



Beyond the Top Talkers

Empirical Correlation of Conficker-C Infected IP Space

Rhiannon Weaver, CERT/NetSA
FloCon 2010
January 12th 2010



NO WARRANTY

THIS MATERIAL OF CARNEGIE MELLON UNIVERSITY AND ITS SOFTWARE ENGINEERING INSTITUTE IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this presentation is not intended in any way to infringe on the rights of the trademark holder.

This Presentation may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

Introduction

Volume statistics and Top-N lists

- Often we prioritize loudest talkers by IP Address
- N has to be small for manual analytic workflow
- NAT, DHCP complicate the picture
- Uncleanliness:
 - Bad guys tend to clump together by administered network
 - Net blocks, responsible parties, and WHOIS, oh my!

What does this look like in a case study?

- Conficker-C botnet
- What does dynamic allocation look like?
- Who do we find with an IP focus? What about /24s?
- Show some pretty pictures

Looking at Conficker-C

Network telescope into
infections:

```
rwfilter --start=2009/03/05:00 --end=2009/03/25:00 /  
--type=in --proto=17 --sport=1024- --pass=stdout | /  
rwfilter -input=stdin --d-conficker --dyn=conficker.so /  
-pass=conCtraffic.rw
```

Conficker-C hosts scan
the internet randomly


Flow rates vary among
IP addresses, /24 blocks

Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks

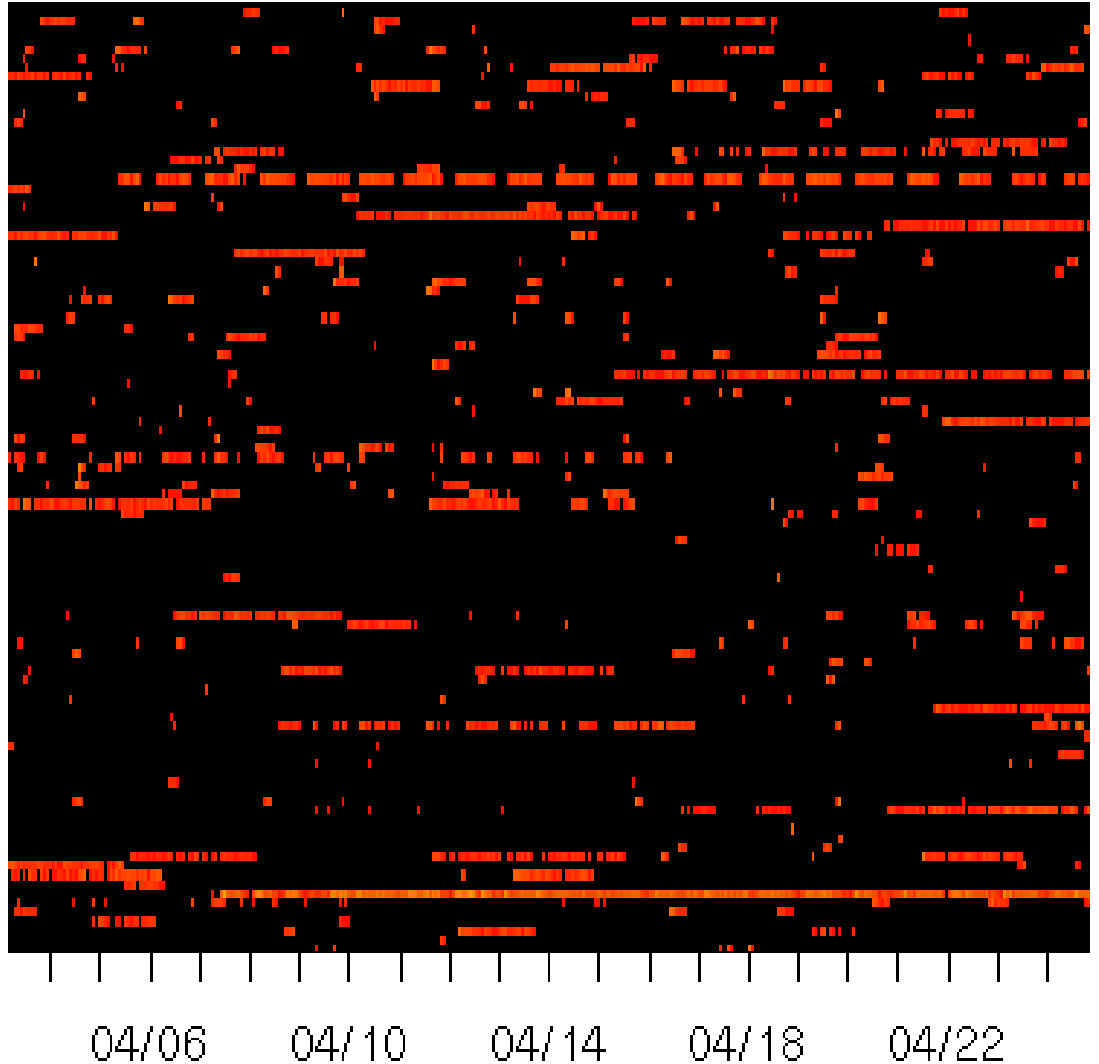


Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks




Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks



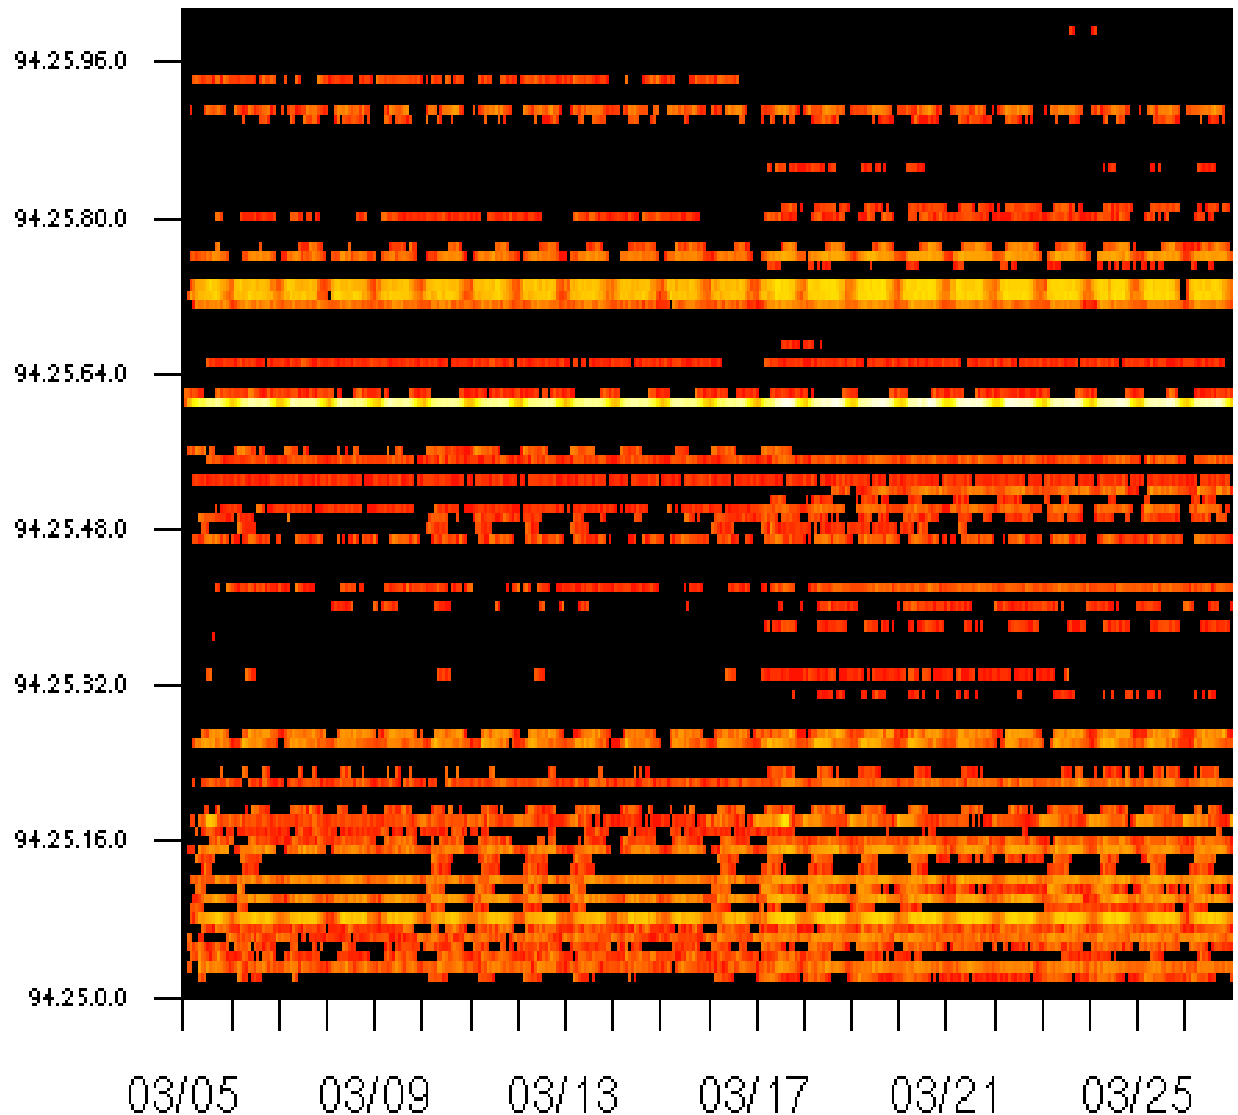
Looking at Conficker-C

Network telescope into infections:

