Cache Pollution (and other trends)

Paul Vixie, ISC

DNS-OARC Prague

2010 May 02

Today's talk

- Cache Pollution
  - SIE noticed that some were taking up more RR space in passive DNS databases than others
  - GaTech was able to quantify and graph it

    Contributions by Manos Antonakakis, GaTech, et.al. (D.Dagon, W.Lee, R.Perdisci, N.Feamster)
- A case for real-time analysis

Number of unique RRs is increasing

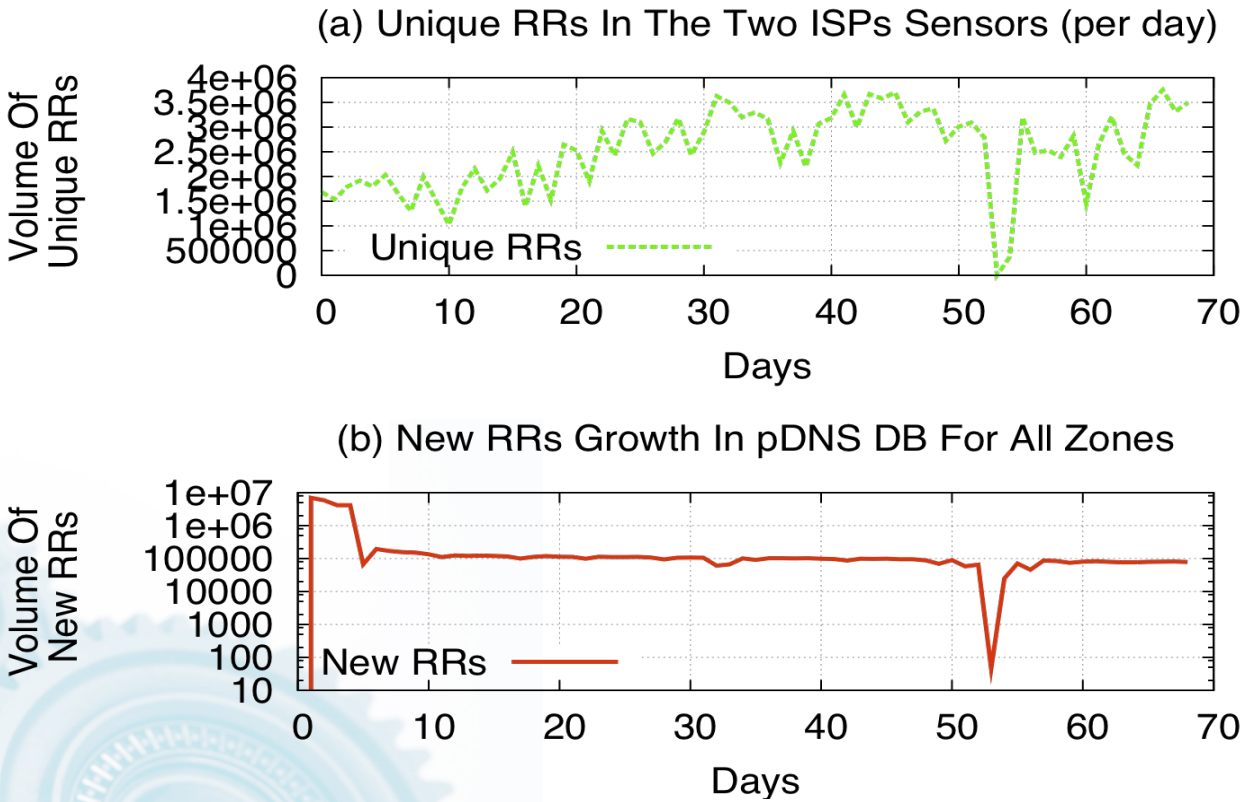

More CDN names and IPs

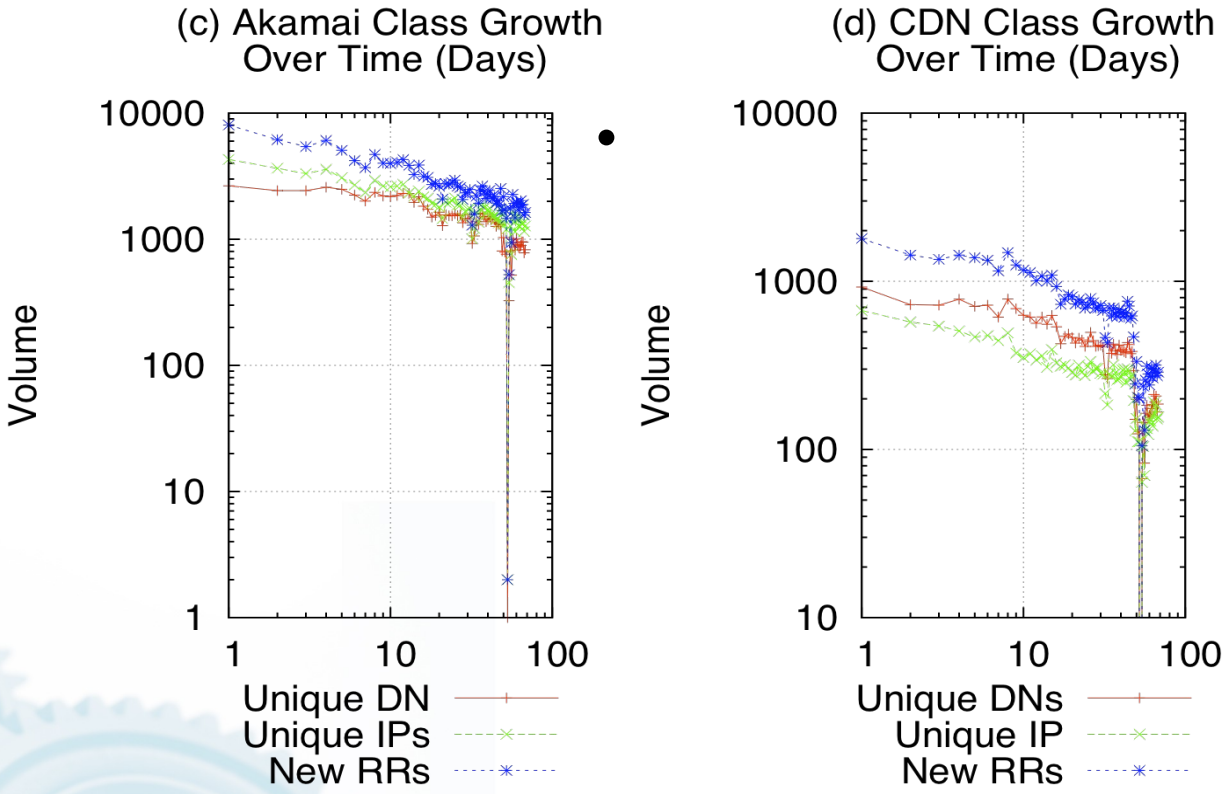

RRs served by large providers and dynamic DNS

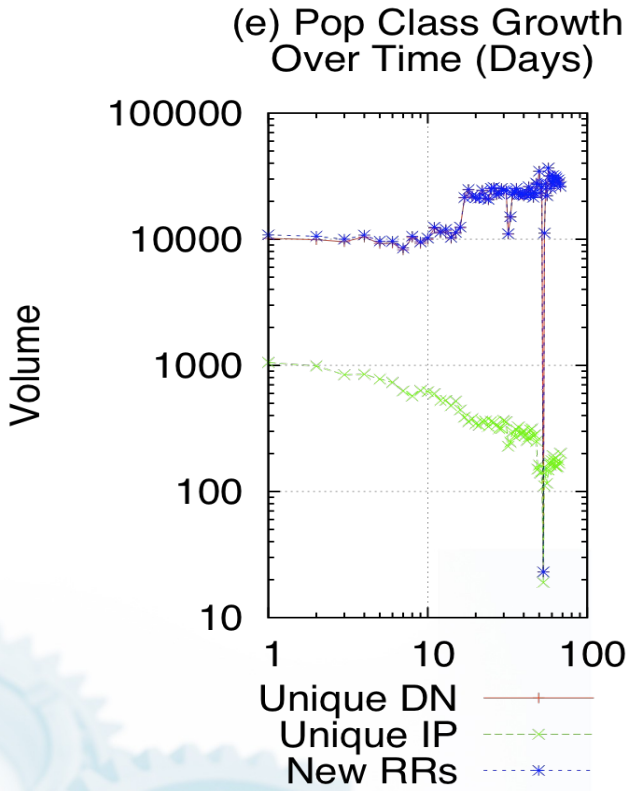

facebook, amazon, google
(note divergence between RRs/IPs)

