

Internet Routing Table Analysis Update



Philip Smith
SANOG 16
21 July 2010
Paro



Motivation

- 1998: No one was publishing any Internet routing table analysis
 - Only CIDR-Report reporting on top 20 contributors to routing table, and top 20 bad aggregators
- With support of APNIC, my weekly reporting report started 23rd February 1999:
 - Routing table size
 - CIDR-Report style reporting on a per-RIR basis
 - ...and many other interesting features

Routing Report 21 July 2010

BGP routing table entries examined:	325968
Prefixes after maximum aggregation:	149997
Deaggregation factor:	2.17
Unique aggregates announced to Internet:	159238
Total ASes present in the Internet Routing Table:	34407
Prefixes per ASN:	9.47
Origin-only ASes present in the Internet Routing Table:	29885
Origin ASes announcing only one prefix:	14473
Transit ASes present in the Internet Routing Table:	4522
Transit-only ASes present in the Internet Routing Table:	103
Average AS path length visible in the Internet Routing Table:	3.6
Max AS path length visible:	26
Max AS path prepend of ASN (41664)	21
Prefixes from unregistered ASNs in the Routing Table:	299
Unregistered ASNs in the Routing Table:	116
Number of 32-bit ASNs allocated by the RIRs:	695
Prefixes from 32-bit ASNs in the Routing Table:	843
Special use prefixes present in the Routing Table:	0
Prefixes being announced from unallocated address space:	163
Number of addresses announced to Internet:	2278316352
Equivalent to 135 /8s, 204 /16s and 89 /24s	
Percentage of available address space announced:	61.5
Percentage of allocated address space announced:	66.3
Percentage of available address space allocated:	92.8
Percentage of address space in use by end-sites:	83.9
Total number of prefixes smaller than registry allocations:	155397

APNIC Region

Prefixes being announced by APNIC Region ASes:	79001
Total APNIC prefixes after maximum aggregation:	27132
APNIC Deaggregation factor:	2.91
Prefixes being announced from the APNIC address blocks:	75900
Unique aggregates announced from the APNIC address blocks:	33485
APNIC Region origin ASes present in the Internet Routing Table:	4119
APNIC Prefixes per ASN:	18.43
APNIC Region origin ASes announcing only one prefix:	1136
APNIC Region transit ASes present in the Internet Routing Table:	636
Average APNIC Region AS path length visible:	3.6
Max APNIC Region AS path length visible:	15
Number of APNIC addresses announced to Internet:	532384032
Equivalent to 31 /8s, 187 /16s and 137 /24s	
Percentage of available APNIC address space announced:	79.3

APNIC AS Blocks	4608-4864, 7467-7722, 9216-10239, 17408-18431
(pre-ERX allocations)	23552-24575, 37888-38911, 45056-46079
	55296-56319, 131072-132095
APNIC Address Blocks	1/8, 14/8, 27/8, 43/8, 58/8 to 61/8,
	110/8 to 126/8, 133/8, 175/8, 180/8, 182/8,
	183/8, 202/8, 203/8, 210/8, 211/8, 218/8 to 223/8

Global per AS prefix count summary

ASN	No of nets	/20 equiv	Max Agg	Description
6389	3890	3718	288	bellsouth.net, inc.
4323	2720	1112	394	Time Warner Telecom
4766	1857	8409	489	Korea Telecom (KIX)
19262	1831	4567	278	Verizon Global Networks
1785	1780	698	129	PaeTec Communications, Inc.
20115	1511	1523	654	Charter Communications
7018	1477	5734	960	AT&T WorldNet Services
4755	1470	297	161	TATA Communications formerly
8151	1463	3049	248	UniNet S.A. de C.V.
7545	1369	232	103	TPG Internet Pty Ltd
17488	1333	145	127	Hathway IP Over Cable Interne
2386	1287	568	909	AT&T Data Communications Serv
6478	1253	251	201	AT&T Worldnet Services
22773	1178	2859	65	Cox Communications, Inc.
17974	1169	285	51	PT TELEKOMUNIKASI INDONESIA
11492	1160	209	78	Cable One
3356	1159	10878	402	Level 3 Communications, LLC
7011	1132	288	651	Citizens Utilities
8452	1102	445	10	TEDATA
18566	1088	296	39	Covad Communications

