# Internet

- Organization
- Addresses
- TCP/IP Protocol stack
- Forwarding




# What defines the Internet ?

1. Use of a globally unique address space based on Internet Addresses
2. Support of the Transmission Control Protocol/Internet Protocol (TCP/IP) suite for communications
3. Offers end-to-end delivery service for applications

# Growth of the Internet

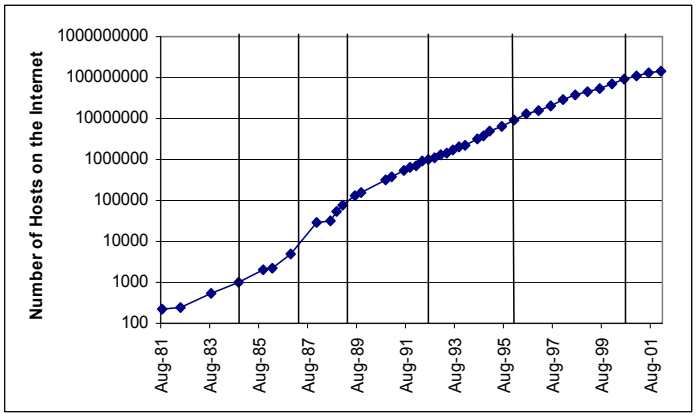


*Source: Internet Software Consortium*





# Internet Infrastructure

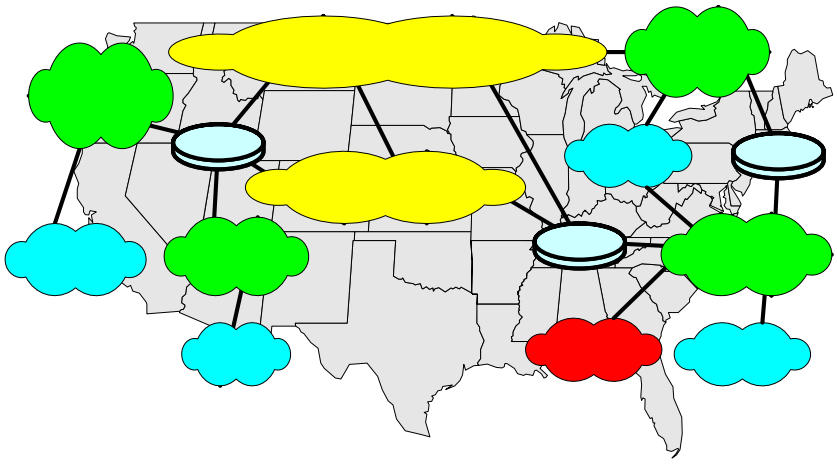

## Internet Infrastructure

- The infrastructure of the Internet consists of a federation of connected networks that are each independently managed ("**autonomous systems**")
  - Note: Each "autononmous system" may consist of multiple IP networks
- Hierarchy of Internet service providers (ISPs)
  - **Tier-1**: Backbone network is a nationwide or worldwide network (US: less than 20)
  - **Tier-2:** regional networks (in US: less than 100)
  - **Tier-3:** local Internet service provider (in US: several thousand)




## Internet Infrastructure

- Location where a network (local or regional ISP, corporate network) gets access to the Internet is called a **Point-of-Presence (POP).**
- Locations (Tier-1 or Tier-2) networks are connected for the purpose of exchanging traffic are called **peering points**.
  - **Public peering**: Traffic is swapped in a specific location, called Internet exchange points (IXPs)
  - **Private peering:** Two networks establish a direct link to each other.

