The Future of IP

Scott Bradner
Harvard University
sob@harvard.edu
5/18/96

Whoda Thunk

- 1974 - TCP/IP by Cerf & Kahn
  datagram wrong answer - should be connection
- 1982 - ARPANET adopts TCP/IP
  telcos - X.25 right answer
- 1988 - T1 NSFnet
  OSI figures it out
- 1991 - T3 NSFnet, 1st ISPs
  GOSIP mandate
- 1992 - getting too big
  GOSIP is the answer
- 1994 - IPv6
  no response

History

- ramp approaching vertical
- > 100% growth / year for most metrics

	'90	'93	'95
hosts	100K	1.2M	7M
domans	1K	20K	120K
.com	500	9K	65K

What is the Internet?

- Separately identifiable data network
- Hype topic
- Wall Street crack
- Security worry
- Reliability worry
- TCP/IP

Future?

- year 2000
  - 250M PCs
  - 200M Internet
- Internet *is* the GII (and NII)
- part of basic daily life
- too damn big?
- TCP/IP!

Background - The Need for An IPng

- August 1990
  - projected exhaustion of Class B space by March 1994
- 32 bit address space can identify 4 billion hosts
  - assignment inefficiency reduces utilization (RFC-1715)
  - use of classful addresses reduces efficiency
- Routing table bloat
  - table space increasing faster than memory technology
  - thus can not just use multiple Class C addresses

Available Timeframe

- Address Lifetime Expectations (ALE) working group
  - Frank Solensky, FTP Software <solensky@ftp.com>
  - Tony Li, Cisco Systems <tli@cisco.com>
- Made prediction at Seattle, Toronto & San Jose IETF meetings
  - 2005 - 2011
- Mixed view of confidence level
  - questions on base data & assumes no paradigm shifts
  - routing tables are still going to be a problem
- CIDR helps
- Projection at Danvers IETF meeting pushes out time

Classless InterDomain Routing (CIDR)

- Aggregate routing information
- Assign addresses in power-of-two chucks
- Advertise power-of-two sized chuck of address space per entry
  - all of provider's customers can be aggregated into one advertisement
  - reduce size & rate of growth of routing table
- Some issues
  - assumes customers renumber to provider address range
  - tends to bind customer to a provider
  - problem with multi-homed customers
- It works, up to a point

Routing Table Size

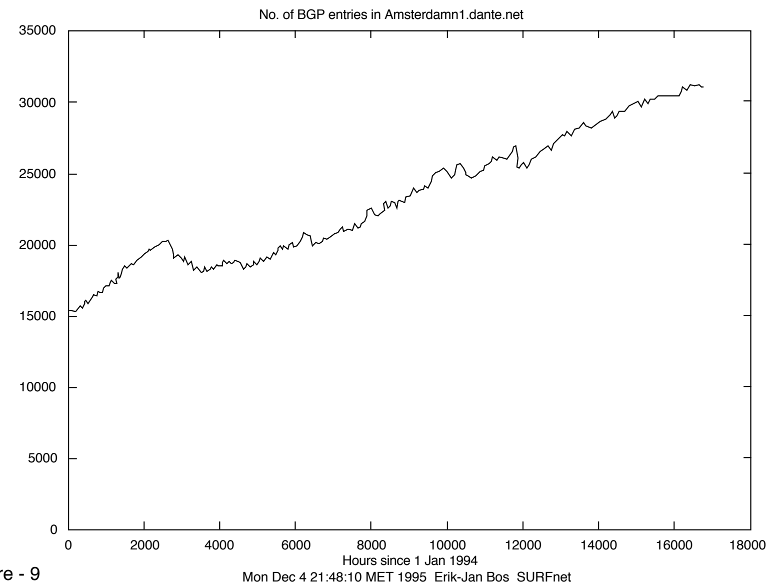

Scope of IPng

- Development, testing & deployment will take time
- Still we seem to have adequate time in IPv4 address space but not excessive

  (excluding paradigm shifts)
- Can do more than "just" fix addresses
- Use requirements process to determine actual scope of IPng effort

