


The Internet

An Engineering Approach to Computer Networking

My how you've grown!

- The Internet has doubled in size every year since 1969
- By 1996, 10 million computers joined the Internet
- Since then:



What does it look like?

- Loose collection of networks organized into a multilevel hierarchy
 - ◆ 10-100 machines connected to a *hub* or a *router*
 - ◆ service providers also provide direct dialup access
 - ◆ or over a wireless link
 - ◆ 10s of routers on a *department backbone*
 - ◆ 10s of department backbones connected to *campus backbone*
 - ◆ 10s of campus backbones connected to *regional service providers*
 - ◆ 100s of regional service providers connected by *national backbone*
 - ◆ 10s of national backbones connected by *international trunks*

Example of message routing

```
# traceroute henna.iitd.ernet.in
traceroute to henna.iitd.ernet.in (202.141.64.30), 30 hops max, 40 byte packets
 1  UPSON2-NP.CIT.CORNELL.EDU (128.84.154.1)  1 ms  1 ms  1 ms
 2  HOL1-MSS.CIT.CORNELL.EDU (132.236.230.189)  2 ms  3 ms  2 ms
 3  CORE1-MSS.CIT.CORNELL.EDU (128.253.222.1)  2 ms  2 ms  2 ms
 4  CORNELLNET1.CIT.CORNELL.EDU (132.236.100.10)  4 ms  3 ms  4 ms
 5  ny-ith-1-H1/0-T3.nysernet.net (169.130.61.9)  5 ms  5 ms  4 ms
 6  ny-ith-2-F0/0.nysernet.net (169.130.60.2)  4 ms  4 ms  3 ms
 7  ny-pen-1-H3/0-T3.nysernet.net (169.130.1.121)  21 ms  19 ms  16 ms
 8  sl-pen-21-F6/0/0.sprintlink.net (144.228.60.21)  16 ms  40 ms  36 ms
 9  core4-hss15-0.WestOrange.mci.net (206.157.77.105)  20 ms  20 ms  24 ms
10  core2.WestOrange.mci.net (204.70.4.185)  21 ms  34 ms  26 ms
11  border7-fddi-0.WestOrange.mci.net (204.70.64.51)  21 ms  21 ms  21 ms
12  vsnl-poone-512k.WestOrange.mci.net (204.70.71.90)  623 ms  639 ms  621 ms
13  202.54.13.170 (202.54.13.170)  628 ms  629 ms  628 ms
14  144.16.60.2 (144.16.60.2)  1375 ms  1349 ms  1343 ms
15  henna.iitd.ernet.in (202.141.64.30)  1380 ms  1405 ms  1368 ms
```

Intranet, Internet, and Extranet

- Intranets are administered by a single entity
 - ◆ e.g. Company or University
- Internet is administered by a coalition of entities
 - ◆ name services, backbone services, routing services etc.
- Extranet refers to a common use pattern
 - ◆ refers to exterior customers who can access privileged Intranet services
 - ◆ e.g. Car manufacturer makes services available to dealers...

What holds the Internet together?

- Addressing
 - ◆ how to refer to a machine on the Internet
- Routing
 - ◆ how to get there
- Internet Protocol (IP)
 - ◆ what to speak to be understood

Example: joining the Internet

- How can people talk to you?
 - ◆ get an IP address from your administrator
- How do you know where to send your data?
 - ◆ if you only have a single external connection, then no problem
 - ◆ otherwise, need to speak a routing protocol to decide next hop
- How to format data?
 - ◆ use the IP format so that intermediate routers can understand the destination address
- If you meet these criteria--you're on the Internet!
- Decentralized, distributed, and chaotic
 - ◆ but it scales (why?)

What lies at the heart?