Conficker-C hosts scan the internet randomly

Flow rates vary among IP addresses, /24 blocks

- Some seen for <24 hours, some for every hour
- Some average 2 to 3 pings/hour, some 1000s



Compiling Top-N lists

Data: Top 1000 talkers from March 3 through April 24, 2009

- by day (53 days) and by hour (1272 hours)
- by IP Address and by /24 net block
- Look at blocks appearing in hourly top 20 IP blacklist, /24 blacklist, or both



Supplementary data: Flows from all /24s seen performing Con-C scans over the 2 month period.

- 1,091,013 blocks.

Summary Information by /24 net block

- TotalIP: Total number of active IP Addresses seen
- Nonzero: Total number of hours observed scanning
- MaxIP: Maximum number of simultaneous IPs per hour
- MaxFL: Maximum number of flows seen per hour
- Mean0FL: Mean number of flows per active hour
- TalkRate: \sim Total volume sent (=Mean0FL * Nonzero).
- Country Code

Top 20 Lists by Net Block




Infected Machines

On average:

approx. 4 flows/hour from
one active infected host

Divide Mean Flows by 4 to
get approximate mean
hosts per hour

(I will affectionately deem
this calculation “cowboy
statistics”)



Who are the “Plague Dogs”?

Persistently seen in both IP and /24 blacklists:

IP Address	#Top20 IP	#Top20 /24	Avg Hosts/HR	Notes
128.125.179.129	1174	1172	591.8	USC, CA children's hospital
75.141.184.3	1153	972	319.5	Charter Communications, FT Worth TX
94.25.61.240-247 (*)	217	916	485.0	JSC Rostelecom Client Ulyanovsk, Russia (/17)
206.113.142.245	1220	582	245.5	MCI Communications Services, Inc. Ashburn, VA
65.122.8.1	1136	556	243.1	Roosevelt School District
206.160.168.34	1147	215	162.7	Sprint, Reston VA
216.115.160.40	1026	212	159.8	Unibase, Utah, US

(*) NAT from the Russian Class B we saw earlier

Other net blocks are US-based, sparse among neighbors.

Sparse Activity

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
75.141.129.0	6	28	1	5	2.25	63.00	us
75.141.130.0	2	22	1	23	3.59	79.00	us
75.141.131.0	2	11	1	11	3.09	34.00	us
75.141.132.0	2	4	1	3	2.00	8.00	us
75.141.133.0	1	8	1	14	5.63	45.00	us
75.141.134.0	1	4	1	6	2.50	10.00	us
75.141.135.0	1	3	1	3	1.67	5.00	us
75.141.136.0	2	5	1	4	2.00	10.00	us
75.141.137.0	2	4	1	5	2.25	9.00	us
75.141.138.0	1	2	1	4	2.50	5.00	us
75.141.139.0	1	27	1	6	2.44	65.99	us
75.141.140.0	1	3	1	5	2.33	7.00	us
75.141.152.0	1	8	1	5	1.88	15.00	us
75.141.184.0	3	1222	2	3283	1278.18	1808926.19	us
75.141.187.0	2	9	1	4	2.11	19.00	us

Characteristics (1 phrase or less): Big NATs in small to mid-sized allocations.

Back to Russia


By eye, this /16 has a lot of activity going on

But the top ranks (barring the plague dog block 94.25.61.0):

IP by Hour: 128
/24 by Hour: 103
IP by Day : 361
/24 by Day: 343

Avg Host/HR = 942

Imagine if someone were trying to hide in this network!



