INTERNET SOFTWARE CONSORTIUM

The Future of IP Addressing

Joe Abley
jabley@isc.org

Why do we need IP addresses?

- means by which one device attached to the internet is distinguished from every other device
- used to direct requests to an appropriate destination (destination address)
- used to indicate where replies should be sent (source address)

Addresses live in the IP Header

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Version|  IHL  |Type of Service|          Total Length         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Identification        |Flags|      Fragment Offset    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Time to Live |    Protocol   |         Header Checksum       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       Source Address                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Destination Address                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Options                    |    Padding    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

(RFC 791)

Address Stability

- Some devices prefer persistent, stable addresses
  - so clients can find a server easily
  - less strain on the DNS
  - renumbering router interfaces can cause instability in the routing system

Address Stability

- Some devices don't mind being renumbered
  - occasionally-connected devices (e.g. dial-up hosts)
  - mobile devices
  - client devices in general
- DHCP, PPP/IPCP

Address Assignment

- There are 2^{32} possible addresses in IPv4
  - (that's 4294967296 addresses)
- How should they be assigned?

Topological Structure

32-bit IP Address			
203	97	2	226
11001011	01100001	00000010	11100010
"network part"			"host part"

- all the hosts on the 203.97.2.0 subnet have the same "network part"
- "host part" is used to identify another device on a connected subnet

Address Classes

- History Lesson! Obsolete Information!
- Initial goals of address structure:
  - simplicity of allocation
  - simplicity of routing
- The Internet is an obscure research project at this stage (1981), so there is no great concern over growth or address conservation

Address Classes

class	class 32-bit IP Address				hosts
A	0xxxxxxx	xxxxxxxx	xxxxxxxx	xxxxxxxx	16,777,216
B	10xxxxxx	xxxxxxxx	xxxxxxxx	xxxxxxxx	65,536
C	110xxxxx	xxxxxxxx	xxxxxxxx	xxxxxxxx	256

- Simple allocation procedure
- Boundary between “network part” and “host part” lies in the same place as the dots in the address

Subnetting

- Expecting a single layer-2 network to contain 16 million devices is clearly silly
  - Split a network up into "sub networks"
  - insert a "subnet identifier" between the "network part" and the "host part"
- Subnets of a particular network are always the same size
  - (there are also other peculiarities)

And Everything Was Fine Until...

- The Internet grew
- The address space was consumed rapidly
- It was widely predicted that there would be no numbers left by 2000

Solutions to the Address Famine

- Short Term:
  - change the allocation policies, since the class-based system is too wasteful
  - encourage address conservation
- Longer Term:
  - Extend IP addresses from 32 bits to something bigger

CIDR, VLSM

- Classless Inter-Domain Routing
- Variable-Length Subnet Masks

<table>
<tr><td colspan="32">32-bit IP Address</td></tr>
<tr><td colspan="8">203</td><td colspan="8">97</td><td colspan="8">2</td><td colspan="8">226</td></tr>
<tr><td>1</td><td>1</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr>
<tr><td colspan="24">“network part”</td><td colspan="8">“host part”</td></tr>
</table>

CIDR, VLSM

- Classless Inter-Domain Routing
- Variable-Length Subnet Masks

<table>
<tr><td colspan="32">32-bit IP Address</td></tr>
<tr><td colspan="8">203</td><td colspan="8">97</td><td colspan="8">2</td><td colspan="8">226</td></tr>
<tr><td>1</td><td>1</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr>
<tr><td colspan="24">“address prefix”</td><td colspan="8">“host part”</td></tr>
</table>

CIDR, VLSM

- No longer restricted to octet boundaries
- Can allocate netblocks more efficiently

<table>
<tr><td colspan="32">32-bit IP Address</td></tr>
<tr><td colspan="8">203</td><td colspan="8">97</td><td colspan="8">2</td><td colspan="8">226</td></tr>
<tr><td>1</td><td>1</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr>
<tr><td colspan="26">“address prefix”</td><td colspan="6">“host part”</td></tr>
</table>

CIDR, VLSM

- No longer restricted to octet boundaries
- Can allocate netblocks more efficiently

<table>
<tr><td colspan="32">32-bit IP Address</td></tr>
<tr><td colspan="8">203</td><td colspan="8">97</td><td colspan="8">2</td><td colspan="8">226</td></tr>
<tr><td>1</td><td>1</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td><td>0</td><td>1</td><td>0</td></tr>
<tr><td colspan="18">“address prefix”</td><td colspan="14">“host part”</td></tr>
</table>

CIDR, VLSM

class	hosts
A	16,777,216
B	65,546
C	256

prefix length	hosts
...	...
19 bits	4096
20 bits	2048
21 bits	1024
22 bits	512
24 bits	256
25 bits	128
...	...