- Two key technical innovations
 - ◆ packets
 - ◆ store and forward

Packets

- Self-descriptive data
 - ◆ packet = data + metadata (header)
- Packet vs. sample
 - ◆ samples are not self descriptive
 - ◆ to forward a sample, we have to know *where* it came from and *when*
 - ◆ can't store it!
 - ◆ hard to handle bursts of data

Store and forward

- Metadata allows us to forward packets when we want
- E.g. letters at a post office headed for main post office
 - ◆ address labels allow us to forward them in batches
- Efficient use of critical resources
- Three problems
 - ◆ hard to control delay within network
 - ◆ switches need memory for buffers
 - ◆ convergence of flows can lead to congestion

Key features of the Internet

- Addressing
- Routing
- Endpoint control

Addressing

- Internet addresses are called IP addresses
- Refer to a *host interface*: need one IP address per interface
- Addresses are structured as a two-part hierarchy
 - ◆ network number
 - ◆ host number

135.105.53	100
------------	-----

An interesting problem


- How many bits to assign to host number and how many to network number?
- If many networks, each with a few hosts, then more bits to network number
- And *vice versa*
- But designer's couldn't predict the future
- Decided three sets of partitions of bits
 - ◆ class A: 8 bits network, 24 bits host
 - ◆ class B: 16 bits each
 - ◆ class C: 24 bits network, 8 bits host

Addressing (contd.)

- To distinguish among them
 - ◆ use leading bit
 - ◆ first bit = 0 => class A
 - ◆ first bits 10 => class B
 - ◆ first bits 110 => class C
 - ◆ (what class address is 135.104.53.100?)
- Problem
 - ◆ if you want more than 256 hosts in your network, need to get a class B, which allows 64K hosts => wasted address space
- Solution
 - ◆ associate *every* address with a *mask* that indicates partition point
 - ◆ *CIDR*

Routing

- How to get to a destination given its IP address?
- We need to know the next hop to reach a particular network number
 - ◆ this is called a *routing table*
 - ◆ computing routing tables is non-trivial
- Simplified example



Default routes

- Strictly speaking, need next hop information for every network in the Internet
 - ◆ > 80,000 now
- Instead, keep detailed routes only for local neighborhood
- For unknown destinations, use a *default* router
- Reduces size of routing tables at the expense of non-optimal paths

Endpoint control

- Key design philosophy
 - ◆ do as much as possible at the edge
 - ◆ dumb network
 - ◆ exactly the opposite philosophy of telephone network
- Layer above IP compensates for network defects
 - ◆ Transmission Control Protocol (TCP)
- Can run over any available link technology
 - ◆ but no quality of service
 - ◆ modification to TCP requires a change at every endpoint
 - ◆ (how does this differ from telephone network?)

Challenges

- IP address space shortage
 - ◆ because of free distribution of inefficient Class B addresses
 - ◆ decentralized control
 - ◆ hard to recover addresses, once handed out
- Decentralization
 - ◆ allows scaling, but makes *reliability* next to impossible
 - ◆ cannot guarantee that a route exists, much less bandwidth or buffer resources
 - ◆ single points of failure can cause a major disaster
 - ◆ and there is no control over who can join!
 - ◆ hard to guarantee security
 - ◆ end-to-end encryption is a partial solution
 - ◆ who manages keys?

Challenges (contd.)

- Decentralization (contd.)
 - ◆ no uniform solution for accounting and billing
 - ◆ can't even reliably identify individual users
 - ◆ no equivalent of white or yellow pages
 - ◆ hard to reliably discover a user's email address
 - ◆ nonoptimal routing
 - ◆ each administrative makes a locally optimal decision

Challenges (contd.).

- Multimedia
 - ◆ requires network to support quality of service of some sort
 - ◆ hard to integrate into current architecture
 - ◆ store-and-forward => shared buffers => traffic interaction => hard to provide service quality
 - ◆ requires endpoint to signal to the network what it wants
 - ◆ but Internet does not have a simple way to identify streams of packets
 - ◆ nor are routers required to cooperate in providing quality
 - ◆ and what about pricing!