Plague Dogs: IP Blacklist

Most visible blocks seen in IP blacklist only

IP Address	#Active in /24	Highest Rank				Avg Host/HR
		Daily IP	Daily /24	Hourly IP	Hourly /24	
77.93.38.7	10	7	43	2	22	129.2
193.239.178.194	13	13	60	7	30	106.2
217.118.82.1	3	14	76	7	30	87.6
195.54.3.58	20	9	61	5	27	87.4
84.22.140.186	2	12	86	6	23	85.6
194.187.148.40	7	11	89	5	44	71.2

- Russian and Ukrainian addresses
- Gaming networks, ISPs
- Allocations /24 through /22

The big NATs are already caught in both IP and /24 blacklists

Plague Dogs: /24 Blacklist

Most visible blocks seen in /24 blacklist only

/24 Net block	#Active in /24	Highest Rank				Avg Host/HR
		Daily IP	Daily /24	Hourly IP	Hourly /24	
195.46.34.0	87	219	1	49	1	886.7
168.8.212.0 (*)	97	41	1	13	1	662.1
125.60.241.0	139	355	3	34	1	423.5
83.234.227.0	16	190	6	64	2	300.5
77.120.128.0	256	1000+	6	198	3	291.4
77.120.129.0	256	1000+	5	612	2	280.6


(*) Showed up in IP list top 20 for 2 out of 1272 hours

- Ukraine, Russia, Philippines, US
- ISPs, Telecom, and the Georgia Board of Education
- More “bang for the buck” than IP lists

