DNS Anycast Stability

Some Early Results

WIDE/CAIDA
2005.03.11

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

<http://rip.psg.com/~randy/050311.wide-anycast.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)

UDP

			Overall		
			UDP		
	Probes	Switches	Failures		
c:	117233282	1893	46652	0.001615%	0.039794%
f:	117233252	1870	59373	0.001595%	0.050645%
i:	117233208	60870	73980	0.051922%	0.063105%
j:	117233194	3202	28063	0.002731%	0.023938%
k:	117233136	2499	25146	0.002132%	0.021450%
m:	115639820	1449	204740	0.001253%	0.177050%
			Planetlab		
			UDP		
	Probes	Switches	Failures		
c:	23458902	186	19444	0.000793%	0.082885%
f:	23458902	181	33160	0.000772%	0.141354%
i:	23458892	158	54386	0.000674%	0.231835%
j:	23458904	285	5467	0.001215%	0.023305%
k:	23458900	191	2735	0.000814%	0.011659%
m:	21868146	62	2960	0.000284%	0.013536%
			Internet		
			UDP		
	Probes	Switches	Failures		
c:	93774380	1707	27208	0.001820%	0.029014%
f:	93774350	1689	26213	0.001801%	0.027953%
i:	93774316	60712	19594	0.064743%	0.020895%
j:	93774290	2917	22596	0.003111%	0.024096%
k:	93774236	2308	22411	0.002461%	0.023899%
m:	93771674	1387	201780	0.001479%	0.215182%

TCP

			Overall		
			TCP		
	Probes	Switches	Failures		
c:	9789330	3791	8221	0.038726%	0.083979%
f:	9789326	3216	4914	0.032852%	0.050198%
i:	9789326	78634	6528	0.803263%	0.066685%
j:	9789326	4674	514473	0.047746%	5.255449%
k:	9789304	3686	206165	0.037653%	2.106023%
m:	9789306	2202	61480	0.022494%	0.628032%
			Planetlab		
			TCP		
	Probes	Switches	Failures		
c:	1987982	107	3903	0.005382%	0.196330%
f:	1987982	99	2957	0.004980%	0.148744%
i:	1987982	104	5408	0.005231%	0.272035%
j:	1987982	143	20517	0.007193%	1.032052%
k:	1987982	98	32679	0.004930%	1.643828%
m:	1987982	58	6588	0.002918%	0.331391%
			Internet		
			TCP		
	Probes	Switches	Failures		
c:	7801348	3684	4318	0.047223%	0.055349%
f:	7801344	3117	1957	0.039955%	0.025085%
i:	7801344	78530	1120	1.006621%	0.014357%
j:	7801344	4531	493956	0.058080%	6.331678%
k:	7801322	3588	173486	0.045992%	2.223803%
m:	7801324	2144	54892	0.027483%	0.703624%