# Topology of a Tier-1 NSP

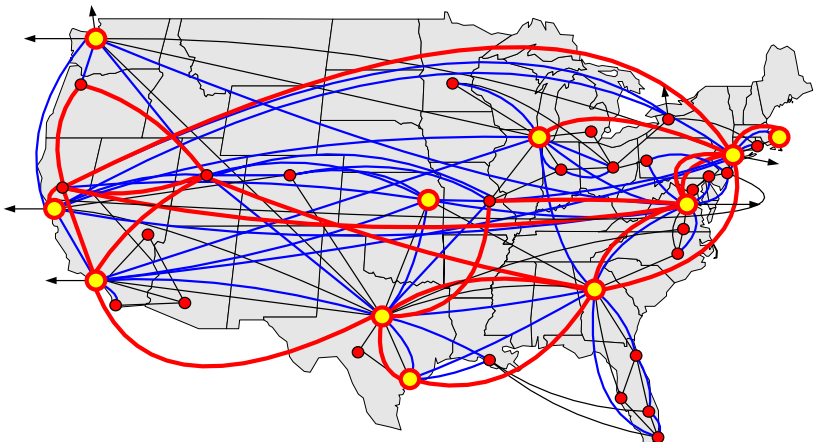





# Organization of a single node in a Tier-1 network

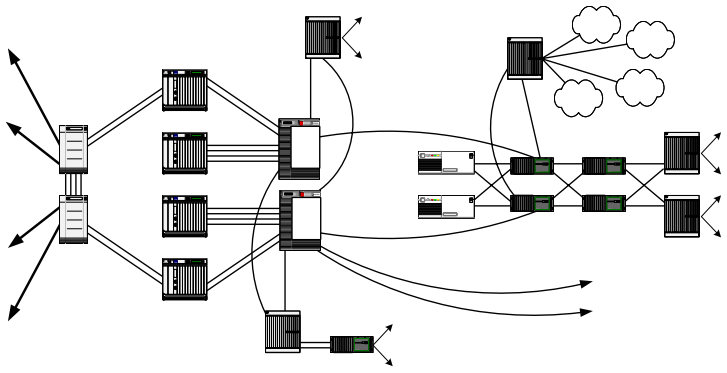

# Addressing in the Internet




# What is an IP Address?

- An IP address is a unique global address for a network interface
  - A computer with 2 network interfaces has 2 IP addresses
- An IP address is a **32 bit long** identifier that encodes a **network number (**or **network prefix)** and a **host number**

## Dotted Decimal Notation

- IP addresses are written in a so-called ***dotted decimal notation***
- Each byte is identified by a decimal number in the range [0..255]:

- **Example:**

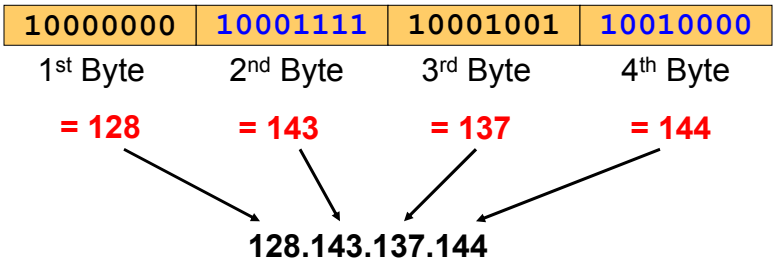





## Network number and Host number

- The network prefix identifies a network and the host number identifies a specific host (actually, interface on the network).
- The part of an IP address that contains the network number is called the **network prefix.**


- **How do we know how long the network prefix is?**
  - The length of the network prefix is often indicated in the prefix notation of an IP address
    - Prefix notation: `128.143.137.144/16`
      → Network has a 16-bit prefix
  - In the past, the network prefix was implicitly defined (**class-based addressing**)

# Example

- **Example**: ellington.cs.virginia.edu


- Prefix notation of IP address of ellington.cs.virginia.edu:
  `128.143.137.144/16`
  - » Network prefix is 16 bits long

- Prefix is: 128.143
- Host number: `137.144`

- IP address of network: `128.143.0.0/16` or 128.143/16
  - → Host number set to zero is the address of the network

- Network mask is: `255.255.0.0` or `ffff000`




# Domain Names and IP Addresses

- Users and applications on the Internet normally do not use IP addresses directly. No one says:

  http://128.143.137.29

- Rather, users and applications use domain names:

  http://www.cs.virginia.edu

- A service on the Internet, called the Domain Name System (DNS) performs the translation between domain names and IP addresses

# Domain Name System (DNS)

- The DNS can be thought of as an Internet-wide database, which translates between Domain names and IP addresses

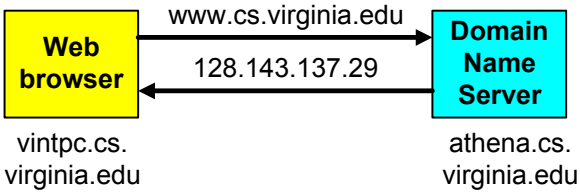


- Many applications can work with both domain names and IP addresses
- The TCP and IP protocols only work with IP addresses




# MAC addresses

- At the data link layer (of local area networks), network interfaces are designated by 6-byte long **MAC addresses**
- Synonyms:
  MAC address
  = physical address
  = hardware address
- Translation between IP addresses and MAC layer addresses is performed by the Address Resolution Protocol (ARP)

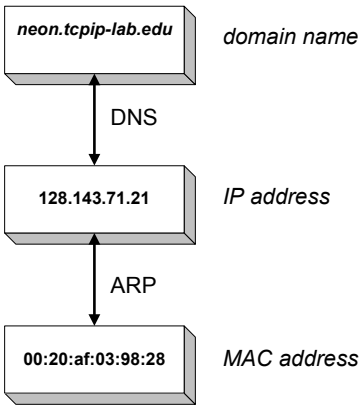

# Finding MAC addresses with ARP

- Router71 needs to find the MAC address of *Neon*.
- **ARP Request** is sent as a broadcast message to all hosts
- Neon responds with an ARP Reply

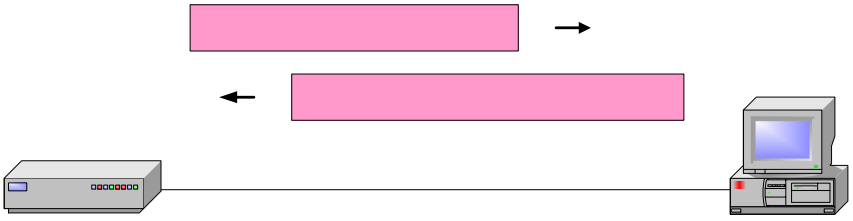





# TCP/IP protocol suite

# TCP/IP Protocol Suite

- The TCP/IP protocol suite is the protocol architecture of the Internet
- The TCP/IP suite has four layers: Application, Transport, Network, and Data Link Layer
- End systems (hosts) implement all four layers. Gateways (Routers) only have the bottom two layers.