East Europe 77.93.x.x

TopIPBlacklist: FL/hr

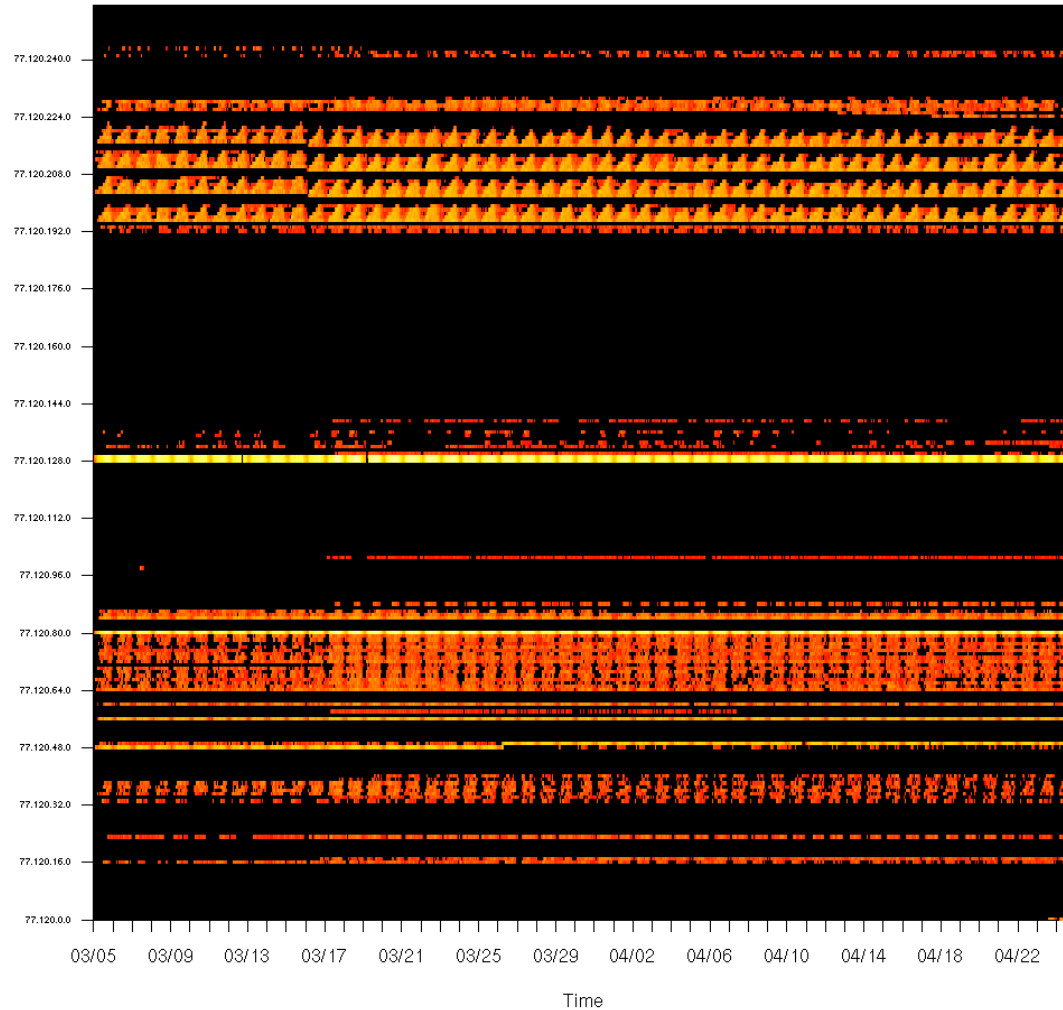
CC
it: 10
ru: 6
ua: 5
cz: 2
lv: 2
ro: 2



Ukraine Datasvit 77.120.x.x

Top24Blacklist: FL/hr

CC
ua: 83



Digging Deeper

“Spread” in summary statistics

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
41.221.16.0	254	1222	243	2036	373.60	402491.16	dz
41.221.17.0	254	1222	242	2051	392.71	421225.87	dz
41.221.18.0	254	1221	229	1889	325.68	354943.85	dz
41.221.19.0	254	1222	224	1550	288.70	318780.14	dz
41.221.20.0	15	1104	12	30	5.92	6311.66	dz
41.221.22.0	1	43	1	31	5.74	246.99	dz
41.221.23.0	5	935	4	85	15.30	12083.11	dz
41.221.24.0	7	1142	6	80	9.19	9745.35	dz
41.221.25.0	1	8	1	13	5.63	45.00	dz
41.221.26.0	254	1222	245	2393	395.43	420475.77	dz
41.221.27.0	128	1207	124	1580	280.46	295347.39	dz
41.221.28.0	2	203	2	14	3.74	757.12	dz
41.221.29.0	2	210	2	17	3.93	828.82	dz

Circled blocks


- ranked between 9-20 in /24 Blacklist
- 70-100 Hosts/HR, each, avg 500 hosts/hr
- none in Top 1000 IPs

Digging Deeper

Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
41.221.160.0	22	491	3	57	7.08	3334.33	ng
41.221.161.0	71	704	10	123	14.76	9832.43	ng
41.221.165.0	41	810	8	88	14.16	8973.12	ng
41.221.166.0	30	340	2	51	6.45	2052.59	ng
41.221.167.0	126	1015	13	90	14.67	14372.03	ng
41.221.168.0	18	387	3	24	6.85	2705.39	ng
41.221.169.0	59	776	8	90	15.84	12033.09	ng
41.221.171.0	28	857	8	70	13.73	10495.70	ng
41.221.172.0	43	243	7	42	7.80	1982.30	ng
41.221.173.0	3	743	2	37	2.67	1936.29	ng
41.221.174.0	16	518	5	40	8.43	4326.92	ng
41.221.175.0	68	437	11	104	13.80	5553.46	ng
41.221.200.0	63	387	3	35	7.24	2798.75	cv
41.221.201.0	59	397	4	27	6.15	2486.93	cv
41.221.202.0	54	346	5	27	6.02	2078.70	cv
41.221.203.0	79	387	5	32	7.70	2927.52	cv
41.221.204.0	76	384	4	69	8.64	3163.38	cv
41.221.205.0	65	384	5	46	6.39	2411.20	cv
41.221.206.0	53	239	4	30	6.40	1453.10	cv
41.221.207.0	41	193	3	23	6.15	1178.31	cv

