



Nameserver Statistics

Bruce Campbell

RIPE NCC

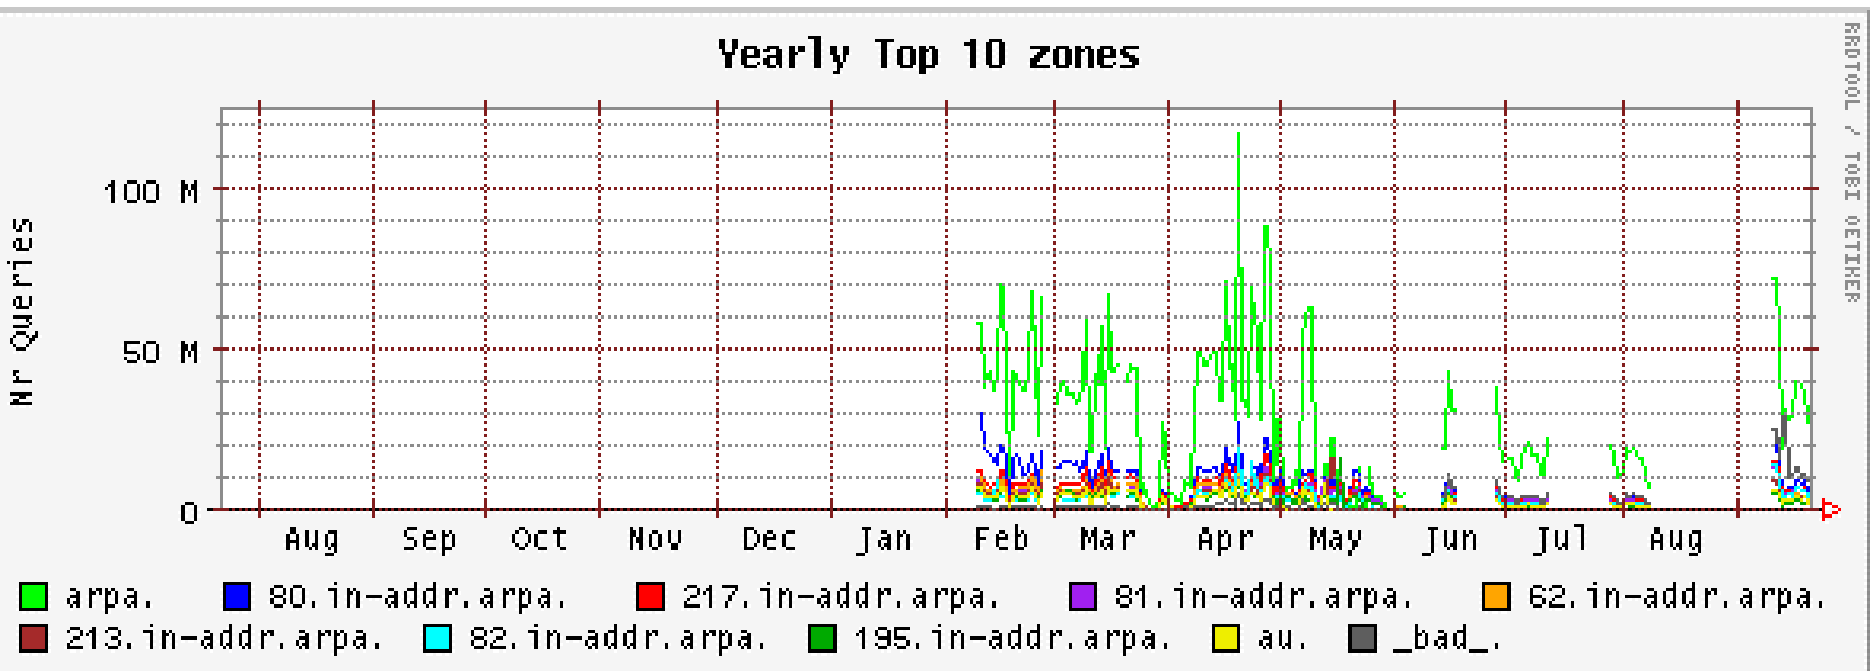
DNS-WG, RIPE49, 23 September
2004




Collection Purpose

- Verification of suspected nameserver traffic.
- Knowing which zone sets need to have dedicated servers.
- Started off as an informal side project (February 2004).

