routing involves many routers working together.

Routers

Routers are devices that forward a packet along a path from one network to another.

Routers build a routing table obtained from information in its own configuration and from updates received from other routers in the network.

Router uses routing algorithms to select the best path to any given destination.

How Does Routing work?

First a router receives a packet inbound on one of its local interfaces.

Then the router looks into the IP header of the given packet to find the destination IP address.

Next the router looks into its routing table and selects the best path to that destination. The best path is associated with an interface.

Finally the router sends the IP packet out of the egress interface to the next router on the path to the final destination.

Route Tables

Routing table is used by routers to lookup the path to the final destination. The route table has the best routes per destination. Each router will have a routing table.

Route table is built in three ways

1. Directly connected interfaces
2. Static routes
3. Dynamic routing protocols

Routing table entries include the following information:

- Destination address - network address of a subnet
- Next hop - interface or IP address of next hop in path
- Egress interface – the interface to the next hop
- Type of route: C- Connected, S - Static, R- RIP, O – OSPF
- Metric - arbitrary number used to help choose the best route

Route Table Example

```
atlanta#sho ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR
Gateway of last resort is not set
     20.0.0.0/24 is subnetted, 1 subnets
S       20.20.20.0 [1/0] via 192.168.0.5
     192.168.4.0/30 is subnetted, 1 subnets
B       192.168.4.4 [20/0] via 192.168.1.6, 00:39:49
     192.168.5.0/30 is subnetted, 1 subnets
O E2    192.168.5.4 [110/1] via 192.168.3.5, 00:01:51, Ethernet0/0
     10.0.0.0/32 is subnetted, 6 subnets
C       10.0.3.1 is directly connected, Loopback0
O E2    10.0.5.5 [110/1] via 192.168.3.5, 00:01:51, Ethernet0/0
B       10.0.4.4 [20/11] via 192.168.1.6, 00:39:49
O       10.0.3.2 [110/11] via 192.168.3.5, 00:01:51, Ethernet0/0
B       10.0.4.3 [20/0] via 192.168.1.6, 00:39:49
     192.168.0.0/30 is subnetted, 1 subnets
O       192.168.0.4 [110/20] via 192.168.3.5, 00:01:51, Ethernet0/0
     192.168.1.0/30 is subnetted, 1 subnets
C       192.168.1.4 is directly connected, Ethernet1/0
     192.168.2.0/30 is subnetted, 1 subnets
B       192.168.2.4 [20/20] via 192.168.1.6, 00:39:50
     192.168.3.0/30 is subnetted, 1 subnets
C       192.168.3.4 is directly connected, Ethernet0/0
```

Routing Decisions

1. Routers match the destination IP address to the network addresses in its routing table based on the most specific entry. The ordering is:
   -Host address (a /32 or individual host route)
   -Subnet
   -Group of Subnets (major network or supernet)
   -Default address (0.0.0.0/0)

2. Administrative distance is the feature used to select the best path when there are more than one route to the same destination from two different protocols.

3. Metric is used by a routing protocol to select the best path to destination subnets.

Default Distance Value Table

The table below lists the administrative distance default values of the protocols that Cisco supports.

Route Source	Default Distance Values
Connected interface	0
Static route*	1
Enhanced Interior Gateway Routing Protocol (EIGRP) summary route	5
External Border Gateway Protocol (BGP)	20
Internal EIGRP	90
IGRP	100
OSPF	110
Intermediate System-to-Intermediate System (IS-IS)	115
Routing Information Protocol (RIP)	120
Exterior Gateway Protocol (EGP)	140
On Demand Routing (ODR)	160
External EIGRP	170
Internal BGP	200
Unknown**	255

Route Decision Example

Route table entries:

- 10.1.0.0/24
- 10.1.0.0/30
- 10.1.0.0/23

1. Which route entry will a packet with destination of 10.1.0.2 use?

   All entries are possible, but 10.1.0.0/30 will be selected because it is the most specific, this entry covers 10.1.0.0 to 10.1.0.3.

