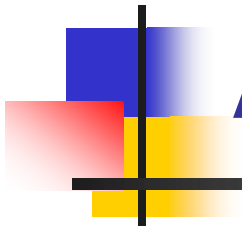

Aggregation (?)

Effect of business practices on the Internet today

Philip Smith

RIPE 50, Stockholm

BGP Report (gih)

- Data taken from:
  - bgp.potaroo.net/as4637/
- 157000 prefixes total announcements
- 108000 prefixes
  - aggregating including full AS PATH info
  - i.e. including each ASN's traffic engineering

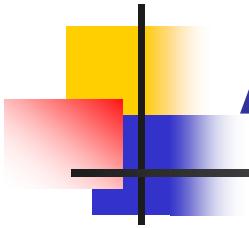

Aggregation/Deaggregation

- 49000 prefixes have no value whatsoever in the Internet Routing Table
  - This view only!
  - They may have value in another view
  - 31% of the Internet Routing table could be discarded with no loss of information

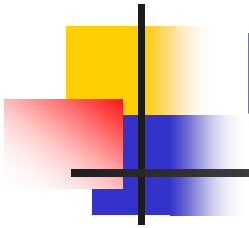

BGP Report (gih)

- 157000 prefixes total announcements
- 108000 prefixes
  - aggregating including full AS PATH info
  - i.e. including each ASN's traffic engineering
- 93000 prefixes
  - aggregating by Origin AS
  - i.e. ignoring each ASN's traffic engineering

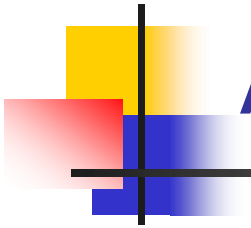

Aggregation/Deaggregation

- 15000 prefixes are caused by ISP traffic engineering
  - This view only!
  - They cause 10% of the Internet Routing Table

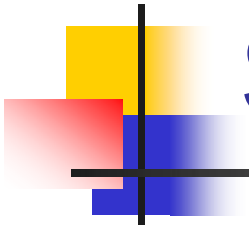

Simplistic Summary

- Deaggregation is a serious problem again
  - 33% of the Internet Routing Table caused by deaggregation
  - 10% of the Internet Routing Table caused by BGP traffic engineering

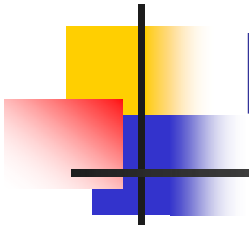

Past Solutions: CIDR Report

- CIDR Report started by Tony Bates in 1995
  - Aim was to encourage ISPs to CIDRise as the Internet moved from classful to classless routing
  - Published top 20 ISPs who could do better at aggregating
  - Weekly mailshot was held in high regard across the industry, and its influence was significant
- Growth of commercial Internet and lack of “clue” reduced the influence

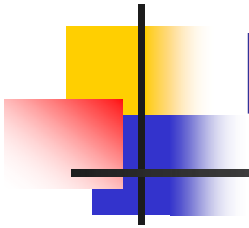

Past Solutions: CIDR Police

- Group of well meaning individuals who in their spare time used my Routing Report and the CIDR Report to encourage ISPs to try and aggregate better
- Were most active in 1999-2002
  - Rampant growth of the Internet Routing Table during the boom years

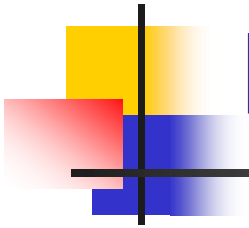

Efforts Today?

- CIDR Report now maintained by Geoff Huston
  - Greatly expanded in scope and available views
  - Web site – www.cidr-report.org
  - Web interface allows any ASN to check on their aggregation effort

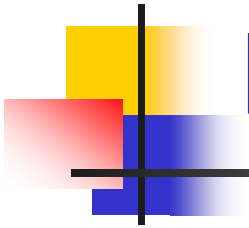

Efforts Today?

