Route Filtering: Handle with Care

Thursday, 26 August 2010

Frank Salanitri, APNIC

Tomoya Yoshida, NTT Communitations

Overview

- Background
- The problem
- APNIC Resource Quality Assurance
- BGP debogon project

Why are IP Addresses Blocked?

IP address can get filtered for various reasons:

- Outdated bogon lists
- Past abusive behaviour
- Blacklist from spamming and DOS attacks
- Security/access policies

IP Filtering Methods

- Route filtering
- Application filtering, esp. Mail
- Firewall filtering

The Problem

- Legitimate Internet traffic fails to reach the destination due to outdated filters and black/bogon lists
- RIR seen as responsible for allocating 'unusable' blocks
- Situation worsens as free pool of IPv4 addresses reaches exhaustion
  - New address blocks attract unwanted levels of traffic from private-use domains, misconfigured equipment, and scanning activity.
  - Prefixes get recycled

What You Can Do

Manage bogon filtering responsibly

- To ensure that addresses are not mistakenly filtered through routers, it is important to keep router ACLs updated

Keep informed about bogon filters and IANA allocations. Visit regularly:

- Team Cymru
- IANA

Resource Quality Assurance

Community awareness campaign

- Build relationships with reputable organizations that maintain bogon/black list
- Education through publications and APNIC training materials
- Keep the Whois Database accurate
  - Actively remind resource holders to update their data

Resource Quality Assurance

APNIC acts to minimize any problems in routability through communication, training, and testing

Testing for new /8 blocks

- NOC mailing lists notification
- Reachability test conducted in conjunction with RIPE NCC
- Collaborative testing

BGP Debogon Project

BGP debogon project

Tomoya Yoshida
NTT Communications
yoshida@nttv6.jp

Your IP Address seen in the world

- My not always reach to every network
  - Even though the ISPs advertise their customers IP blocks in stable of course…
  - In case of the new IP allocation in particular

- We often encounter “(BGP) bogon filtering issue”

Bogon, Bogon Route, Bogon Filtering

- Bogon
  - Originated by the word bogus(False, Fake etc)
- Bogon Route
  - prefix which is not advertised or must not to be advertised usually
- Bogon filtering
  - Filtering Bogon Route by ISP's border GW router generally, including contents filtering at server side

Present use	Address Block
Private Address (RFC1918)	10.0.0.0/8、172.16.0.0/12、192.168.0.0/16
Loopback Address	127.0.0.0/8
Link Local Address	169.254.0.0/16
TEST-NET	192.0.2.0/24
Benchmark Test Address	198.18.0.0/15
Multicast Address	224.0.0.0/3
IANA Reserve	**Now /8 x14**

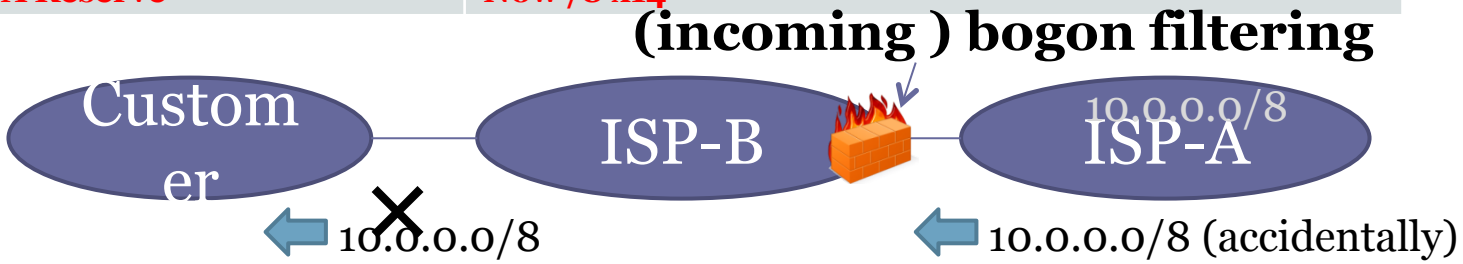

ESTA Problem

https://esta.cbp.dhs.gov/

2 years ago…

Bogon filtering issue

One week advertisement of 14/8, 223/8

- Recently whole x/8 advertisement is observed more often just after the IANA allocated to the RIRs those blocks
  - Investigation 1/8 pollution at first
  - Other x/8s are also investigated for the situations and checking the trend
- Overview of Investigation
  - Period：19th Apr 2010 ~ 26th (1 week)
    - Allocation from IANA to APNIC：10th Apr 2010
  - Prefixes：14/8, 223/8 from AS38639(NTTCom)
  - Packet collecting way：tcpdump + netFlow sampling (Samurai)
  - Reachability check for those two blocks using routeview (router server)

Per Protocol

Normally 30Mbps ~ 50Mbps, it is like a normal traffic curve

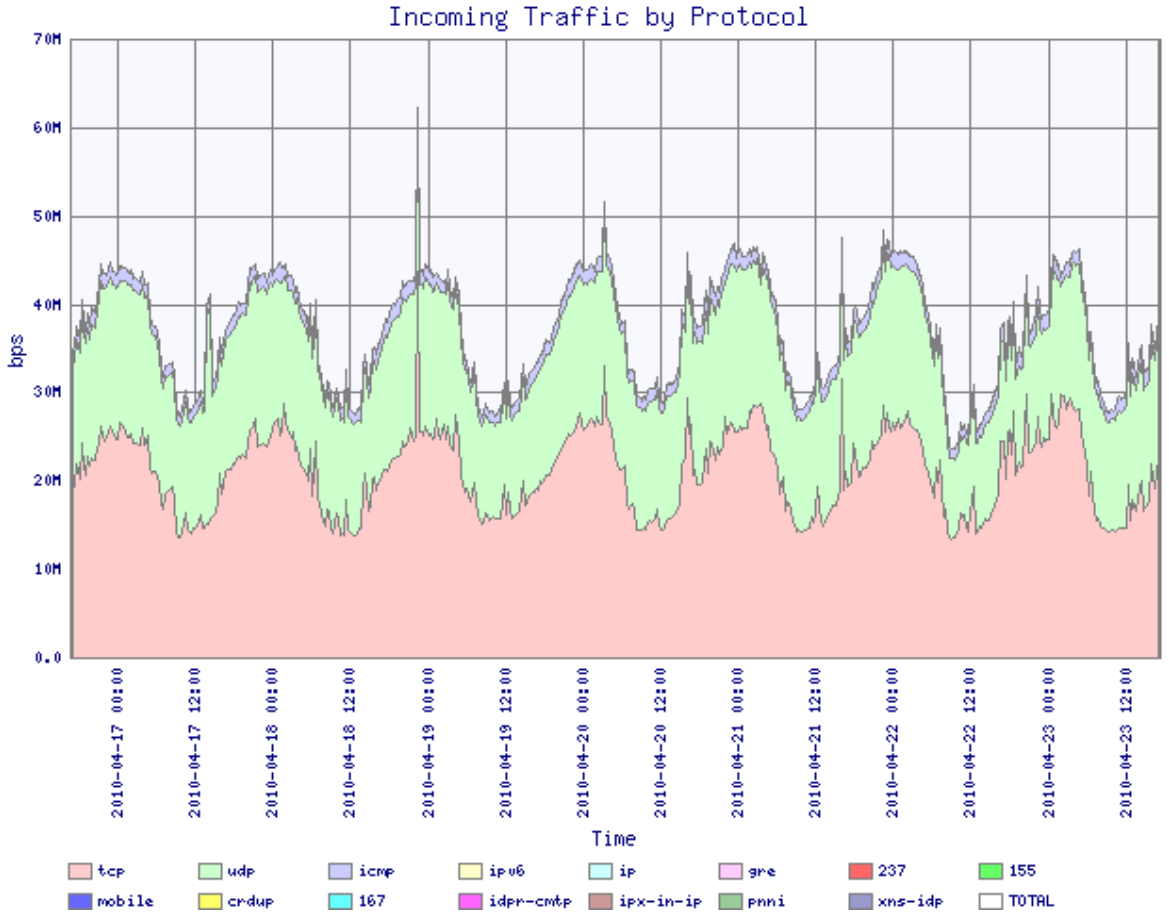

Per Protocol and Port

A half is tcp/445(Conficker , Downadup), second udp/1434(sql-slammer)