2. Which route entry will a packet with destination of 10.1.0.5 use?

   10.1.0.0/24 and 10.1.0.0/23 are possible, 10.1.0.0/30 is not possible because it covers 10.1.0.0 to 10.1.0.3. 10.1.1.0/24 will be selected because it is the most specific.
3. Which route entry will a packet with destination of 10.1.1.5 use?

   10.1.0.0/23 is the only one possible, this covers 10.1.0.0 to 10.1.1.255.

Connected Routes

The first type of routes that a router inserts into its routing table are connected to its interfaces. Based on the IP addresses and subnet masks of its own interfaces, the router inserts route statements of the network addresses and the next hop as the configured interface. No additional configuration of routing or routing protocols is necessary. The administrative distance of a connected route is zero. The interface must be up.

Connected Route Table Entry

C 10.0.0.0/8 is directly connected, FastEthernet 1/0

Directly Connected Routes

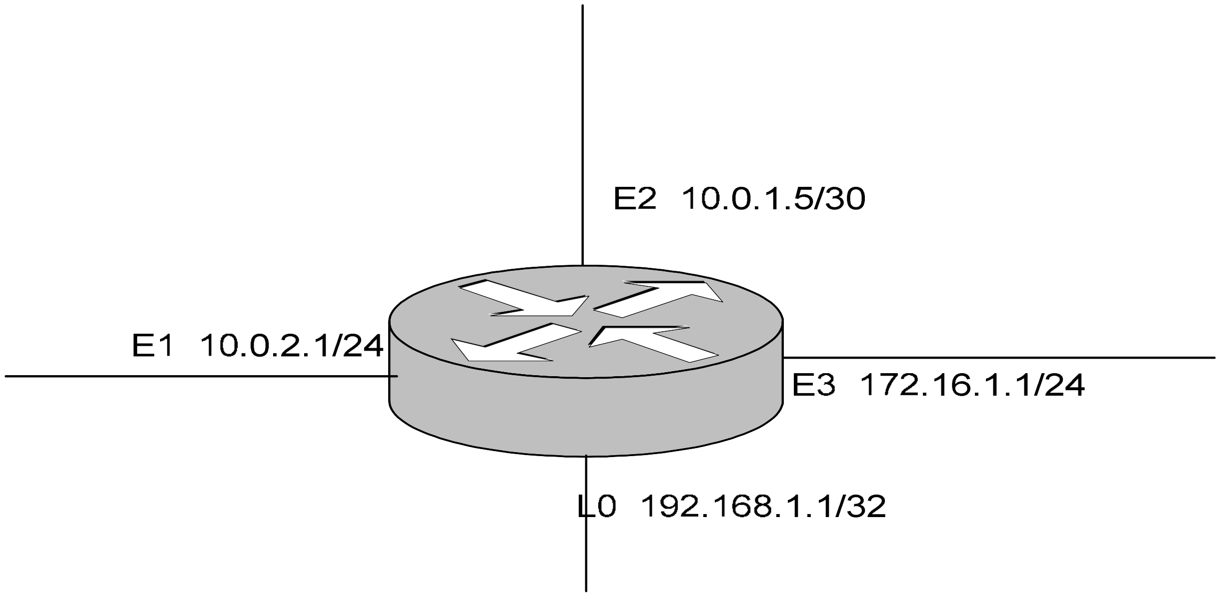

Route Table

- 10.0.1.4/30 Directly Connected E2
- 10.0.2.0/24 Directly Connected E1
- 172.16.1.0/24 Directly Connected E3
- 192.168.1.1/32 Directly Connected L0

Static Routes

Static routes are manually configured on a router by the network administrator.

A static route has 2 parts,

1. Destination network address with subnet mask
2. Pointer to the destination network.
   - An IP address already in the routing table
   - A specified interface on the router

The interface must be up for the router to put the static route pointing to it into its routing table. This is also called the egress interface. If the interface goes down, the route is removed.

The next hop IP address must be in the routing table to put the static route pointing to it into its routing table. If the route table statement for the next hop address is removed, the static route is also removed from the route table.

The static route can point to both next hop IP and egress interface to provide more granularity and control.

Static Routes Cont.