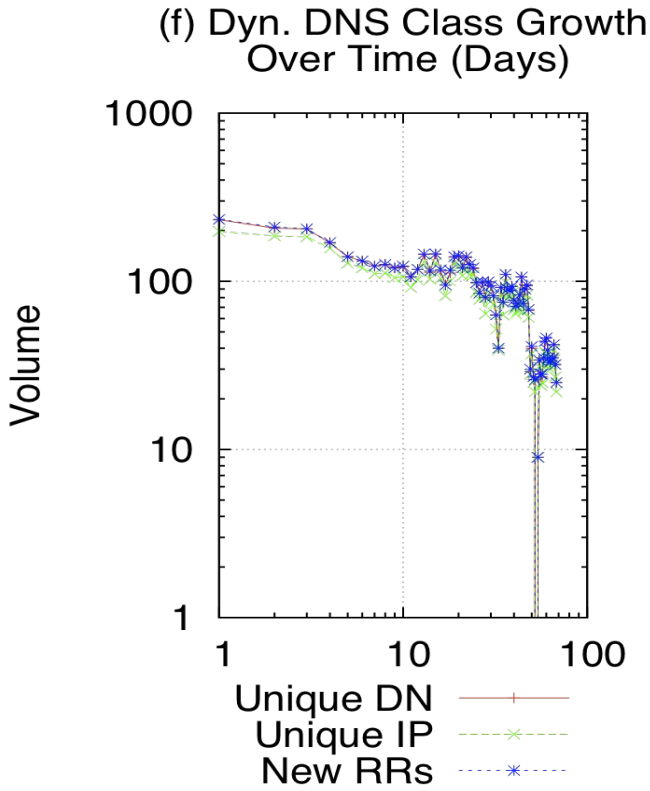

no-ip, dyndns

Less growth in “normal” sites, while large growth in CDNs and popular sites like Facebook