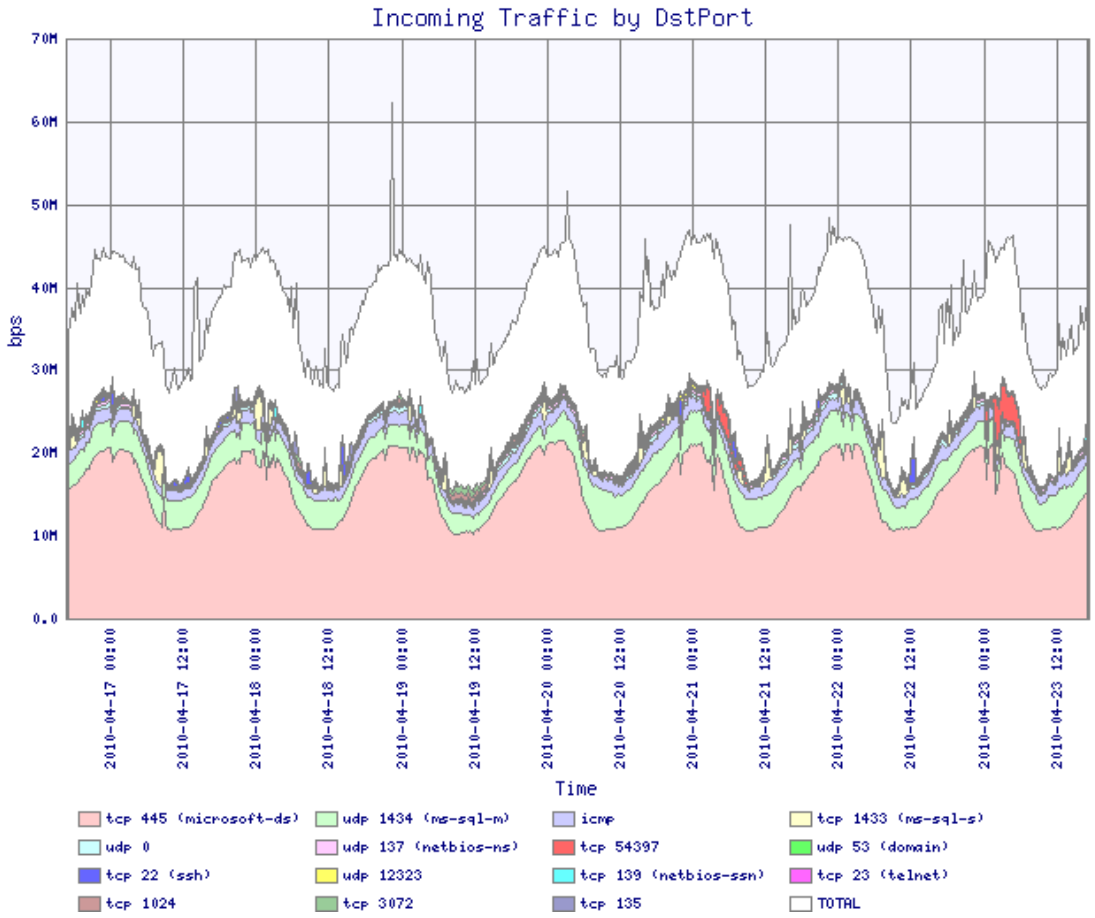

Per Origin_AS

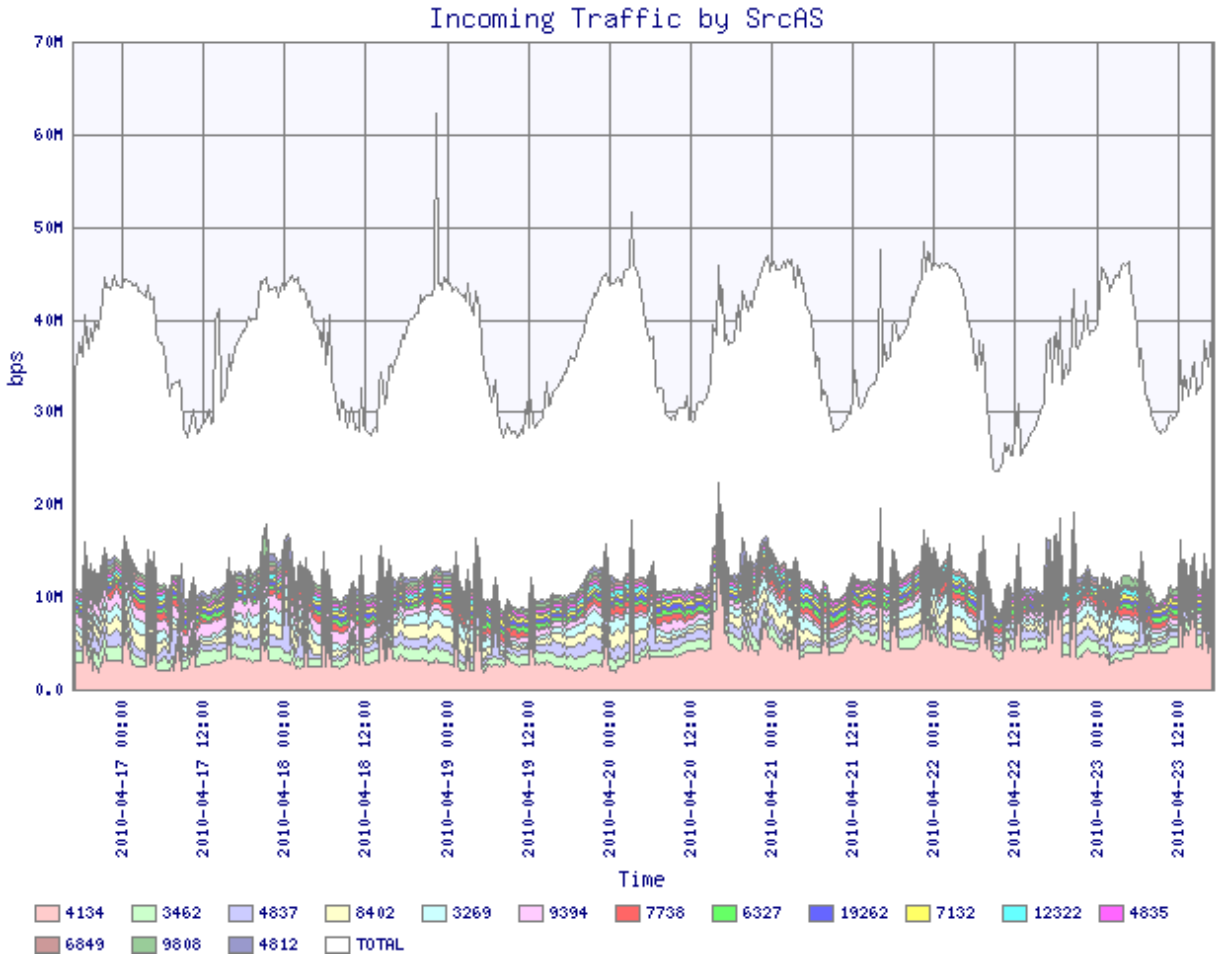

AS4134:ChinaNet
AS3462:Hinet
AS4837:CNCG
AS8402:Corbina Tel
AS3269:Telecom Italy

Per Destination IP

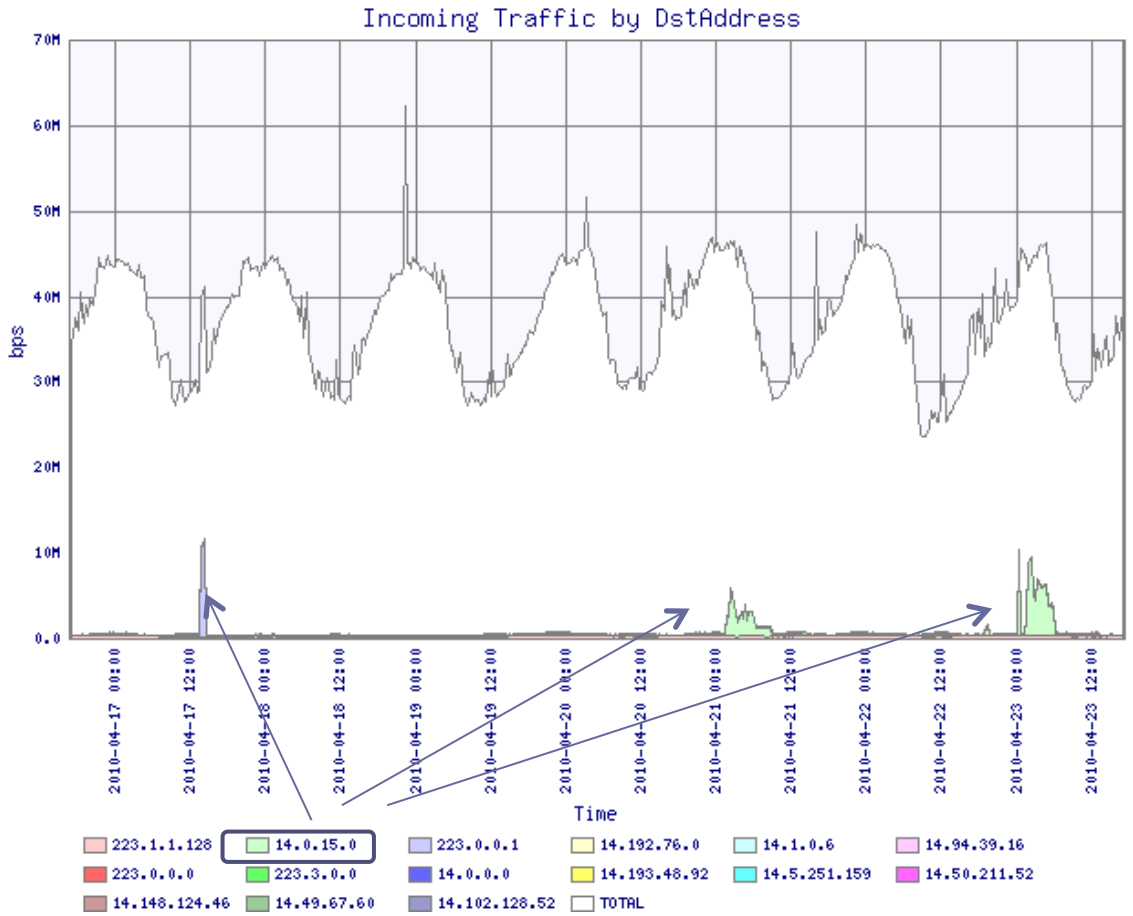

Per Source IP （bps）

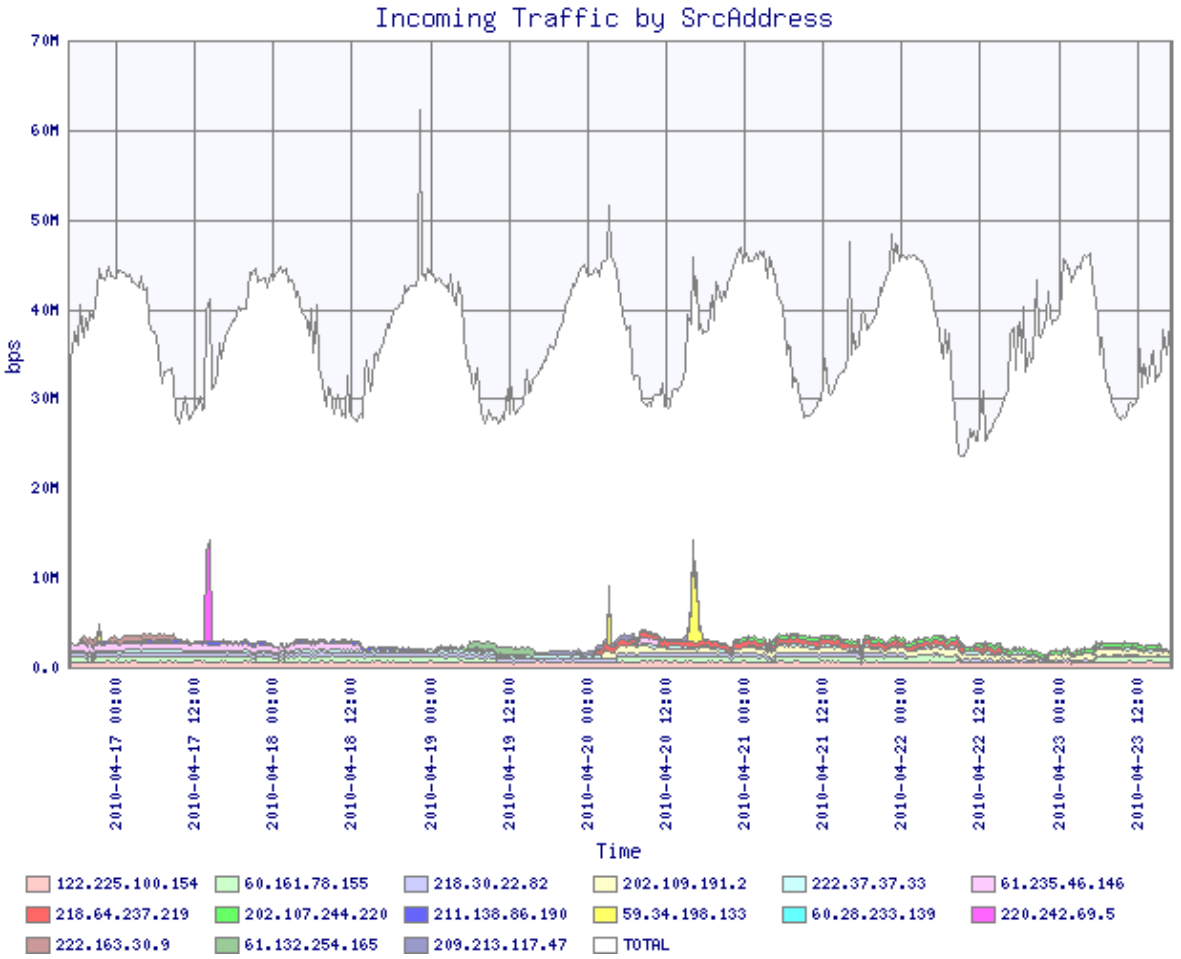

Per Source IP （pps）

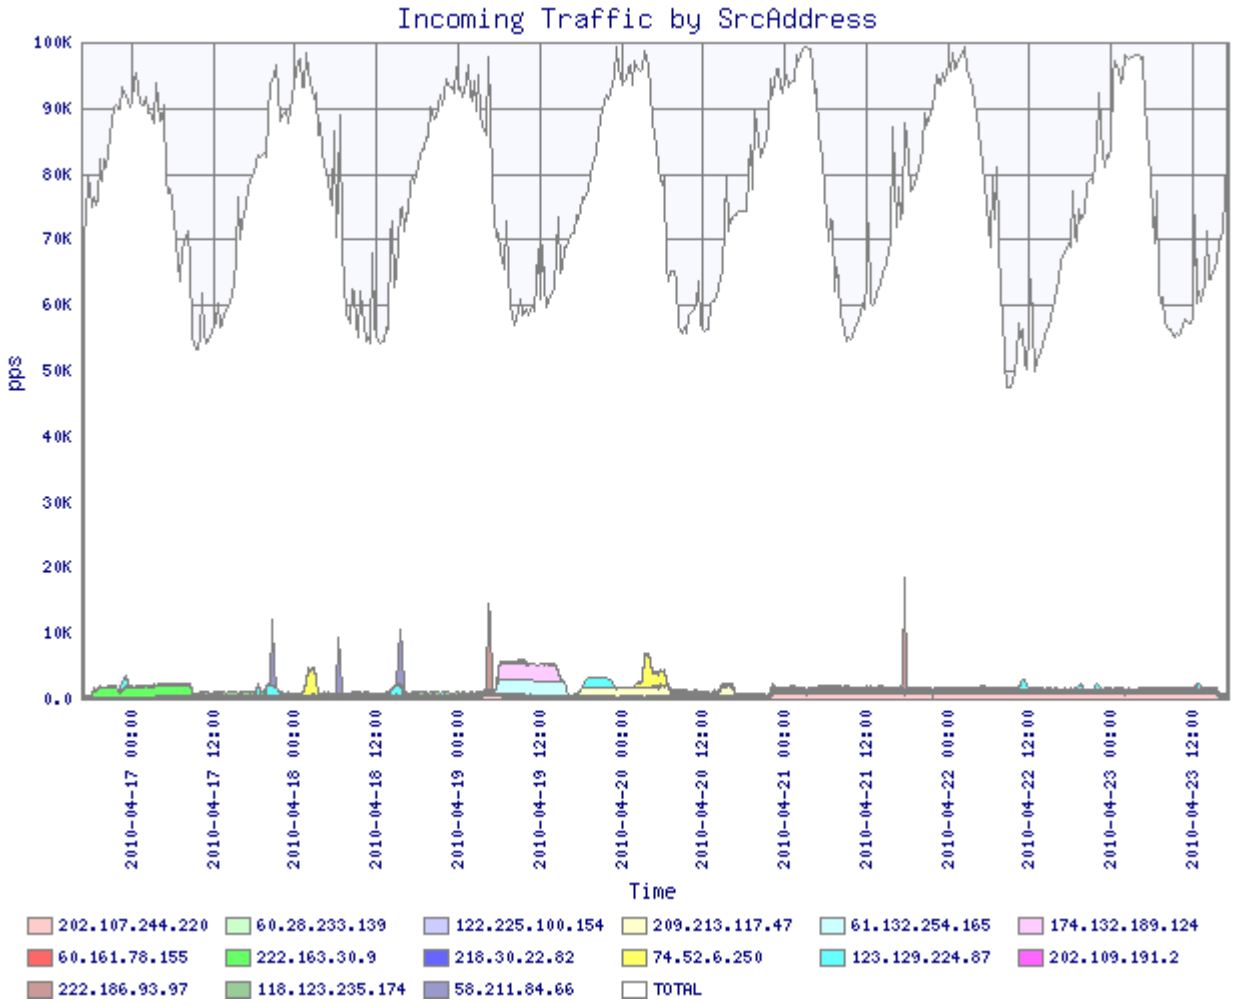

Some Specific Packet (e.g.)