- And that's all
- CIDR Police have "retired"
  - Harder times, more to do, less time to do it
  - "Charity" is the first to suffer
- No one seems to care about size of Internet Routing Table
  - "Problem solved! Vendors make routers with fast CPUs and large memory"

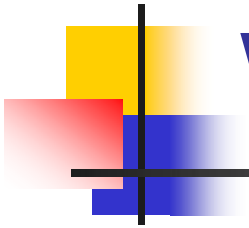

What's going wrong?

- Internet has bigger reach
  - All countries in the world are connected
  - Has everyone been trained on the requirements of being an Internet Service Provider?
- Education system is STILL teaching classful routing 10 years after its obsolescence
  - New engineers are still thinking Class A, Class B and Class C...
  - ...and configure BGP as such

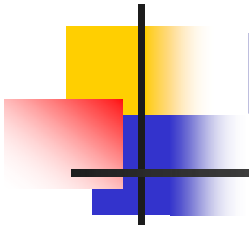

Now?

- RIRs request that address allocations made to ISPs are announced as such
  - Some protest that the RIRs are telling them how to run their networks! ☹
  - Other people only understand Class As, Class Bs, and Class Cs, so announce their networks as /16s or /24s, rather than aggregates

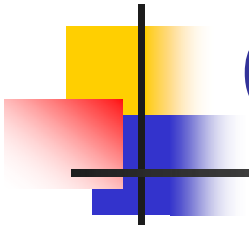

Commercial Pressures

- ISPs deliberately deaggregating
  - To avoid “DoS attacks” from other ISPs falsely announcing their deaggregated address space
  - Oft used excuse but published evidence of these events?
- Such miscreant behaviour encourages others to do likewise with impunity
  - We should all route /32s and be done with it (!)
    - Routed address space span is 1,383,395,136 /32s
  - Even announcing /24s makes this 5.4 million prefixes

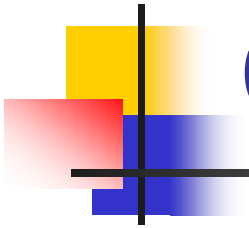

Commercial Pressures

- Network engineers:
  - Paid less (so job rotation is significant)
  - Untrained (training costs money)
  - Have less time (expected to do everything)
  - Participate less in NOGs, if at all
  - Smaller NetEng teams
- Results:
  - Cookbook “knowledge”
  - Mailing list myths and bad/wrong advice
  - Temporary hacks become permanent solutions

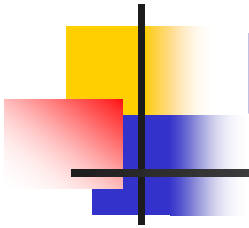

Multihoming

- Multihoming is a basic requirement
  - Improves redundancy and operational reliability
  - Commercial service ⇒ SLAs + non-stop operation
  - “BGP Traffic Engineering”
- Lack of knowledge on what to do
  - Deaggregation is common solution
- Myths of Multihoming:
  - Big router with lots of memory FALSE
  - Need the full routing table FALSE

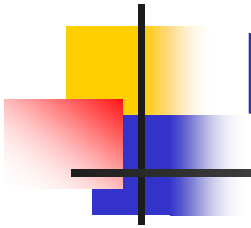

Multihoming

- Lack of training on current multihoming solutions for IPv4
- Lack of agreement between experts on how to implement multihoming
  - One size does NOT fit all
- Lack of clear concise documentation on how to multihome
  - Elusive because solutions are often particular to specific situations

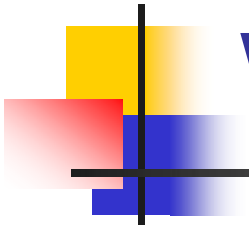

What next?

- Proposal to introduce a Routing WG work item
- Aim: Aggregation Recommendations for ISPs
  - Spin off would be improved multihoming solutions – maybe even best practice documentation