White Papers

Adamson, B.	Tactical Radio Frequency Communication Requirements for IPng, RFC 1677
Bellovin, S.	On Many Address per Host, RFC 1681
Bellovin, S.	Security Concerns for IPng, RFC 1675
Bound, J.	IPng BSD Host Implementation Analysis, RFC 1682
Brazdziunas, C.	IPng Support for ATM Services, RFC 1680
Britton, E. et.al.	IPng Requirements for Large Corporate Networks, RFC 1678
Brownlee, J.	Accounting Requirements for IPng, RFC 1672
Carpenter, B.	IPng White Paper on Transition and Other Considerations, RFC 1671
Chiappa, J.N.	IPng Tech. Req. of the Nimrod Routing and Addressing Architecture, RFC 1753
Clark, R. et.al	Multiprotocol Interoperability In IPng, RFC 1683
Curran, J.	Market Viability as an IPng Criteria, RFC 1669
Estrin, D. et.al.	United Routing Requirements for IPng, RFC 1668
Fleischman, E.	A Large Corporate User's View of IPng, RFC 1687
Green, D. et.al.	HPN Working Group Input to the IPng Requirements Solicitation, RFC 1679
Ghisselli, A., et.al.	INFN Requirements for an IPng, RFC 1676
Heagerty, D.	Input to IPng Engineering Considerations, RFC 1670
Simpson, W.	IPng Mobility Considerations, RFC 1688
Skelton, R.	Electric Power Research Institute Comments on IPng, RFC 1673
Symington, S. et.al.	Modeling and Simulation Requirements for IPng, RFC 1667
Taylor, M.	A Cellular Industry View of IPng, RFC 1674
Vecchi, M.	IPng Requirements: A Cable Television Industry Viewpoint, RFC 1686

IPng Technical Requirements

- IPng requirements process

  Frank Kastenholz, FTP Software <kasten@ftp.com>

  Jon Crowcroft, UCL <J.Crowcroft@cs.ucl.ac.uk>
- RFC1550 request for white papers
- Requirements document

  based on Frank Kastenholz/Craig Partridge draft

  criteria, discussion & time frame
- RFC 1726

IPng Criteria - Guessing IP’s Future

- At least 10^9 networks, 10^{12} end-systems
  safer goal 10^{12} nets, 10^{15} end-systems
- Conservative routing schemes
- Topology flexible
- High performance
- Straightforward transition plan from IPv4
- Robust service
- Media independent
- Datagram service
- Autoconfiguration

IPng Criteria, cont.

- Secure operation
- Globally unique names
- Access to standards
- Support multicasting
- Extensible
- Support service classes
- Support mobility
- Include control protocol (ping, etc.)
- Support for private networks (tunneling)

Address Length

- Hotly discussed issue
- Four basic views
  - 8 bytes is enough, more is inefficient
  - 16 bytes is about right, 8 is not enough
  - use 20 byte NSAPs, provide global harmonization
  - variable length gives best safety and efficiency
- Many detailed arguments
- Consensus is that 16 bytes is enough

IPv6 Overview

- Expanded from IPv4 addressing capability (16 byte addresses)
- Simple header
- Support for extension headers and options
- Support for authentication and privacy
- Support for autoconfiguration
- Support for source routed
- Simple and flexible transition from IPv4
- Flow ID

IPng Addresses

- Propose mapping algorithms from and to other environments
- Where addresses are globally unique and assigned with regard to network topology
- IETF should work with other organizations for development of such mappings
- Common addresses facilitate transition to IPng
- Goal to provide a 1:1 mapping between address types (e.g. IPX, NSAP, E164)

IPv6 Header

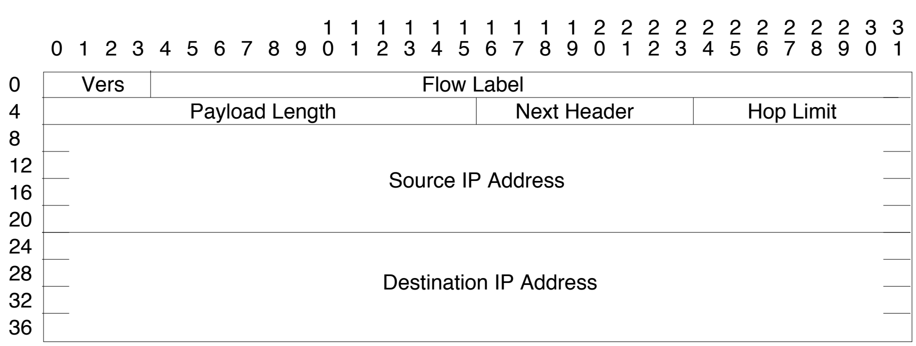

IPv4 Header

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	Vers				Hlen				Pre			D	T	R	C		Total Length															
4	Identifier																	N	M	Fragment Offset												
8	TTL								Protocol								Header Checksum															
12	Source IP Address																															
16	Destination IP Address																															
20	Options																								Padding							