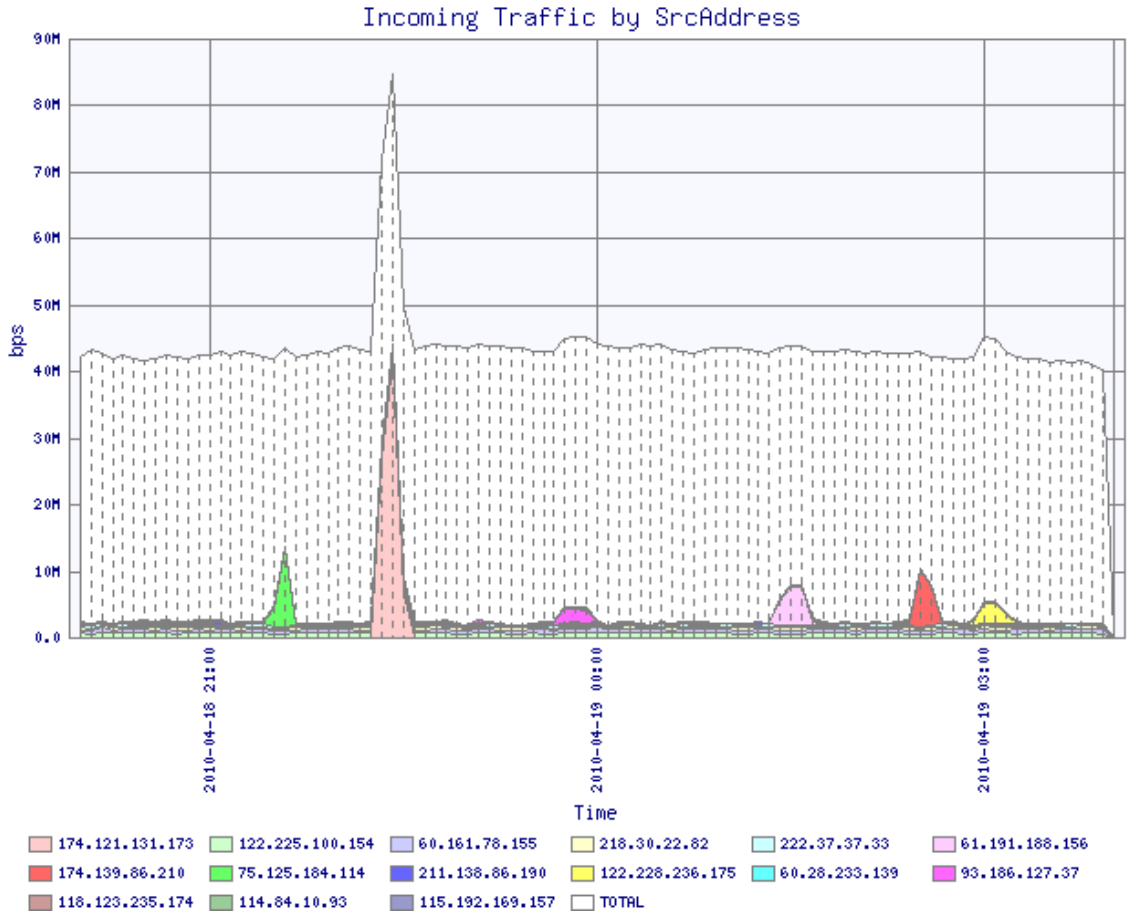

Some Specific Packet (e.g.)

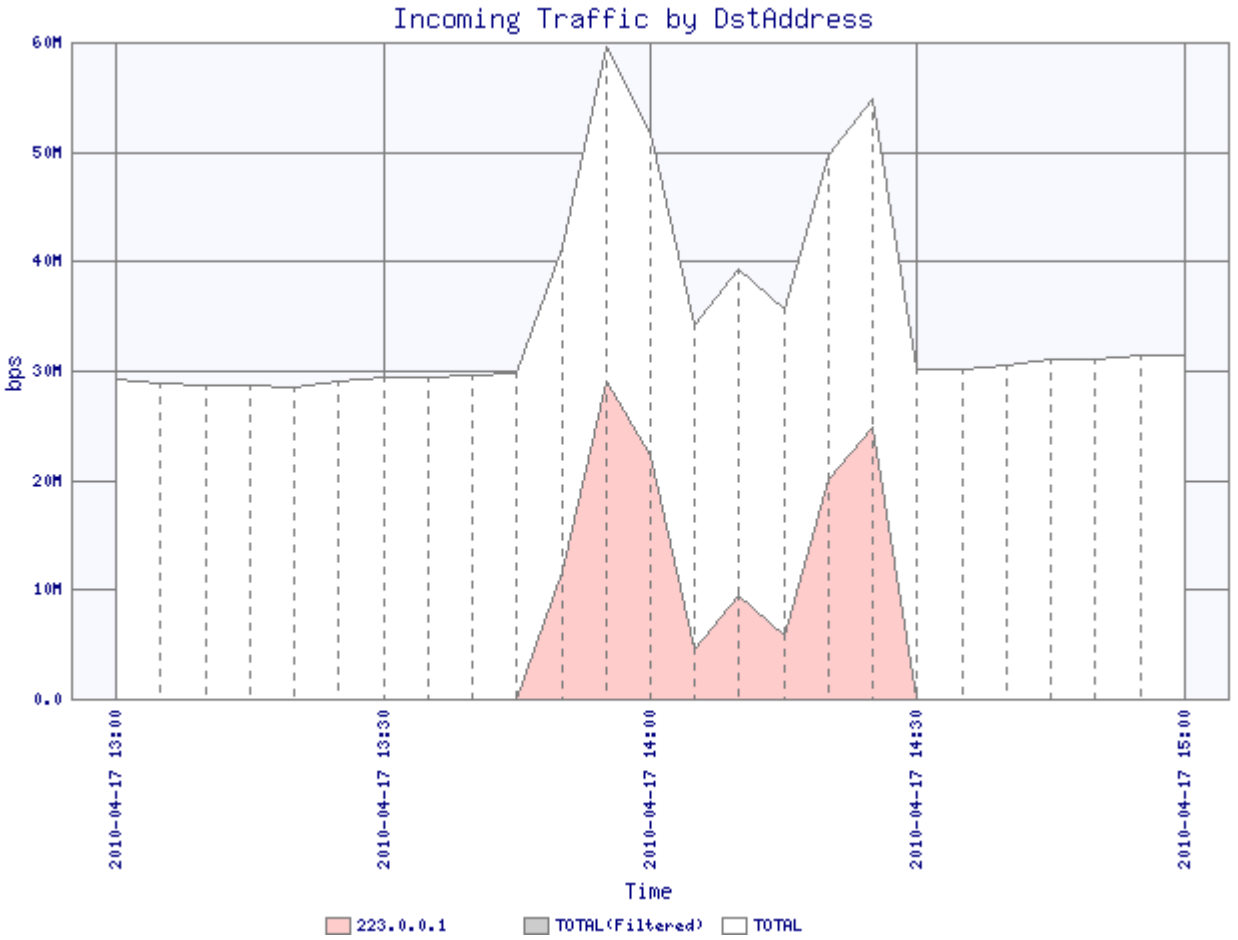

Some Specific Packet (e.g.)

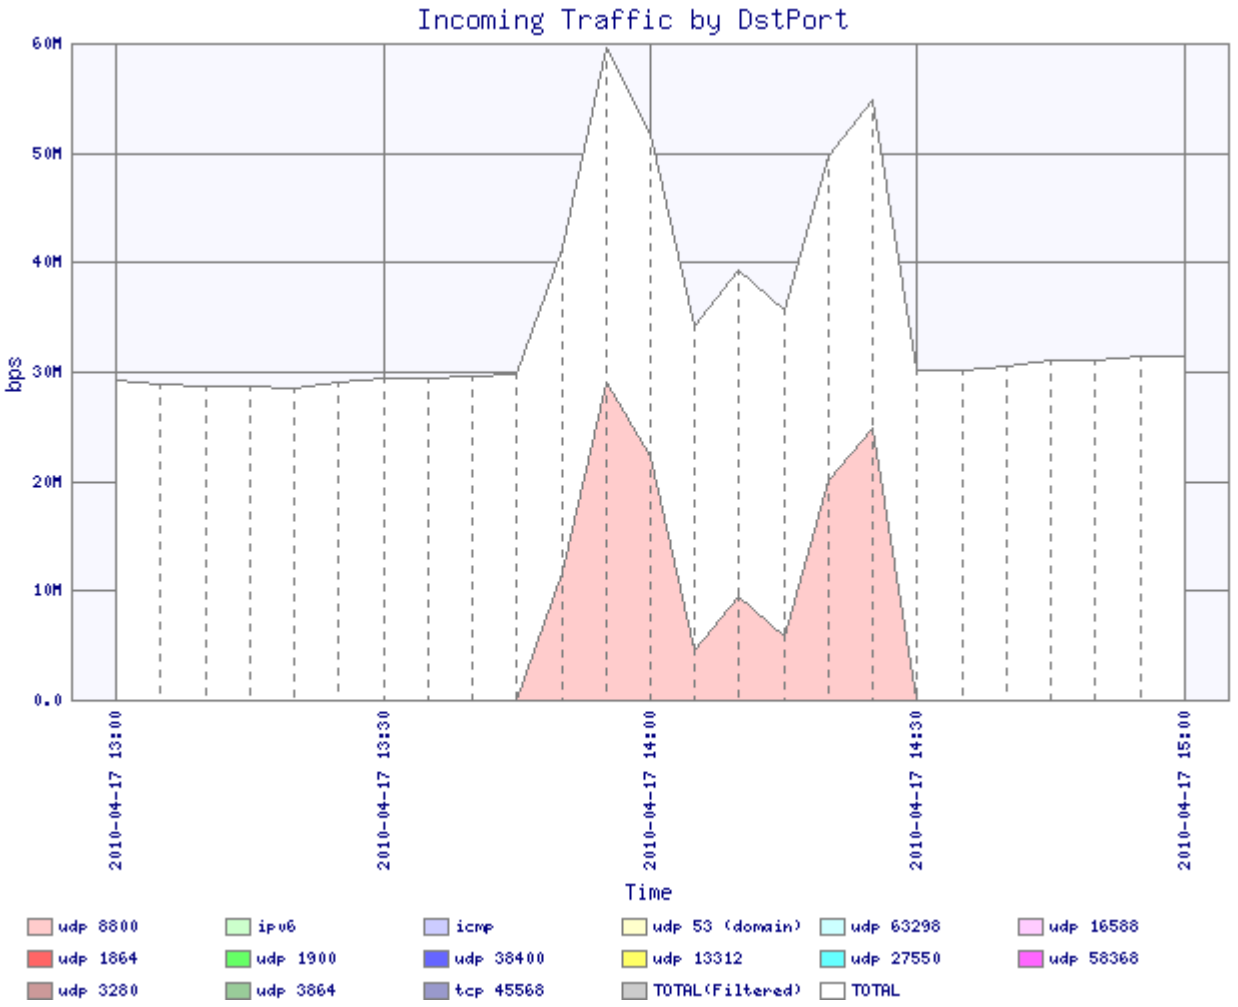

Some Specific Packet (e.g.)

