



BGP Aggregation & The Deaggregation Report

Philip Smith

PacNOG 3
16-23 June 2007
Rarotonga



Route Aggregation Recommendations

- LINX attempted aggregation policy for members
 - It failed even though most members voted for policy
- RIPE Routing Working Group work item from early 2006
 - Based on early LINX concept
 - Authored by Philip Smith, Mike Hughes (LINX) and Rob Evans (UKERNA)



Route Aggregation Recommendations

- RIPE Document — RIPE-399
 - <http://www.ripe.net/ripe/docs/ripe-399.html>
- Discusses:
 - History of aggregation
 - Causes of de-aggregation
 - Impacts on global routing system
 - Available Solutions
 - Recommendations for ISPs



History:

- Classful to classless migration
 - Clean-up efforts in 192/8
- CIDR Report
 - Started by Tony Bates to encourage adoption of CIDR & aggregation
 - Mostly ignored through late 90s
 - Now part of extensive BGP table analysis by Geoff Huston
- Introduction of Regional Internet Registry system and PA address space



Deaggregation: Claimed causes (1):

- Routing System Security
 - “Announcing /24s means that no one else can DOS the network”
- Reduction of DOS attacks & miscreant activities
 - “Announcing only address space in use as rest attracts ‘noise’”
- Commercial Reasons
 - “Mind your own business”



Deaggregation: Claimed causes (2):

- Leakage of iBGP outside of local AS
 - eBGP is NOT iBGP – how many ISPs know this?
- Traffic Engineering for Multihoming
 - Spraying out /24s hoping it will work
 - Rather than do any **real engineering**
- Legacy Assignments
 - “All those pre-RIR assignments are to blame”
 - In reality it is both RIR and legacy assignments



Impacts (1):

- Router memory
 - Shortens router life time as vendors underestimate memory growth requirements
 - Depreciation life-cycle shortened
 - Increased costs for ISP and customers
- Router processing power
 - Processors are underpowered as vendors underestimate CPU requirement
 - Depreciation life-cycle shortened
 - Increased costs for ISP and customers



Impacts (2):

- Routing System convergence
 - Larger routing table → slowed convergence
 - Can be improved by faster control plane processors — see earlier
- Network Performance & Stability
 - Slowed convergence → slowed recovery from failure
 - Slowed recovery → longer downtime
 - Longer downtime → unhappy customers



Solutions (1):

- CIDR Report
 - Global aggregation efforts
 - Running since 1994
- Routing Table Report
 - Per RIR region aggregation efforts
 - Running since 1999
- Filtering recommendations
 - Training, tutorials, Project Cymru,...
- “CIDR Police”



Solutions (2):

- BGP Features:
 - NO_EXPORT Community
 - NOPEER Community
 - RFC3765 — but no one has implemented it
 - AS_PATHLIMIT attribute
 - Still working through IETF IDR Working Group
 - Provider Specific Communities
 - Some ISPs use them; most do not



RIPE-399 Recommendations:

- Announcement of initial allocation as a single entity
- Subsequent allocations aggregated if they are contiguous and bit-wise aligned
- Prudent subdivision of aggregates for Multihoming
- Use BGP enhancements already discussed
- (Oh, and all this applies to IPv6 too)



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
- New **Deaggregation Factor**:
 - Measure of Routing Table size/Aggregated Size
 - Global value has been increasing slowly and steadily since “records began”



“Original Internet” — 2007/06

Total Prefixes

- Global BGP Table
 - 223k prefixes
- North America
 - 106k prefixes
- Europe & Middle East
 - 47k prefixes