Number of prefixes announced by prefix length

/1:0	/2:0	/3:0	/4:0	/5:0	/6:0
/7:0	/8:21	/9:10	/10:25	/11:68	/12:198
/13:409	/14:710	/15:1294	/16:11175	/17:5323	/18:9150
/19:18487	/20:23055	/21:23044	/22:29962	/23:29612	/24:170203
/25:1052	/26:1311	/27:771	/28:28	/29:47	/30:6
/31:0	/32:7				

July 2010 ↑

July 2009 ↓


Number of prefixes announced by prefix length

/1:0	/2:0	/3:0	/4:0	/5:0	/6:0
/7:0	/8:19	/9:10	/10:24	/11:58	/12:169
/13:350	/14:612	/15:1165	/16:10533	/17:4754	/18:8224
/19:17157	/20:20506	/21:20376	/22:26165	/23:26077	/24:152659
/25:930	/26:1031	/27:560	/28:152	/29:8	/30:7
/31:1	/32:8				


Prefixes Smaller than Registry Allocations

ASN	No of nets	Total ann.	Description
6389	2491	3890	bellsouth.net, inc.
4766	1487	1857	Korea Telecom (KIX)
4323	1392	2720	Time Warner Telecom
1785	1241	1780	PaeTec Communications, Inc.
17488	1078	1333	Hathway IP Over Cable Interne
18566	1069	1088	Covad Communications
11492	1068	1160	Cable One
8452	996	1102	TEDATA
10620	988	1072	TVCABLE BOGOTA
7018	877	1477	AT&T WorldNet Services
24560	865	966	Bharti Airtel Ltd., Telemedia
9583	853	1001	Sify Limited
7011	832	1132	Citizens Utilities
19262	815	1831	Verizon Global Networks
4755	795	1470	TATA Communications formerly
3356	789	1159	Level 3 Communications, LLC
17974	745	1169	PT TELEKOMUNIKASI INDONESIA
5668	702	986	CenturyTel Internet Holdings,
22773	670	1178	Cox Communications, Inc.
7545	624	1369	TPG Internet Pty Ltd