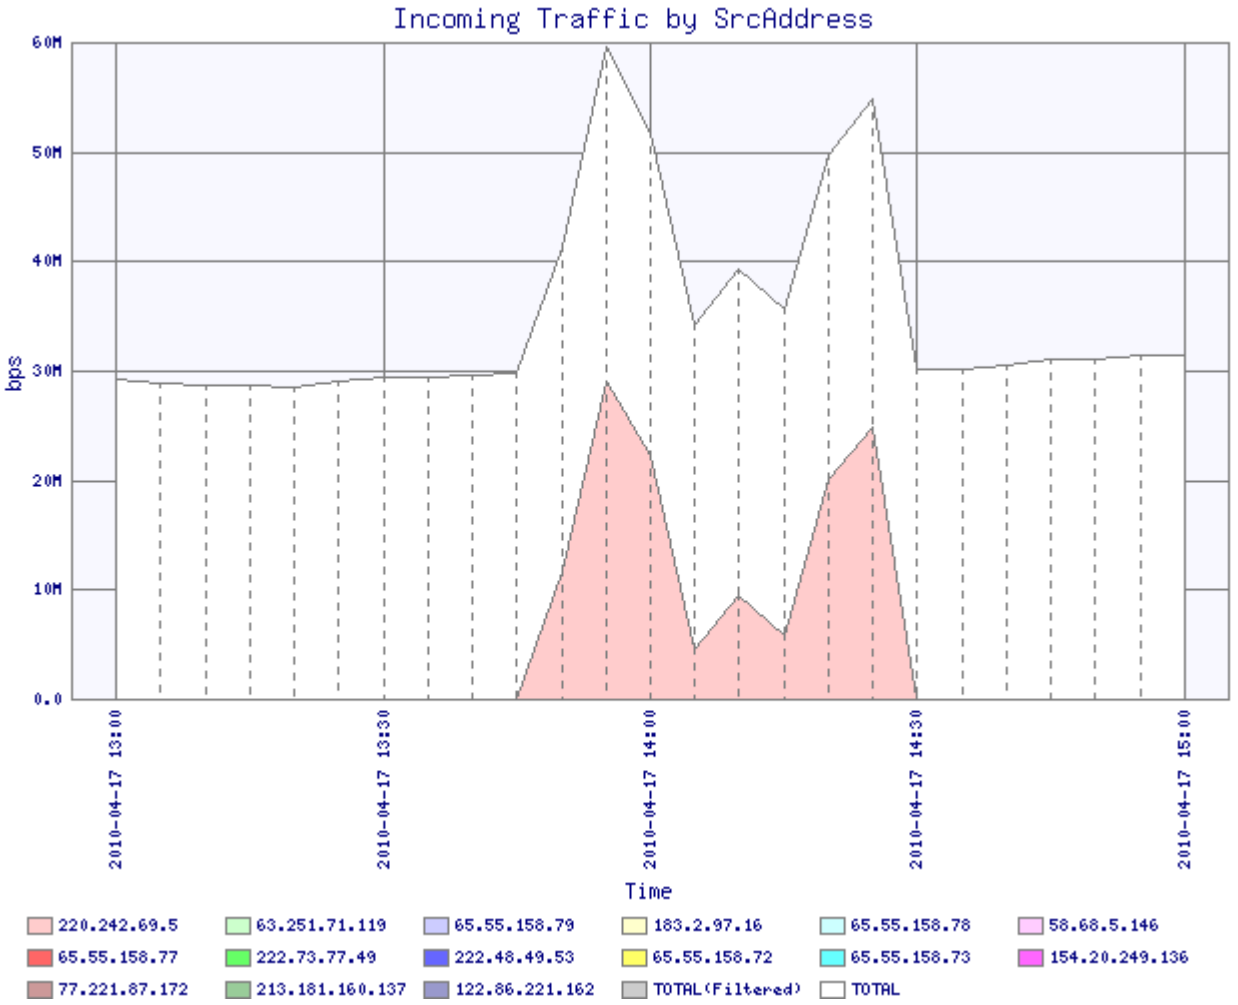

Some Specific Packet (e.g.)