Failures

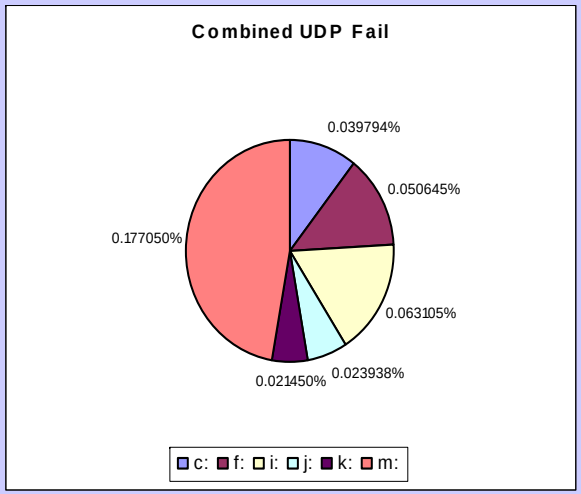

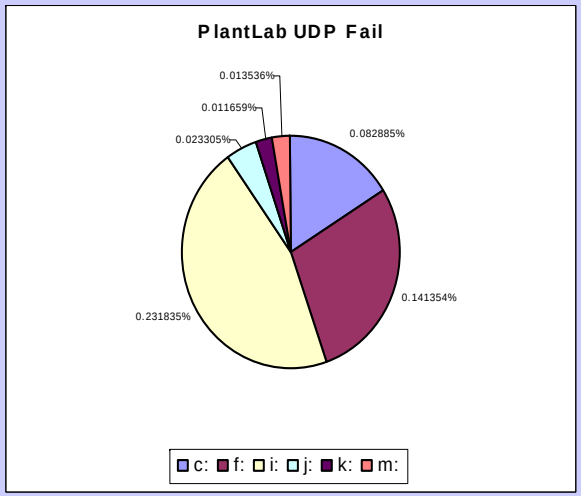

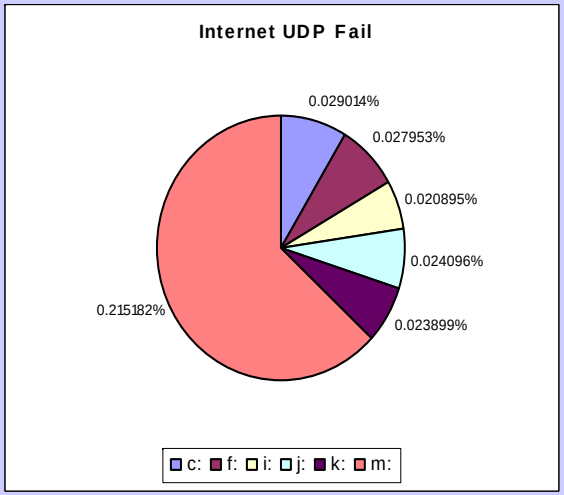

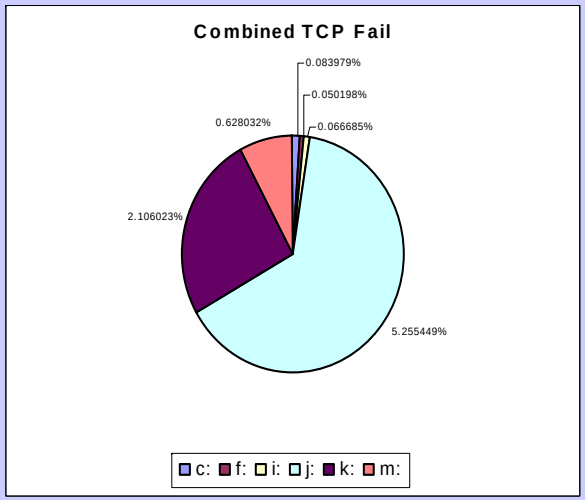

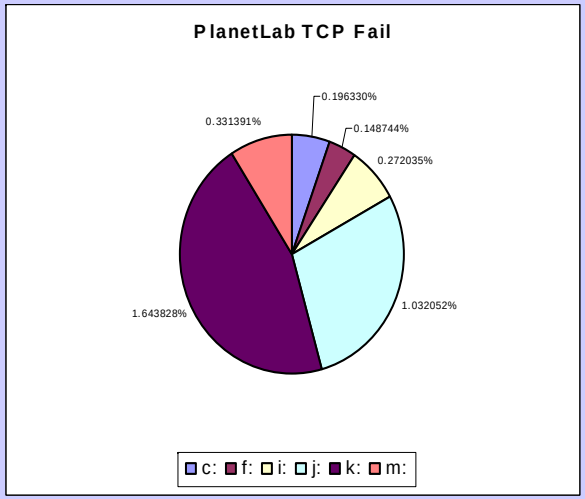

