

Chapter 1 – Introduction

Historical perspective

Evolved from DARPA funded work on Packet Switched Networks in 1960's.
Current architecture and protocols emerged in 1977 – 1980 time frame as the ARPANET
Research into supporting attachment of other networks began about this time.
Use of TCP/IP and evolution to the internet began about 1980.
Berkeley Unix led to wide–spread adoption by Universities.
Domain Name System (DNS) added in the middle 80's.

Administration

1979 – ICCB (Internet Control and Configuration Board) founded by ARPA
1983 – IAB (Internet Architecture Board)
 Provide focus and coordination for protocol development
 Comprised of the heads of (about 10) ITFs.
1989 – Reorganized IAB based on transition from *research* to *production* system.
 New subgroups
 IRTF – (Research task force) SG heads + IRTF head
 IRSG (Steering group)
 Research groups
 IEFT – (Engineering task force)
 IESG (Steering group) – IETF Head + Area mgrs)
 Areas 1 -- n (~10)
 Working groups
1992 – ISOC (Internet Society) emerged as an umbrella administrative agency

Web resources

Internet Society	www.isoc.org
Internet Architecture Board	www.iab.org
Internet Engineering Task Force	www.ietf.org
Internet Research Task Force	www.irtf.org
Internet Assigned Numbers Authority	www.iana.org
Internet Corp. for Assigned Names and Numbers	www.icann.org
American Registry for Internet Numbers	www.arin.net
Asia Pacific Network Information Centre	www.apnic.net
Resaux IP Europeen	www.ripe.net
Latin–American and Caribbean IP Address Registry	www.lacnic.net

Conclusion

Bureaucracies create new layers as least as fast as network architectures!

IETF Areas and Working Groups

Applications Area Working Groups:

acap	Application Configuration Access Protocol
calsch	Calendaring and Scheduling
crisp	Cross Registry Information Service Protocol
ediint	Electronic Data Interchange–Internet Integration
fax	Internet Fax
ftptext	Extensions to FTP
geopriv	Geographic Location/Privacy
imapext	Internet Message Access Protocol Extension
impp	Instant Messaging and Presence Protocol
ipp	Internet Printing Protocol

General Area Working Groups:

ipr	Intellectual Property Rights
nomcom	Operation of the IESG/IAB Nominating and Recall Committees
problem	Problem Statement

Internet Area Working Groups:

Area Director(s):

[Thomas Narten <narten@us.ibm.com>](mailto:narten@us.ibm.com)
[Erik Nordmark <erik.nordmark@sun.com>](mailto:erik.nordmark@sun.com)

atommib	AToM MIB
dhc	Dynamic Host Configuration
dnsext	DNS Extensions
eap	Extensible Authentication Protocol
ifmib	Interfaces MIB
ipcdn	IP over Cable Data Network
ipoib	IP over InfiniBand
iporpr	IP over Resilient Packet Rings
ipv6	IP Version 6 Working Group

Operations and Management Area Working Groups:

aaa	Authentication, Authorization and Accounting
adslmib	ADSL MIB
bmwg	Benchmarking Methodology
bridge	Bridge MIB
disman	Distributed Management
dnsop	Domain Name System Operations
entmib	Entity MIB

Routing Area Working Groups:

bgmp	Border Gateway Multicast Protocol
forces	Forwarding and Control Element Separation
idmr	Inter-Domain Multicast Routing
idr	Inter-Domain Routing
isis	IS-IS for IP Internets

Security Area Working Groups:

idwg	Intrusion Detection Exchange Format
inch	Extended Incident Handling
ipsec	IP Security Protocol
ipseckey	IPSEC KEYing information resource record
ipsp	IP Security Policy
kink	Kerberized Internet Negotiation of Keys
krb-wg	Kerberos WG
msec	Multicast Security