CIDR, VLSM

- Some implications on routing algorithms
  - actually simplified them
  - “longest prefix wins”
- New notation
  - address/prefix length
  - 203.97.2.226/28

The Rise of the NAT

- Network Address Translators provide a boundary between addressing domains
  - "private-use" addresses within an enterprise, mapped to globally unique addresses at the addressing domain boundary
  - RFC1918 describes private-use addresses

The Short-Term Fix

- Engineering fixes
  - CIDR, VLSM, NAT
- Procedural fixes
  - address conservation policies in RIRs
- Return to sanity
  - dot-bomb explosion

The Short-Term Fix

We are not running out of IPv4 addresses any time soon.

The sky is still falling, just not very quickly.

The Long-Term Fix

- IPv6
  - 128-bit addresses
  - (340,282,366,920,938,463,463,374,607,431,768,211,456 of them)
    - this is a big number
  - IPv6 deployment is happening

IPv6 Addresses

- IPv6 addresses are weird-looking
  - **2001:4f8:3:bb:200:f8ff:fe02:7fe7**
- Each of those numbers is a 16-bit value, represented in hexadecimal
- There are a lot of numbers to choose from
  - internal address assignment policy?

ISC IPv6 Addressing Scheme

ARIN		Site	VLAN	Host			
2001	04f8	0003	00bb	0200	f8ff	fe02	7fe7
"ISC"		SQL1	"main"	00-00-f8-02-7f-e7			

- The "host" portion is a "universal" EUI-64 address, which is constructed from a 48-bit IEEE 802 address (RFC 2464)
- So, we still have to think up numbers for new sites and new VLANs, but not new hosts

Panic Over

- So, we have enough IPv6 addresses to last us for a long time
- Managing those new addresses is not that hard\
- CIDR, VLSM, NAT and exploding dot-coms have given us breathing space to finish the migration

Uh Oh

More Problems!