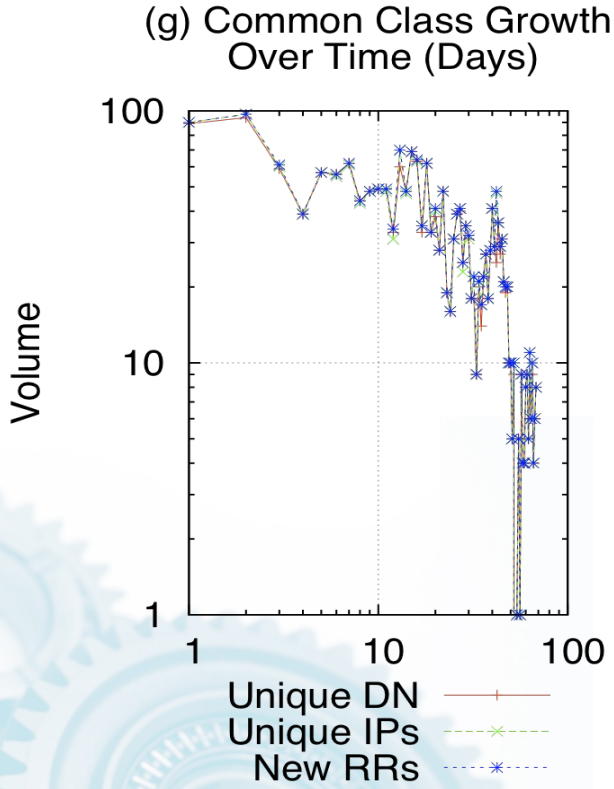

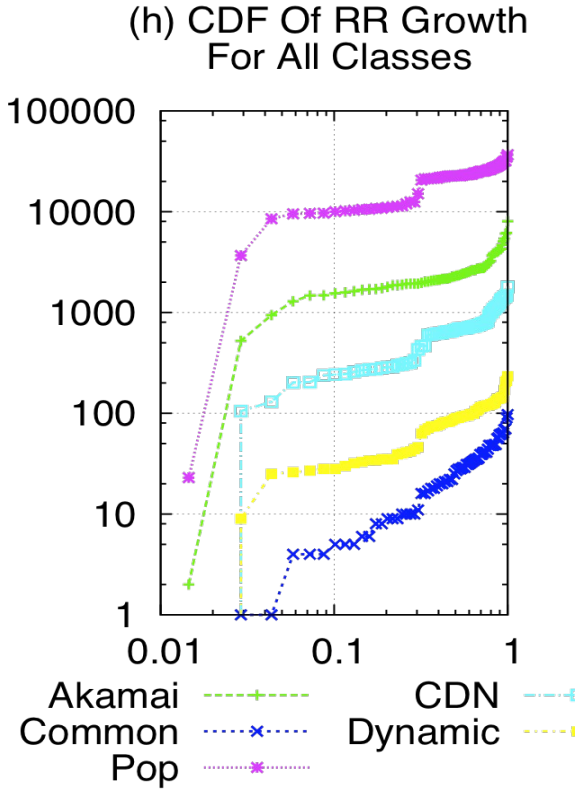

Akamai vs Facebook: RR growth

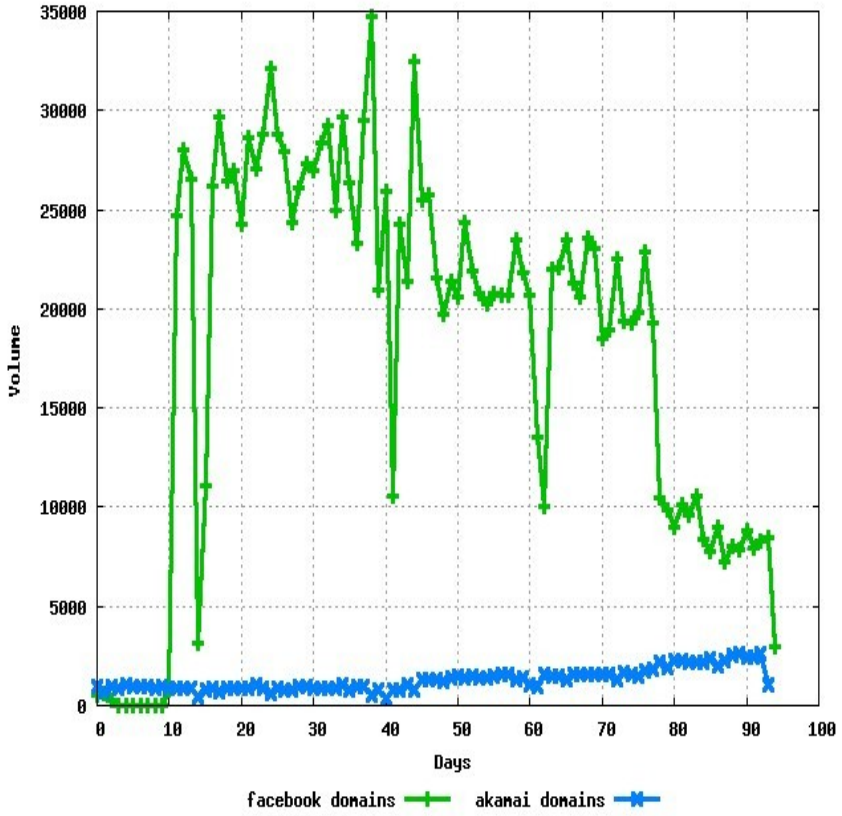

- Facebook generates large number of domains per day
- Akamai generates a lot but nowhere close to Facebook.