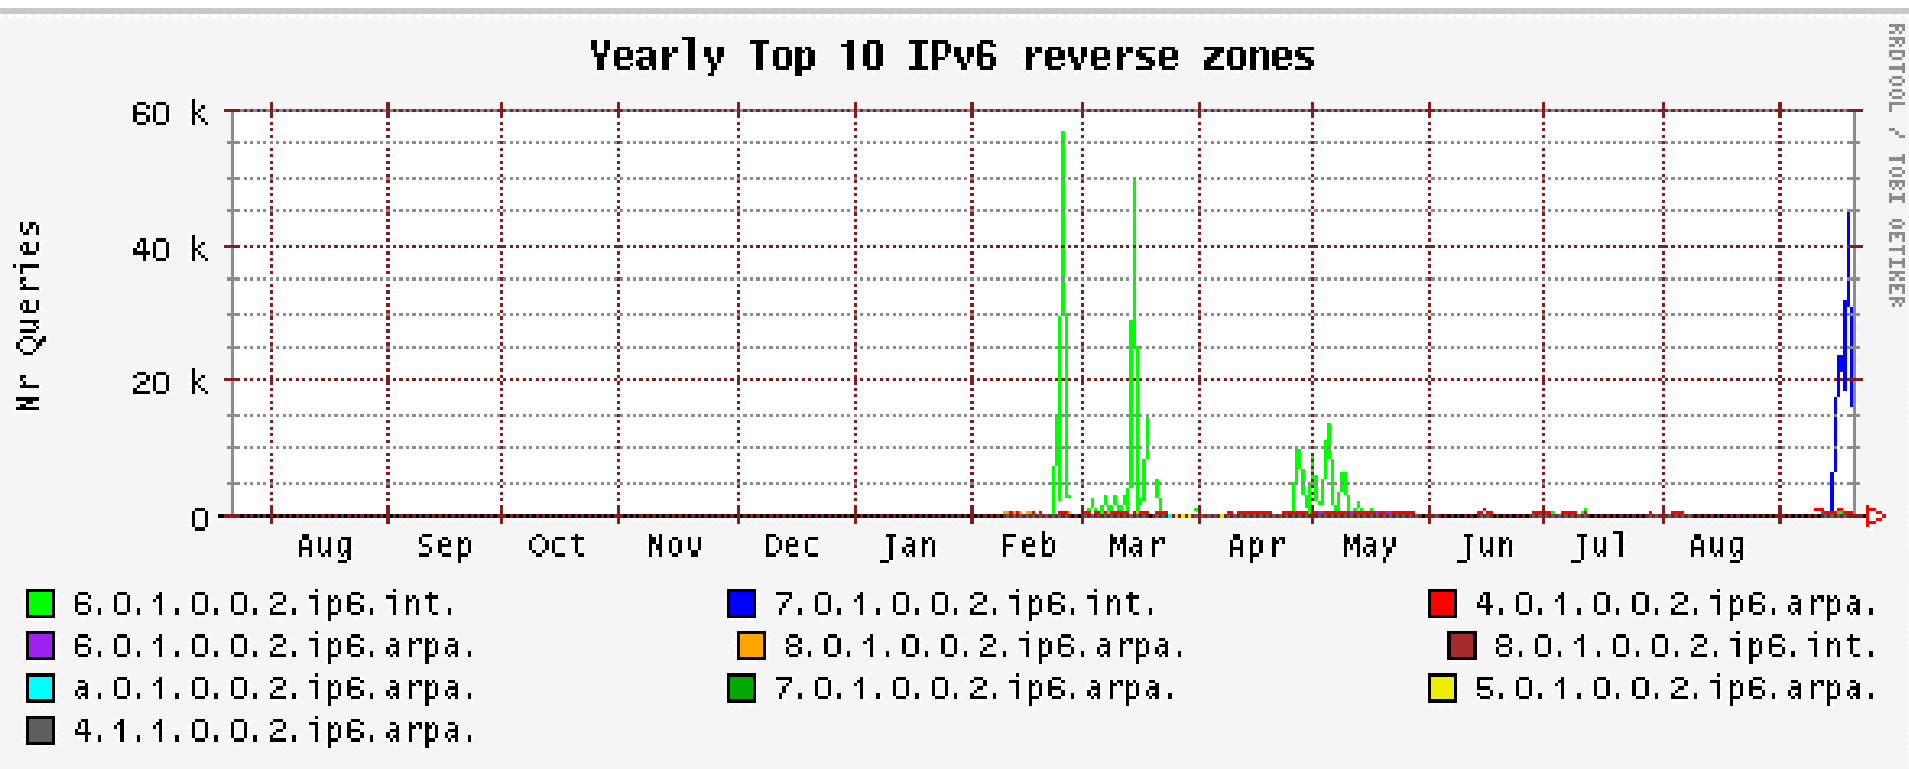
Top 10 Zones




Top 10 non-reverse zones



Top 10 IPv6 Reverse Zones



Top 10 Query Types



Over long term, A, AAAA, A6 and PTR have nearly same rate? Granularity of graphing?



Data Collection Method

- Used switch monitor port.
- Tcpdump of all DNS queries.
- Summarisation script on tcpdumps.
- Totals queries into known zones and query types.



Problems with Collection (or, Just plain Weird stuff)

- Summarisation took longer than collection.
 - Perl pattern matching. ‘nuff said.
- Disk intensive:
 - Have killed one set of disks already.
- Switch monitoring not only used for nameserver queries (monitored ports sometimes changed)
- Inexplicable queries.

Conclusions

- Several classes of zone identified and moving processes initiated.
 - (Primary, ccTLDs, Policy Secondaries, Secondaries)
- ‘We get asked a lot of stuff’
- A6 vs AAAA ? They’re about the same.

Questions, Discussion