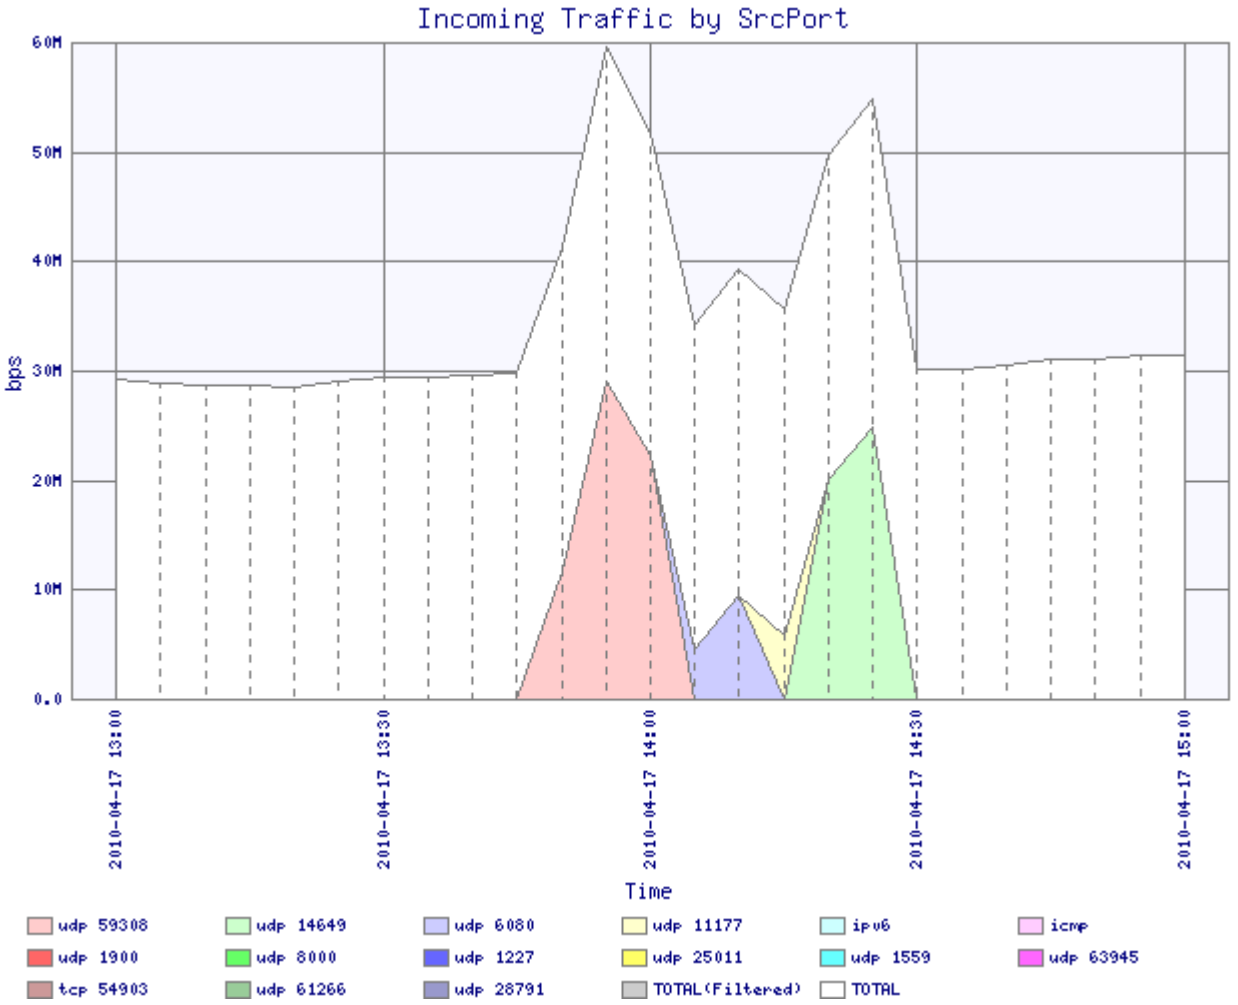

14/8 Traffic to AS38639

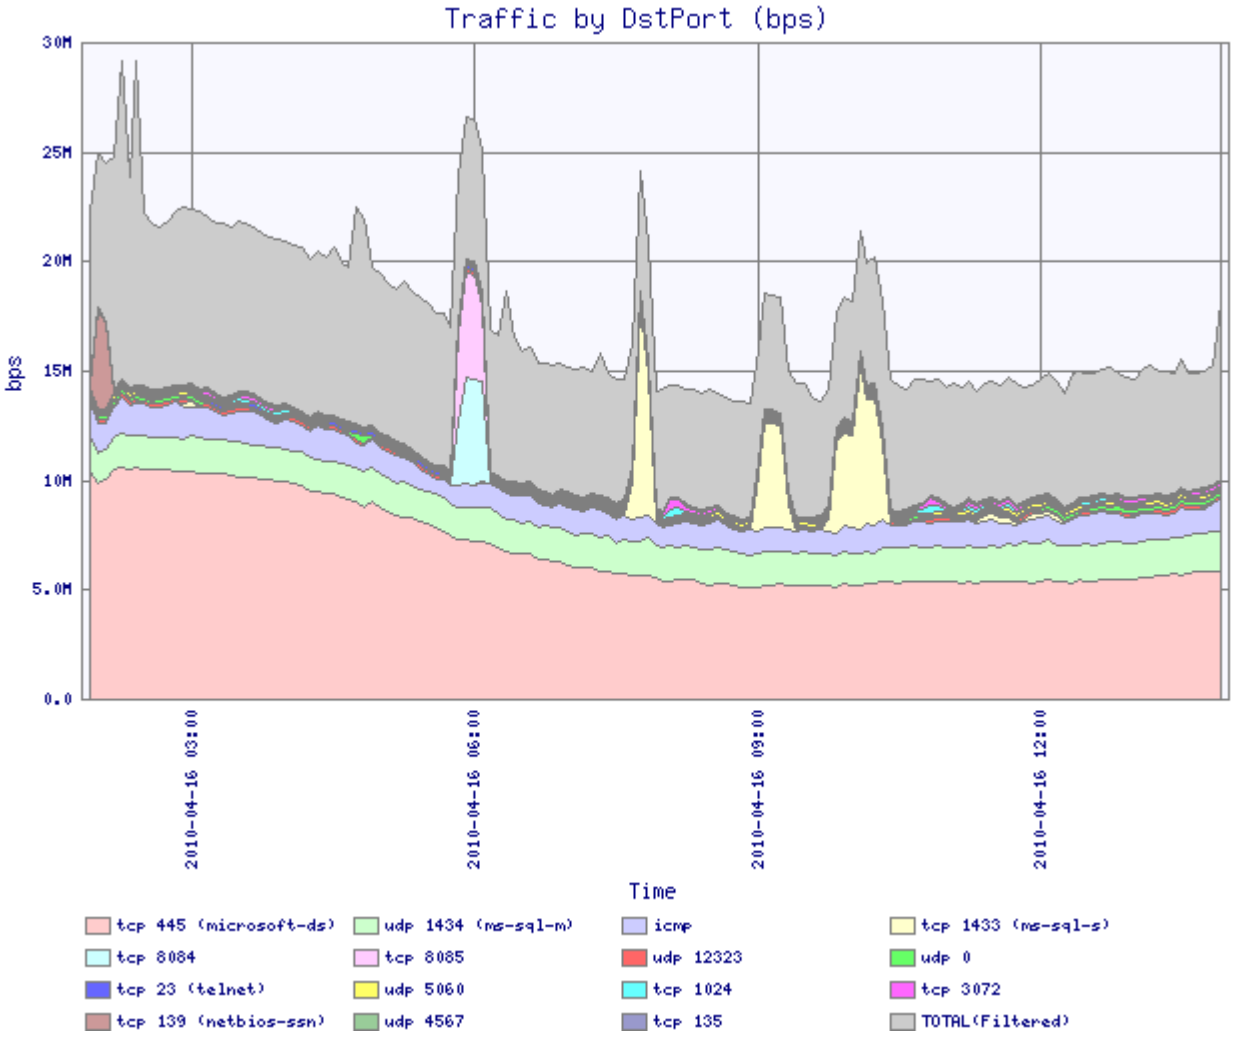

223/8 Traffic to AS38639

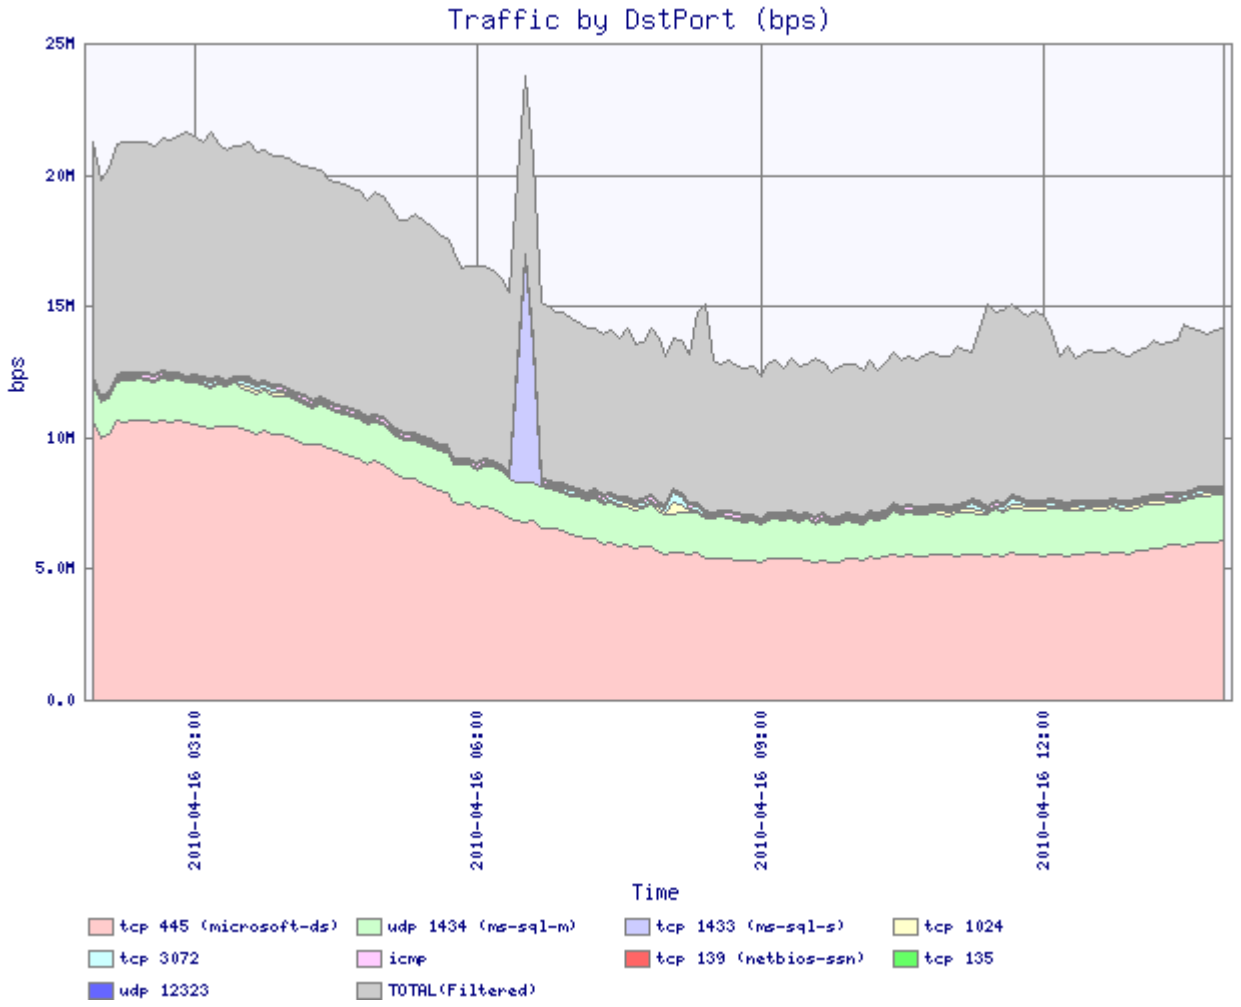

{ 27, 14, 223 } /8 reachability investigation (25th Apr, 2010)

- Approximately 20-30% are not reachable from new allocation IP immediately after new allocation from the IANA to the APNIC
- There are differences between 14/8 and 223/8 even though the same allocation timing

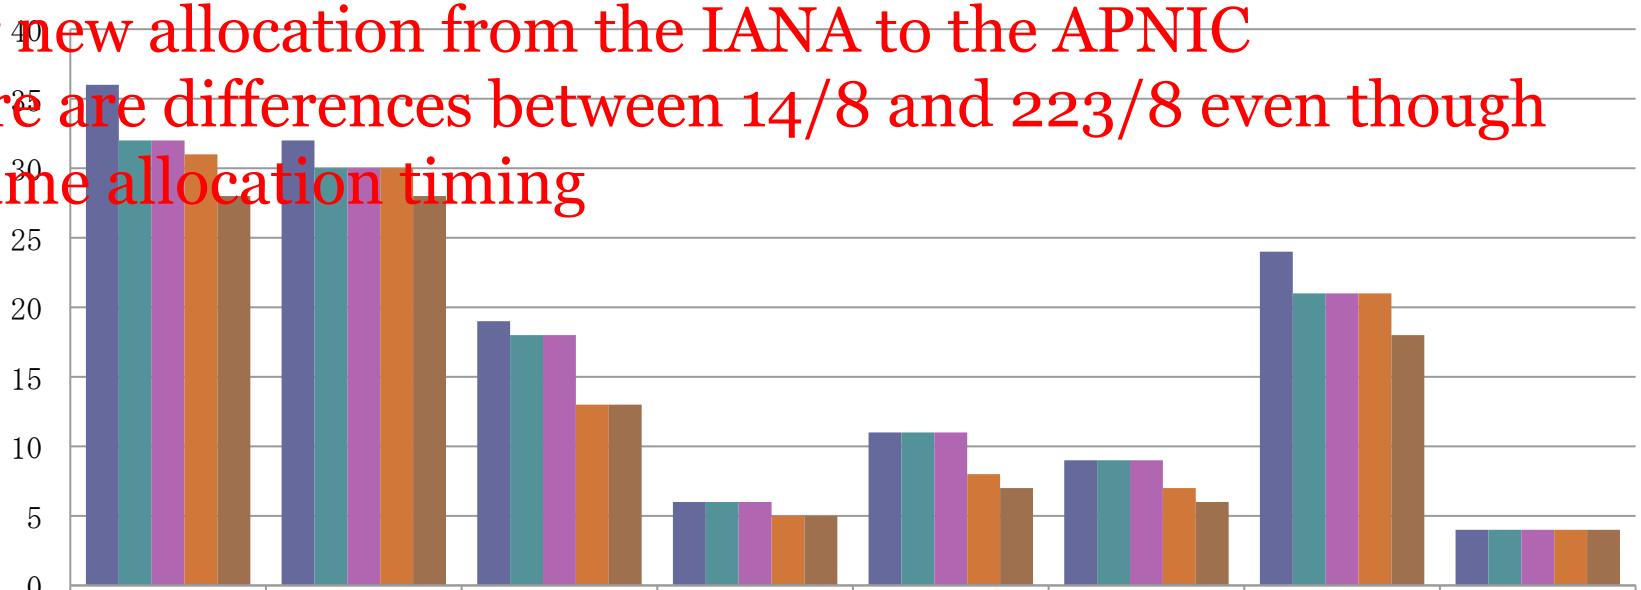

	route-views.routeviews.org	route-views2.routeviews.org	route-views3.routeviews.org	route-views4.routeviews.org	route-views.eqix.routeviews.org	route-views.isc.routeviews.org	route-views.linx.routeviews.org	route-views.wide.routeviews.org
115.69.224.0/21	36	32	19	6	11	9	24	4
27.0.1.0/24	32	30	18	6	11	9	21	4
27.50.8.0/22	32	30	18	6	11	9	21	4
14.0.0.0/8	31	30	13	5	8	7	21	4
223.0.0.0/8	28	28	13	5	7	6	18	4

IPv4 address space recently allocated by the IANA to the RIRs

- 10 /8 allocations to the RIRs,
  to APNIC, ARIN, RIPE, LACNIC (not AfriNIC)
  - 20100119 #APNIC 001/8, 027/8
  - 20100212 #ARIN 050/8, 107/8
  - 20100411 #APNIC 014/8, 223/8
  - 20100511 #RIPE 031/8, 176/8
  - 20100603 #LACNIC 177/8, 181/8
  - 20100805 #APNIC 049/8, 101/8

http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.txt

1/8 reachability investigation （7th Jun, 2010）

Approximately 10% are not reachable even though 5 months later

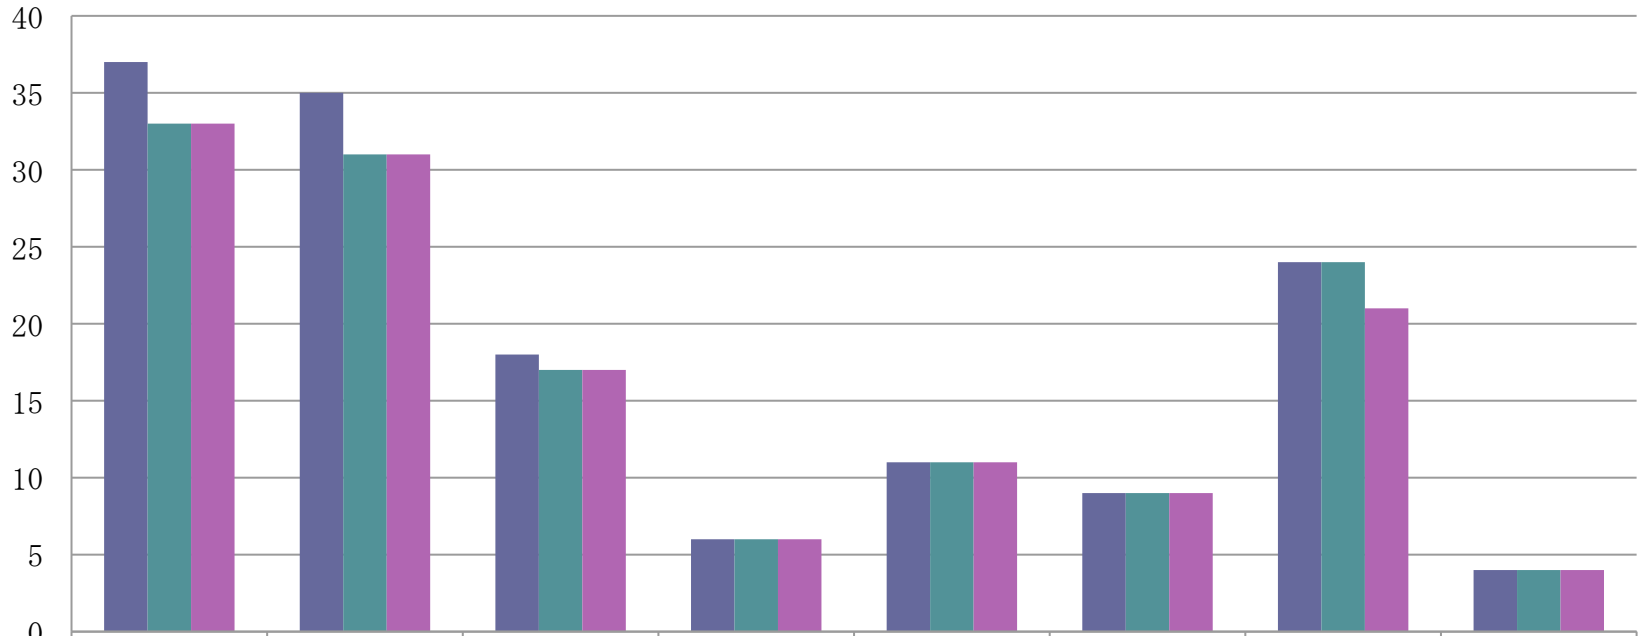

	route-views.routeviews.org	route-views2.routeviews.org	route-views3.routeviews.org	route-views4.routeviews.org	route-views.eqix.routeviews.org	route-views.isc.routeviews.org	route-views.linx.routeviews.org	route-views.wide.routeviews.org
115.69.224.0/21	37	35	18	6	11	9	24	4
1.50.8.0/24	33	31	17	6	11	9	24	4
1.200.0.0/22	33	31	17	6	11	9	21	4

1/8 reachability investigation (8th Jul, 2010)