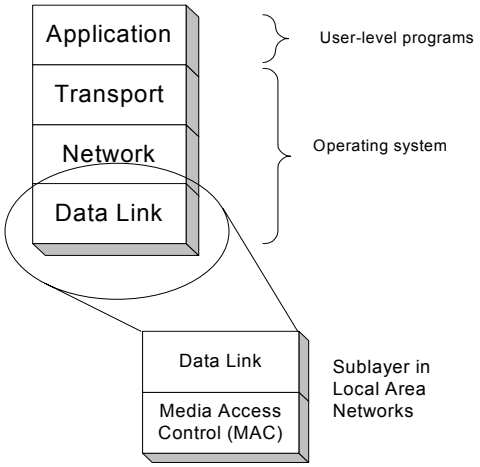





# Functions of the Layers

- Data Link Layer:
  - Service: Reliable transfer of frames over a link<br>Media Access Control on a LAN
  - Functions: Framing, media access control, error checking
- Network Layer:
  - Service: Move packets from source host to destination host
  - Functions: Routing, addressing
- Transport Layer:
  - Service: Delivery of data between hosts
  - Functions: Connection establishment/termination, error control, flow control
- Application Layer:
  - Service: Application specific (delivery of email, retrieval of HTML documents, reliable transfer of file)
  - Functions: Application specific

## Layers in routers and hosts

- IP is the highest layer protocol which is implemented at both routers and hosts

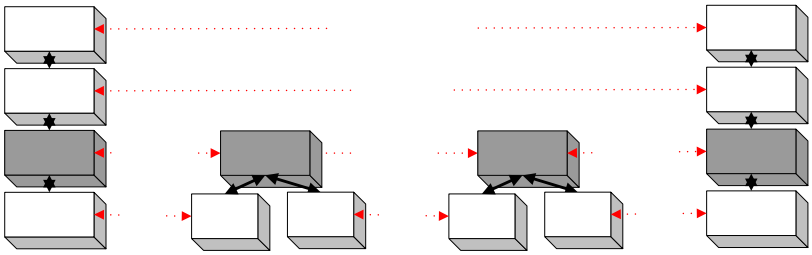





## IP: The waist of the hourglass

- **IP is the waist of the hourglass of the Internet protocol architecture**
- Multiple higher-layer protocols
- Multiple lower-layer protocols
- Only one protocol at the network layer.

