DNS Anycast Stability

Some Early Results

APNIC / Kyoto
2005.02.23

Peter Boothe <peter@cs.uoregon.edu>
Randy Bush <randy@psg.com>

<http://rip.psg.com/~randy/050223.anycast-apnic.pdf>

Why do This?

- Verisign presentation "Life and Times of J-Root" <http://www.nanog.org/mtg-0410/pdf/kosters.pdf>
- Foils 27 to 29, reported non-trivial routing jitter and therefore suggested "DO NOT RUN anycast with stateful transport."
- But for almost a decade, there have been reports of successful delivery of stateful services over anycast
- Was their measurement from an abnormal vantage point, or are there other things going on?

Experimental Method

- Volunteers on hundreds of hosts around the world ran a multi-day script
- Every two seconds it probed the known anycast root servers

```
dig @X.root-servers.net. hostname.bind chaos txt
```

- Both UDP and TCP queries
- Results were collected at a central server

```
42.666.7.11
Sat Nov 13 01:29:29 UTC 2004 f UDP pao1e.f.root-servers.org
Sat Nov 13 01:29:29 UTC 2004 i UDP s1.lnx
Sat Nov 13 01:29:30 UTC 2004 j UDP jns4-kgtld.j.root-servers.net
Sat Nov 13 01:29:30 UTC 2004 k UDP k1.linx
Sat Nov 13 01:29:30 UTC 2004 m UDP M-d3
Sat Nov 13 01:29:30 UTC 2004 c TCP lax1a.c.root-servers.org
Sat Nov 13 01:29:31 UTC 2004 f TCP pao1c.f.root-servers.org
Sat Nov 13 01:29:31 UTC 2004 i TCP s1.lnx
Sat Nov 13 01:29:31 UTC 2004 j TCP jns1-kgtld.j.root-servers.net
Sat Nov 13 01:29:32 UTC 2004 k TCP k1.linx
Sat Nov 13 01:29:33 UTC 2004 m TCP M-d3
Sat Nov 13 01:29:35 UTC 2004 c UDP lax1a.c.root-servers.org
Sat Nov 13 01:29:35 UTC 2004 f UDP pao1e.f.root-servers.org
Sat Nov 13 01:29:35 UTC 2004 i UDP s1.lnx
Sat Nov 13 01:29:35 UTC 2004 j UDP jns3-kgtld.j.root-servers.net
Sat Nov 13 01:29:36 UTC 2004 k UDP k1.linx
Sat Nov 13 01:29:36 UTC 2004 m UDP M-d3
...
```

Warning

- This is about routing
- Not root server performance
- The effects you are about to see are likely caused by
  - Inter-ISP eBGP
  - Intra-ISP iBGP
  - Intra-ISP IGP (OSPF or IS-IS)