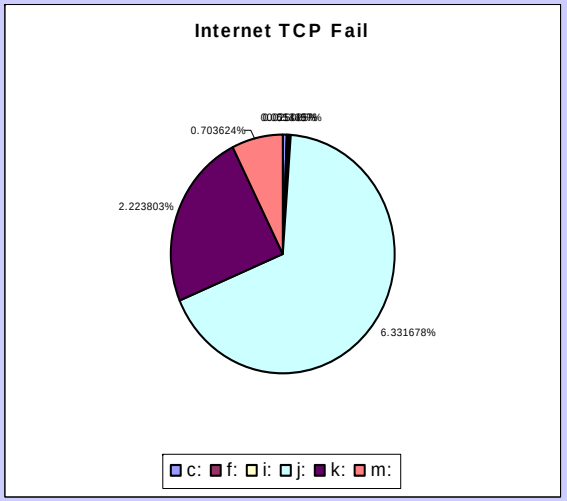

Switches

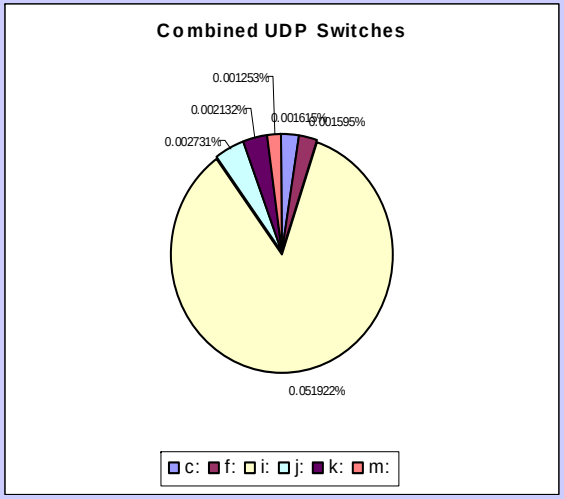

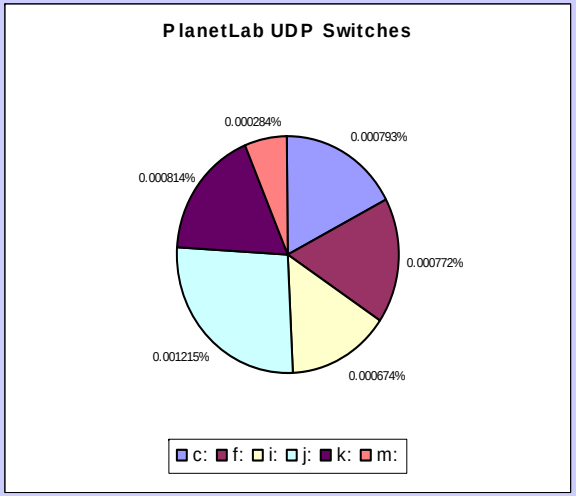

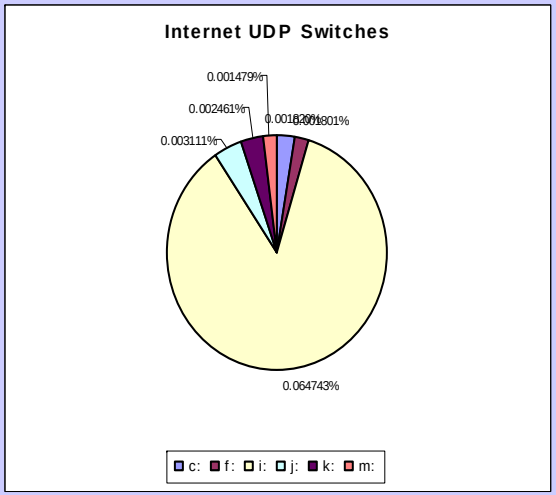