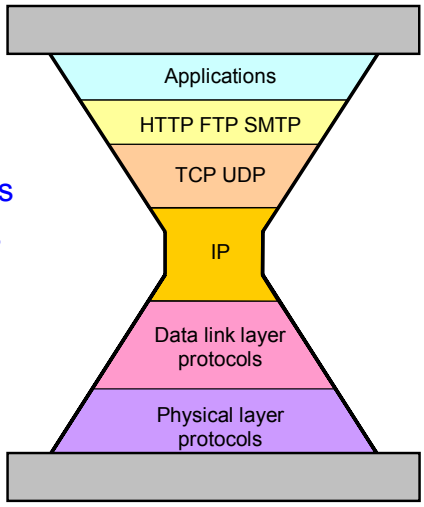






ation

TCP

# Assignment of Protocols to Layers

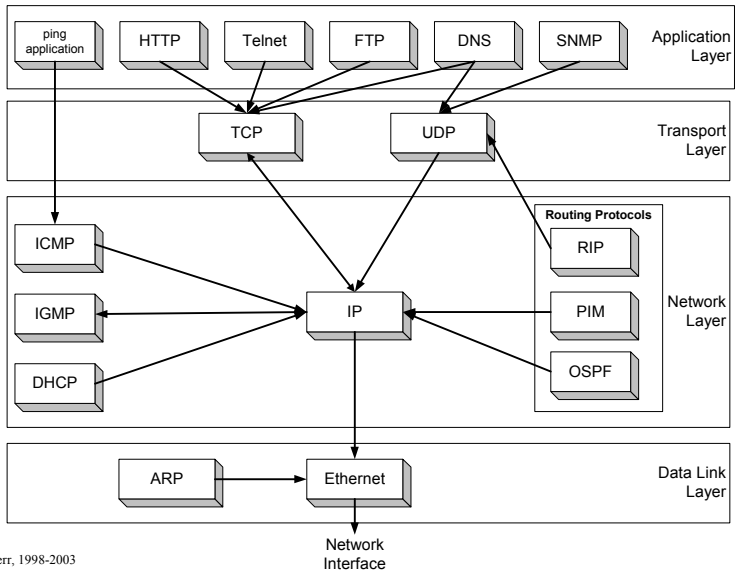






# Sending a packet

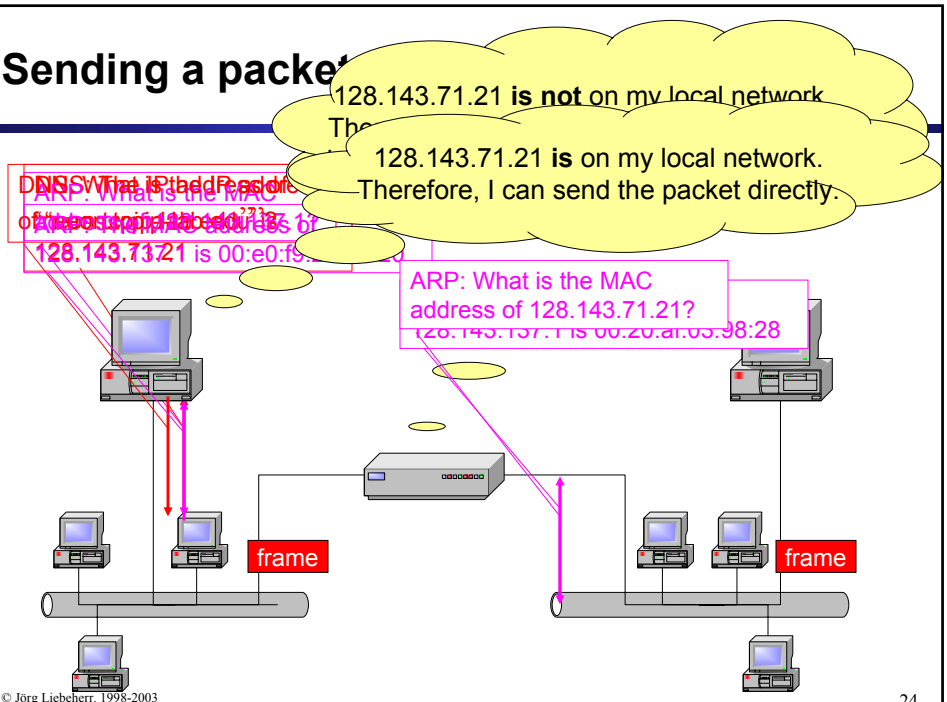

## Layers in the Example

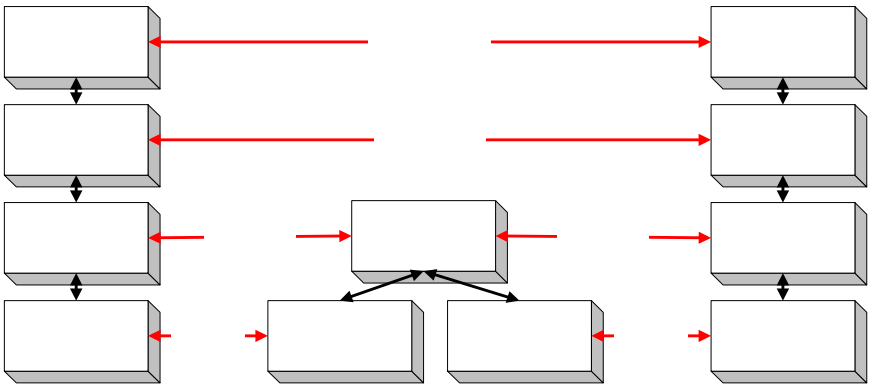





## Layers in the Example

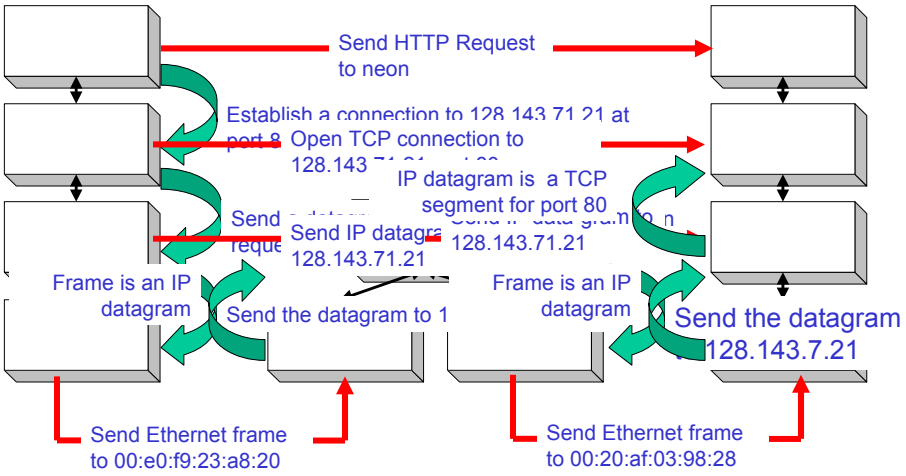






TP

CP

IP

# Encapsulation and Demultiplexing

- As data is moving down the protocol stack, each protocol is adding layer-specific control information