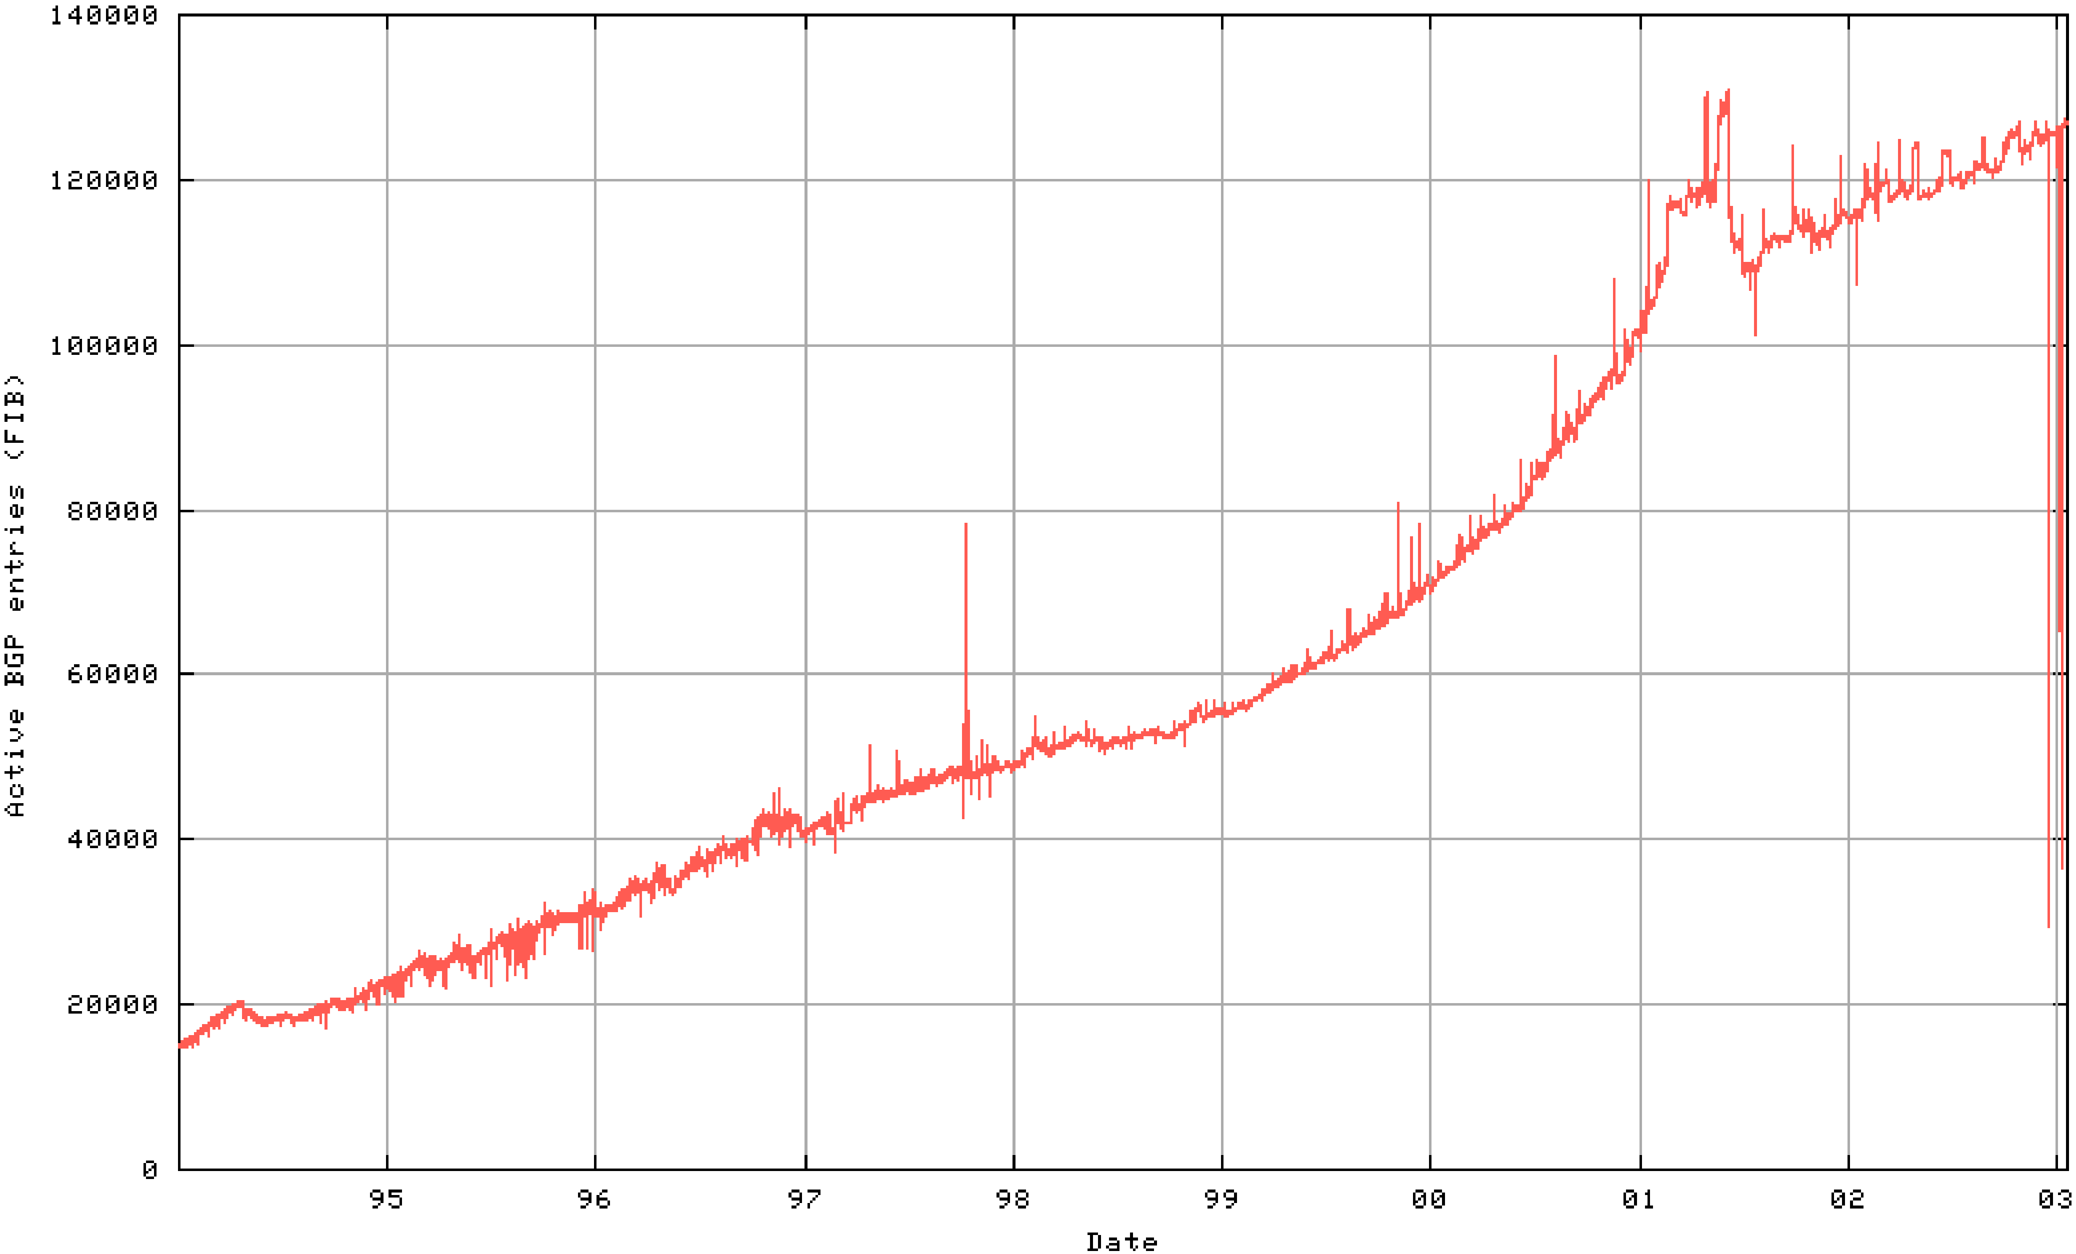

measurement by Geoff Huston of route-views data
http://bgp.potaroo.net/as6447/

Routing System Complexity

- The same growth that raised the spectre of address exhaustion has also had an impact on the routing system
  - growing number of routing entries
  - growing number of paths
  - increasing turbulence

Routing System Complexity

- A strained routing system is an unhappy routing system
  - increasing convergence times
  - routing loops
  - persistent oscillations

Information Hiding

- If every router on the internet needs to know the intimate details of the entire network in order to work, then we are doomed
- Routers need to be deterministically ignorant of other peoples' networks in order for the routing system to scale

Aggregation

- Replace large amounts of fine-grained routing information with smaller amounts of coarse-grained routing information
  - leave the detail in the routing system where it is needed, and summarise everwhere else

Aggregation

- Aggregation works best when addresses have been allocated to be aggregatable
  - topologically-close networks should use addresses which are numerically similar
  - if the topology changes, networks should renumber

Aggregation

- ISP A is allocated 203.97.0.0/17, and makes assignments from within that block for customers
- ISP B is allocated 203.97.128.0/17, and makes customer assignments from that block
- ISP A and ISP B are customers of ISP C

Aggregation

ISP	ISP sees	ISP announces
A	203.97.1.0/24 203.97.2.0/29 ...	203.97.0.0/17
B	203.97.128.0/28 203.97.128.16/28 ...	203.97.128.0/17