Sub-IP Area Working Groups:

ccamp	Common Control and Measurement Plane
gsmg	General Switch Management Protocol
ipo	IP over Optical
mpls	Multiprotocol Label Switching
ppvnp	Provider Provisioned Virtual Private Networks
tewg	Internet Traffic Engineering

Transport Area Working Groups:

avt	Audio/Video Transport
dccp	Datagram Congestion Control Protocol
enum	Telephone Number Mapping
ieprep	Internet Emergency Preparedness
ippm	IP Performance Metrics
ips	IP Storage
iptel	IP Telephony

IRTF Active Research Groups

- * Anti-Spam
- * Authentication Authorisation Accounting Architecture
- * Crypto Forum
- * Delay-Tolerant Networking
- * End-to-End
- * Group Security
- * Internet Measurement
- * Network Management
- * NameSpace
- * Peer-to-Peer
- * Routing
- * Searchable Internet Resource Names
- * Services Management

IP Address allocation

Both IPv4 and IPv6 addresses are assigned in a delegated manner.

Users are assigned IP addresses by Internet service providers (ISPs).

ISPs obtain allocations of IP addresses from a local Internet registry (LIR) or national Internet registry (NIR), or from their appropriate Regional Internet Registry (RIR):

APNIC (Asia Pacific Network Information Centre) – Asia/Pacific Region

ARIN (American Registry for Internet Numbers) – North America and Sub-Saharan Africa

LACNIC (Regional Latin-American and Caribbean IP Address Registry) – Latin America and some Caribbean Islands

RIPE NCC (Réseaux IP Européens) – Europe, the Middle East, Central Asia, and African countries located north of the equator

The IANA's role is to allocate IP addresses from the pools of unallocated addresses to the RIRs according to their established needs. When an RIR requires more IP addresses for allocation or assignment within its region, the IANA makes an additional allocation to the RIR.

The Internet Corporation for Assigned Names and Numbers (ICANN) is the non-profit corporation that was formed to assume responsibility for the IP address space allocation, protocol parameter assignment, domain name system management, and root server system management functions previously performed under U.S. Government contract by IANA and other entities.

Growth of the Internet

Date	Hosts		Date	Hosts	Networks	Domains
12/69	4		07/89	130,000	650	3,900
06/70	9		10/89	159,000	837	
10/70	11		10/90	313,000	2,063	9,300
12/70	13		01/91	376,000	2,338	
04/71	23		07/91	535,000	3,086	16,000
10/72	31		10/91	617,000	3,556	18,000
01/73	35		01/92	727,000	4,526	
06/74	62		04/92	890,000	5,291	20,000
03/77	111		07/92	992,000	6,569	16,300
12/79	188		10/92	1,136,000	7,505	18,100
08/81	213		01/93	1,313,000	8,258	21,000
05/82	235		04/93	1,486,000	9,722	22,000
08/83	562		07/93	1,776,000	13,767	26,000
10/84	1,024		10/93	2,056,000	16,533	28,000
10/85	1,961		01/94	2,217,000	20,539	30,000
02/86	2,308		07/94	3,212,000	25,210	46,000
11/86	5,089		10/94	3,864,000	37,022	56,000
12/87	28,174		01/95	4,852,000	39,410	71,000
07/88	33,000		07/95	6,642,000	61,538	120,000
10/88	56,000		01/96	9,472,000	93,671	240,000
01/89	80,000		07/96	12,881,000	134,365	488,000
			01/97	16,146,000		828,000
			07/97	19,540,000		1,301,000

Month	1998	1999	2000
Jan	30.3514	44.2211	70.2063
Feb	31.3885	48.0524	72.9320
Mar	32.4240	50.4706	75.0822
Apr	33.5683	53.4577	77.1310
May	34.5121	55.7970	80.1703
Jun	35.4462	57.3369	82.8540
Jul	36.4815	59.2785	85.8201
Aug	37.2117	61.3473	88.6047
Sep	38.5580	63.3193	91.5382
Oct	39.9277	65.6249	94.2421
Nov	41.5121	67.3009	96.5888
Dec	43.0256	68.7849	
Africa		18.469	2388.86
Asia		7173.72	54912.0
Europe		5854.3	67547.7
Oceania		823.8	14409.4
Central America		412.34	1334.75
South America		1015.95	13185.5
North America		42857.7	118694.