Akamai vs Facebook: IP growth

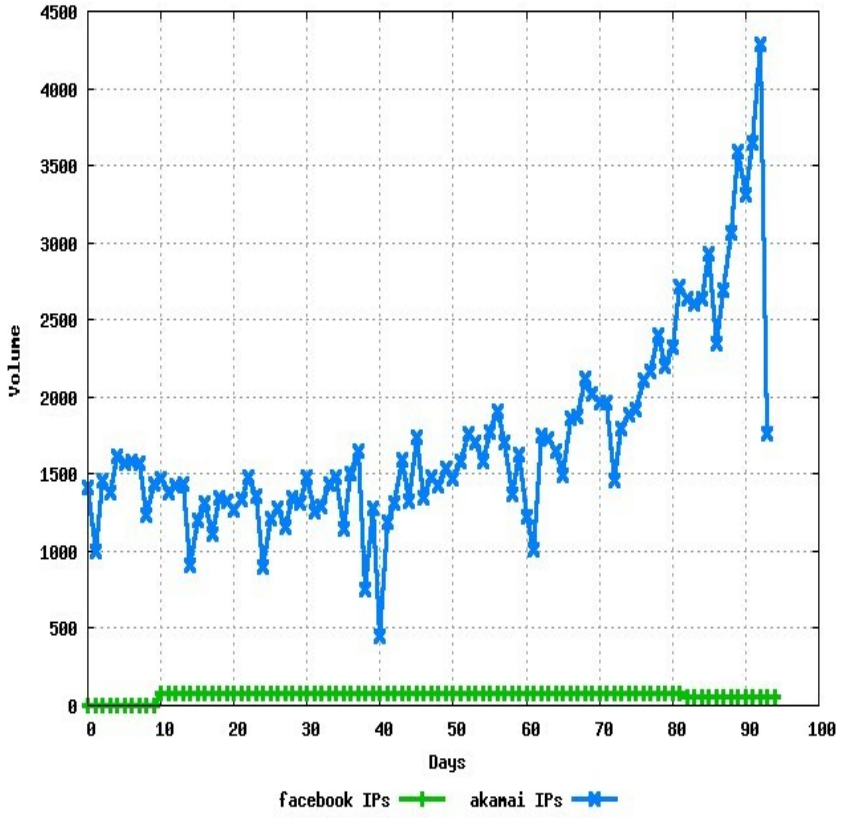

- Akamai maps new RRs to new IPs each day
- Facebook uses a relatively static IP address space into which all new RRs are pointed

facebook

```
# nmsgtool -C ch202 -o - -c 5000 | egrep “channel.*facebook” | fgrep -v \;
053822xxxx.channel53.facebook.com. 3600 IN A 69.63.178.123
0199595xxxx.channel66.facebook.com. 3600 IN A 69.63.178.136
0253010xxxx.channel11.facebook.com. 3600 IN A 69.63.176.171
06670xxxx.87.channel.facebook.com. 3600 IN A 69.63.180.44
0289776xxxx.channel07.facebook.com. 3600 IN A 69.63.176.167
0285474xxxx.channel05.facebook.com. 3600 IN A 69.63.176.165
47771xxxx.92.channel.facebook.com. 3600 IN A 69.63.180.45
0381825xxxx.channel35.facebook.com. 3600 IN A 69.63.176.195
173621xxxx.channel08.facebook.com. 3600 IN A 69.63.176.168
0257487xxxx.channel12.facebook.com. 3600 IN A 69.63.176.172
0241726xxxx.106.channel.facebook.com. 3600 IN A 69.63.180.47
040887xxxx.channel11.facebook.com. 3600 IN A 69.63.176.171
```

Akamai

a1394.g.akamai.net. 20 IN A 70.167.151.137
a1394.g.akamai.net. 20 IN A 70.183.191.114
a1492.g.akamai.net. 20 IN A 70.167.151.190
a1492.g.akamai.net. 20 IN A 70.183.191.113
a1339.b.akamai.net. 20 IN A 98.174.28.121
a1339.b.akamai.net. 20 IN A 98.174.28.123
a1108.da1.akamai.net. 20 IN A 70.167.151.182
a1108.da1.akamai.net. 20 IN A 70.167.151.195
a1493.g.akamai.net. 20 IN A 70.183.191.114
a1493.g.akamai.net. 20 IN A 70.167.151.134
a1593.g.akamai.net. 20 IN A 70.167.151.196
a1593.g.akamai.net. 20 IN A 70.167.151.171

low ttl at least

Take-aways

- Need to research impact of CDNs and Facebook on caching resolver behavior
- How many similar services can large-population resolvers take?
- When inserting lots of RRs, perhaps lowering TTL can help mitigate effect?
- We need continued real-time monitoring
  - Every day is a DITL
  - SIE is here to help