Europe/Africa 41.221.x.x

AlgerianISP: FL/hr



Digging Deeper

“Spread” in summary statistics

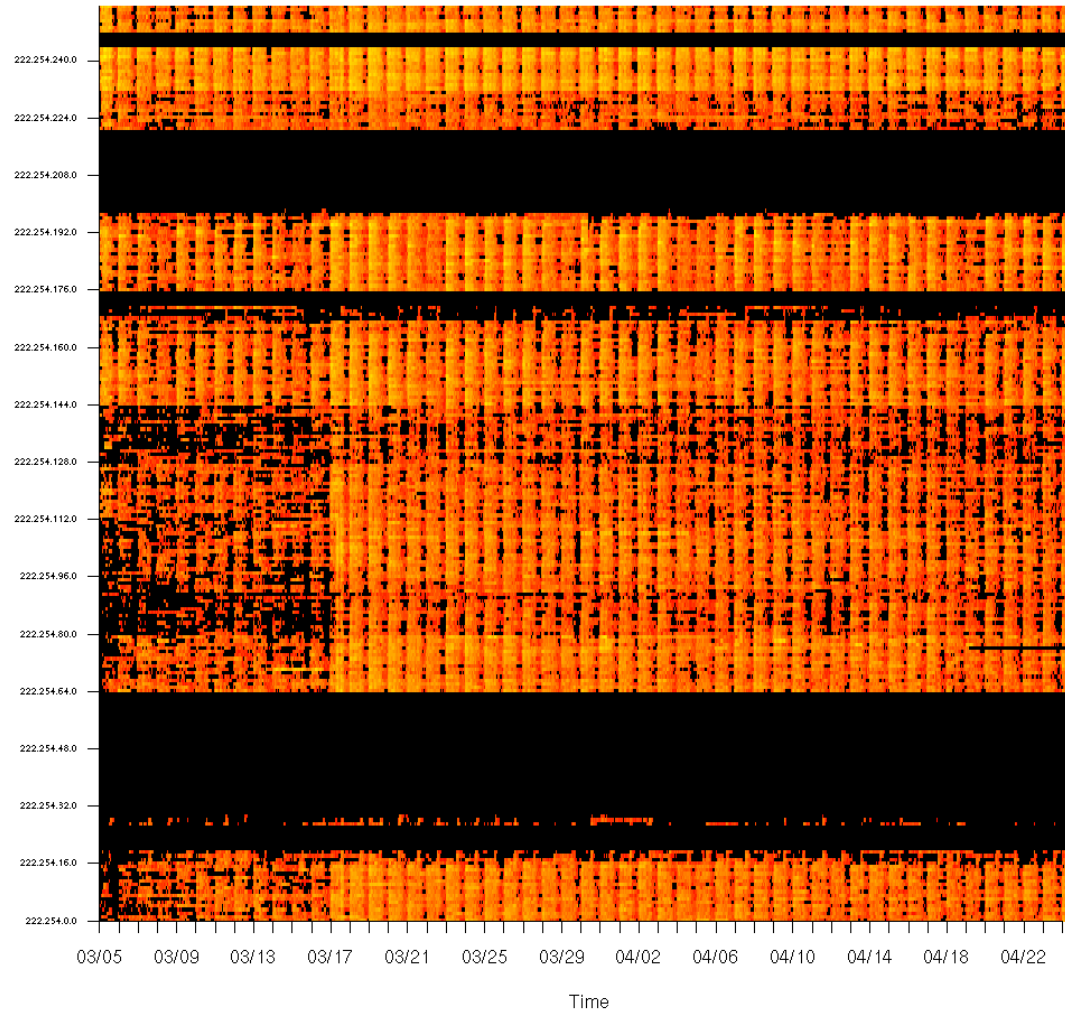
Slash24	TotalIP	Nonzero	MaxIP	MaxFL	Mean0FL	TalkRate	CC
222.254.180.0	229	1132	19	497	35.68	34702.36	vn
222.254.181.0	225	1149	18	550	30.91	31987.85	vn
222.254.185.0	230	1149	16	458	34.88	35522.45	vn
222.254.188.0	246	1108	24	508	44.17	40827.85	vn
222.254.189.0	244	1078	20	549	45.17	39800.97	vn
222.254.190.0	240	1120	19	747	33.12	30691.78	vn
222.254.191.0	243	1153	22	535	48.32	44407.26	vn
222.254.192.0	238	1078	16	477	34.85	34866.95	vn
222.254.194.0	230	1146	16	597	47.48	50603.01	vn
222.254.195.0	232	1045	19	423	44.75	41512.70	vn

All blocks in this grid

- IP address appeared briefly in at least Top 20 hourly IP rank
- 7-12 Hosts/Hr
- none in Top 20 /24 Blacklist


Viet Nam 222.254.x.x (3 Telecoms)

VietNamISP: FL/hr




Smoothing Across Net Blocks

- Highlights contiguous blocks with similar behavior
- Use variable bandwidth for multiple views
- Max # IPs; Bandwidth=75




UKR Telecom 94.179.x.x

UKtelecom2: FL/hr




UKR telecom 94.178.x.x

UKtelecom: FL/hr




UKR telecom 94.178.x.x

?????