Review of Network Principles and Terminology

Network architecture

A collection of service, interface, and protocol definitions that provide a means for communication among elements of a distributed system.

Service classes

Connectivity

Connection-oriented (a.k.a. virtual circuit) – Like a telephone

Connection-less (a.k.a. datagram) – Like mail

Reliability

Unreliable

Loses packets

Duplicates packets

Delivers packets with inverted bits

Reorders the delivery of packets

Reliable

Doesn't do any of the above.

Protocols

Rules that define the format, meaning, and sequencing of messages on a network

Layered protocol design

Several competing layered designs exist

Objective is improved modularization of network software for

– faster development

– more reliability

– easier to add new functionality

Elements in each layer

Communicate over logical channels

With counterparts in the same layer at other nodes

Using *Peer Protocols*

TCP/IP

A term used as a name for a network architecture (like SNA or ISO/OSI)
Derived from the names of its two most well known protocols
Consists of 4 layers

Link

Physical encoding, transmission and receiving of bits of information
May include error detecting or correcting
Includes adapter card, microcode, and any device drivers.
Protocols include ARP, RARP, NDIS, IEEE 802.2, IEEE 802.3
Absence of a specification here was an important factor in TCP/IP dominance

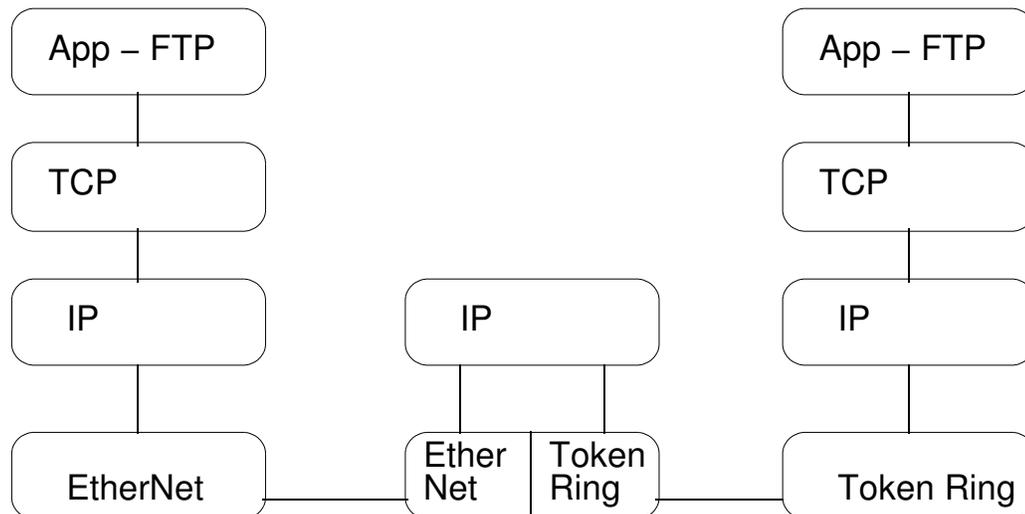
Network

Routing of packets and congestion control.
Fragmentation and reassembly.
Includes IP, ICMP, and IGMP (Internet Group Management Protocol)
IP provides unreliable connectionless service
Simplicity of the specification here was a key factor in TCP/IP dominance

Transport

Provides the network API (application level programming interface).
Reliable (TCP) end to end transport of data.
Break messages into packets and reassemble.
Handle setup and take down of sessions.
TCP provides reliable connection oriented service
UDP provides unreliable connection less service

Application



Header encapsulation and removal

Link sees

Ethernet header	14	bytes	(RFC 894)
IP header	20(+)	bytes	
TCP header	20(+)	bytes	
Application data	?	bytes	
Ethernet trailer	4	bytes	

Net sees

IP header	20(+)	bytes	
TCP header	20(+)	bytes	
Application data	?	bytes	
Ethernet trailer	4	bytes	

etc.