No big changes one more month later…

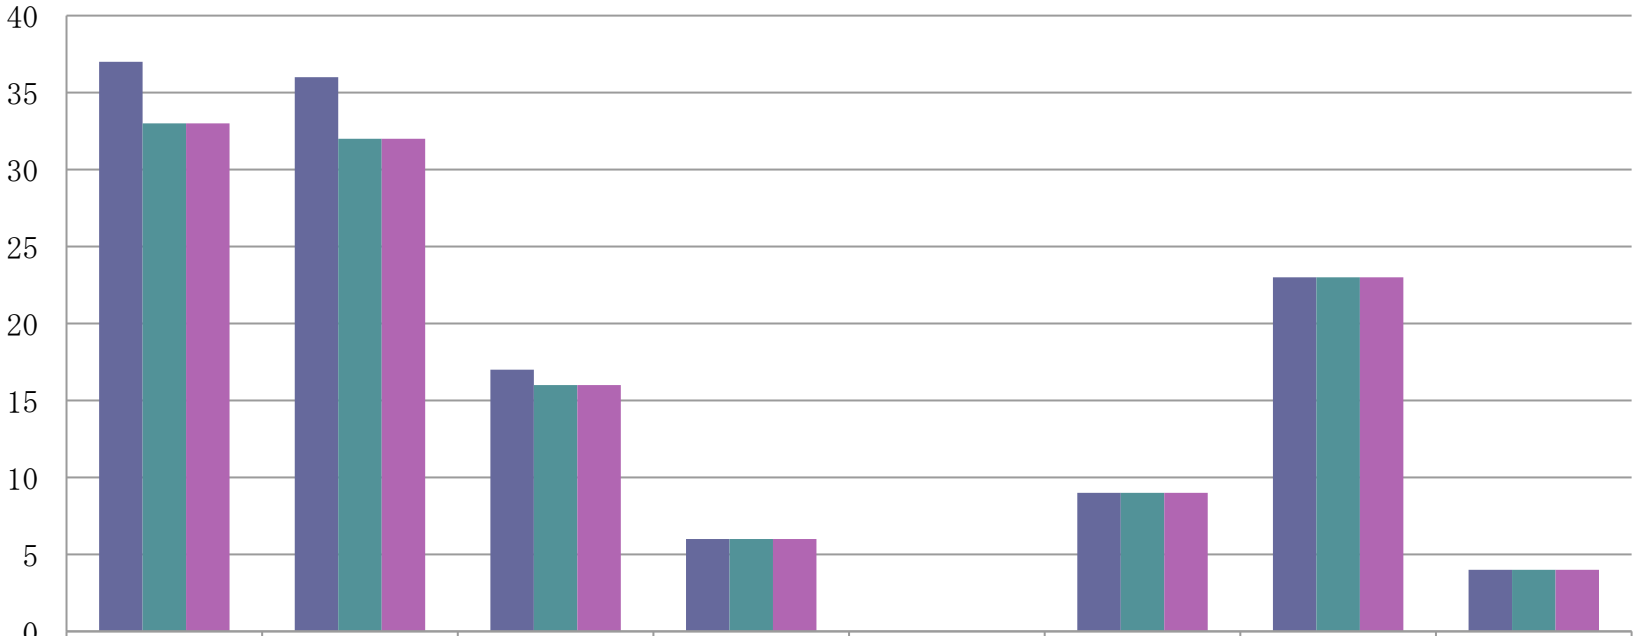

	route-views.routeviews.org	route-views2.routeviews.org	route-views3.routeviews.org	route-views4.routeviews.org	route-views.eqix.routeviews.org	route-views.isc.routeviews.org	route-views.linx.routeviews.org	route-views.wide.routeviews.org
115.69.224.0/21	37	36	17	6		9	23	4
1.50.8.0/24	33	32	16	6		9	23	4
1.200.0.0/22	33	32	16	6		9	23	4