Flow Label

1	3	24
T	TC	flow ID

- T - time sensitivity
  - 0 = yes
  - 1 = no
- TC - Traffic Class
  - type of flow
- Flow ID - random, unique-to-source value
  - combined with source address to identify traffic flow

IPv6 Extension Headers

- Less used functions moved to Extension Headers
- Only present when needed
- Only looked at by node with address in Destination Address (except Hop-by-Hop Options)
- Extensible
  - Hop-by-Hop Options
  - Routing
  - Fragment
  - Authentication
  - Privacy
  - Destination Options

Hop-by-Hop and Destination Options

- Contain one or more options

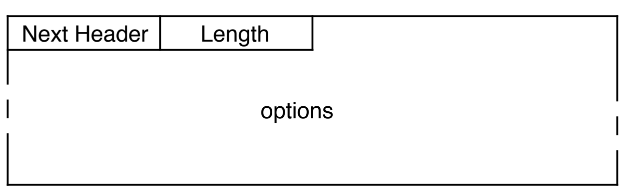

- Pad options
  - options header must be multiple of 8 bytes

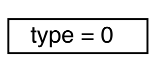

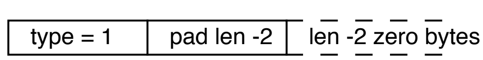

Header Option Handling

AIU	C	Option ID

- AIU - action to be taken if unknown option
  - 00 - skip this option
  - 01 - discard the packet
  - 10 - discard the packet & send ICMP error message
  - 11 - same as 10 except send message only if destination was not a multicast address

  eases introduction of new options
- C - set if option data changes en-route
  - (Hop-by-Hop Options only)

  include option in the Authentication integrity assurance computation

Jumbogram Option

- If header length field in IPv6 header = 0
- Use jumbogram option in Hop-by-Hop options header to find actual datagram length

type = 194	opt dat len=4	datagram payload length

Fragment Header

- Path MTU discovery recommended
- If required, Fragment Header can be used by packet source (routers do not fragment)

Next Header	Length	Fragment Offset	00M
Packet Identifier			

M = More Fragments

Routing Header

Next Header	Length	Routing type	Segments left
reserved	loose/strict bit mask		
address 0			
address 1			

Address Types

- Unicast (single destination)
  - global
  - compatible (IPv4, IPX, NSAP, X.121...)
  - site-local
  - link-local
- Multicast (multiple destinations)
- Anycast (nearest destination)
  - prefix with trailing zeros

IPv6 Address Prefixes

Allocation	Prefix (binary)	Fraction
reserved	0000 0000	1/256
reserved	0000 0001	1/256
NSAP Allocation	0000 001	1/128
IPX Allocation	0000 010	1/128
reserved	0000 011	1/128
reserved	0000 1	1/32
reserved	0001	1/16
reserved	001	1/8
provider-based unicast	010	1/8
reserved	011	1/8
reserved for geographic	100	1/8
reserved	101	1/8
reserved	110	1/8
reserved	1110	1/16
reserved	1111 0	1/32
reserved	1111 10	1/64
reserved	1111 110	1/128
local use address	1111 1110	1/256
multicast address	1111 1111	1/256

IPv6 Transition Goals

- Allow incremental upgrade from IPv4 hosts to IPv6
- Few sequence dependencies
- Support what vendors will do
- Allow IPv4 only hosts to talk to IPv6-only hosts
- Finish before IPv4 addresses run out

IPv6 Transition Techniques

- dual stack
- IPv4 compatible addresses
  IPv4 address embedded in IPv6 address
- IPv6 in IPv4 encapsulation
  tunnel IPv6 across IPv4 topology
- IPv4 <-> IPv6 header translation
  optional

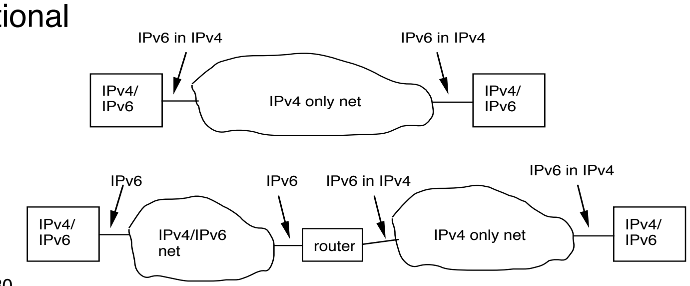

Address Autoconfiguration

- Two types of autoconfiguration
  - server-less
  - state-full server
- DHCPng deals with state-full server
- Security policy an issue
- Trying for *plug & play* in dentist's office
- Autoconfiguration support required in IPv6