Addressing in the Internet

A number of address classes are relevant

Link:	Every Ethernet adapter has a 48 bit address
Network:	Hosts may support one <i>or more</i> network level or IP addresses (at least one per network to which host is attached.)
Transport:	Applications connect to 16 bit port addresses

The *one* address which must be unique within the world is the IP address

IP addresses

32 bits — To be expanded in IPV6 to 128 bits

Written in dotted decimal notation 130.127.48.118

IP addresses were originally class oriented (A, B, C...)

A – 0	7 bit netid	24 bit hostid
B – 10	14 bit netid	16 bit hostid
C – 110	21 bit netid	8 bit hostid
D – 1110	28 bit multicast id	
E – 11110	27 bits reserved for future use	

Wide area routing is now classless..

CIDR (network number/prefix length)

Special network ids

<i>Net part</i>	<i>Host part</i>	
All	zeros	This host (source only)
All zeros	hostid	Host on this net
All	ones	Local broadcast
Netid	all ones	Directed b'cast
127	anything (1)	Loopback

Weaknesses in original TCP/IP addressing scheme

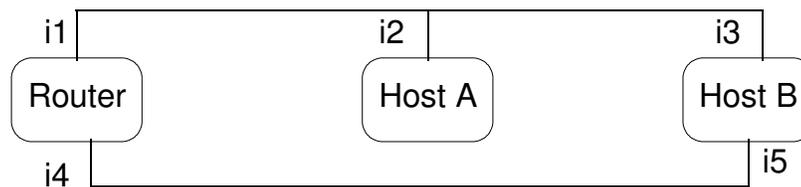
Running out of addresses because routing / address space interaction constrains address usage (now largely resolved via CIDR and NAT).

This interaction also means if a host changes networks => its IP address must change

Have to change all addresses in net when migrating from class C to class B.

Since addresses are associated with *interfaces not hosts* loss of a link can make it impossible to reach a multihomed host.

Example:



If Host A apps normally talk to Host B apps over i3 and i3 fails, recovery may require *application level* or *DNS* changes

App makes a connection using the name Host_B

DNS maps the name HostB to an IP address bound to i3

B may have another DNS name Host_B_i5

DNS maps the name HostB_i5 to an IP address bound to i5

So the app can change names or someone can change DNS bindings

Internet level routing is done by (Network number/Prefix length)

A prefix consists of the n most significant bits of a 32 bit IP address
Longest prefixing matching a given destination address wins
Core routers have large tables of the form (*prefix, pfxlen, next-hop*)

Network/Autonomous System level routing

Routing within a specific network is the responsibility of that network.
Subnetworks are often used in *intra* network routing.
Partitioning of `hostid` into `subnet:host` is responsibility of the subnet.

Local network (LAN) level routing

Mapping of IP addresses to MAC addresses is responsibility of the local net.
The *ARP* Protocol performs this function.

Transport layer routing

Packets are actually routed to applications using

Destination IP address
Protocol id (e.g. TCP or UDP)
Port number
Sender IP address and sender port number
e.g. Telnet Server listens on TCP port 23

Assigned ports now lie in range 0 – 1023 and are managed by IANA (or ICANN)
Some ports above this range are also allocated (see RFC 1340)

Extracted from RFC 1340:

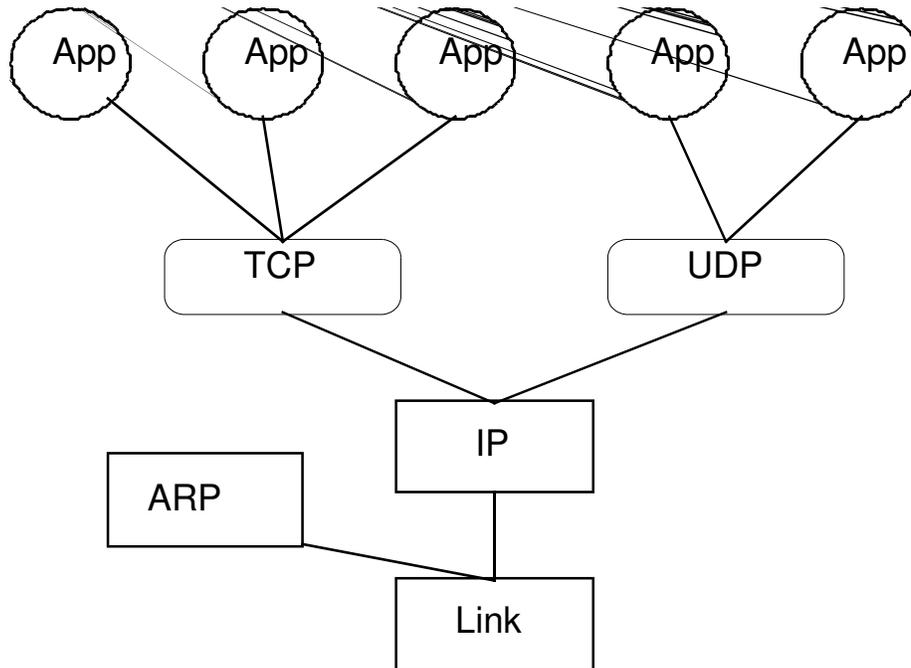
ftp-data	20/tcp	File Transfer [Default Data]
ftp-data	20/udp	File Transfer [Default Data]
ftp	21/tcp	File Transfer [Control]
ftp	21/udp	File Transfer [Control]
	22/tcp	Unassigned
	22/udp	Unassigned
telnet	23/tcp	Telnet
telnet	23/udp	Telnet
	24/tcp	any private mail system
	24/udp	any private mail system
smtp	25/tcp	Simple Mail Transfer
smtp	25/udp	Simple Mail Transfer
	26/tcp	Unassigned
	26/udp	Unassigned
nsw-fe	27/tcp	NSW User System FE
nsw-fe	27/udp	NSW User System FE

Other port numbers are "registered" but not permanently assigned:
Registered numbers also appear in RFC 1340

Keyword	Decimal	Description
-----	-----	-----
blackjack	1025/tcp	network blackjack
blackjack	1025/udp	network blackjack
hermes	1248/tcp	
hermes	1248/udp	
bbn-mmc	1347/tcp	multi media conferencing
bbn-mmc	1347/udp	multi media conferencing
bbn-mmx	1348/tcp	multi media conferencing
bbn-mmx	1348/udp	multi media conferencing
sbook	1349/tcp	Registration Network Protocol
sbook	1349/udp	Registration Network Protocol
editbench	1350/tcp	Registration Network Protocol
editbench	1350/udp	Registration Network Protocol
equationbuilder	1351/tcp	Digital Tool Works (MIT)
equationbuilder	1351/udp	Digital Tool Works (MIT)
lotusnote	1352/tcp	Lotus Notes
lotusnote	1352/udp	Lotus Notes

Multiplexing and demultiplexing

Within a host specific instances of the protocol stack actually form a tree



Multiplexing occurs as packets travel toward root.

Demultiplexing as they arrive at destinations

<i>Level</i>	<i>Based on</i>
Link	Frame type in ethernet
IP	Protocol type in IP header
TCP/UDP	(dest port, source port, source IP, dest IP)

Naming in the internet

Names provide a "user friendly form of address"

Names are managed in a distributed data base using the DNS protocol

Domain name services

.edu

.com

.net

.in

Exercises:

1.1, 1.2, 1.3, 1.4, 1.6, 1.8