RIPE Debogon Project

Checking reachability per /8s in case of new allocation based on RIPE

http://www.ris.ripe.net/debogon/

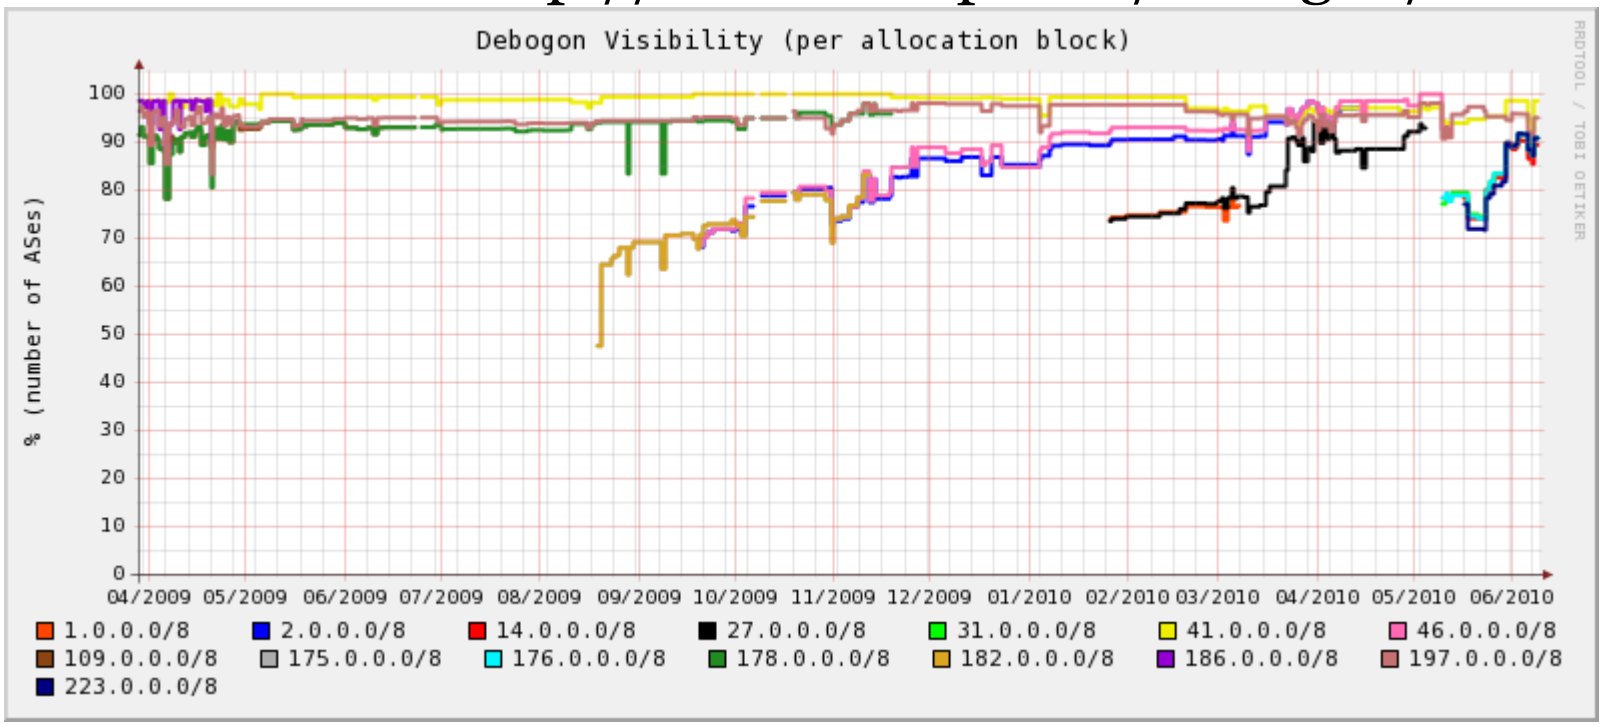

1.50.0.0/22 Reachability（Feb.-Mar 2010）

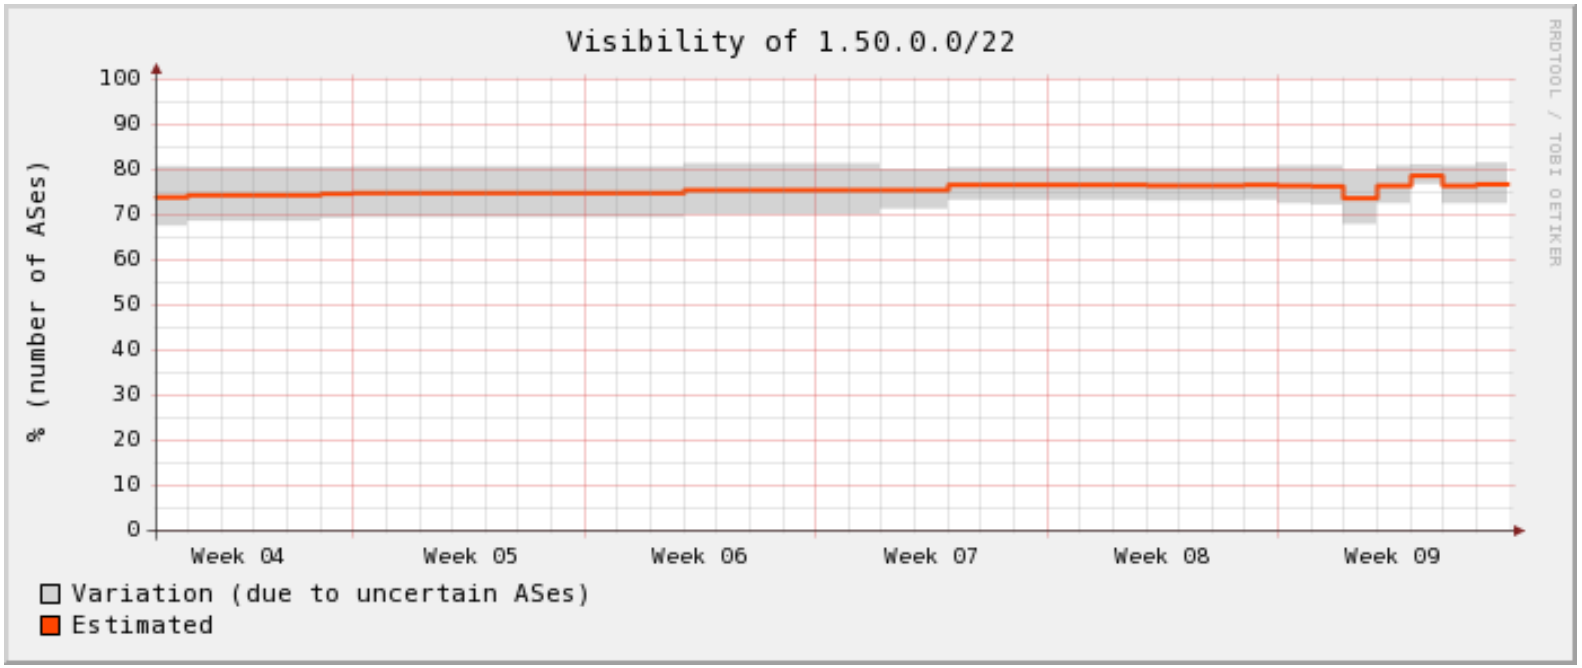

27.50.0.0/22 Reachability（Feb.- Apr. 2010）

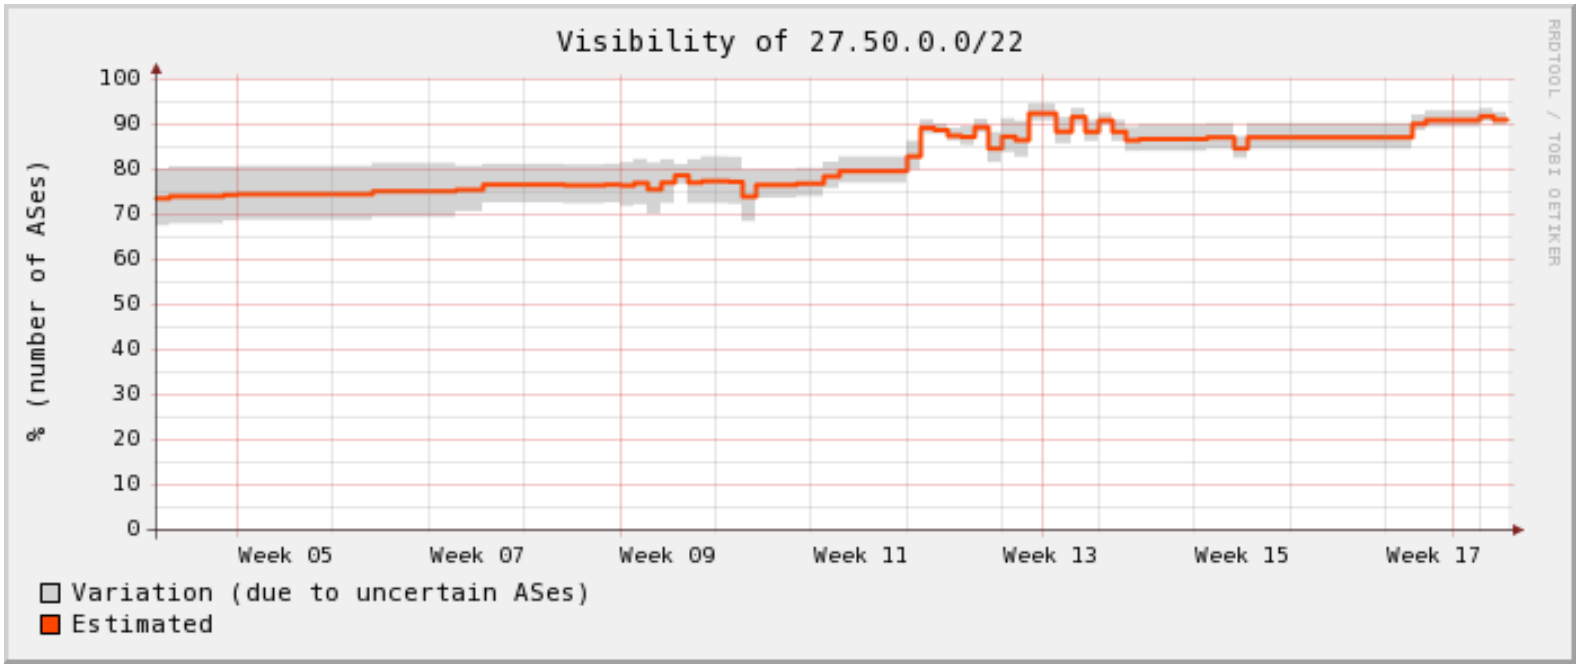

May 2010

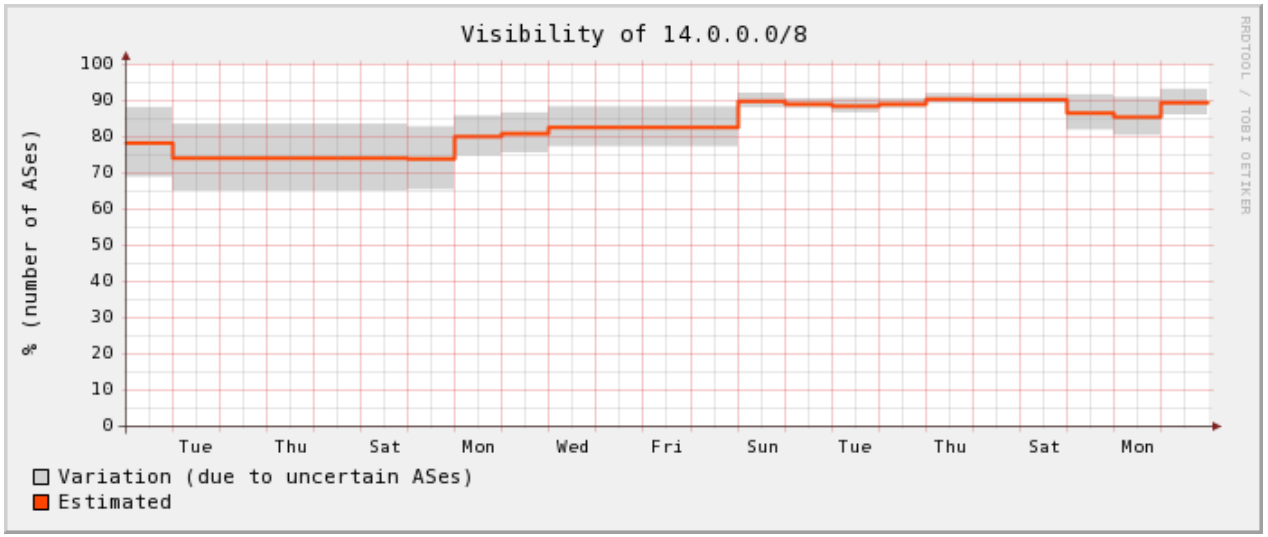

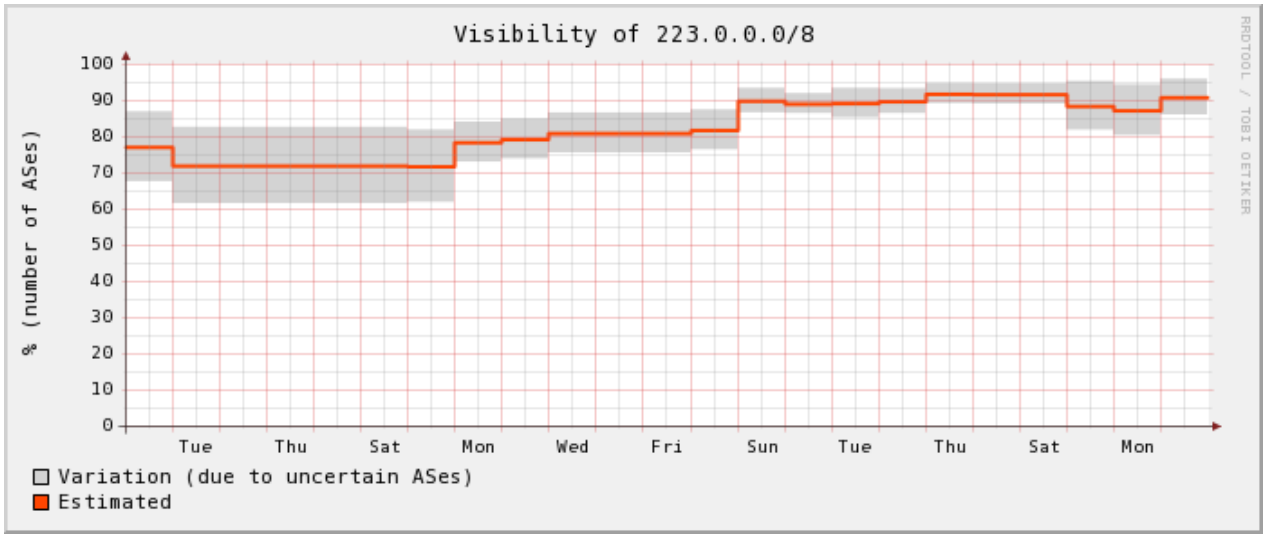

http://www.ris.ripe.net/debogon/2010/05/index.shtml

{1,27}/8 reachability investigation from NTTCom AS38639

- Checking 22581 Ases by ping reachability check
  - Results：Approximately 10% are unreachable
    - It's similar to Routeview、RIPE debogon results
  - When assignment is begun to LIR, improving gradually

	15th Apr. 2010		8th Jun. 2010	
	27/8(new)	203/8(old)	1/8(new)	203/8(old)
# of Dest AS	22581	22581	22581	22581
# of Ping OK AS	20086	22167	19787	21177
Diff	2495	414	2794	1404
% of NG	11%	2%	12%	6%

Probably now it's less than 10%

Checking (a part of) Blacklist

esta.cbp.dhs.gov
www.2ch.net (bulletin board)
www.mlb.com
www.bbc.co.uk
www.americanairlines.jp
...

Approximately 50% is not reachable

27.0.0.0/8

38/72

2010/04/15 investigation

1.0.0.0/8

39/72

2010/06/05 investigation

1slash8 NW for 7/8-9 janog26

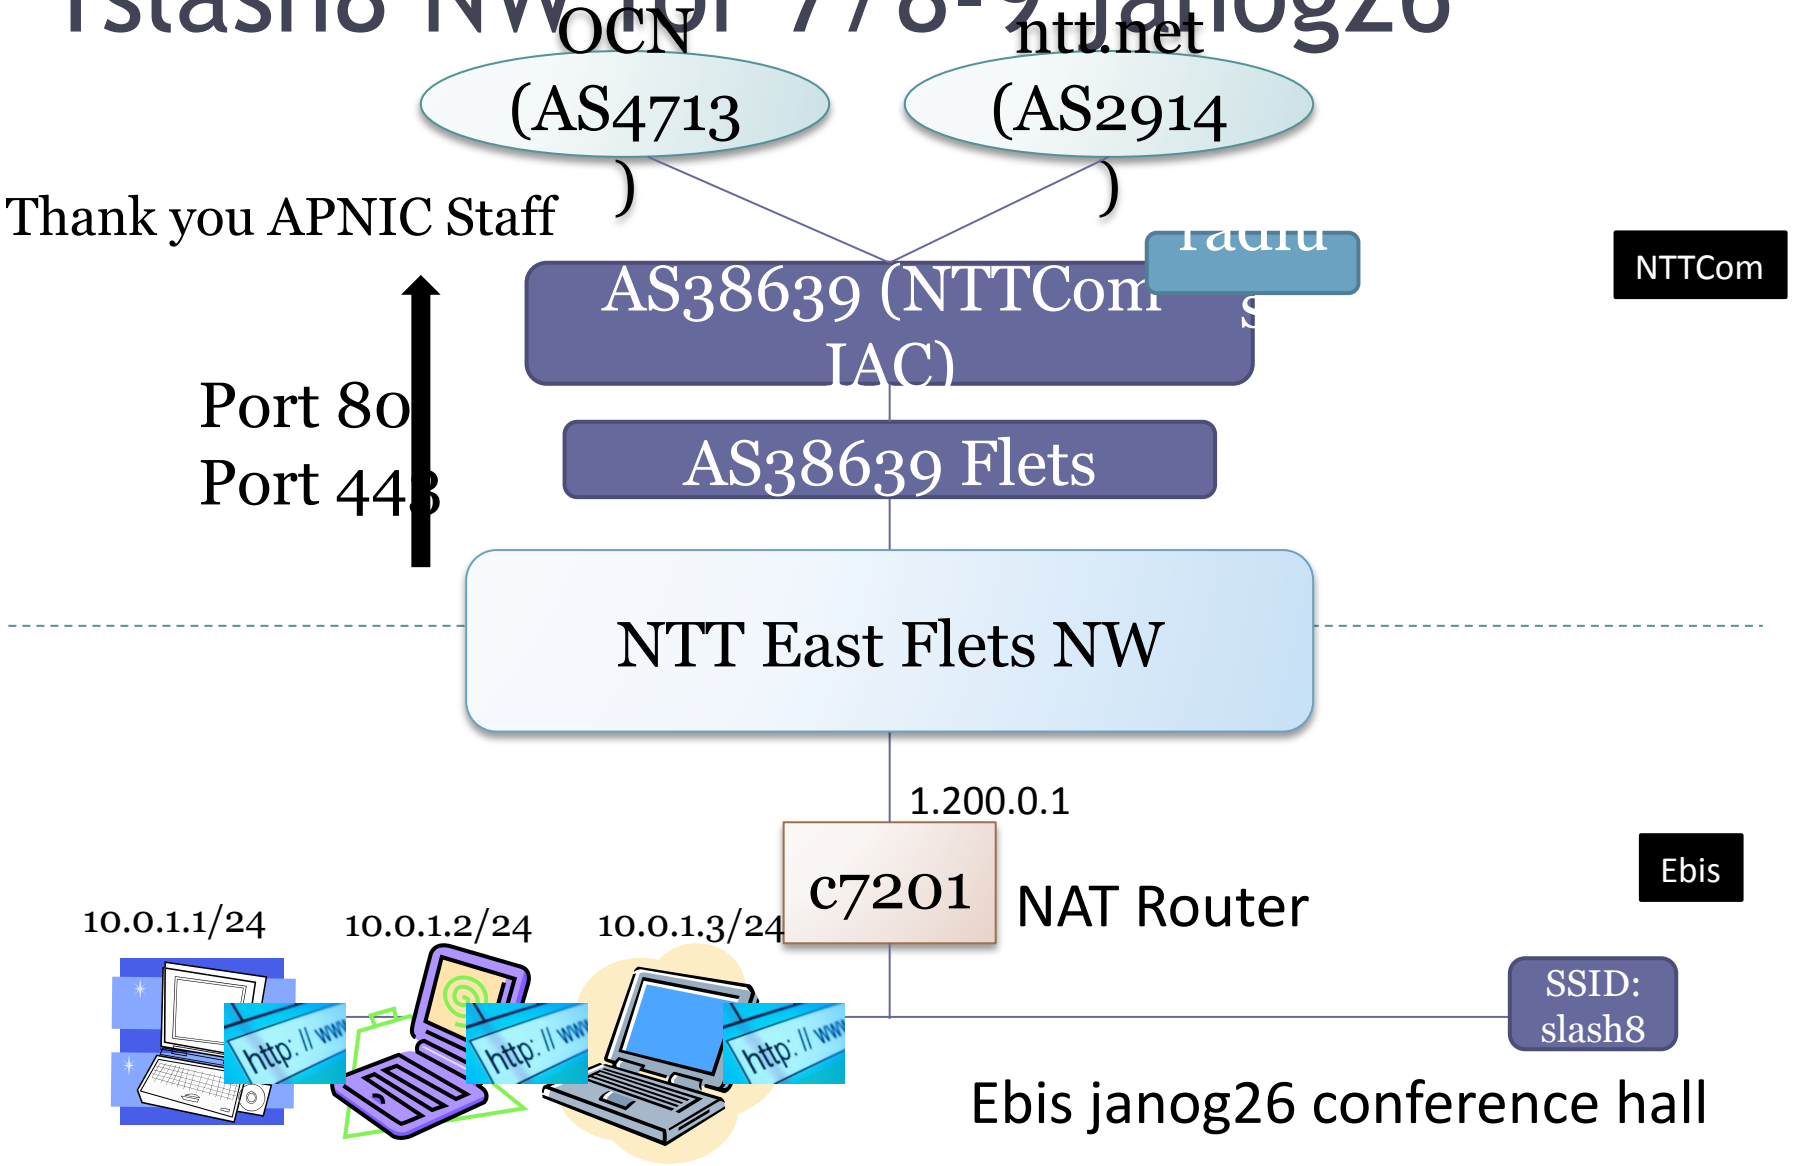

(A part of) reachability NG list from 1/8 netowrk only @ janog26 meeting

- http://www.metro.tokyo.jp/
- http://www.sangiin.go.jp/
- http://www.lottehotel.com/
- http://www.xn--w22as22a.com/
- http://www.mizuho-tb.co.jp/
- http://www.mizuhocbk.co.jp/
- http://www.alaxala.co.jp/
- http://www.admission.jp/
- http://www.clarion.com/
- http://chizu-route-susumu.jp/
- http://metacafe.com/
- http://softonic.com/
- http://gougou.com/
- http://www.bbc.co.uk/
- http://www.e-tokyo.lg.jp/
- http://www.ebookjapan.jp/
- http://www.nta.go.jp/
- http://www.ietf.org/ (tools.ietf.org OK)