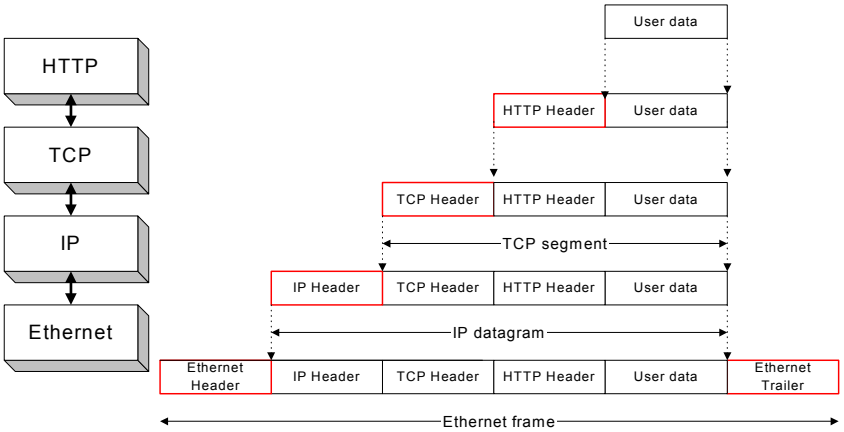





# Encapsulation and Demultiplexing in our Example

- Let us look in detail at the Ethernet frame between Argon and the Router, which contains the TCP connection request to Neon.

- This is the frame in hexadecimal notation.

```
00e0 f923 a820 00a0 2471 e444 0800 4500 002c
9d08 4000 8006 8bff 808f 8990 808f 4715 065b
0050 0009 465b 0000 0000 6002 2000 598e 0000
0204 05b4
```

