


Classful Address Allocation

Surprisingly not as evil as you think.

- <http://bill.herrin.us/network/201010-classful.ppt>

What was Classful addressing?



- All allocations in 0.0.0.0/1 were exactly **/8** (24M addresses).
- All allocations in 128.0.0.0/2 were exactly **/16** (65k addresses).
- All allocations in 192.0.0.0/3 were exactly **/24** (256 addresses).

Traffic engineering was not a problem with Classful

145.6.0.0/16

145.6.7.0/24 –
Traffic engineering

- 145.6.7.0/24 must be traffic engineering – filterable.

167.8.0.0/14

- 167.8.0.0/14 is lie!

178.9.0.0/16

- 178.9.0.0/16 is unambiguously a whole allocation.

TE in CIDR costs you money!

- CIDR Report 9/24/2010

- | ASnum | NetsNow | NetsAggr | NetGain | % | |
|---------|---------|----------|---------|-------|-----------|
| Table | 337009 | 208022 | 128987 | 38.3% | All ASes |
| AS6389 | 3779 | 282 | 3497 | 92.5% | BellSouth |
| AS4323 | 4468 | 1924 | 2544 | 56.9% | twtelecom |
| AS19262 | 1819 | 285 | 1534 | 84.3% | Verizon |

- Four cents per BGP route per router per year.
- $\$0.04 \times 128,987 \times \text{your } 100 \text{ routers} = \text{half a million bucks}$
- $\$0.04 \times 128,987 \times 150,000 \text{ BGP routers worldwide} = \$775 \text{ million per year to carry all that TE.}$

Traffic Engineering in CIDR – The Cutout Problem

207.199.128.0/18 - ISP

207.199.128.0/20 –
TE preference to
receive via a
particular upstream

207.199.136.0/23 –
multihomed customer

Must carry the /18 route

The /20 is discardable, particularly if the ISP is distant. System remains reachable regardless.

The /23 must be carried – you won't always and may never be able to reach it via the /18.

But how do you configure YOUR router to tell the difference between THEIR /20 and /23?

TE in CIDR: The Variable Subnet Problem

207.199.128.0/19 - ISP

207.199.128.0/22 -
Traffic engineering

This /22 is discardable
As are any other routes
inside this /19.

207.199.160.0/22 – User 2

207.199.164.0/22 – User 3

These /22s are not
discardable.

How do you configure YOUR router to tell the difference
between the /22 announcements it receives?

But Classful failed, remember?

- Not enough B's for everyone who needed more than one C. Not enough A's period.
 - So we caused a routing table explosion by allocating many C's per org.
- We routed cutouts anyway for folks who changed ISPs – so much for filterable TE.
- Rigid class structure wasted address space.

Classful and CIDR

- CIDR created the TE disaggregation problem.
- We're stuck with it for IPv4 – the addresses have already been allocated.
- But Classful can offer us fresh insights for IPv6 address allocation.

Classifying IPv6 Allocation Size By Pool

2000::/3 – IPv6 Internet

3210::/12 – One IANA allocation

3210::/16 – ARIN /28 pool

3211::/16 – ARIN /32 pool

3212::/16 – ARIN /40 pool

3213::/16 – ARIN /48 pool

- Only allocate the exact sizes.
- Fixes the TE disaggregation problem.
- Doesn't suffer IPv4 Classful's problems.
- Plenty of /32's for anyone who needs more than a /40.

Classifying IPv6 Allocations

2000::/3 – IPv6 Internet

3210::/12 – One IANA allocation

3210::/14 – Multihomed pool

3211::/16 – /32 pool

3212::/16 – /48 pool

3214::/14 – Singlehomed

3215::/16 – /32 pool

3216::/16 – /48 pool


- Why classify only by address allocation size?
- Classify by any important criteria.

Classiful IPv6

Good Today
Better Tomorrow

Future Tech: map-encap

Send to: 321F:1234::1



Future Tech: map-encap

- IRTF Routing Research Group has found a class of technologies called “Map-Encap” that allows cheap provider-independent addressing.
- Needs registered addresses
- Preferably addresses classified as matching the given map-encap protocol and not accepted into the BGP table.

Conclusion

- A hybrid CIDR/classful addressing scheme could yield technically superior IPv6 addressing practices.
- <http://bill.herrin.us/network/201010-classful.ppt>
- <http://bill.herrin.us/network/bgpcost.html>