1/8 usage （by Routing Info）

2010/7/8

Already used(routable) one third space in 1/8 but still not good conditions

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

1.X/16 (X:0-255)

- 100% used
- less than 100% used
- Ping source block

Recent Allocation to AS4713

27/8 allocation to APNIC : Jan. 2010

Network Information:

[Network Number]	**27.114.0.0/17**
[Network Name]	
[Organization]	NTT COMMUNICATIONS CORPORATION
[Administrative Contact]	AY1361JP
[Technical Contact]	KK551JP
[Technical Contact]	TS19037JP
[Technical Contact]	TT10660JP
[Abuse]	abuse@ocn.ad.jp
[Allocated Date]	2010/07/12
[Last Update]	2010/07/12 15:43:21(JST)

Recent Allocation to AS4713

223/8 allocation to APNIC : Jan. 2010

Network Information:

[Network Number]	**223.216.0.0/14**
[Network Name]	
[Organization]	NTT COMMUNICATIONS CORPORATION
[Administrative Contact]	AY1361JP
[Technical Contact]	KK551JP
[Technical Contact]	TS19037JP
[Technical Contact]	TT10660JP
[Abuse]	abuse@ocn.ad.jp
[Allocated Date]	2010/07/12
[Last Update]	2010/07/12 15:43:21(JST)

New Allocation IP's reachability investigation (16th Jul, 2010)

115/8 > 27/8 > 1/8, 14/8 223/8

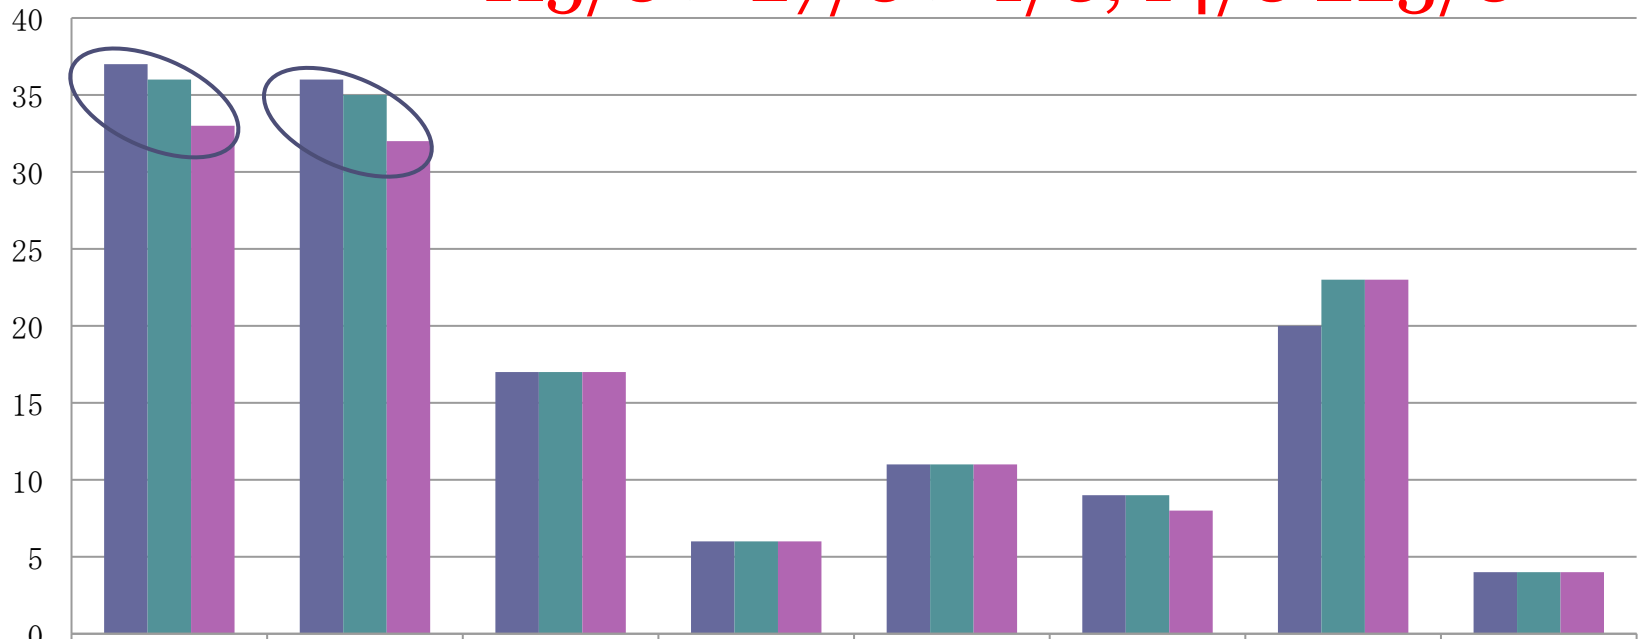

	route-views.routeviews.org	route-views2.routeviews.org	route-views3.routeviews.org	route-views4.routeviews.org	route-views.eqix.routeviews.org	route-views.isc.routeviews.org	route-views.linx.routeviews.org	route-views.wide.routeviews.org
115.69.224.0/21	37	36	17	6	11	9	20	4
27.114.0.0/17	36	35	17	6	11	9	23	4
223.216.0.0/14	33	32	17	6	11	8	23	4

Google search for “1.200.0.1”

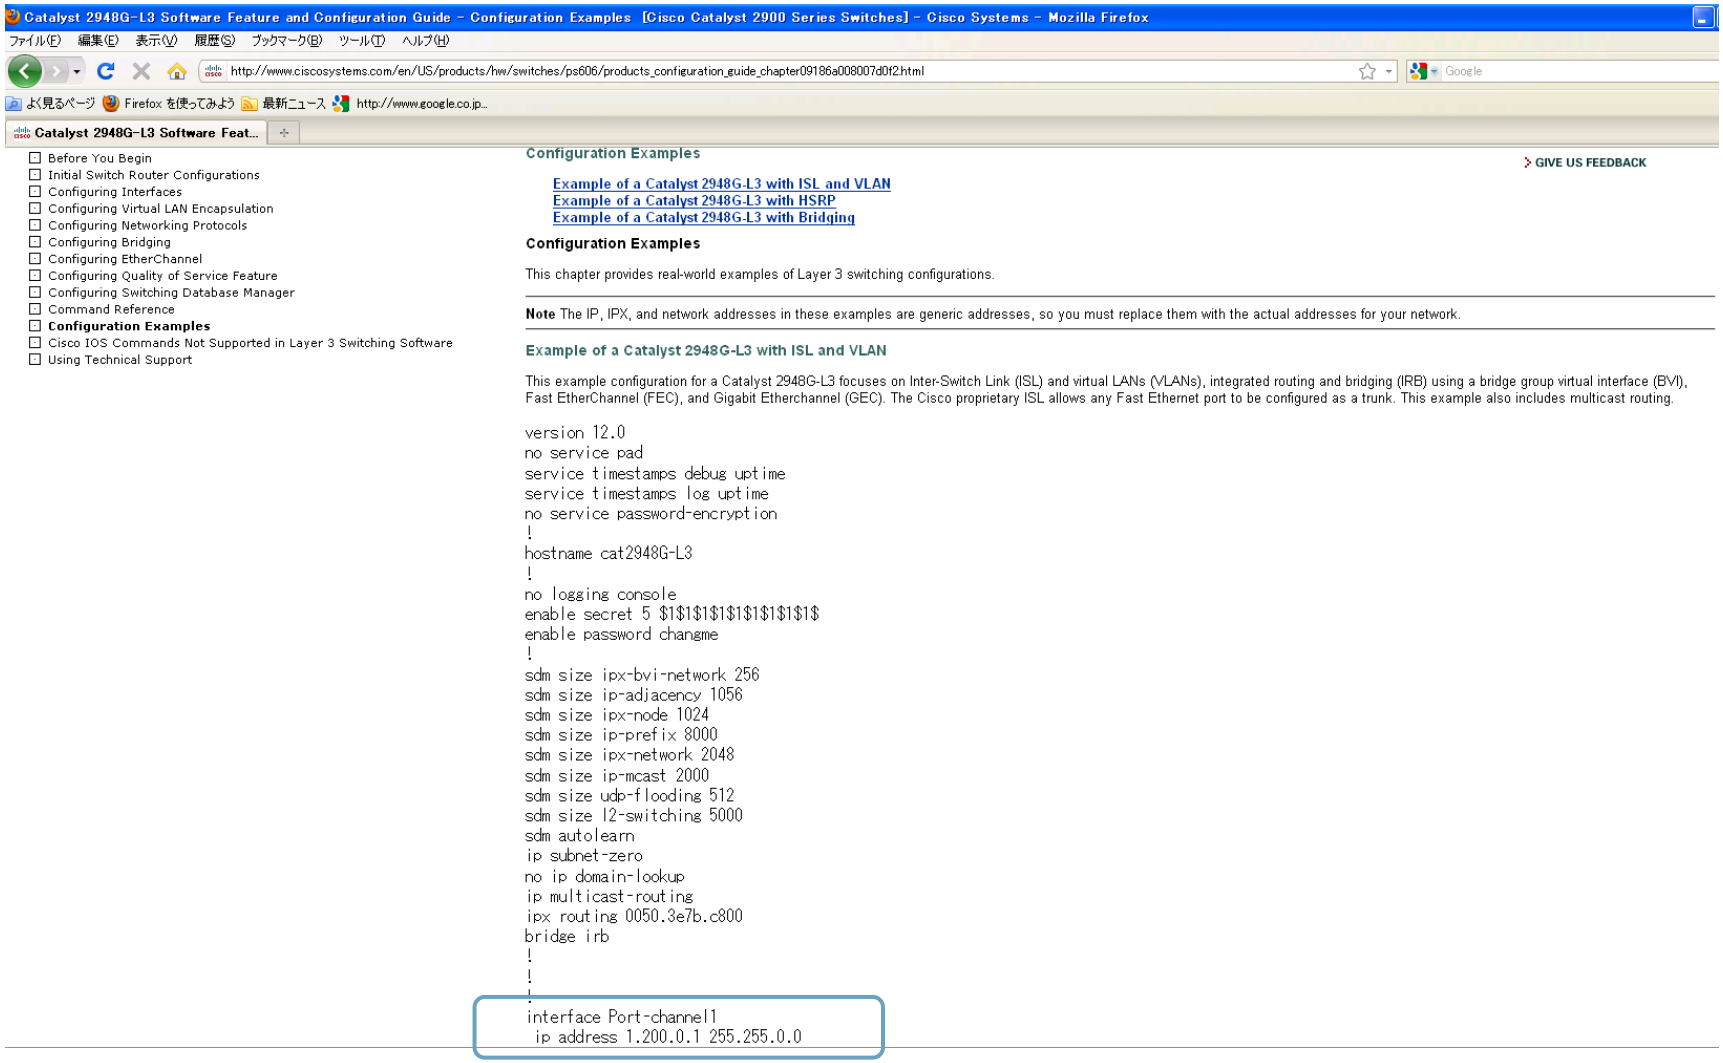

Thank You

frank@apnic.net
yoshida@nttv6.jp