# Encapsulation and Demultiplexing

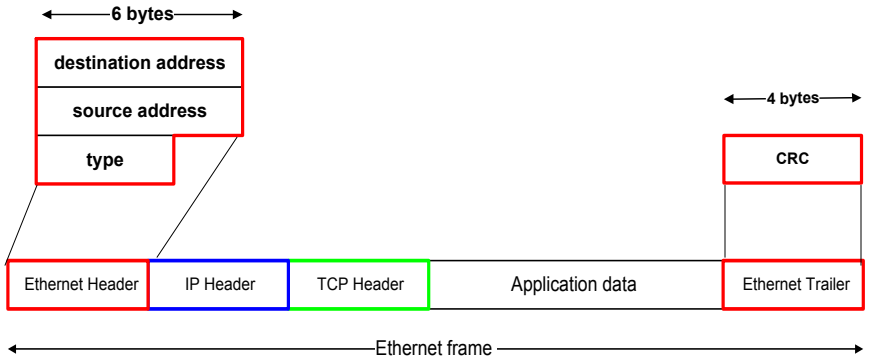





# Encapsulation and Demultiplexing: Ethernet Header

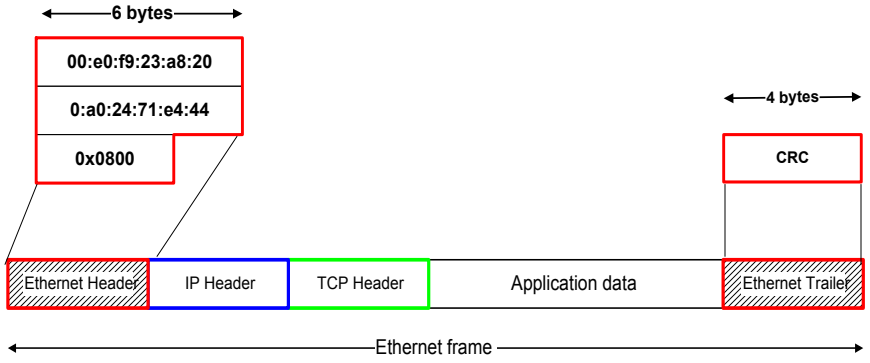

# Encapsulation and Demultiplexing: IP Header

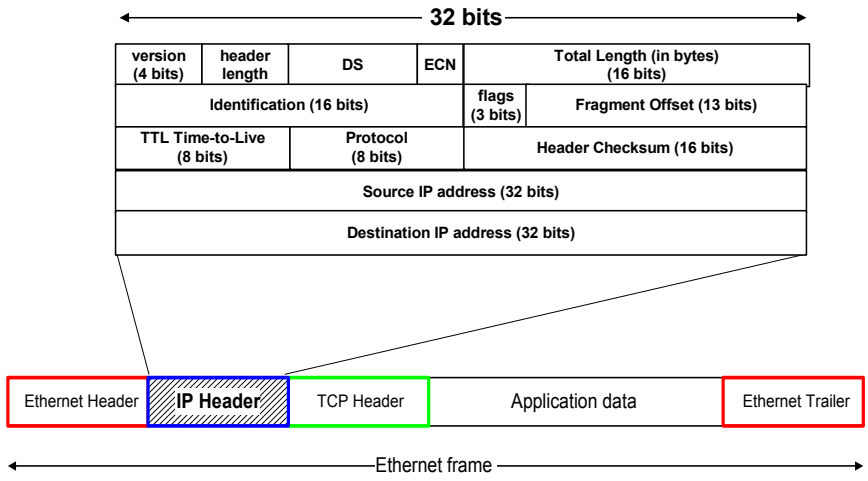






# Encapsulation and Demultiplexing: IP Header

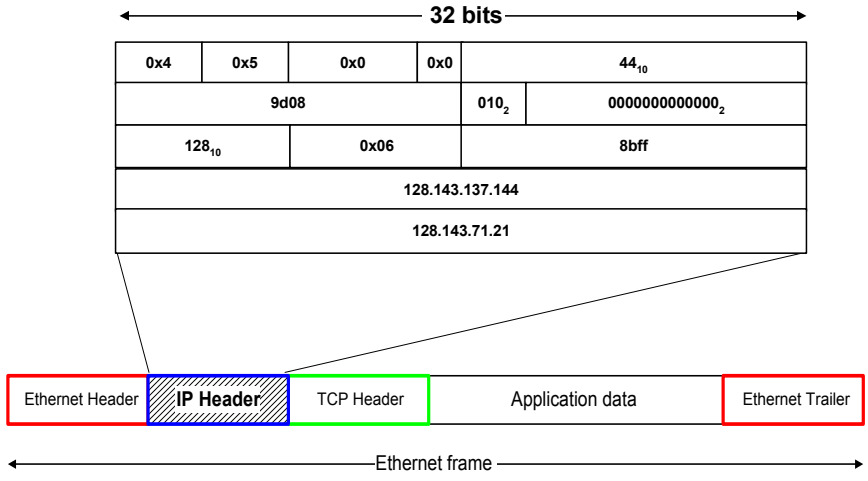

# Encapsulation and Demultiplexing: TCP Header

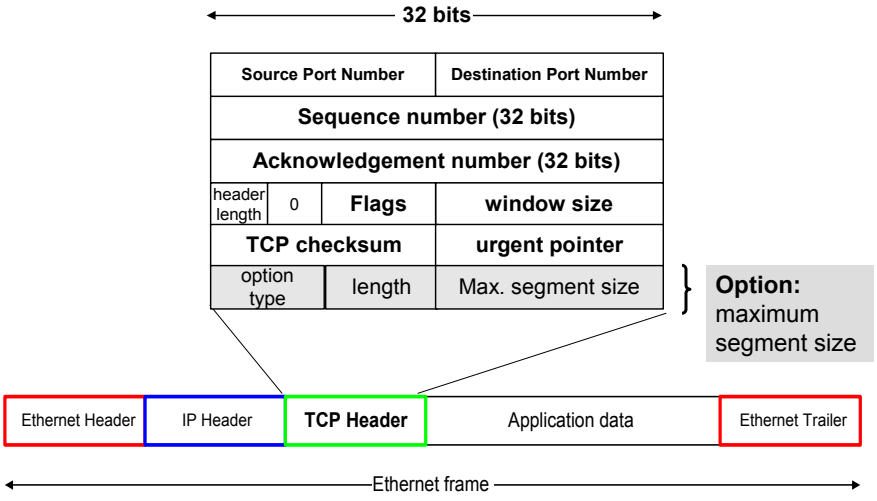






# Encapsulation and Demultiplexing: TCP Header

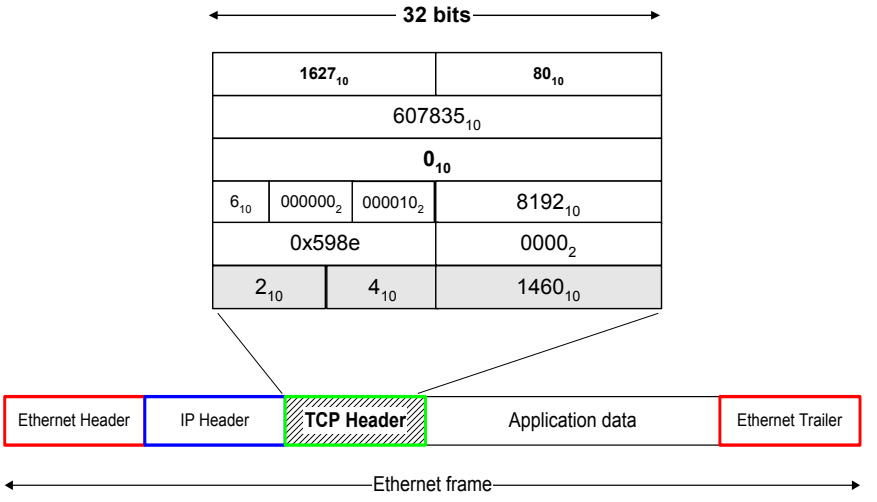

# Encapsulation and Demultiplexing: Application data

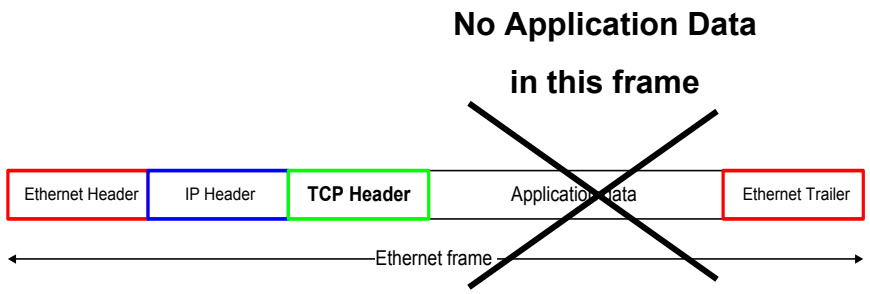





# Forwarding IP datagrams

# IP Forwarding

- Internet is a collection of IP networks
- IP provides an end-to-end delivery service for IP datagrams between hosts
- The delivery service is realized with the help of IP routers

- The forwarding service of IP is:
  - Best effort
  - Connectionless
  - Unreliable

**Important: Don't confuse "collection of IP networks" with "collection of autonomous systems" !**

**The two concepts are very different.**




# Delivery of an IP datagram

- View at the data link layer layer:
  - Internetwork is a collection of LANs or point-to-point links or switched networks that are connected by routers

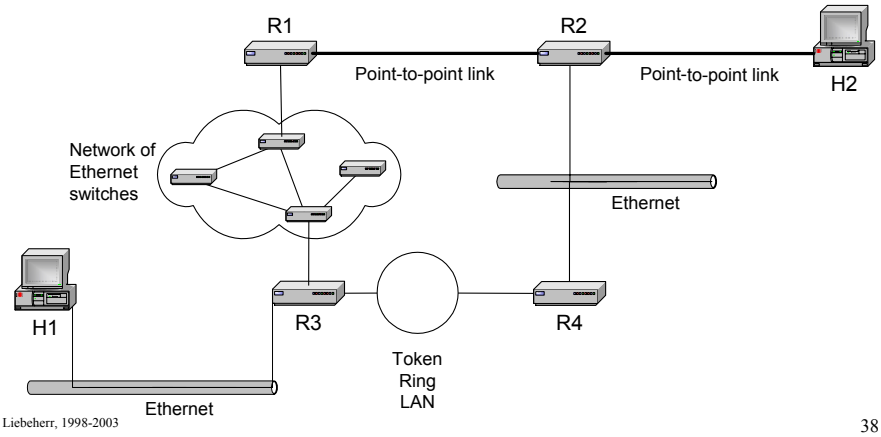

# Delivery of an IP datagram

- View at the IP layer:
  - An IP network is a logical entity with a network number
  - We represent an IP network as a “cloud”
  - The IP delivery service takes the view of clouds, and ignores the data link layer view

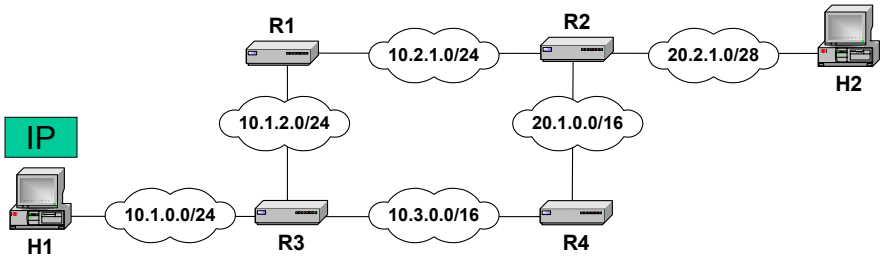





# Tenets of end-to-end delivery of datagrams

The following conditions must hold so that an IP datagram can be successfully delivered

1. The network prefix of an IP destination address must correspond to a unique data link layer network (=LAN or point-to-point link or switched network).
   (The reverse need not be true!)
2. Routers and hosts that have a common network prefix must be able to exchange IP dagrams using a data link protocol (e.g., Ethernet, PPP)
3. Every data link layer network must be connected to at least one other data link layer network via a router.

