



BGP Aggregation & The Deaggregation Report

Philip Smith

MENOG 1
3-5 April 2007
Bahrain

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 being sparing
- 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
 - 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” — 2007/02

Total Prefixes

- Global BGP Table
 - 217k prefixes
- North America
 - 105k prefixes
- Europe & Middle East
 - 45k prefixes

Deaggregation Factor

- Global Average
 - 1.87
- North America
 - 1.71
- Europe & Middle East
 - 1.53



“Newer Internet” — 2007/02


Total Prefixes

- Global BGP Table
 - 217k prefixes
- Asia & Pacific
 - 50k prefixes
- Africa
 - 3k prefixes
- Latin America & Caribbean
 - 14k prefixes

Deaggregation Factor

- Global Average
 - 1.87
- Asia & Pacific
 - 2.48
- 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	208	201	TEDATA
6713	143	132	Itissalat Al-MAGHRIB
15475	128	124	Nile Online
5536	123	108	Internet Egypt Network
33783	113	106	EEPAD TISP TELECOM & INTERNET
24835	85	79	RAYA Telecom - Egypt
3741	291	59	The Internet Solution
2561	61	59	Egyptian Universities Network
15706	55	51	Sudatel Internet Exchange Aut
23889	49	36	MAURITIUS TELECOM
33766	33	32	Nyala Communications Pty Ltd
2018	138	25	Tertiary Education Network
8524	30	24	AUCEGYPT Autonomous System
21280	27	23	Swift Global Kenya Ltd. Is an
33776	27	22	Starcomms Nigeria Limited
12455	25	22	Jambonet Autonomous system
15804	23	22	AS of The Way Out Internet So
21491	23	21	UTL On-line is RF broadband I
33774	42	19	AS Number for Telecom Algeria
14988	16	15	Botswana Telecommunications C

Asia & Pacific Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
4755	1173	1096	Videsh Sanchar Nigam Ltd. Aut
4134	1286	1006	CHINANET-BACKBONE
9498	968	898	BHARTI BT INTERNET LTD.
9583	1103	726	Sify Limited
17488	626	594	Hathway IP Over Cable Interne
7545	609	532	TPG Internet Pty Ltd
18101	538	511	Reliance Infocom Ltd Internet
4668	502	492	LG-EDS Systems Inc.
4766	788	474	Korea Telecom (KIX)
17676	503	438	Softbank BB Corp.
9443	453	380	Primus Telecommunications
4812	451	379	China Telecom (Shanghai)
17974	341	327	PT TELEKOMUNIKASI INDONESIA
9829	315	301	BSNL National Internet Backbo
9929	298	255	China Netcom Corp.
9394	259	252	CHINA RAILWAY Internet (CRNET)
10139	257	250	Meridian Telekoms
17849	284	242	Telecommunications Technology
4780	267	232	Digital United Inc.
9800	278	225	CHINA UNICOM

North America Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
11492	1017	1005	Cable One
18566	999	990	Covad Communications
6478	1080	926	AT&T Worldnet Services
4323	1267	915	Time Warner Telecom
22773	695	653	Cox Communications, Inc.
5668	581	564	CenturyTel Internet Holdings,
19262	721	559	Verizon Global Networks
6197	1032	532	BellSouth Network Solutions,
19916	568	514	OLM LLC
855	565	494	Canadian Research Network
7029	564	483	Alltel Information Services,
15270	516	482	PaeTec.net -a division of Pae
3602	518	419	Sprint Canada, Inc.
33588	431	404	Bresnan Communications, LLC.
6517	432	395	Yipes Communications, Inc.
20115	810	385	Charter Communications
2386	1089	358	AT&T Data Communications Serv
11139	368	343	Cable & Wireless Dominica
3464	362	337	Alabama SuperComputer Network
721	596	323	DLA Systems Automation Center

Latin America Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
8151	1061	850	UniNet S.A. de C.V.
11830	483	464	Instituto Costarricense de El
16814	361	353	NSS, S.A.
22047	307	296	VTR PUNTO NET S.A.
11172	361	295	Servicios Alestra S.A de C.V
7303	265	241	Telecom Argentina Stet-France
11556	237	233	Cable-Wireless Panama
6471	258	228	ENTEL CHILE S.A.
14117	242	228	Telefonica del Sur S.A.
6147	229	205	Telefonica Del Peru
10481	181	173	Prima S.A.
20299	181	157	NEWCOM AMERICAS
10620	164	142	TVCABLE BOGOTA
14522	147	142	SatNet S.A.
18822	146	140	TELEFONICA MANQUEHUE
6503	213	129	AVANTEL, S.A.
7910	149	116	ANDINET ON LINE
7738	138	116	Telecomunicacoes da Bahia S.A
19169	138	112	Telconet
8163	118	112	METROTEL REDES S.A.

EU & Middle East Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
24863	339	308	LINKdotNET AS number
5416	254	242	BATELCO-BH
20858	199	195	EgyNet
12479	185	179	Uni2 Autonomous System
8551	206	176	Bezeq International
3352	195	159	Ibernet, Internet Access Netw
9121	172	149	TTnet Autonomous System
3215	239	147	France Telecom Transpac
3269	215	143	TELECOM ITALIA
5486	159	141	Euronet Digital Communication
6830	169	130	UPC Distribution Services
9116	143	119	Goldenlines main autonomous s
30890	205	112	SC Kappa Invexim SRL
8866	132	110	Bulgarian Telecommunication C
702	522	106	UUNET - Commercial IP service
9051	148	101	INCONET Autonomous System
3300	183	95	AUCS Communications Services
15471	148	90	SNR - Societatea Nationala de
12302	122	89	MobiFon S.A.
8584	109	89	Barak AS



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