One AS's View

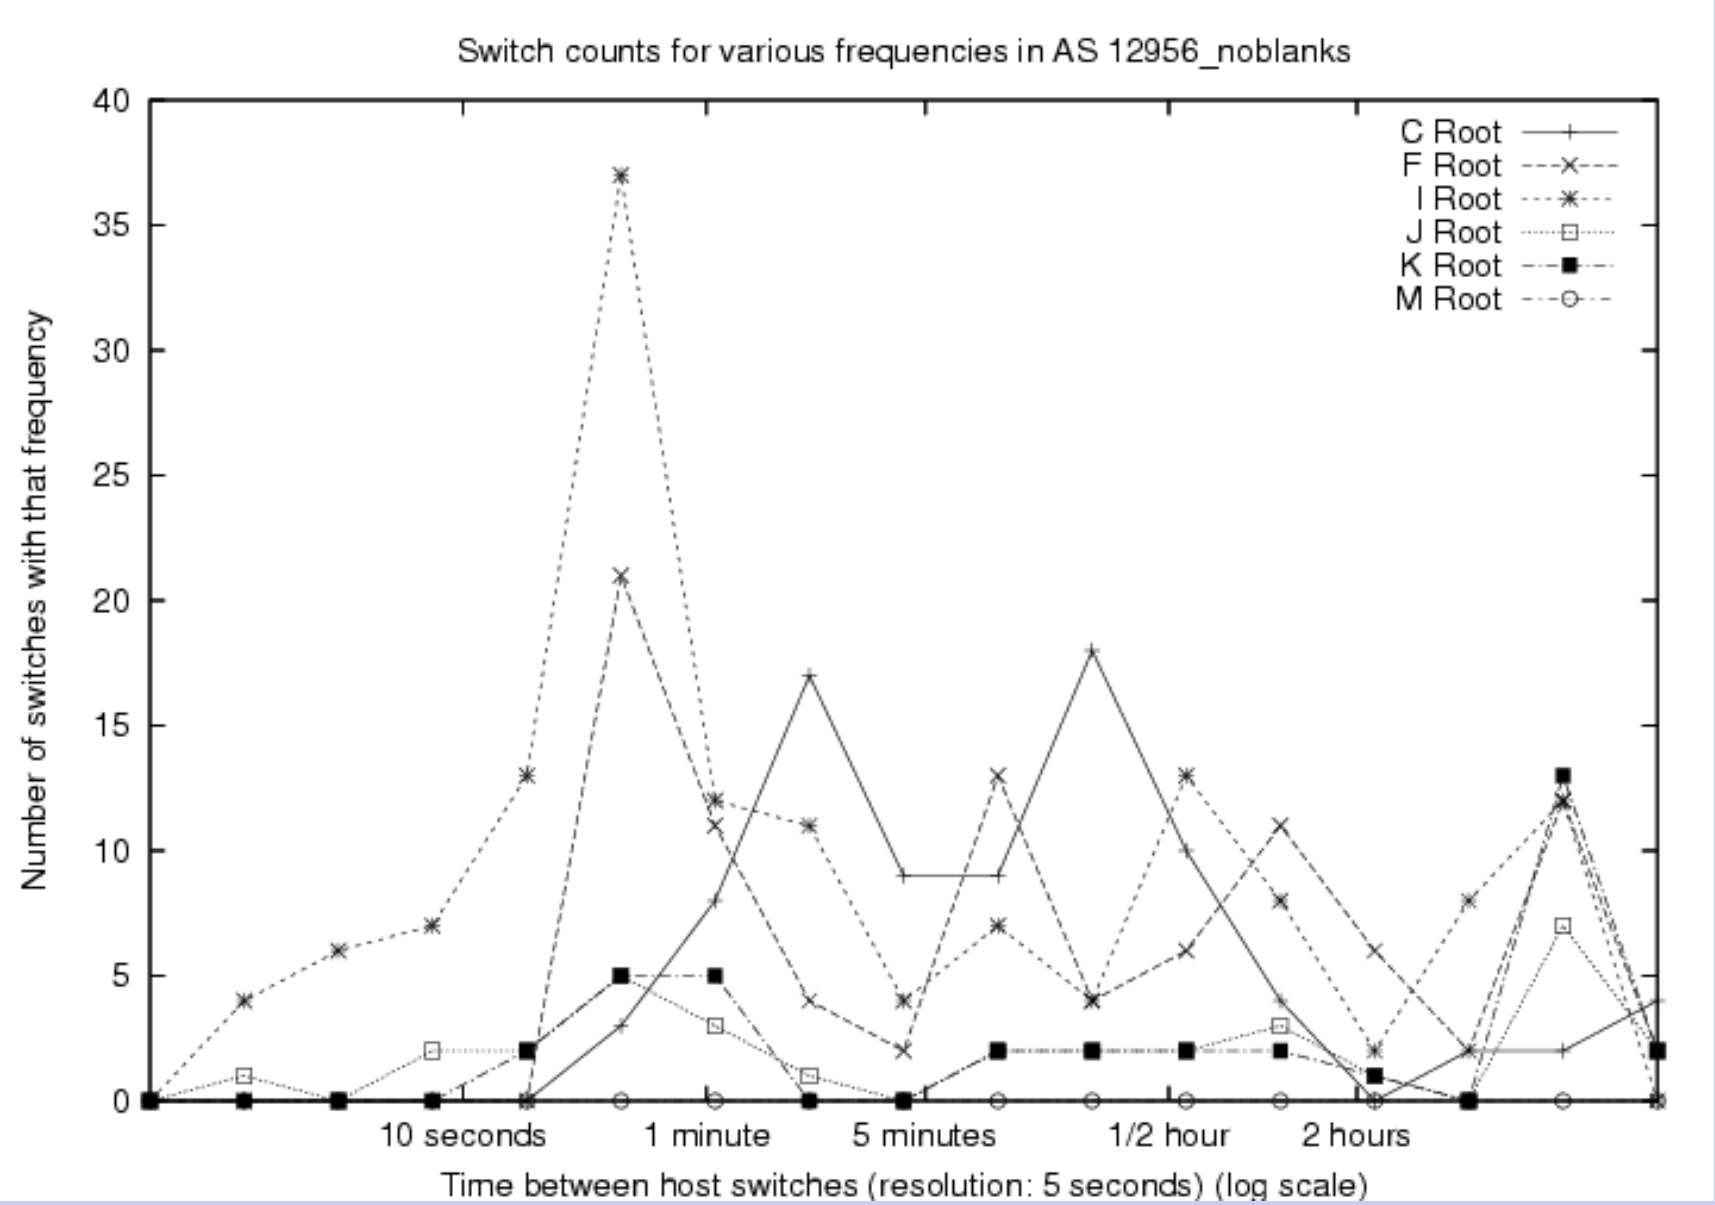

The Good

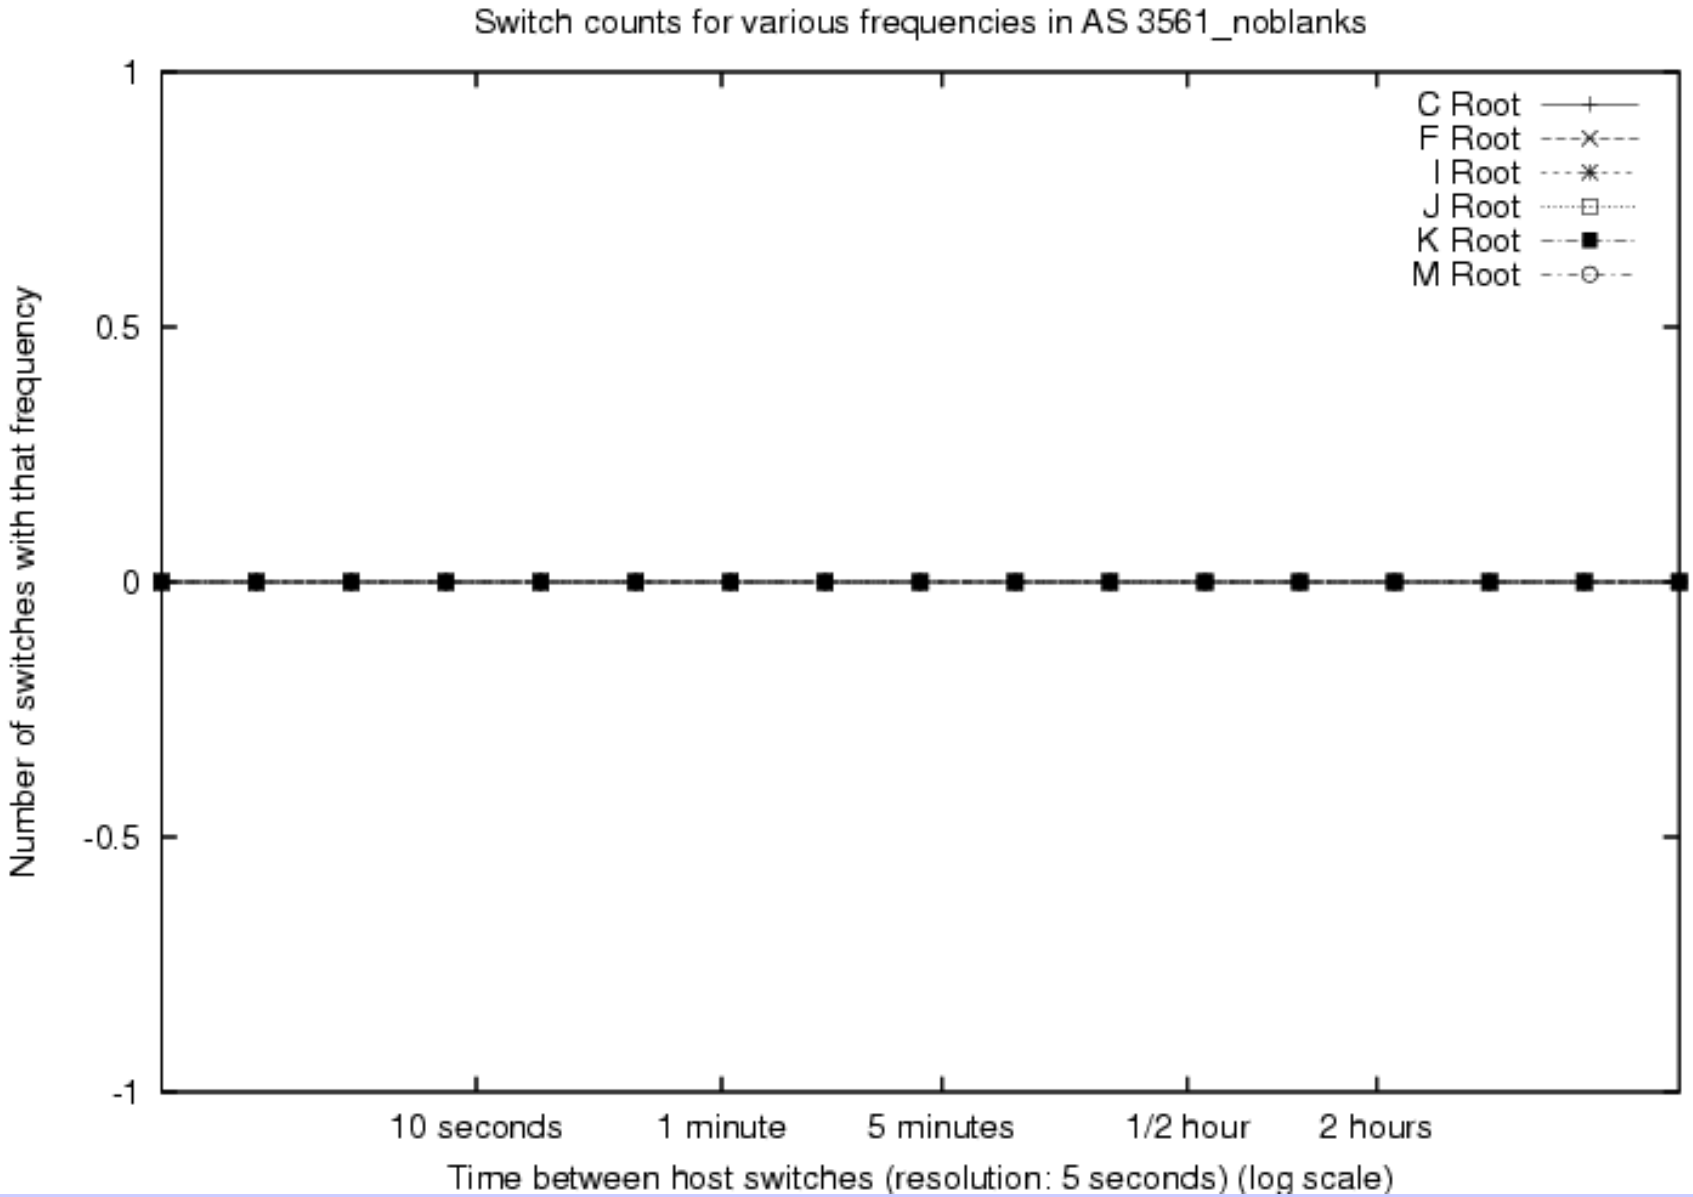