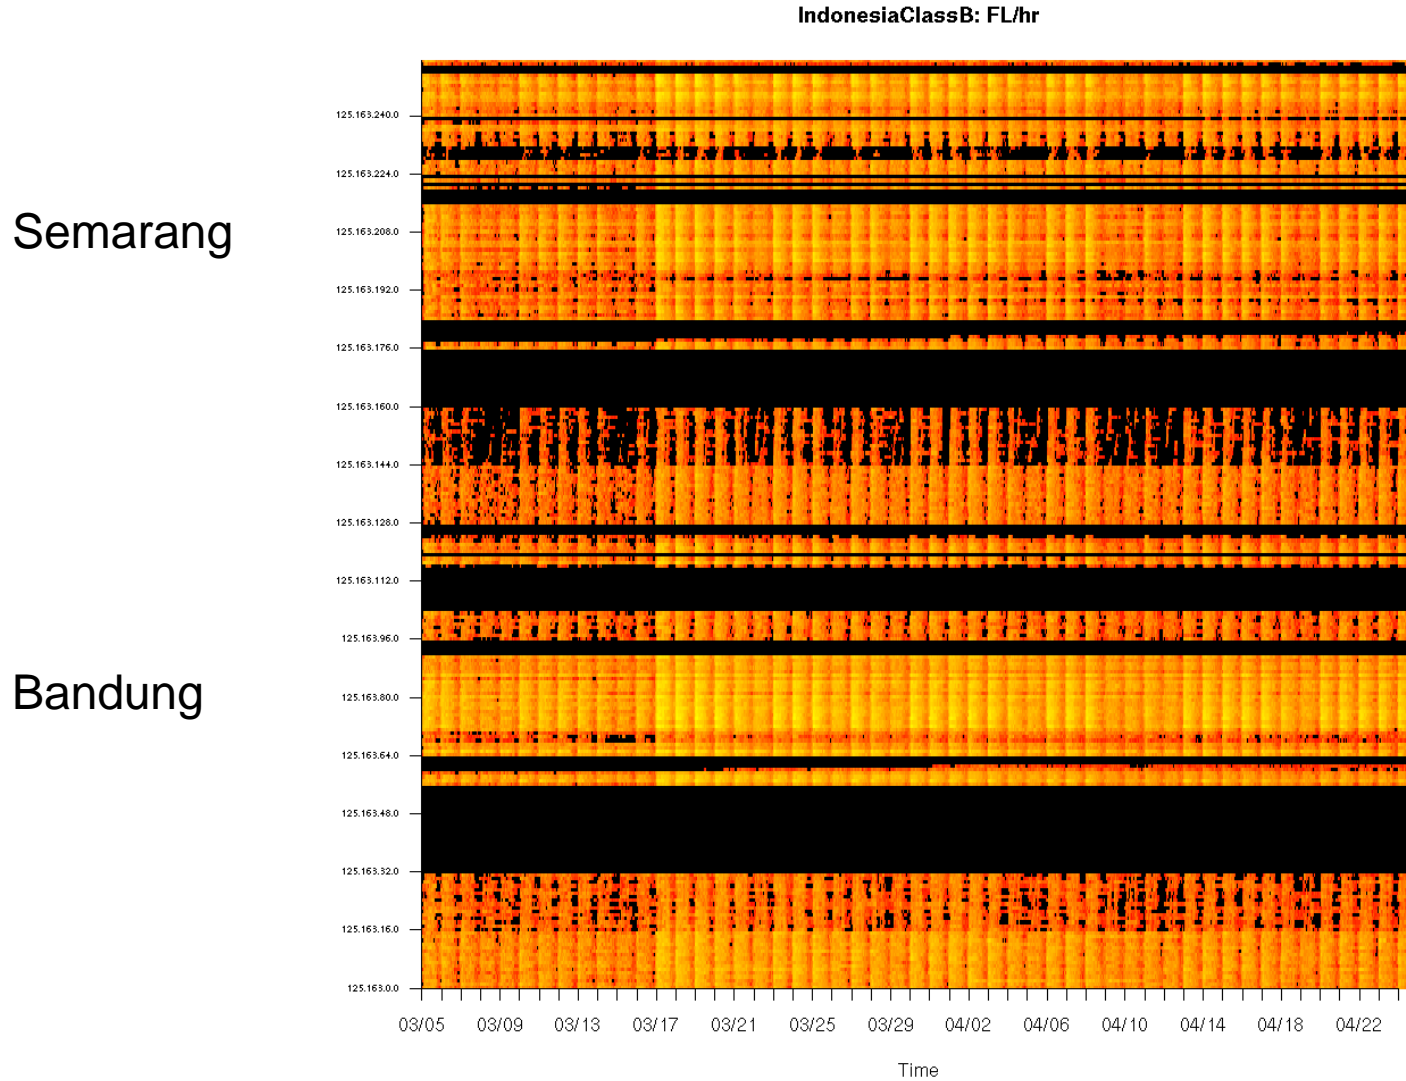


Viet Nam 118.71.x.x (1 Telecom)

VietNamClassB: FL/hr



Indonesia 125.163.x.x (2 cities)



Who do we see with Top 20?

Class B (/16)	#Active /24s	#Top 20 IP	#Top 20 /24	Highest Rank				Avg Host/HR
				Daily IP	Daily /24	Hourly IP	Hourly /24	
118.71.0.0	215	8	103	163	4	5	1	2294.2
125.163.0.0	183	2	1	683	270	20	18	2136.6
222.254.0.0	185	29	1	174	172	8	19	1049.9
94.179.0.0	236	0	0	609	669	129	194	760.0
94.178.0.0	256	0	0	693	1000+	216	345	500.4

Conficker Attribution

Who is behind Conficker?