Deaggregation Factor

- Global Average
 - 1.89
- North America
 - 1.72
- Europe & Middle East
 - 1.55



“Newer Internet” — 2007/06

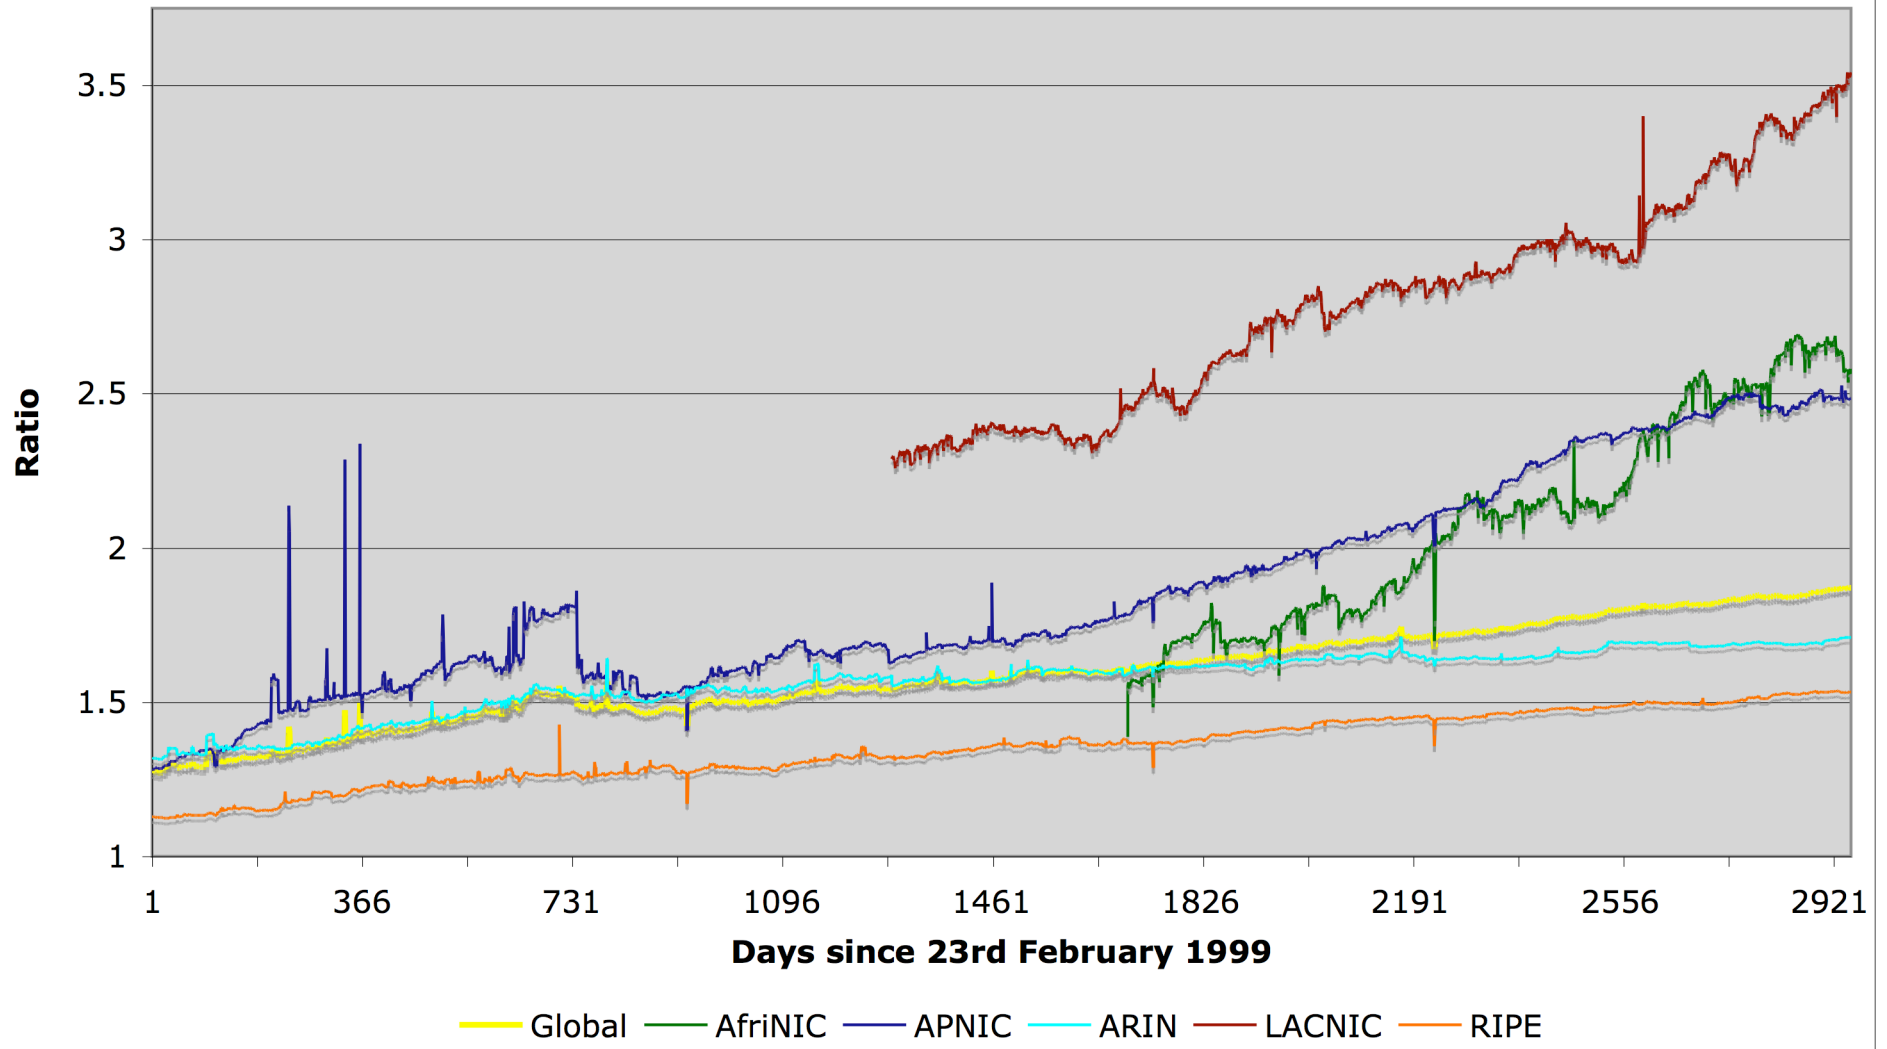
Total Prefixes

- Global BGP Table
 - 223k prefixes
- Asia & Pacific
 - 52k prefixes
- Africa
 - 3k prefixes
- Latin America & Caribbean
 - 15k prefixes

Deaggregation Factor

- Global Average
 - 1.89
- Asia & Pacific
 - 2.55
- Africa
 - 2.56
- Latin America & Caribbean
 - 3.54

Deaggregation: RIR Regions vs Global



Africa Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
8452	225	218	TEDATA
6713	143	132	Itissalat Al-MAGHRIB
15475	129	125	Nile Online
5536	123	108	Internet Egypt Network
33783	117	108	EEPAD TISP TELECOM & INTERNET
24835	85	79	RAYA Telecom - Egypt
3741	291	59	The Internet Solution
2561	61	59	Egyptian Universities Network
15706	51	47	Sudatel Internet Exchange Aut
23889	55	42	MAURITIUS TELECOM
33766	29	28	Nyala Communications Pty Ltd
2018	137	25	Tertiary Education Network
8524	30	24	AUCEGYPT Autonomous System
33776	30	23	Starcomms Nigeria Limited
21280	26	22	Swift Global Kenya Ltd.Is an
21491	23	21	UTL On-line is RF broadband I
33774	43	19	AS Number for Telecom Algeria
15804	19	18	AS of The Way Out Internet So
14988	16	15	Botswana Telecommunications C
25163	16	15	21 st Century Technologies Li