The Bad

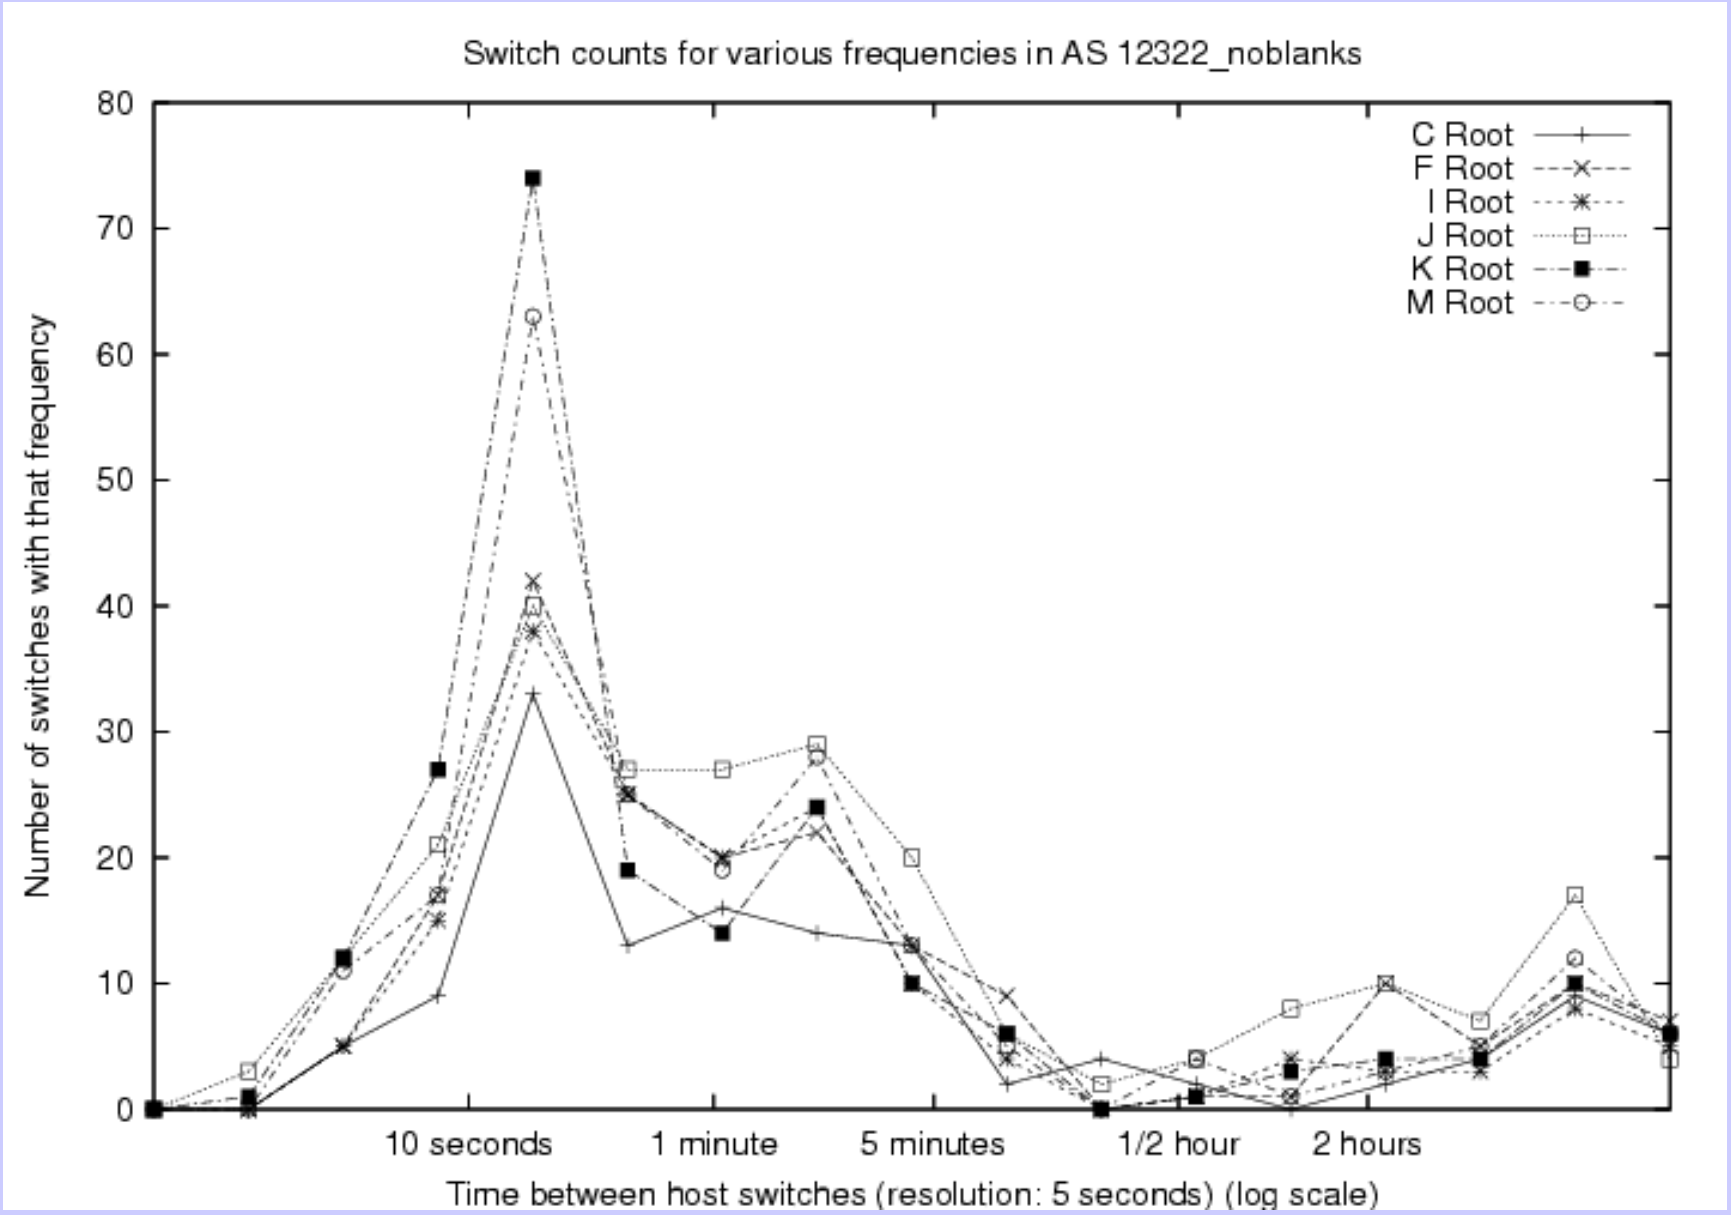

And the Ugly