- Conficker A would shut itself down if it detected a Ukranian keyboard setup
- Two IPs were able to interact with both Conficker.B and Conficker.A hosts
 - 200.68.xxx.xx Alternativagratis.com – Argentina
 - 81.23.xxx.xxx Kyivstar.net - Kiev, Ukraine
- Rogue AV Product source is Baka Software (Kiev, UK)
- Two Kiev based ISPs with large netblocks run under the radar
 - Con-C bootstraps a peer list, so it is in the interest of the controllers to have peers available

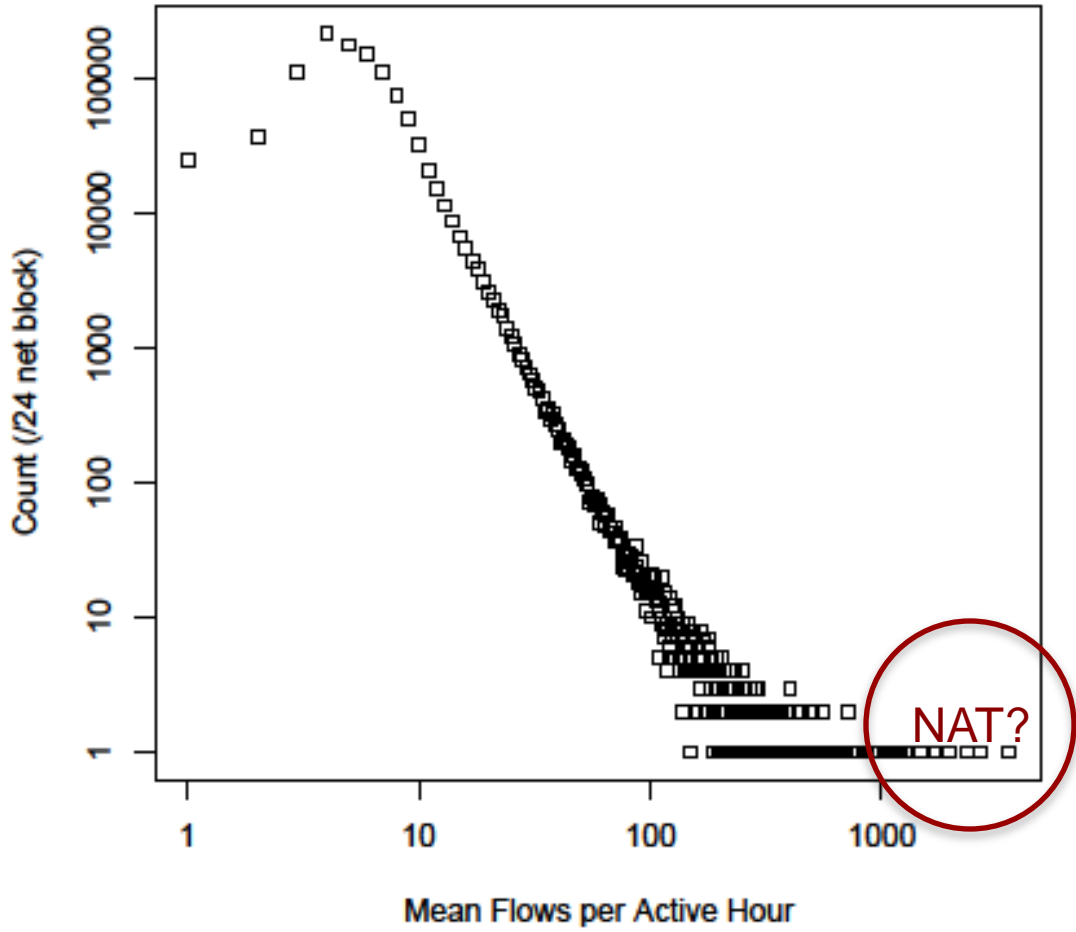
Take-home points/discussion

- Top-N has to go pretty far down to find all the interesting stuff
- Let correlation in summary statistics help us find the big blocks
- Big blocks seem a bit more international (network conventions?)
- Are evil-doers really trying to hide (zippy “bullet-proof” networks), or is it just large scale DHCP?
- Telecom ISPs have abuse contacts, but how useful are they?

Thank You!

Extra slides

Measurement



Measurement