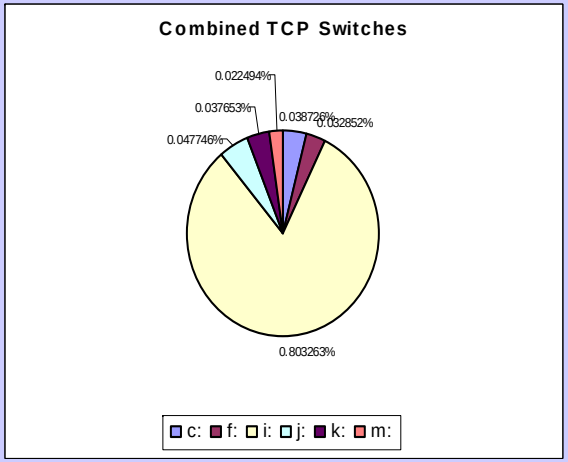

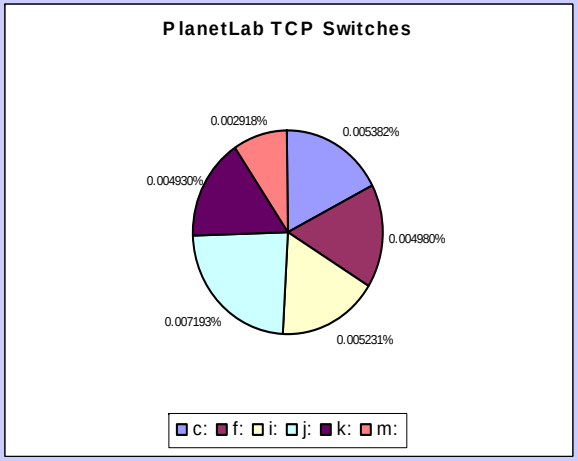

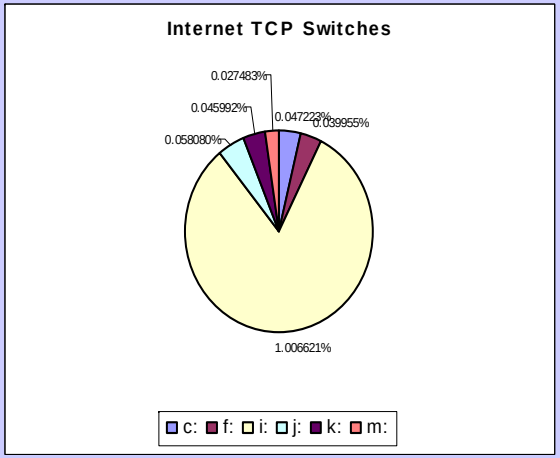

Switches v Failures - UDP

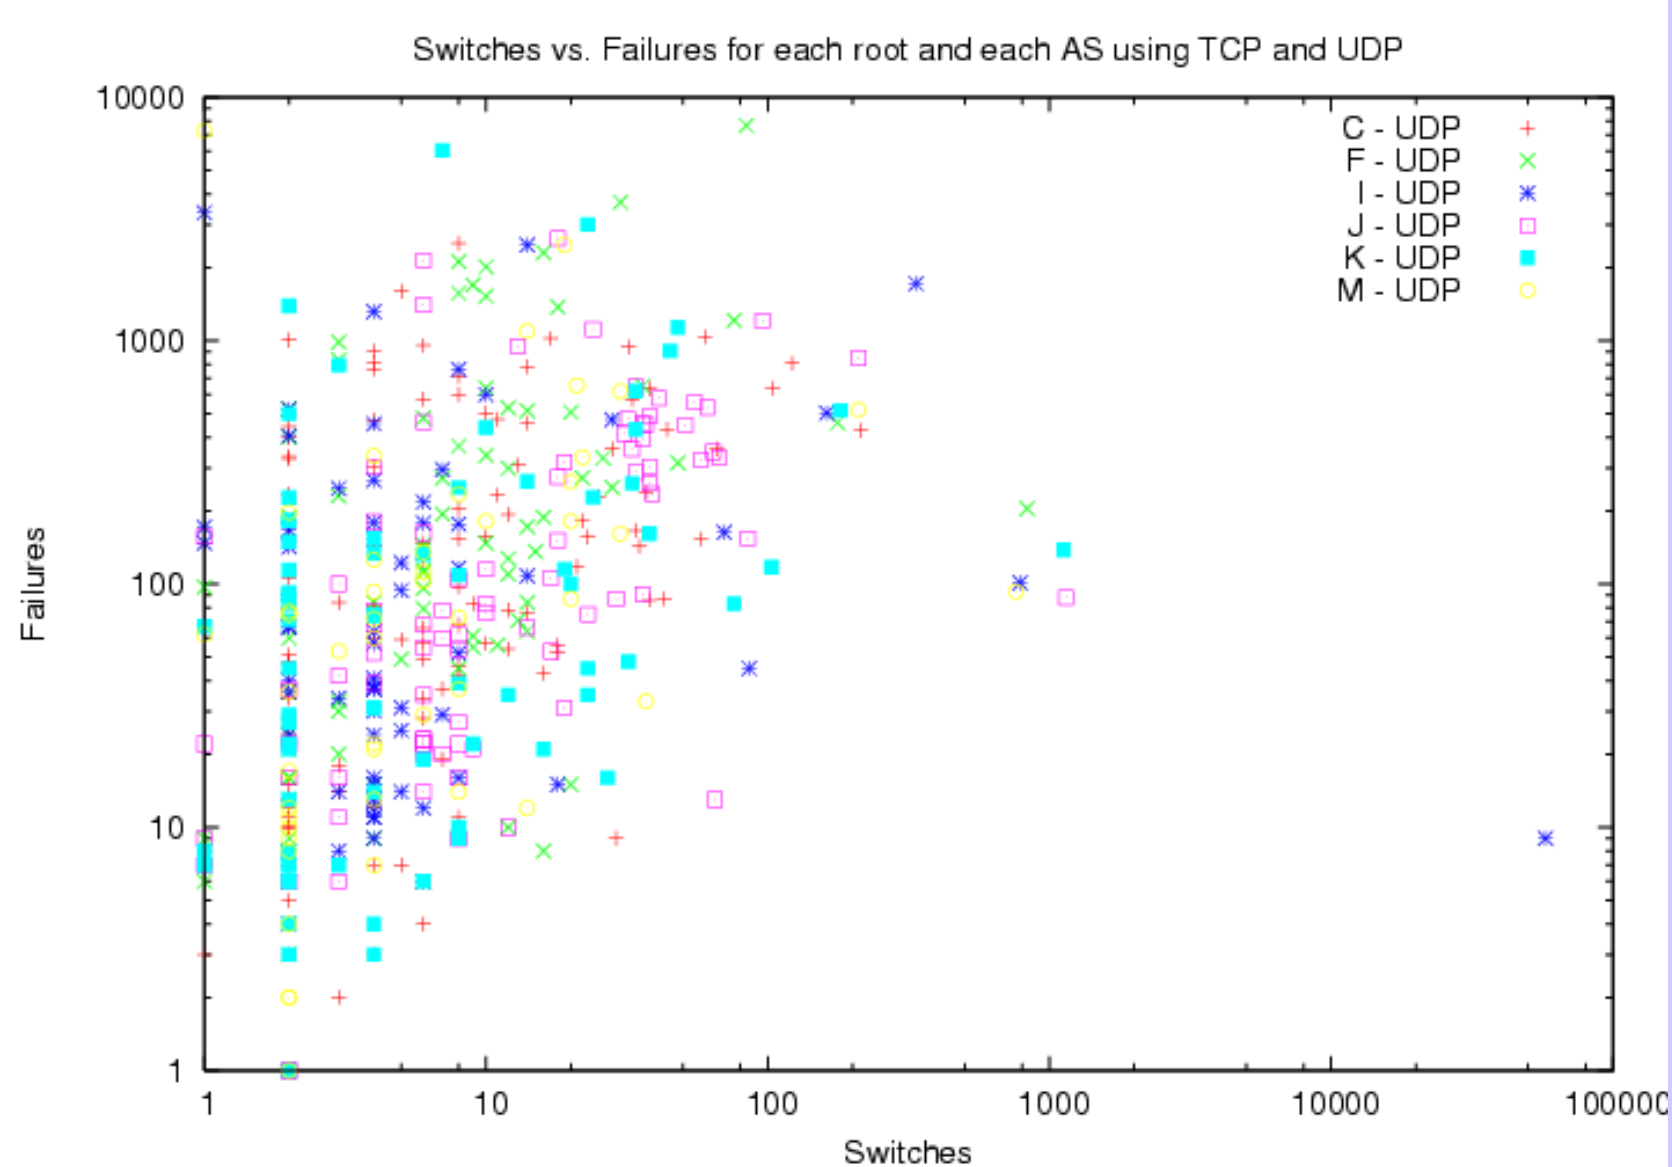