Asia & Pacific Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
4755	1202	1130	Videsh Sanchar Nigam Ltd. Aut
4134	1278	1007	CHINANET-BACKBONE
9498	974	900	BHARTI BT INTERNET LTD.
9583	1099	718	Sify Limited
17488	666	630	Hathway IP Over Cable Interne
7545	624	544	TPG Internet Pty Ltd
18101	529	503	Reliance Infocom Ltd Internet
4668	507	497	LG-EDS Systems Inc.
4766	782	470	Korea Telecom (KIX)
17676	504	439	Softbank BB Corp.
4812	463	388	China Telecom (Shanghai)
9443	458	384	Primus Telecommunications
9829	358	344	BSNL National Internet Backbo
17974	346	334	PT TELEKOMUNIKASI INDONESIA
9394	262	255	CHINA RAILWAY Internet (CRNET)
9929	294	253	China Netcom Corp.
10139	258	251	Meridian Telekoms
17849	286	243	Telecommunications Technology
4780	270	233	Digital United Inc.
9800	280	228	CHINA UNICOM

North America Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
11492	1043	1022	Cable One
18566	1001	992	Covad Communications
6478	1090	941	AT&T Worldnet Services
4323	1270	910	Time Warner Telecom
22773	707	667	Cox Communications, Inc.
5668	595	578	CenturyTel Internet Holdings,
19262	724	551	Verizon Global Networks
6197	1036	533	BellSouth Network Solutions,
19916	568	517	OLM LLC
855	568	497	Canadian Research Network
7029	575	494	Alltel Information Services,
15270	527	492	PaeTec.net -a division of Pae
33588	437	406	Bresnan Communications, LLC.
6517	442	404	Yipes Communications, Inc.
20115	822	396	Charter Communications
2386	1127	370	AT&T Data Communications Serv
3602	444	345	Sprint Canada, Inc.
721	615	342	DLA Systems Automation Center
3464	367	342	Alabama SuperComputer Network
11139	367	342	Cable & Wireless Dominica

Latin America Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
8151	1097	886	UniNet S.A. de C.V.
11830	483	464	Instituto Costarricense de El
16814	362	354	NSS, S.A.
11172	365	301	Servicios Alestra S.A de C.V
22047	311	299	VTR PUNTO NET S.A.
7303	271	243	Telecom Argentina Stet-France
11556	235	232	Cable-Wireless Panama
6471	259	228	ENTEL CHILE S.A.
14117	242	228	Telefonica del Sur S.A.
6147	220	196	Telefonica Del Peru
20299	192	178	NEWCOM AMERICAS
10481	181	173	Prima S.A.
10620	185	162	TVCABLE BOGOTA
14522	146	141	SatNet S.A.
6503	211	126	AVANTEL, S.A.
18822	132	125	TELEFONICA MANQUEHUE
7910	153	122	ANDINET ON LINE
7738	140	118	Telecomunicacoes da Bahia S.A
8163	124	118	METROTEL REDES S.A.
23216	157	115	RAMtelecom Telecomunicaciones

EU & Middle East Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
24863	343	312	LINKdotNET AS number
20858	196	192	EgyNet
12479	185	179	Uni2 Autonomous System
8551	206	176	Bezeq International
3352	200	167	Ibernet, Internet Access Netw
9116	188	164	Goldenlines main autonomous s
3215	240	148	France Telecom Transpac
9121	171	145	TTnet Autonomous System
3269	214	142	TELECOM ITALIA
5486	160	142	Euronet Digital Communication
6830	169	130	UPC Distribution Services
30890	205	110	SC Kappa Invexim SRL
8866	130	110	Bulgarian Telecommunication C
702	520	106	UUNET - Commercial IP service
9051	147	101	INCONET Autonomous System
15471	160	98	SNR - Societatea Nationala de
3300	183	95	AUCS Communications Services
12302	122	89	MobiFon S.A.
15611	86	84	Iranian Research Organisation
8764	115	79	LIETUVOS-TELEKOMAS Autonomous



Observations

- Huge gulf in operational good practices between “older” and “newer” Internet
 - Could threaten the Internet as we know it
- 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



Conclusion

- “Newer” Internet is growing rapidly
 - As is the deaggregation there
- RIPE-399 now exists
- Make it your BGP good practice document