A static pointing to an interface routing table entry looks like:

S 172.16.0.0/12 is directly connected via FastEthernet 1/0

A static point to an IP address routing table entry looks like:

S 172.16.0.0/12 via 10.0.0.1

Static routes, along with dynamic routing protocols, are needed to reach networks not directly connected to the router.

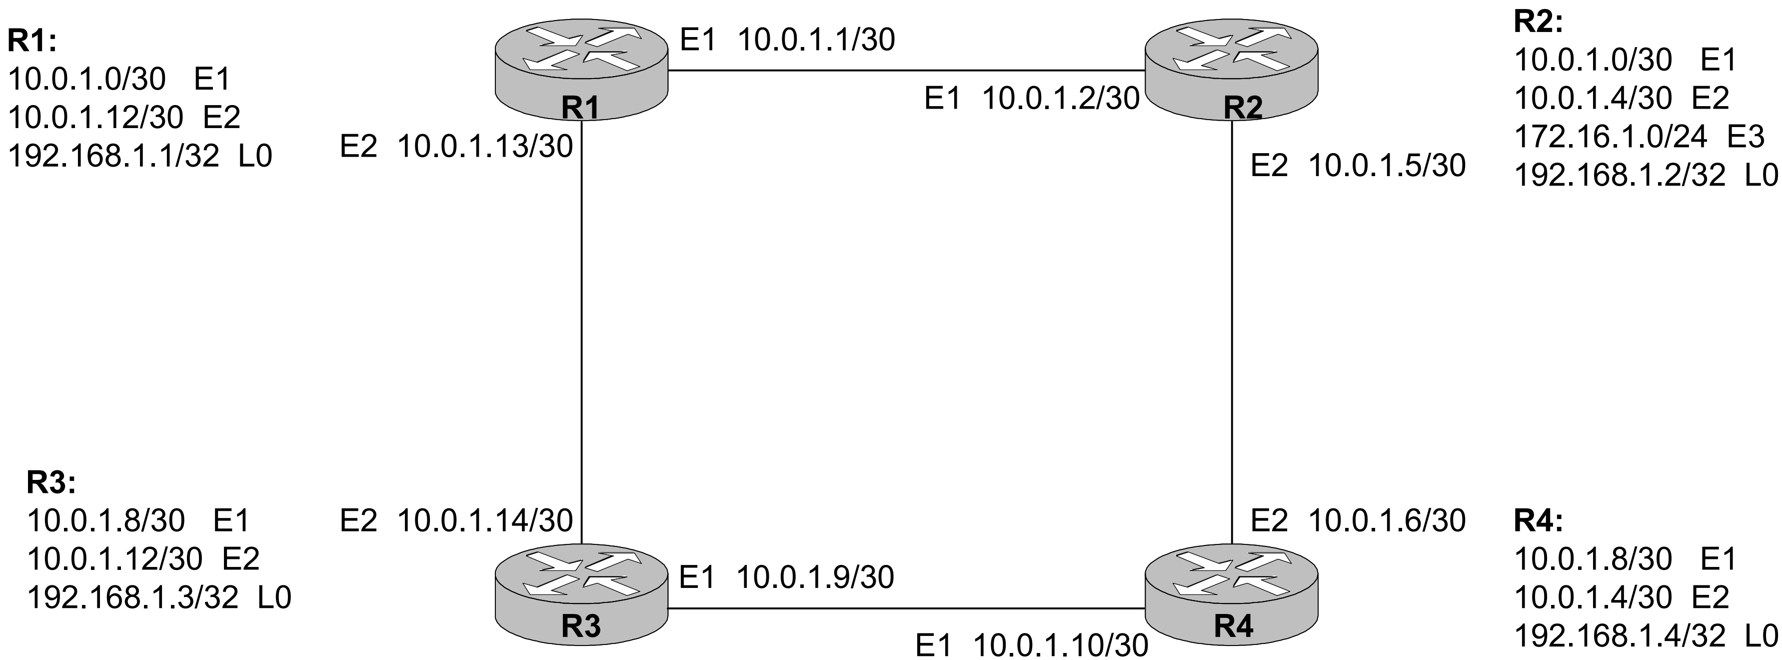

1. what can R1 get to?