Switches v Failures - TCP

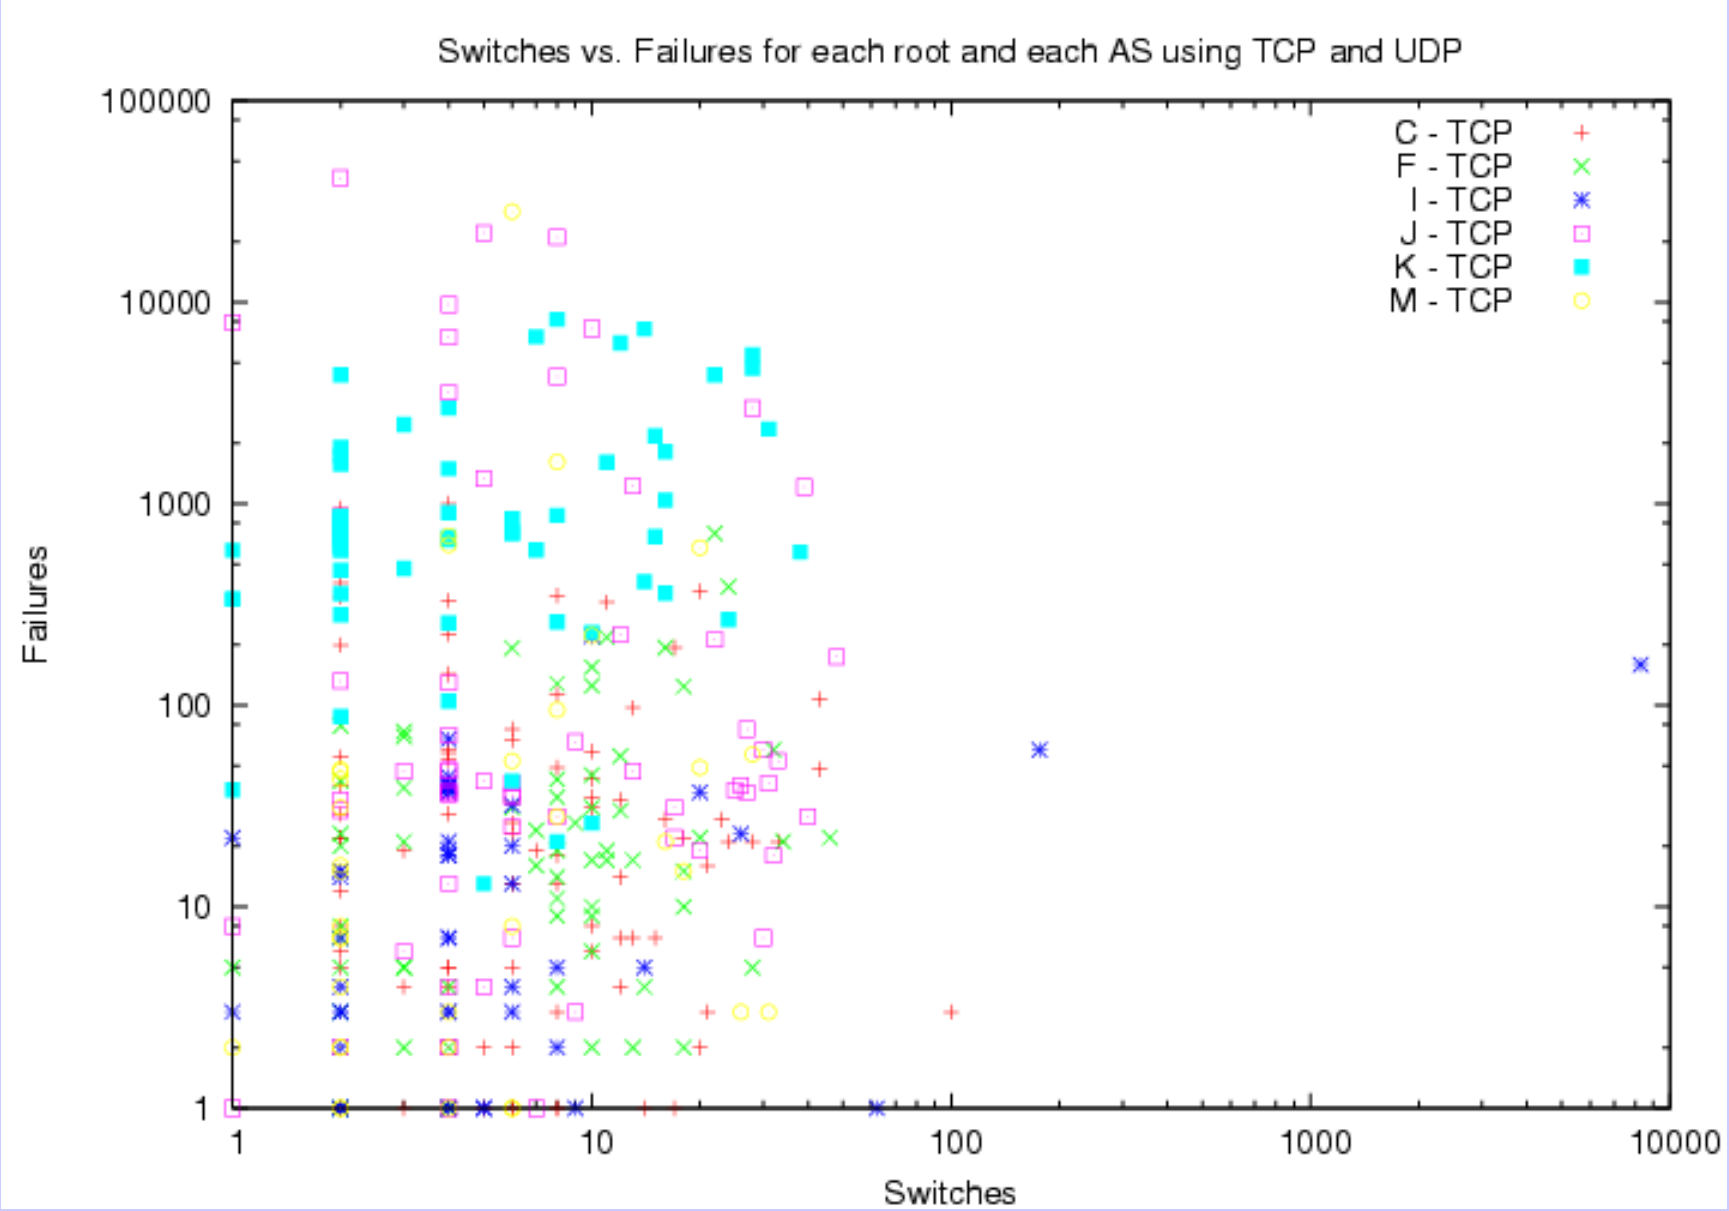

One AS's View

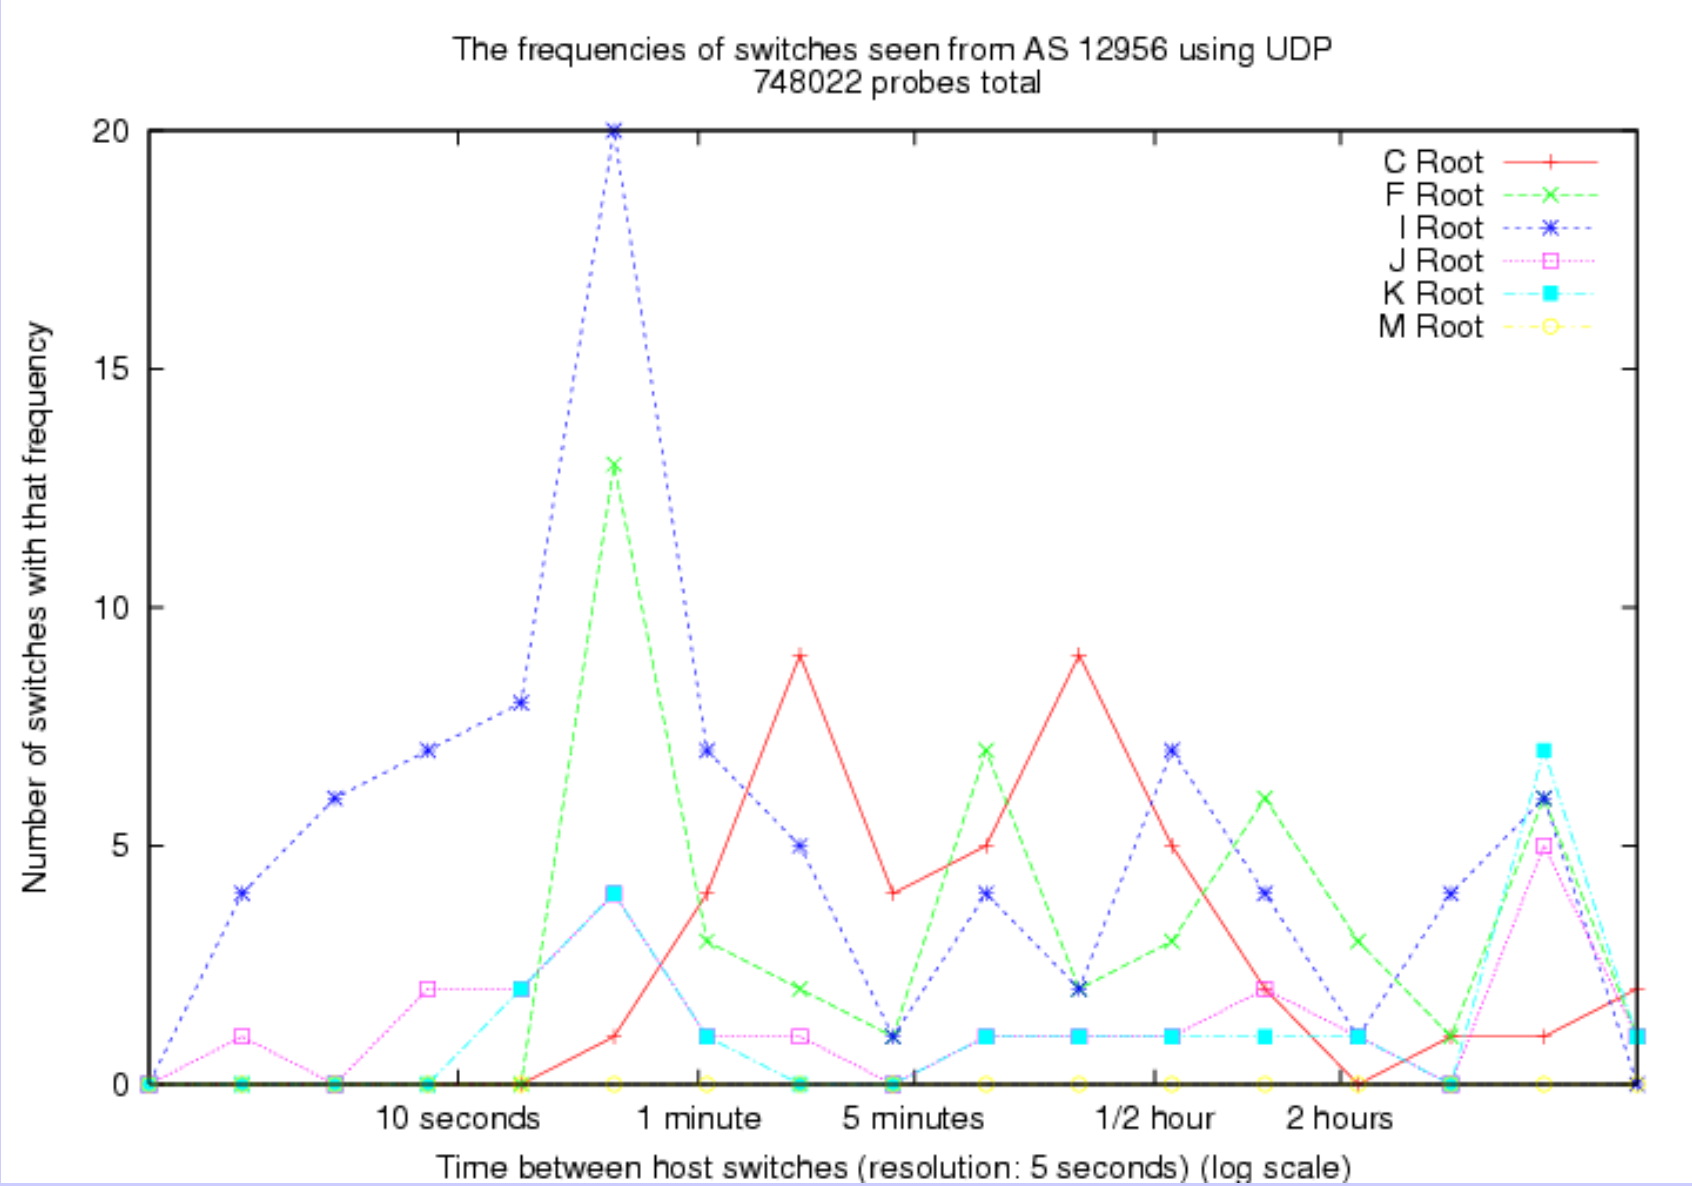

The Good

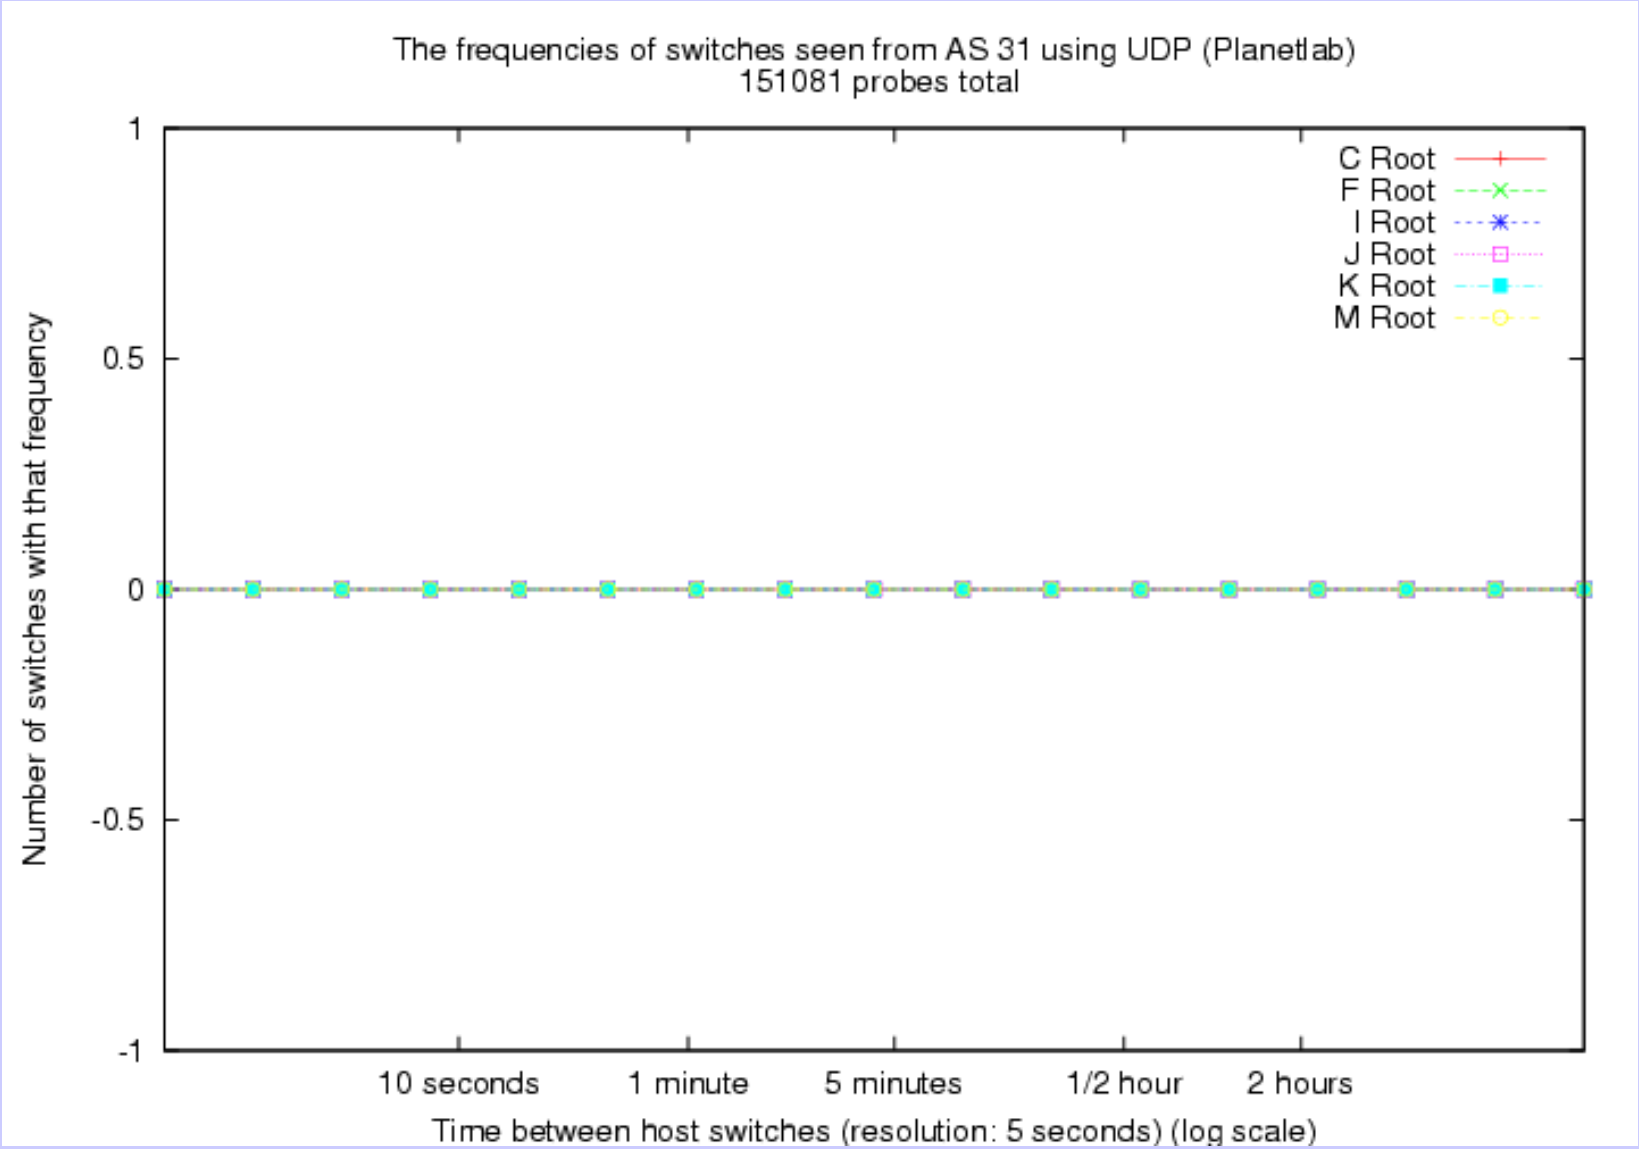

The Bad

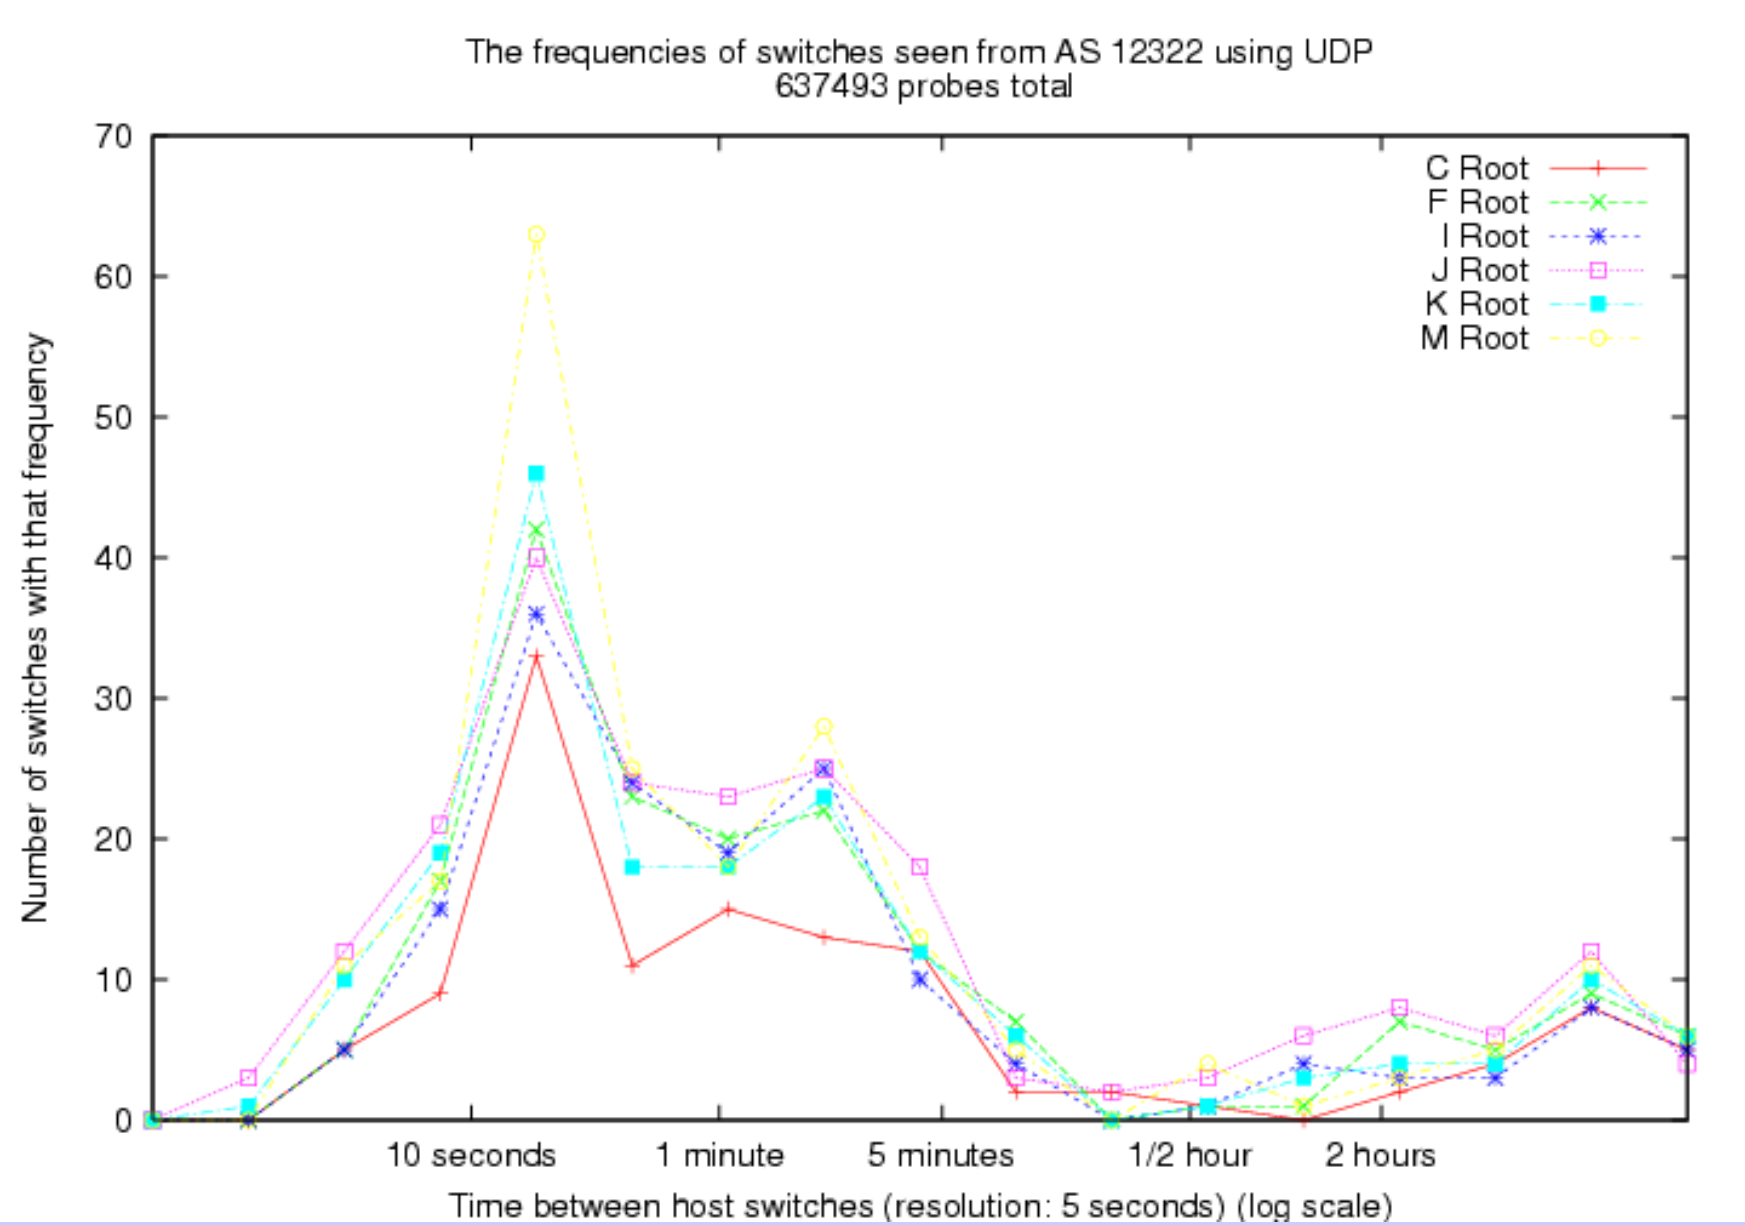