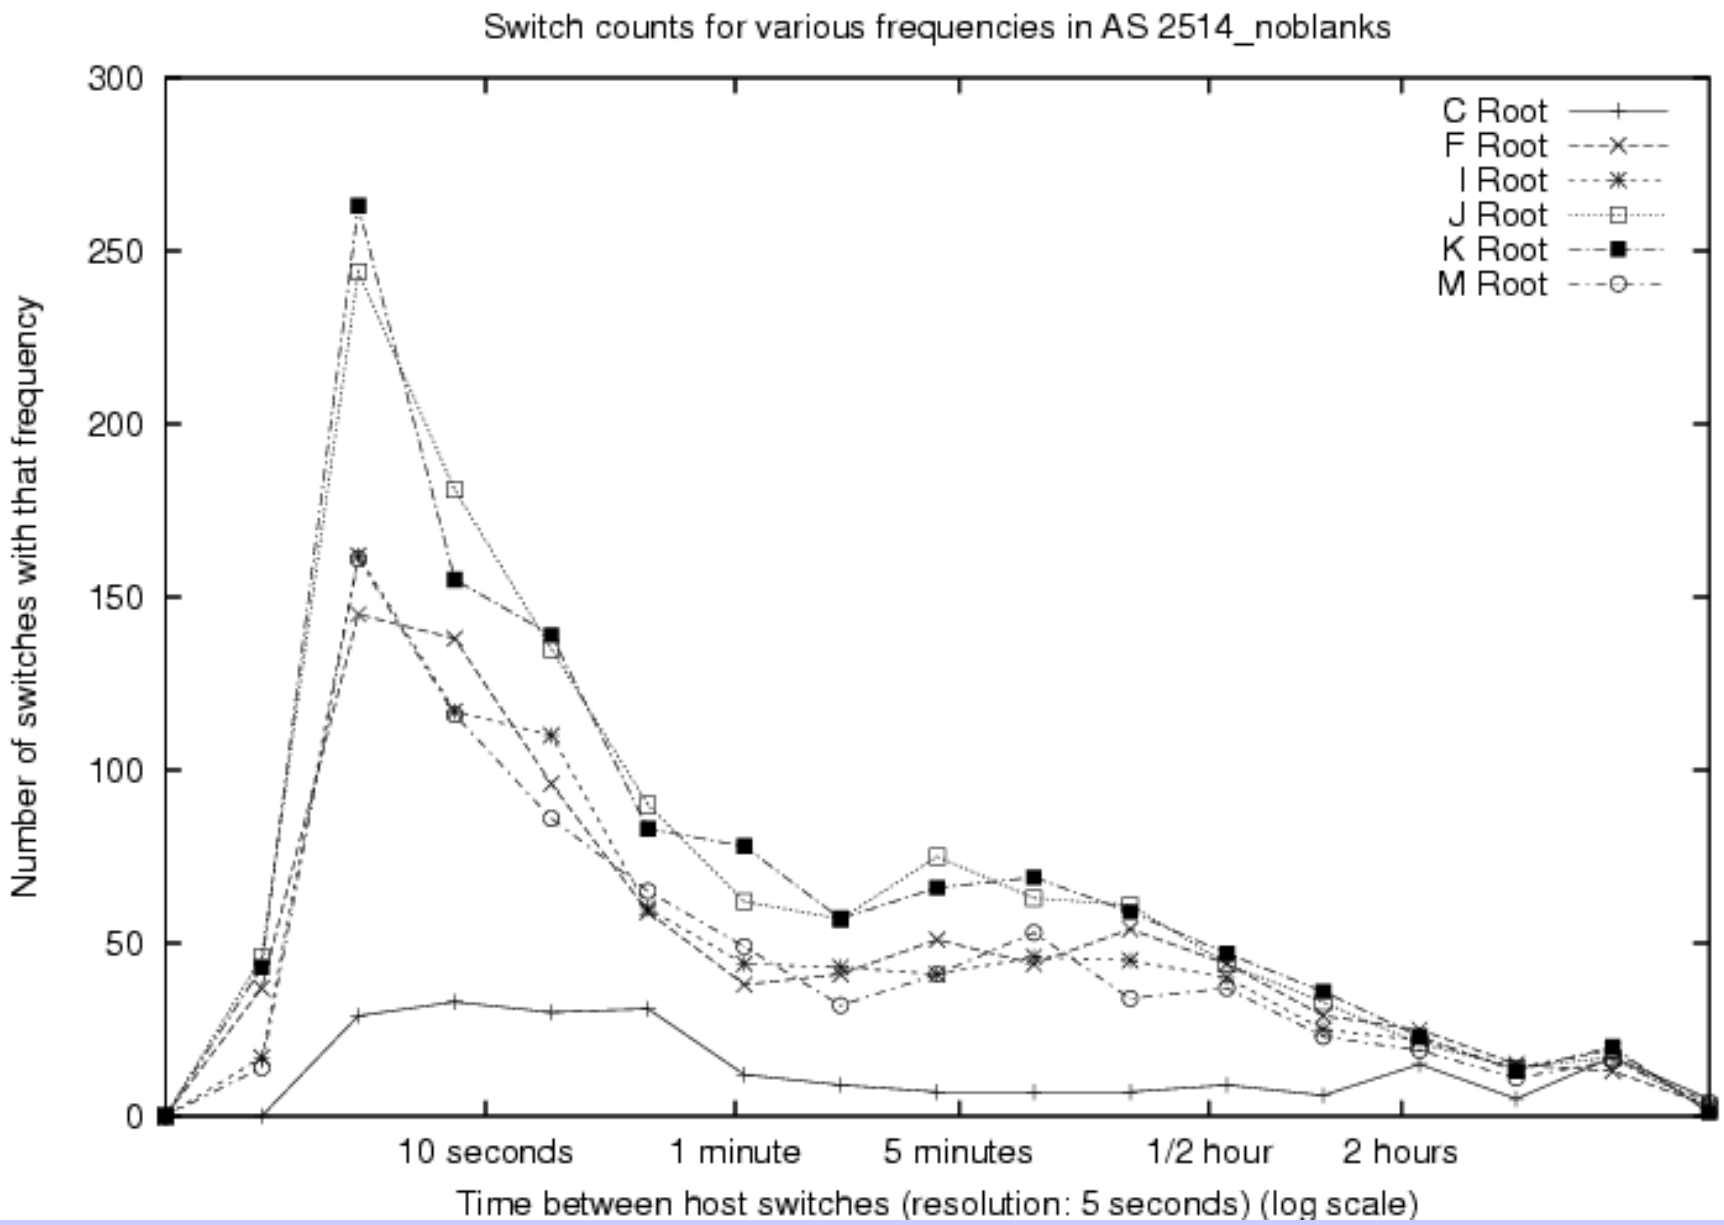

Even Uglier

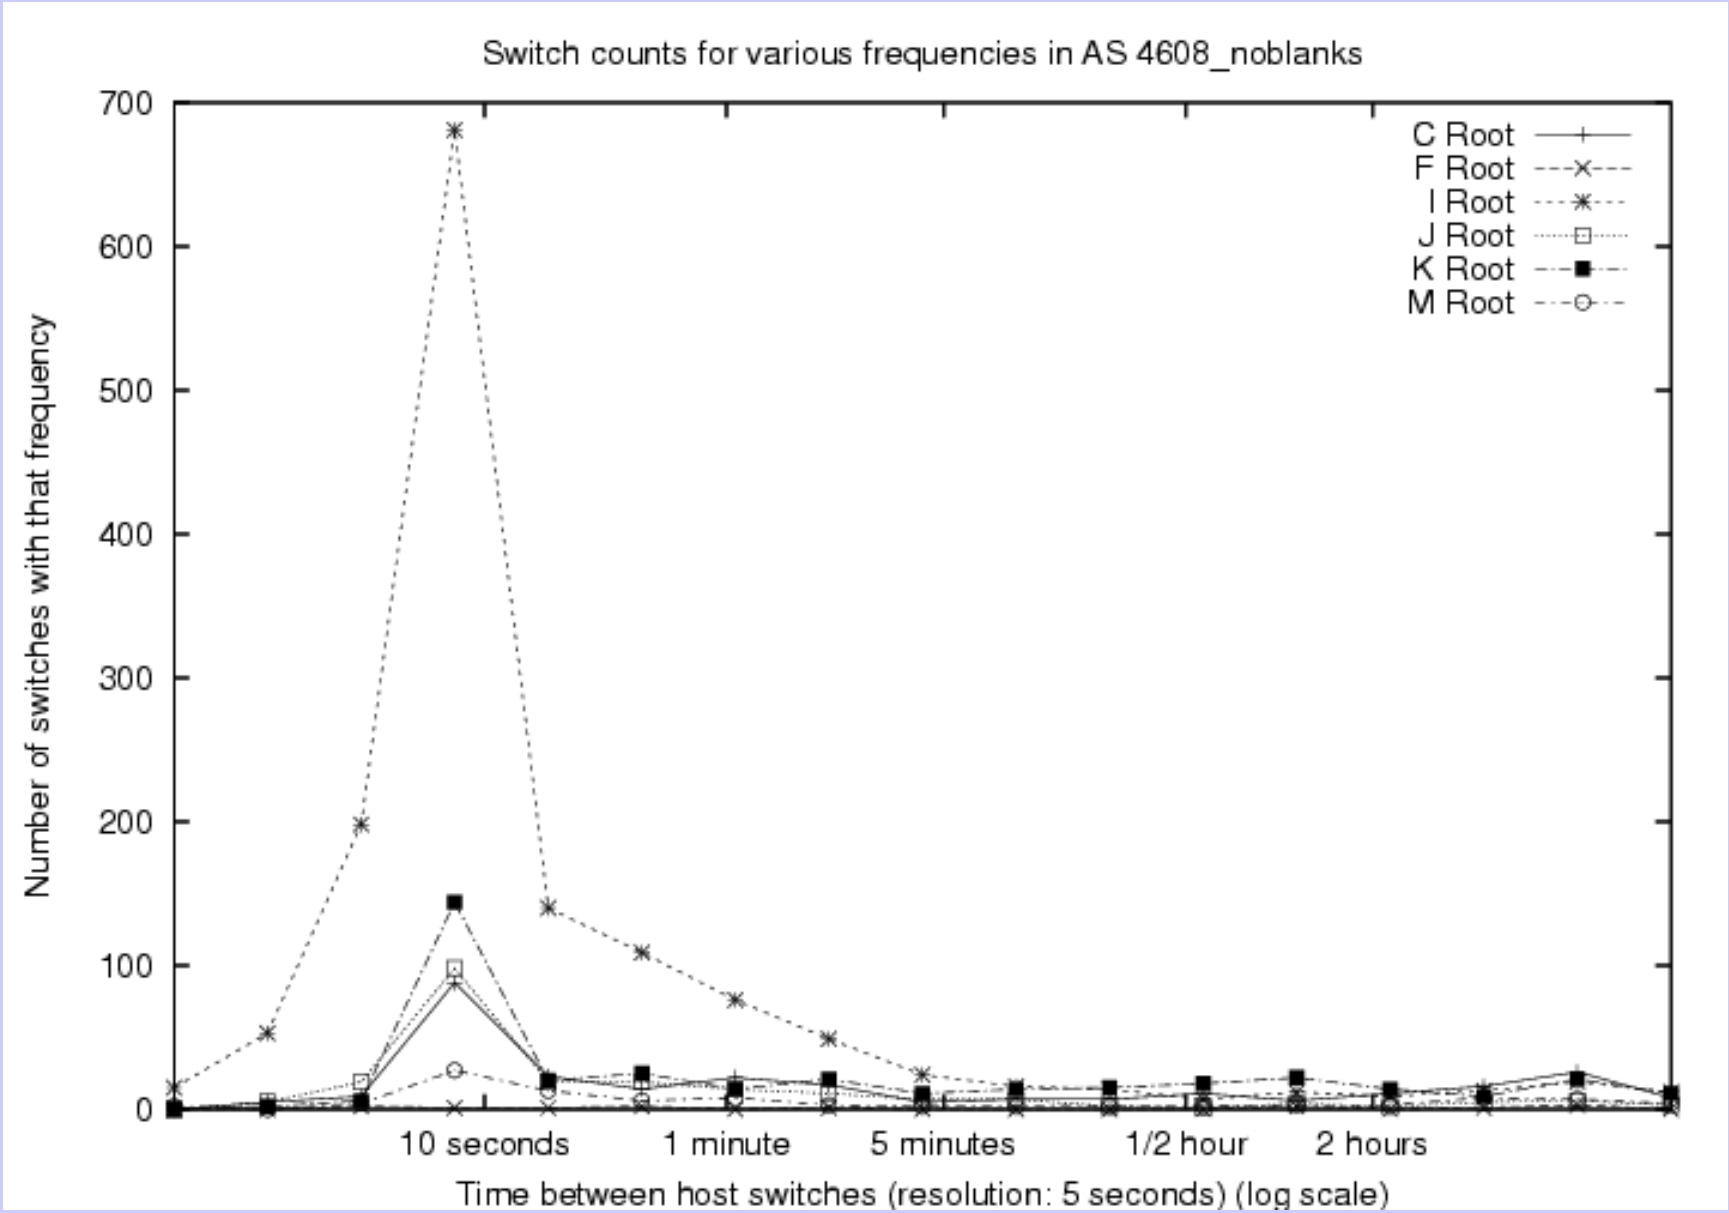