BGP Routing Table



AS Growth




Max Aggregation vs Unique Prefixes



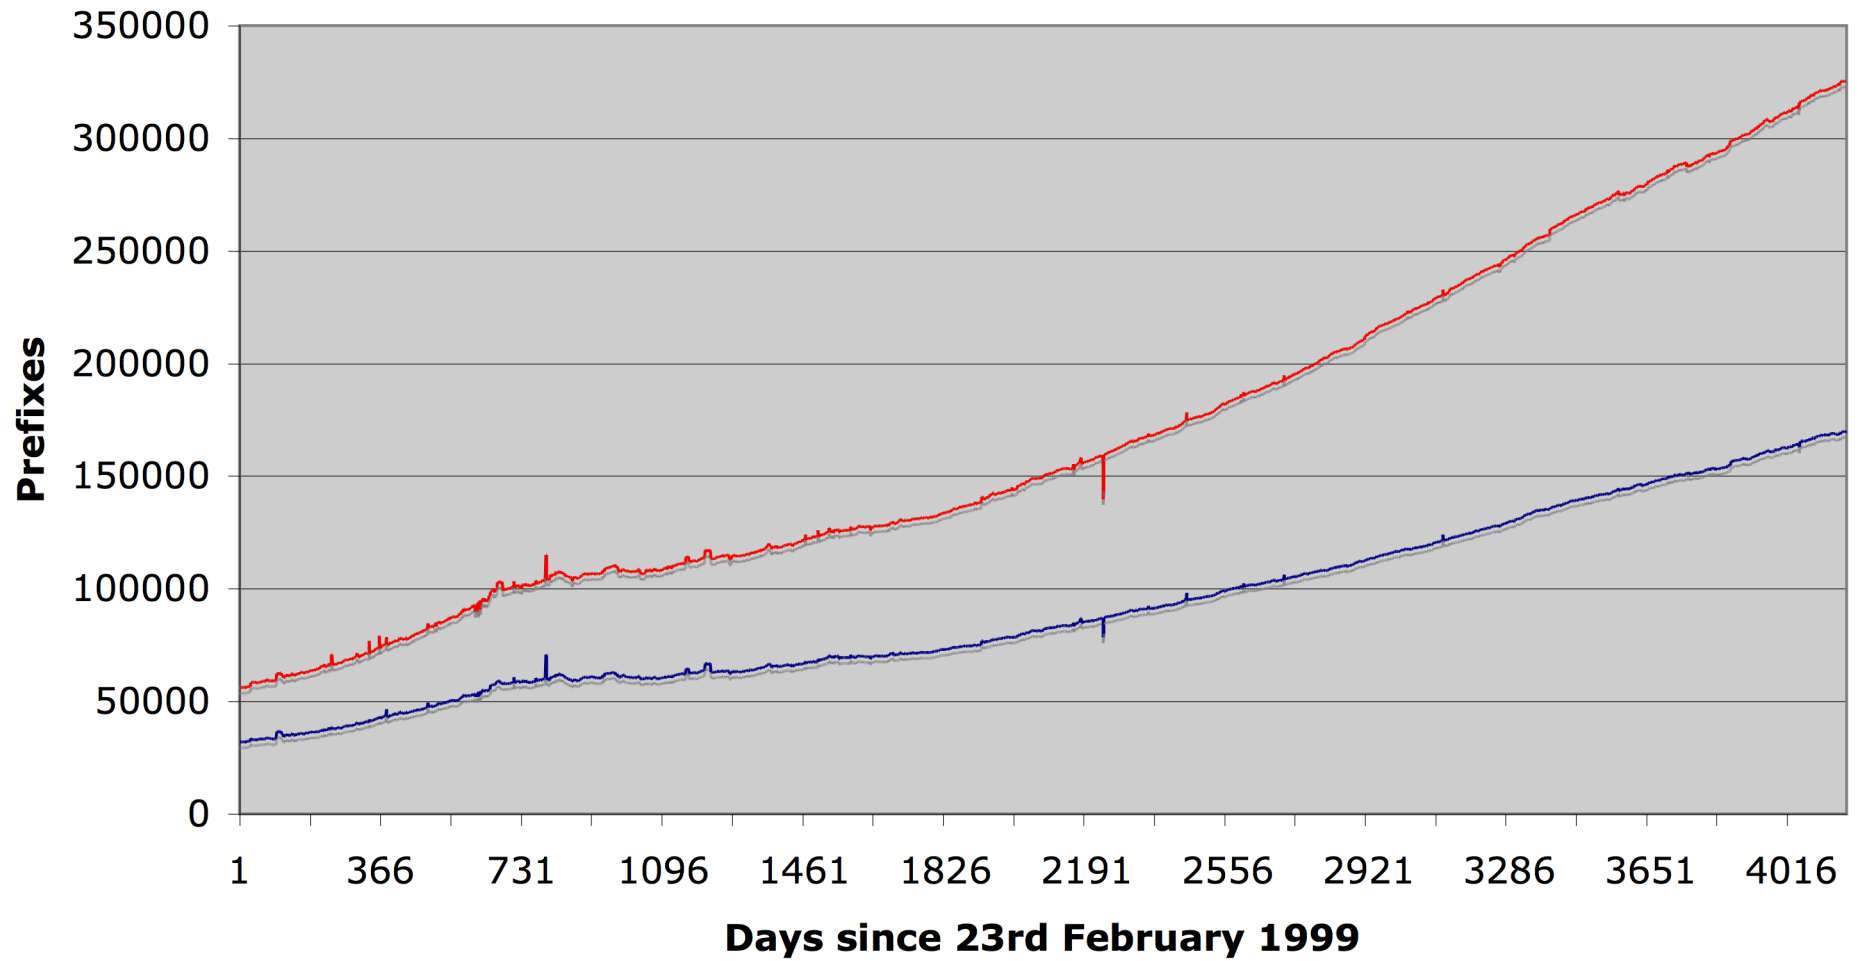
— Max Aggregation — Unique Prefixes — Global BGP Table

Prefix sizes announced




— /19 — /20 — /21 — /22 — /23

/24s announced




— /24s — Global BGP Table

Address Space announced




AS Announcements




— Total ASNs — Origin-only ASNs — ASN providing Transit & Origin — Transit-only ASNs