And the Ugly

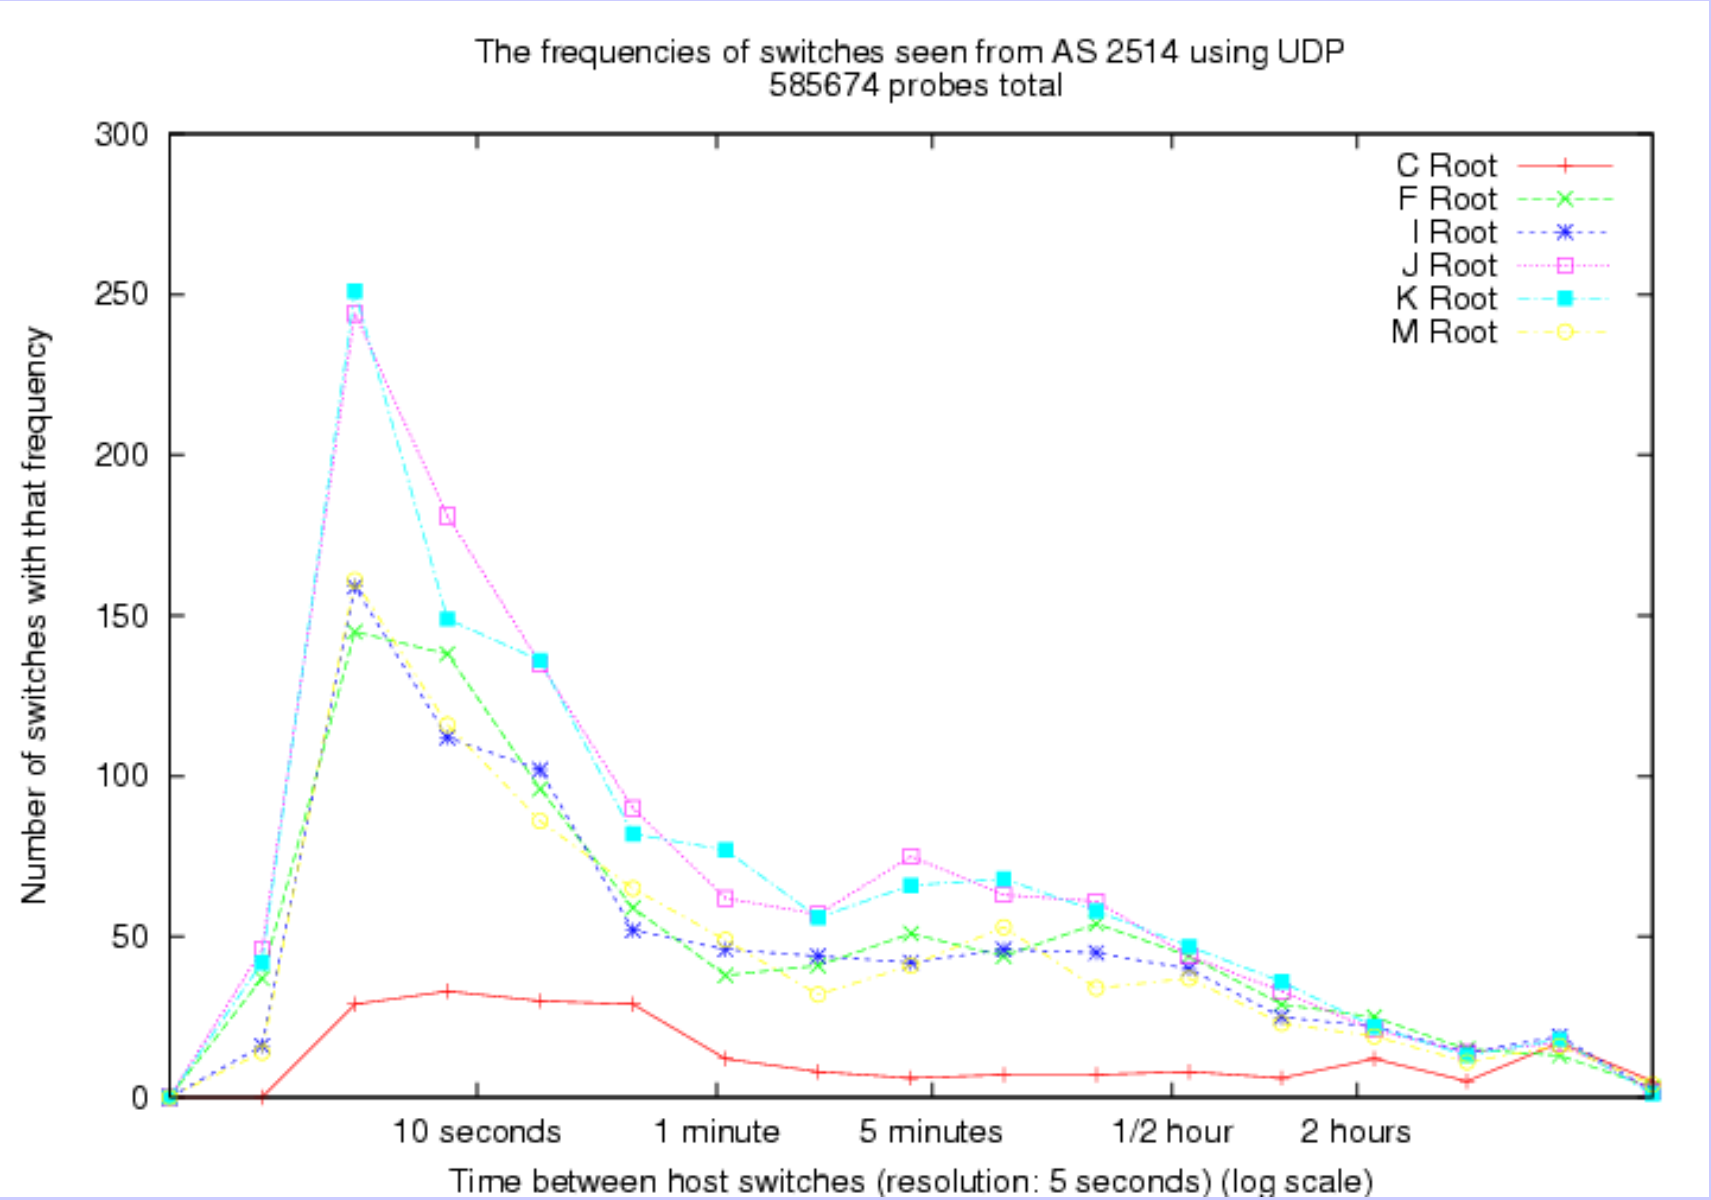

Even Uglier

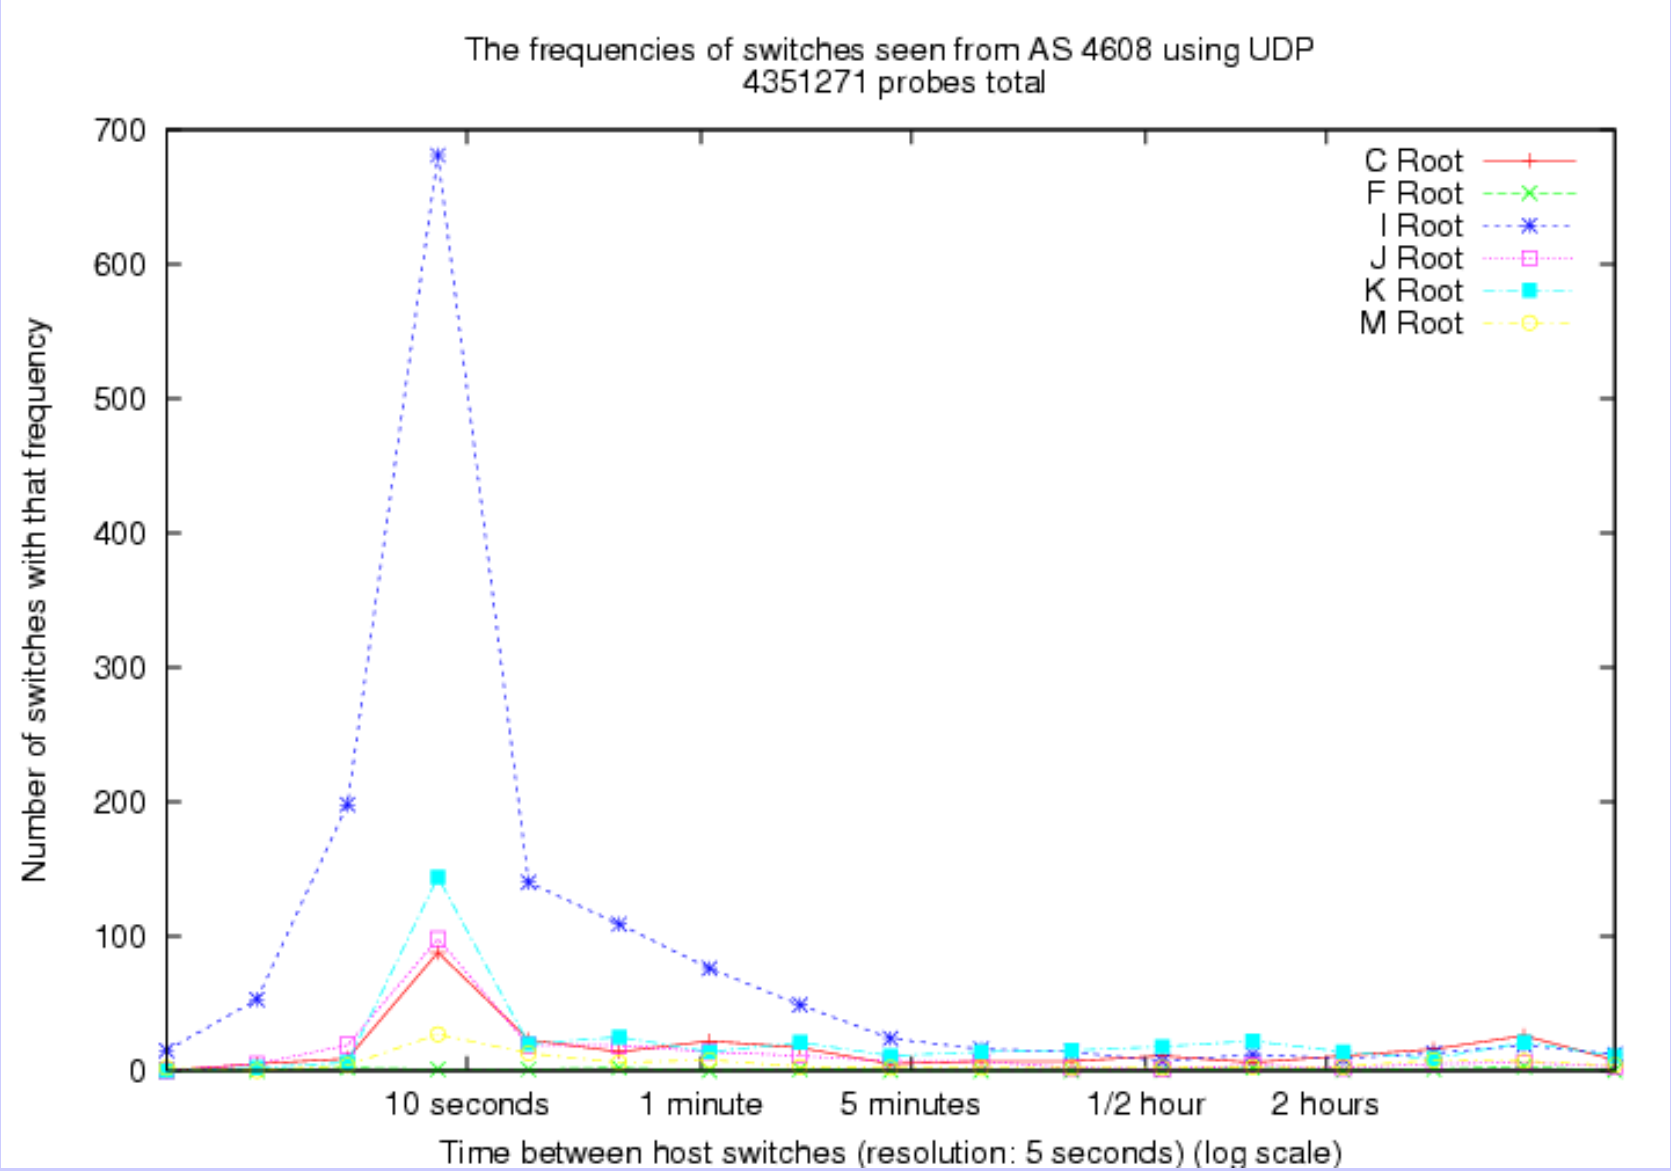

And it Goes On and On

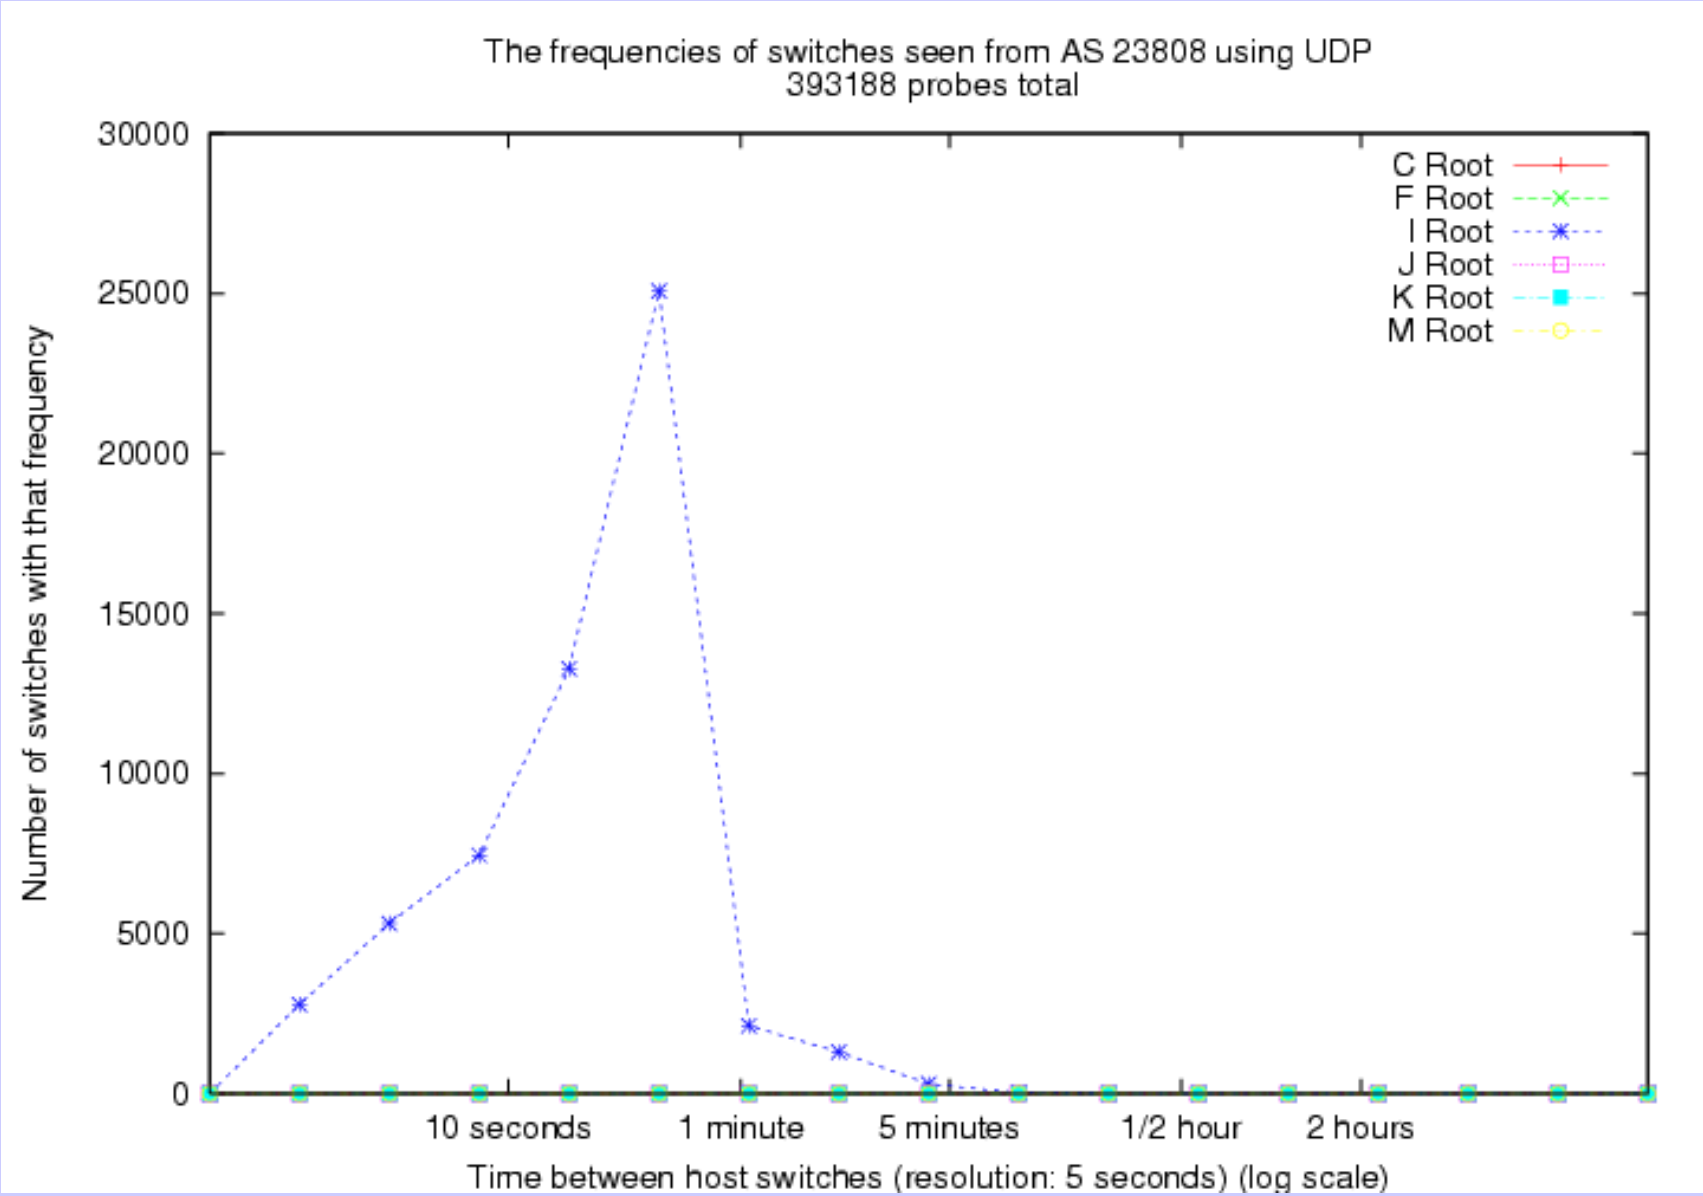