And it Goes On and On

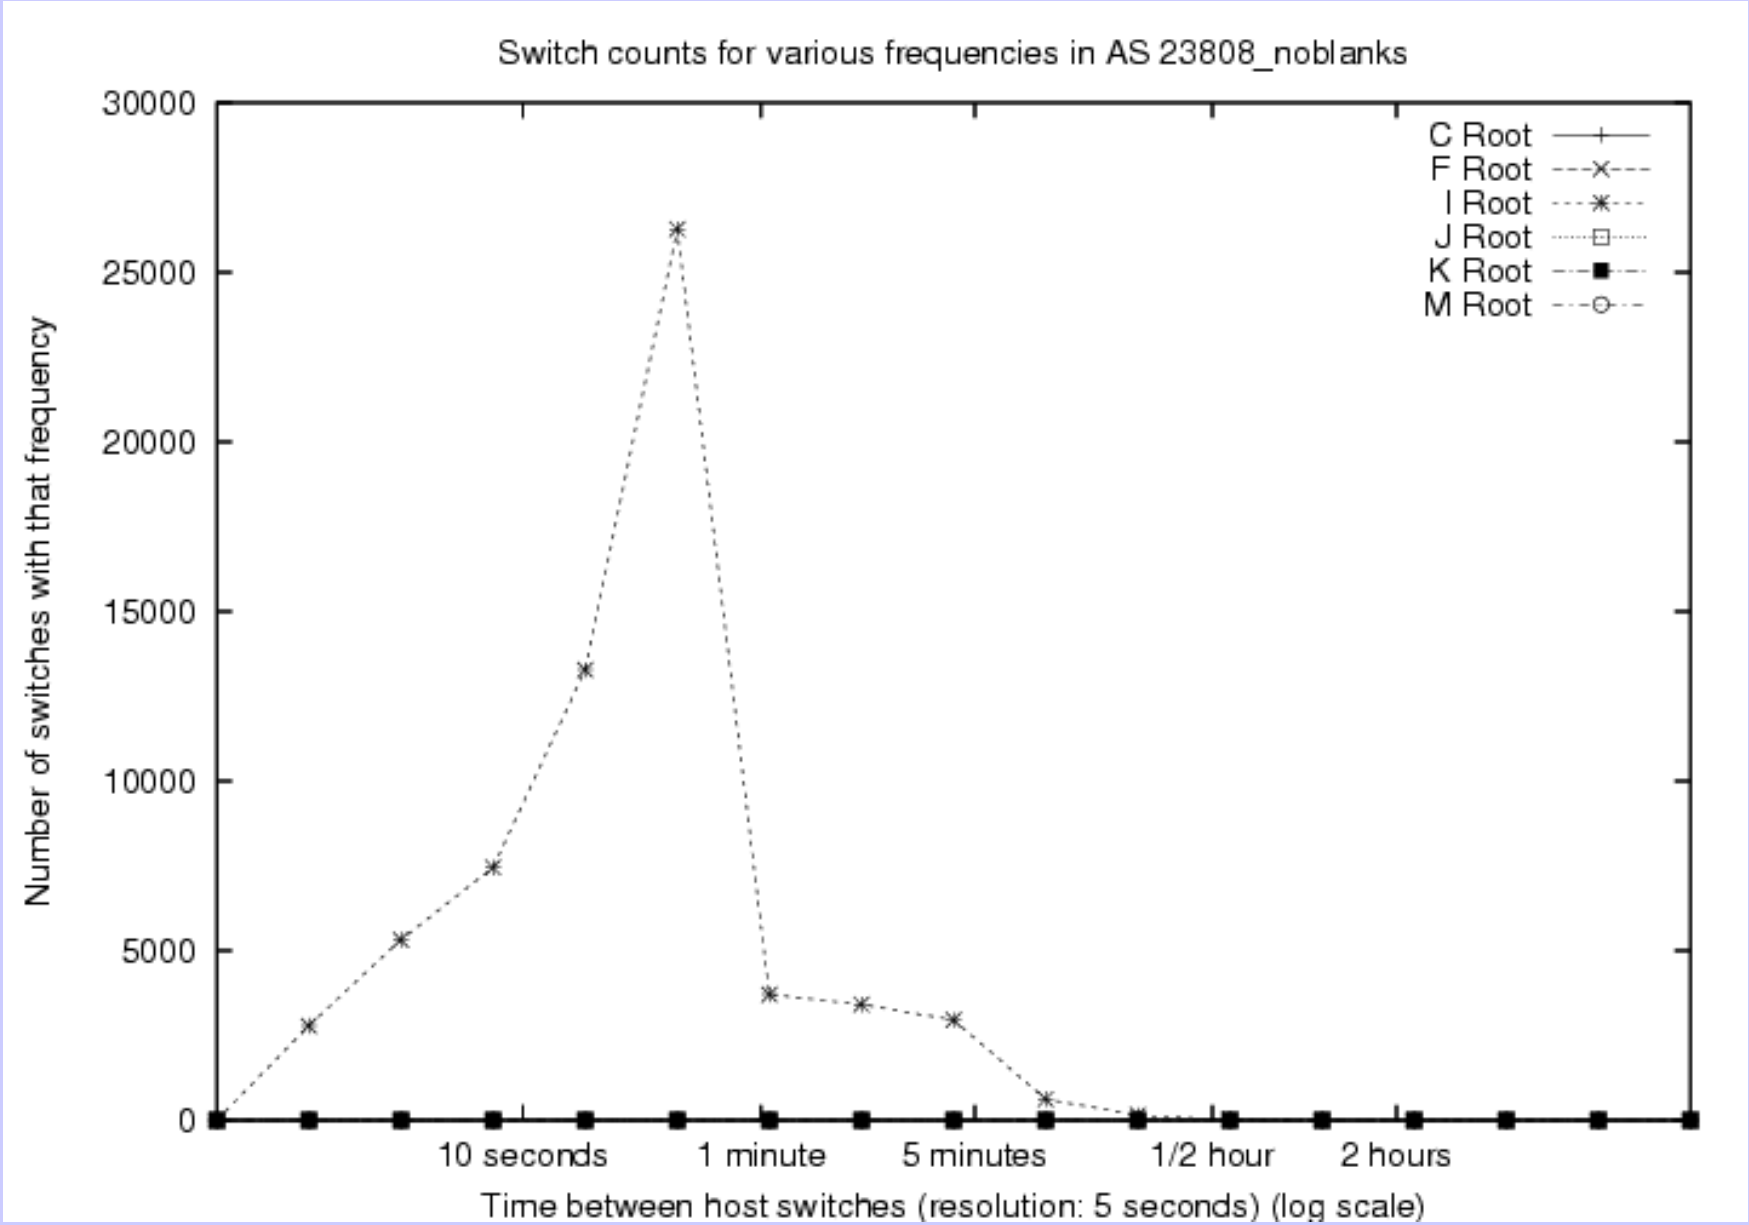

No-Response

Root	Probes	Drops	Percent Drop
c	63529640	27318	0.043000%
f	63529623	17534	0.027600%
i	63529602	14865	0.023399%
j	63529595	14588	0.022963%
k	63529553	1114427	1.754187%
m	62732897	18041	0.028758%