Maximum AS Path Length



Growth in IPv4 Address Space Allocations





Looking at Deaggregation

- CIDR Report
 - www.cidr-report.org
 - Encourages aggregation following CIDRisation of Internet
 - Today: extensive suite of reports and tools covering state of BGP table
- Routing Report
 - BGP table status on per RIR basis
 - Original CIDR Report and a whole lot more



Deaggregation Factor

- Routing Report
 - One summary takes BGP table and aggregates prefixes by origin AS
 - Called “Max Aggregation” in report
 - Global and per RIR basis
 - <http://thyme.apnic.net/current/>
- Calculates **Deaggregation Factor**:
 - Measure of Routing Table size/Aggregated Size
 - Global value has been increasing slowly and steadily since “records began”



July 2010


Total Prefixes

- Global BGP Table
 - 326k prefixes
- Europe & Middle East
 - 75k prefixes
- North America
 - 135k prefixes
- Asia & Pacific
 - 79k prefixes
- Africa
 - 7k prefixes
- Latin America & Caribbean
 - 29k prefixes

Deaggregation Factor

- Global Average
 - 2.17
- Europe & Middle East
 - 1.72
- North America
 - 1.94
- Asia & Pacific
 - 2.91
- Africa
 - 3.91
- Latin America & Caribbean
 - 4.15

Deaggregation: RIR Regions vs Global



Global AfriNIC APNIC ARIN LACNIC RIPE

Asia Pacific Aggregation Savings Summary

ASN	No of Nets	Savings	Description
4766	1857	1368	Korea Telecom (KIX)
4755	1470	1309	TATA Communications formerly
7545	1369	1266	TPG Internet Pty Ltd
17488	1333	1206	Hathway IP Over Cable Interne
17974	1169	1118	PT TELEKOMUNIKASI INDONESIA
24560	966	788	Bharti Airtel Ltd., Telemedia
9829	813	781	BSNL National Internet Backbo
9498	693	651	BHARTI Airtel Ltd.
7552	652	643	Vietel Corporation
4780	692	619	Digital United Inc.
4808	828	609	CNCGROUP IP network: China169
18101	691	568	Reliance Infocom Ltd Internet
9583	1001	516	Sify Limited
17676	580	516	Softbank BB Corp.
9808	515	502	Guangdong Mobile Communicatio
17908	544	487	Tata Communications
23700	479	473	Indonesia Network Information
9443	544	468	Primus Telecommunications
17803	539	406	BSES TeleCom Limited
4804	459	391	Microplex PTY LTD

<http://thyme.apnic.net/current/data-CIDRnet-APNIC>



Importance of Aggregation

- Size of routing table
 - Memory is no longer a problem
 - Routers can be specified to carry 1 million prefixes
- Convergence of the Routing System
 - This is a problem
 - Bigger table takes longer for CPU to process
 - BGP updates take longer to deal with
 - BGP Instability Report tracks routing system update activity
 - <http://bgpupdates.potaroo.net/instability/bgpupd.html>

The BGP Instability Report

The BGP Instability Report is updated daily. This report was generated on 14 July 2010 06:12 (UTC+1000)

50 Most active ASes for the past 7 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	30890	260803	14.67%	443	588.72	EVOLVA Evolva Telecom s.r.l.
2	24400	62801	3.53%	12	5233.42	CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
3	9808	58522	3.29%	510	114.75	CMNET-GD Guangdong Mobile Communication Co.Ltd.
4	50075	38768	2.18%	66	587.39	INBAR-BUSINESS-CORPORATION INBAR BUSINESS CORPORATION S.R.L.
5	5416	30112	1.69%	137	219.80	BATELCO-BH
6	2018	25065	1.41%	210	119.36	TENET-1
7	44475	20317	1.14%	32	634.91	TIADOLI-AS SC Tiadoli Company SRL
8	39543	18337	1.03%	25	733.48	TENNET-AS SC TENNET TELECOM SRL
9	11981	17498	0.98%	9	1944.22	SPECIALSYSTEMS - Special Systems Inc.
10	48838	16945	0.95%	33	513.48	EUROLAN-ASN Eurolan Solutions SRL
11	35931	15479	0.87%	6	2579.83	ARCHIPELAGO - ARCHIPELAGO HOLDINGS INC
12	25620	14829	0.83%	176	84.26	COTAS LTDA.
13	37204	14571	0.82%	10	1457.10	TELONE
14	8452	13369	0.75%	1173	11.40	TEDATA TEDATA
15	41852	13238	0.74%	19	696.74	EXPERTNET-AS S.C. EXPERTNET S.R.L.
16	32528	12831	0.72%	8	1603.88	ABBOTT Abbot Labs
17	34916	10834	0.61%	18	601.89	DNET-AS SC Digital Construction Network SRL
18	9829	10776	0.61%	813	13.25	BSNL-NIB National Internet Backbone
19	14420	10199	0.57%	499	20.44	CORPORACION NACIONAL DE TELECOMUNICACIONES - CNT EP
20	49829	8716	0.49%	9	968.44	CAPITAL-COM SC Capital Communication Investment SRL
21	35048	8300	0.47%	12	691.67	NETZONE-AS SC Power Netzone SRL
22	5800	8140	0.46%	185	44.00	DNIC-ASBLK-05800-06055 - DoD Network Information Center
23	33588	7976	0.45%	647	12.33	BRESNAN-AS - Bresnan Communications, LLC.