Failures are Sometimes Poisson

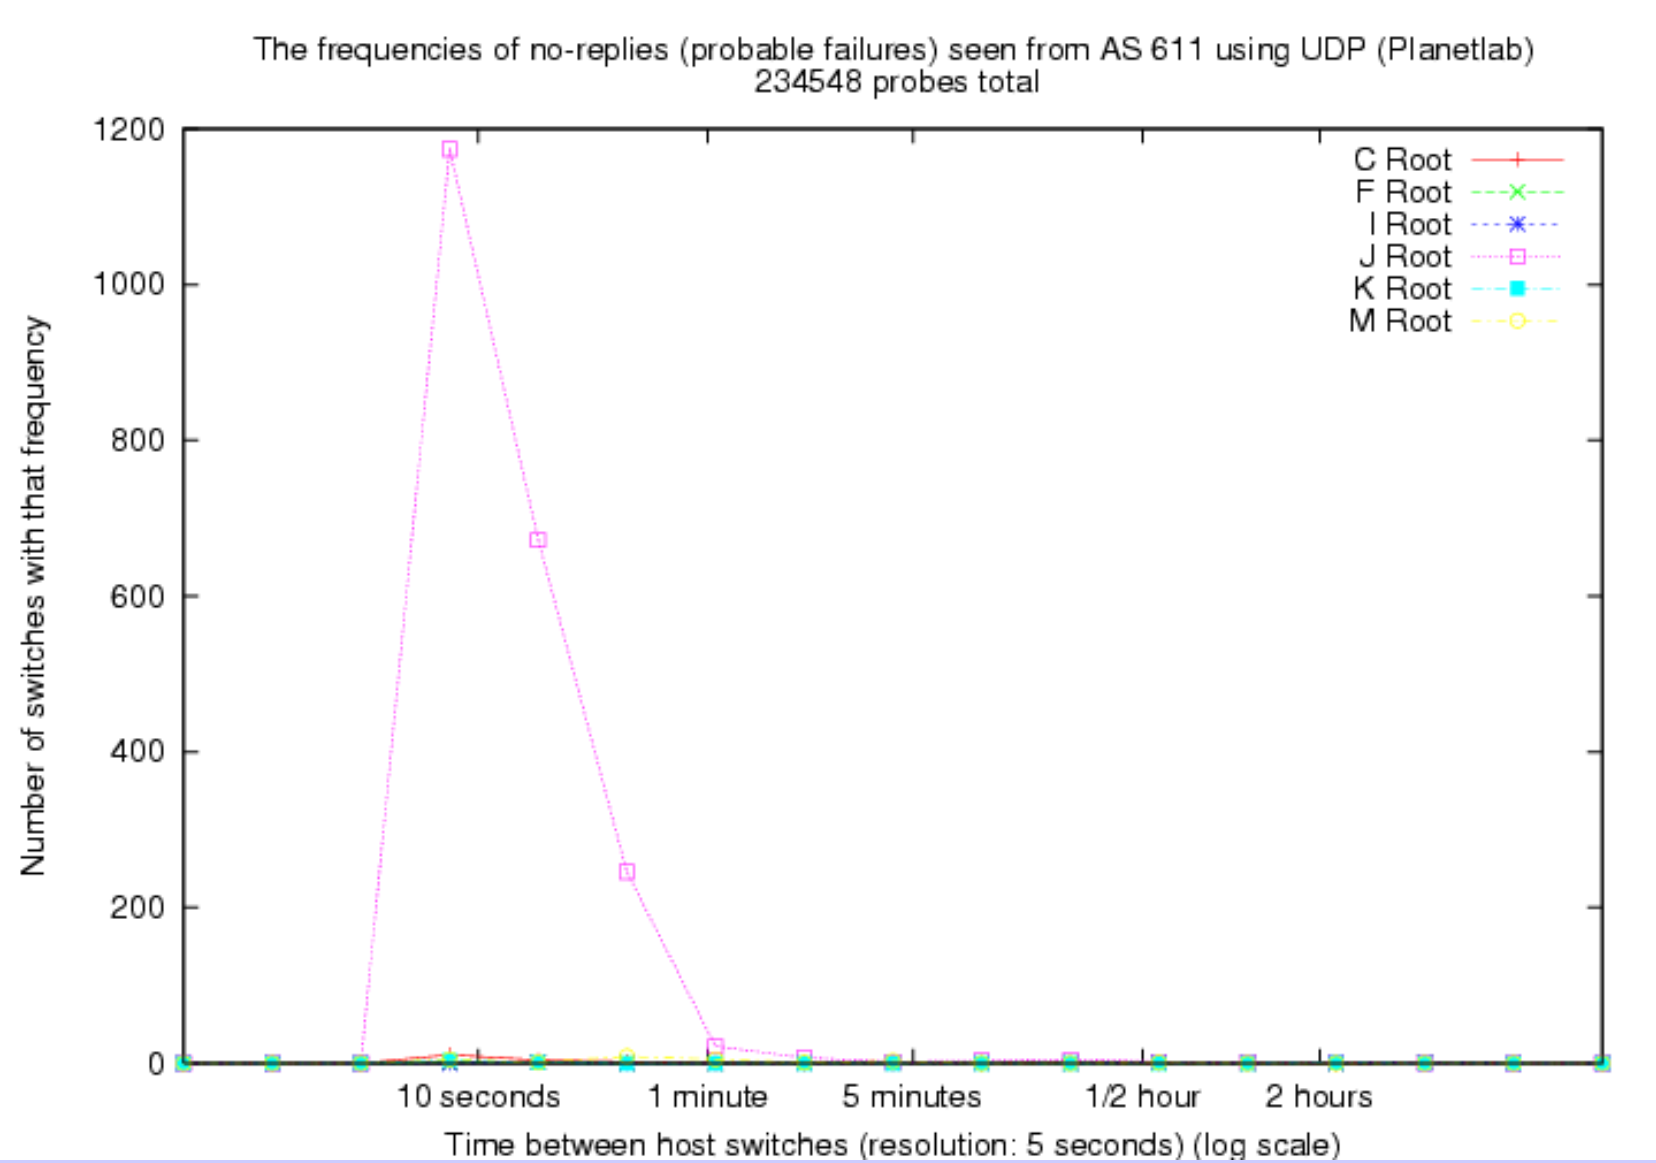

Thoughts

- This likely is more about an AS's routing than about root server stability per se
- Anycast may amplify/exacerbate routing tweakiness, inter- and/or intra-AS
- How much is 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
- PlanetLab is not the same as the Internet

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)