No-Response is Often Poisson

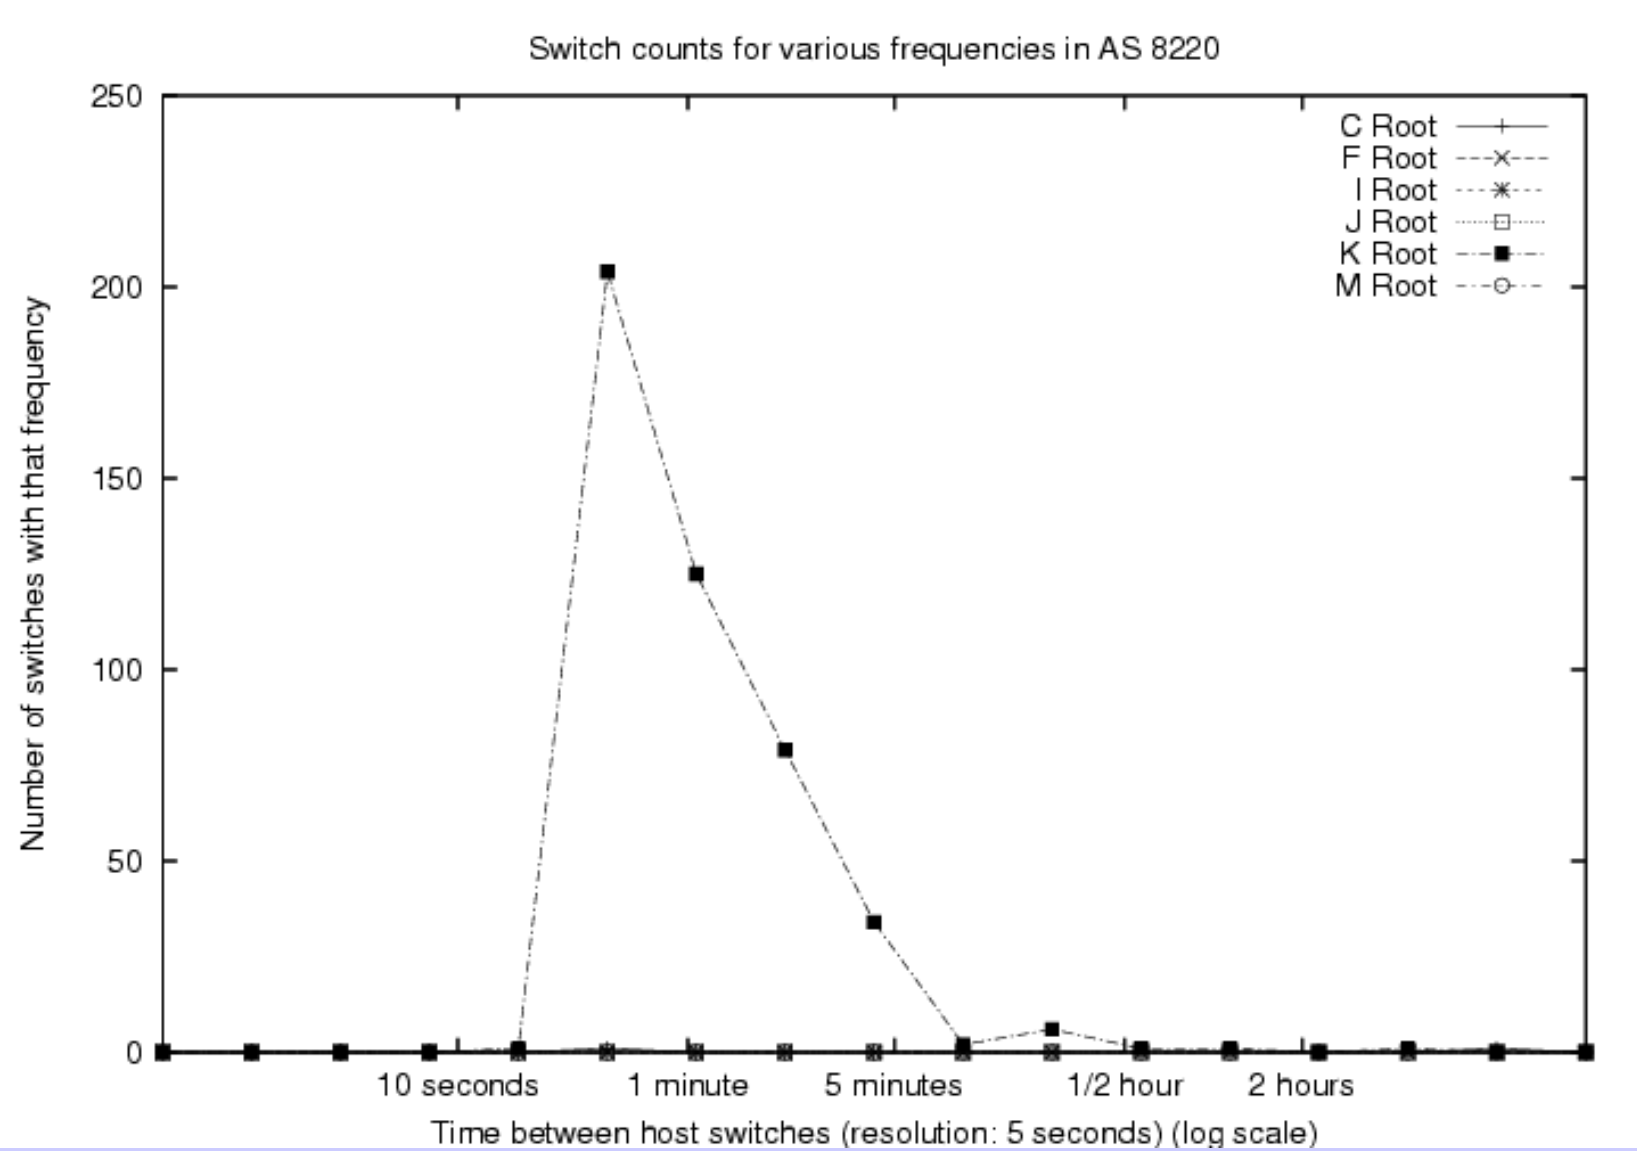

Thoughts

- This may be more about an AS's routing than about root server stability
- Is it due to AS's IGP?
- But it makes no difference to the user They still can not tell which server gave a DNS resolution problem
- Resiliency has been gained, but seemingly at the expense of transparency and debugability
- Intra-AS anycast may be more stable

Thanks To

- Mark Kosters, who asked a question
- The many volunteer host admins who ran the experiment
- NSF via award ANI-0221435
- The University of Oregon (Peter, Lucy, Joel)
- Internet Initiative Japan (Randy)