# Routing tables

- Each router and each host keeps a **routing table** which tells the router how to process an outgoing packet
- Main columns:
  1. **Destination address:** where is the IP datagram going to?
  2. **Next hop or interface:** how to send the IP datagram?
- Routing tables are set so that datagrams gets closer to the its destination

Routing table of a host or router

IP datagrams can be directly delivered ("direct") or is sent to a router ("R4")

| Destination | Next Hop |
|---|---|
| 10.1.0.0/24 | direct |
| 10.1.2.0/24 | direct |
| 10.2.1.0/24 | R4 |
| 10.3.1.0/24 | direct |
| 20.1.0.0/16 | R4 |
| 20.2.1.0/28 | R4 |





# Delivery with routing tables

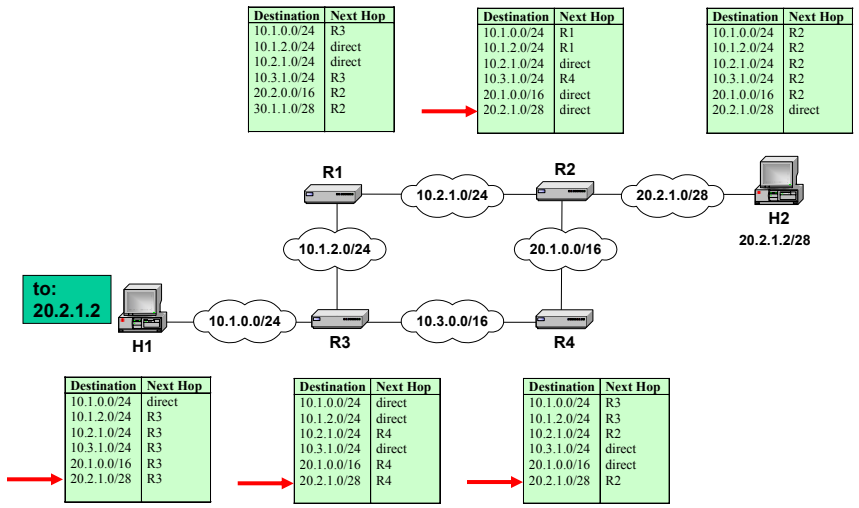

# Delivery of IP datagrams

- There are two distinct processes to delivering IP datagrams:
  1. **Forwarding:** How to pass a packet from an input interface to the output interface?
  2. **Routing:** How to find and setup the routing tables?




# Processing of an IP datagram in IP

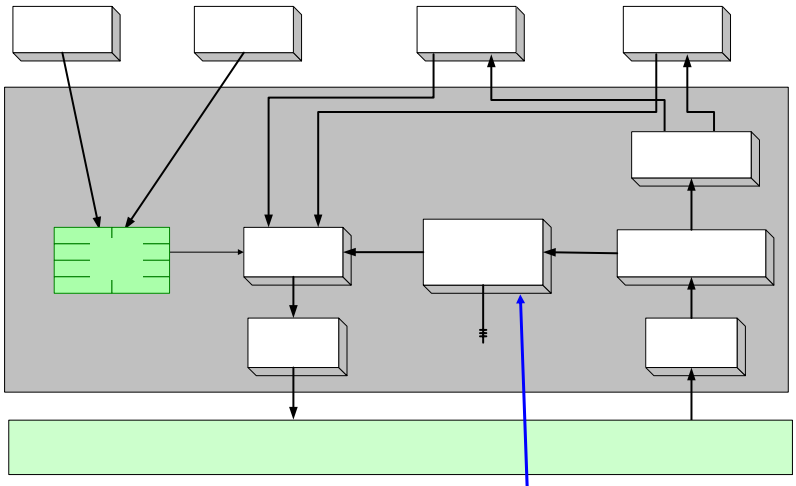

IP router: IP forwarding enabled
Host: IP forwarding disabled

# Routing table lookup

- When a router or host need to transmit an IP datagram, it performs a routing table lookup
- **Routing table lookup:** Use the IP destination address as a key to search the routing table.
- Result of the lookup is the IP address of a next hop router, or the name of a network interface

| Destination address | Next hop |
|---|---|
| network prefix<br>*or*<br>host IP address<br>*or*<br>loopback address<br>*or*<br>default route | IP address of next hop router<br><br>*or*<br><br>Name of a network interface |




# Longest Prefix Match

**Longest Prefix Match:** Search for the routing table entry that has the longest match with the prefix of the destination IP address

1. Search for a match on all 32 bits
2. Search for a match for 31 bits

   …..

32. Search for a mach on 0 bits

Host route, loopback entry
→ 32-bit prefix match

Default route is represented as 0.0.0.0/0
→ 0-bit prefix match

**128.143.71.21**

| Destination address | Next hop |
|---|---|
| 10.0.0.0/8 | R1 |
| 128.143.0.0/16 | R2 |
| 128.143.64.0/20 | R3 |
| 128.143.192.0/20 | R3 |
| 128.143.71.0/24 | R4 |
| 128.143.71.55/32 | R3 |
| default | R5 |

**The longest prefix match for 128.143.71.21 is for 24 bits with entry 128.143.71.0/24**

**Datagram will be sent to R4**

## Design Principles of the Internet Architecture

- Run over multiple (link layer) network technologies
- Distributed organization and control
- Robustness to partial failures

- **End-to-end principle**: Keep functionality inside the network (layer 3) simple. Implement complex functions outside the network




## Internet talking points

- Address depletion → IPv6, Network address translation
- Enhancing the basic best effort service → Internet QoS
- Convergence of routing protocols → BGP convergence
- Security → DDoS, Encryption