Neighbor Discovery

- Router send out advertisements
  - lists prefix(es) for link
  - say if host can use prefix to create global address
  - say if host can use prefix to determine "on-link"
  - say if host must use DHCPng to get address
- If host can use prefix to create global address
  - host appends "MAC" address to prefix
  - checks for duplicate addresses
- Host MAC addresses resolved with ARP-like request/response procedure
  - sent to multicast address formed from dest IP address

Renumbering

- Renumbering hosts
  - advertise prefix for provider **A**
  - connect to provider **B**
  - advertise prefix for provider **A** and provider **B**
    - use prefix **B** for new connections
  - advertise prefix for provider **B** only - do not use **A**

IPv6 Routing

- Hierarchical addresses used in IPv6
- 1st version "provider based" hierarchy
- Working on geographic based
- Address assignment a concern from the start
- Easy(er) renumbering may be important in maintaining efficient use of routing table space

Routing Paradigm

- Longest-match routing will be used
- Existing routing protocols will be modified for IPv6
  - RIPv2
  - OSPF
  - IS-IS
  - IDRP
- Also source routing - ERP header
  - provider section
  - reduce per packet processing

From the IPng Recommendation

"We feel that an improvement in the basic level of security in the Internet is vital to its continued success. Users must be able to assume that their exchanges are safe from tampering, diversion and exposure. Organizations that wish to use the Internet to conduct business must be able to have a high level of confidence in the identity of their correspondents and communications. **The goal is to provide strong protection as a matter of course throughout the Internet.**"

IPng Security Recommendations

- Support for the Authentication Header be required in all compliant IPv6 implementations
- Support for a specific authentication algorithm by required
- Support for the Privacy Header be required in all complaint IPv6 implementations
- Support for a privacy authentication algorithm be required

Authentication Header

Next Header	Length	Reserved
Security Parameters Index (SPI)		
Authentication Data		

- Destination Address + SPI = security association
  - identifies algorithm, key etc.
- Used to authenticate all fields I packet that do not change en-route
- Keyed MD5 is the required default algorithm

Encapsulating Security Payload

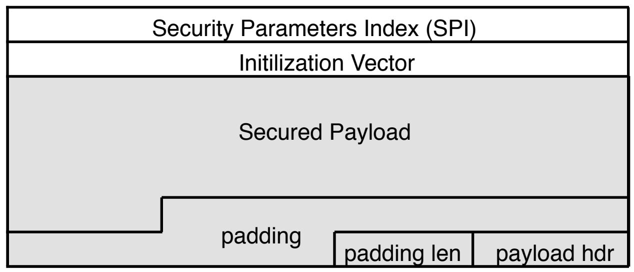

- DES-CBC is required default algorithm
- Must be last non-encrypted header
- Can encapsulate part or full packet

Pretenders to TCP/IP Crown

- X.25
- OSI
- SNA / APPN
- IPX
- ATM

- prediction - IP will remain the bearer service
  - hide media from apps

Post WWW

- the web filled a hole we did not know we had
- lowered the Internet entry requirements
  - mom can surf
- all too-ubiquitous client
  - the world is not all nails
- what additional holes are out there?

Future Applications

- prediction -- we only know a few of the apps that will be important in 5 years
  - email
  - www
  - ftp
  - remote access (diminshing)
  - “buy” button
- how will these map into ‘top 10’ apps is unclear

Will the Structure Hold?

- traffic (both bits and routing info) are stressing existing equipment
- don't know what will be on the ends of the glass in a year
- can't see the shape of what will be needed in the year 2000
  - can't guess the important apps / functions
- business advantage to those who can figure it out

Are Voice & Video Special?

- broadcast -- no
  - semi-reliable data to client within a few seconds
  - memory is cheap
- interactive -- yes
  - QoS (as in latency & latency variation) very important
  - new technologies needed (e.g. RSVP)
  - prediction -- ATM will not help much
    - not ATM end to end, desktop lost to Ethernet

Shifting Viewpoints #1

- original Internet
  - lots of little dumb boxes
  - lots of bright hosts
- telco net
  - fewer but bright (& big) boxes
  - lots of dumb hosts (phones)
- prediction
  - combination

Shifting Viewpoints #2

- move from

data is a special case of voice
to
voice is a special case of data

- what other special cases?
- billing will be fun

Threat vs. Promise

- this data network will be both threat & promise
  - just like the telephone
  - just like the auto

- it will succeed in being both