50 Most active Prefixes for the past 7 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	111.10.4.0/24	11531	0.62%	9808 -- CMNET-GD Guangdong Mobile Communication Co.Ltd.
2	111.10.0.0/24	11531	0.62%	9808 -- CMNET-GD Guangdong Mobile Communication Co.Ltd.
3	111.10.1.0/24	11529	0.62%	9808 -- CMNET-GD Guangdong Mobile Communication Co.Ltd.
4	111.10.3.0/24	11528	0.62%	9808 -- CMNET-GD Guangdong Mobile Communication Co.Ltd.
5	111.10.2.0/24	11526	0.62%	9808 -- CMNET-GD Guangdong Mobile Communication Co.Ltd.
6	117.131.0.0/17	9283	0.50%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
7	120.204.0.0/16	9179	0.49%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
8	117.136.8.0/24	8901	0.48%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
9	117.135.128.0/18	8693	0.47%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
10	117.135.0.0/17	8684	0.46%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
11	198.140.43.0/24	8532	0.46%	35931 -- ARCHIPELAGO - ARCHIPELAGO HOLDINGS INC
12	212.51.128.0/19	7720	0.41%	13030 -- INIT7 Init Seven AG, Zurich, Switzerland
13	63.211.68.0/22	6947	0.37%	35931 -- ARCHIPELAGO - ARCHIPELAGO HOLDINGS INC
14	130.36.34.0/24	6393	0.34%	32528 -- ABBOTT Abbot Labs
15	130.36.35.0/24	6392	0.34%	32528 -- ABBOTT Abbot Labs
16	190.65.228.0/22	5574	0.30%	3816 -- COLOMBIA TELECOMUNICACIONES S.A. ESP
17	221.181.64.0/18	5136	0.27%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
18	207.254.176.0/20	4548	0.24%	19174 -- CNC-USA - China Netcom (USA) Operations Ltd.
19	41.34.29.0/24	3640	0.19%	8452 -- TEDATA TEDATA
20	211.136.160.0/19	3266	0.17%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
21	211.136.96.0/19	3266	0.17%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
22	211.136.128.0/19	3266	0.17%	24400 -- CMNET-V4SHANGHAI-AS-AP Shanghai Mobile Communications Co.,Ltd.
23	206.184.16.0/24	3064	0.16%	174 -- COGENT Cogent/PSI
24	64.86.26.0/23	3042	0.16%	37204 -- TELONE
25	143.138.107.0/24	2897	0.16%	747 -- TAEGU-AS - Headquarters, USAISC
26	91.212.23.0/24	2858	0.15%	48754 -- SOBIS-AS SOBIS SOLUTIONS SRL
27	196.29.32.0/21	2676	0.14%	37204 -- TELONE



Observations

- Range of operational “practices” between RIR regions
 - Deaggregation by newer ISPs & developing regions is growing rapidly
 - Is harming the **entire** Internet
- RIPE-399 is only a recommendation
 - Hopefully all the RIRs will include pointers with each address allocation
 - Hopefully more ISPs will pay attention to it
 - Training is there — most ISPs choose to ignore it

Internet Routing Table Analysis Update



Questions?