2. I want to connect to R2 loopback 0

3. I want to connect to hosts on R2 E2

4. I want to connect to R3 E1 (10.0.1.9)

5. I want to telnet to R4 loopback,

will it work?

What's needed to make it work?

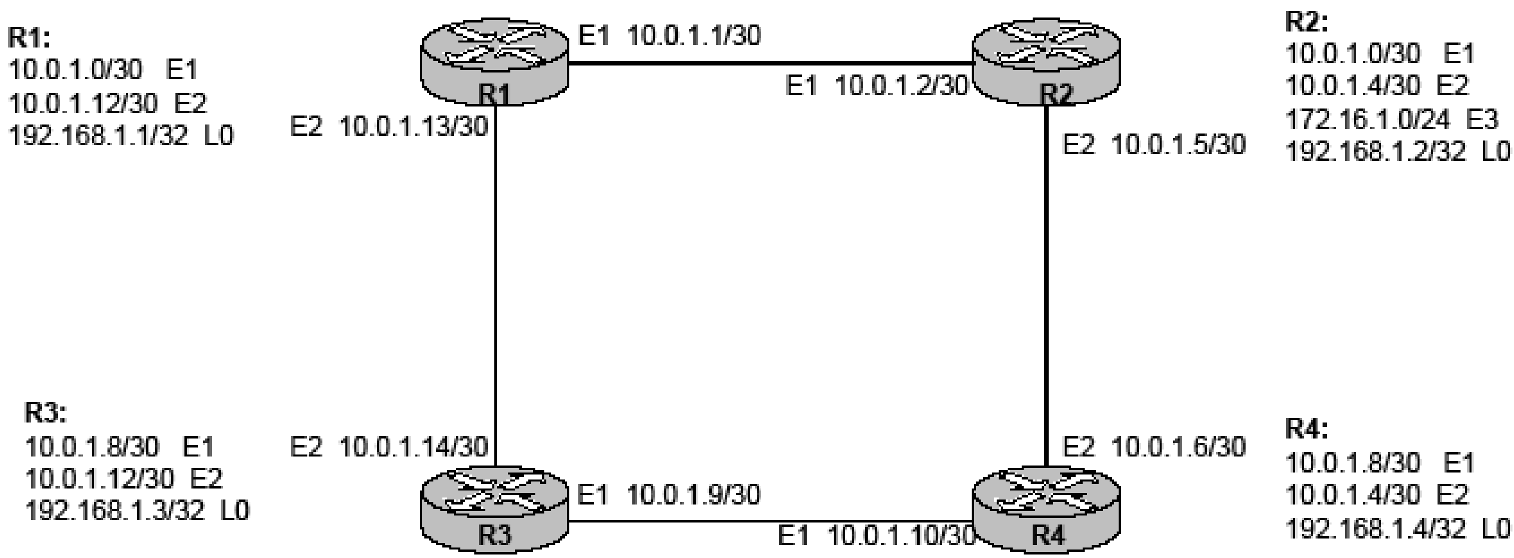

1. what can R1 get to?
10.0.1.0/30 -> .1 and .2
10.0.1.12/30 -> .13 and .14
192.68.1.1/32
These are all directly connected networks

2. I want to connect to R2 loopback 0
Do I need static route? YES, R2 loopback is not directly connected to R1
Ip route 192.168.1.2/32 E1 and/or 10.0.1.2

3. I want to connect to hosts on R2 interface E2
Ip route 10.0.1.4/30 E1 and/or 10.0.1.2

4. I want to connect to 10.0.1.9 on R3 interface E1
IP route 10.0.1.8/30 E2 and/or 10.0.1.14

5. I want to telnet to R4 loopback 0
IP route 192.168.1.4/32 E1/10.0.1.2 OR E2/10.0.1.14 (both paths are equal)
5A. *Will it work?*
NO, R2 and R3 need static route to R4 loopback, and static routes for return trip are needed also.

Summary

- IP addressing
- Routing decisions
- Connected Routes
- Static Routes
- Our lab will cover router configuration, directly connected routes and static routes. Please make sure you read the lab instructions prior to lab
- Reading assignments
  - Chapter 1
  - Chapter 3 – skip pages 90-94, 100-103, and 116-122 (Avoid IPv6).