C	203.97.0.0/17 203.97.128.0/17	203.97.0.0/16

Problems

- Multihoming
  - what if a customer of ISP A also wants to have connectivity through ISP B?
- Customers find renumbering difficult
  - hence there is market pressure not to force customers to renumber
- Inter-Domain Traffic Engineering

Multihoming

- Increasing numbers of customers want to multi-home, since the cost of doing so is low and the cost of experiencing downtime is growing
- It is becoming cost-effective for individual customers to multi-home

Inter-Domain Traffic Engineering

- Sometimes the information that is hidden by aggregation is necessary in order to manage traffic flow
- It is not easy to tell whether aggregation is appropriate unless you are the actual operator of a range of addresses
- Proxy Aggregation is hard

What is the Answer?

- There is no good answer right now, although it is an area of active research
  - ptomaine, multi6 working groups at the IETF
- Answers that are good for the routing system are often bad in other ways
  - operational complexity
  - software incompatibility

The Future

- Maybe service identifiers and network addresses will be separate and different
- Maybe network addresses will change frequently and transparently
- Maybe more of the decision-making in the routing system will move to the edge, away from the core
  - maybe even to the host

Predictions

- IPv6 deployment will happen
- IPv4 will go away
- The Internet will continue to grow
- The numbering system and the routing system will change to accommodate that growth
  - and users will not notice
- New Zealand will win the Americas Cup