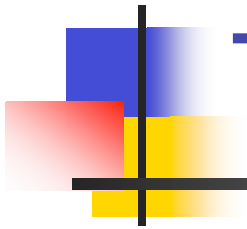

BGP Aggregation & The Deaggregation Report

Philip Smith

APRICOT 2007
Bali, Indonesia

Route Aggregation Recommendations

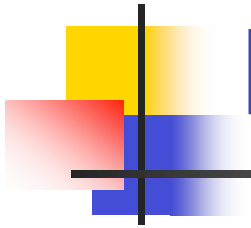

- LINX attempted aggregation policy for members
  - It failed even though most members voted for policy
- RIPE Routing Working Group work item from early 2006
  - Based on early LINX concept
  - Authored by Philip Smith, Mike Hughes (LINX) and Rob Evans (UKERNA)

Route Aggregation Recommendations

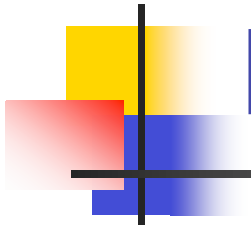

- RIPE Document — RIPE-399
  - http://www.ripe.net/ripe/docs/ripe-399.html
- Discusses:
  - History of aggregation
  - Causes of de-aggregation
  - Impacts on global routing system
  - Available Solutions
  - Recommendations for ISPs

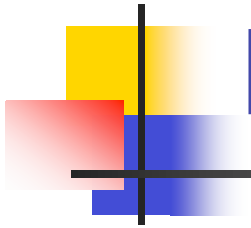

History:

- 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

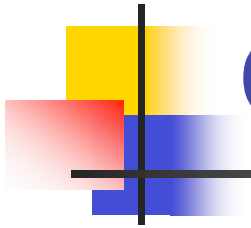

Deaggregation: Claimed causes (1):

- 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"

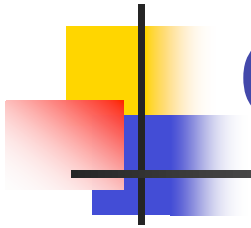

Deaggregation: Claimed causes (2):

- 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

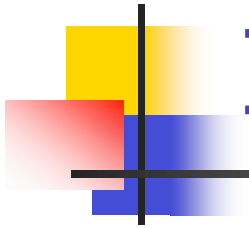

Impacts (1):

- 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

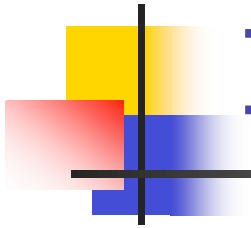

Impacts (2):

- 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

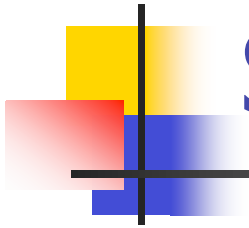

Solutions (1):

- 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”

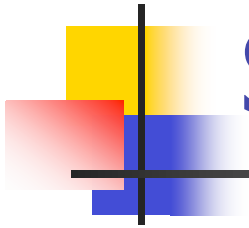

Solutions (2):

- 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

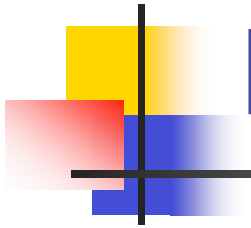

RIPE-399 Recommendations:

- 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)

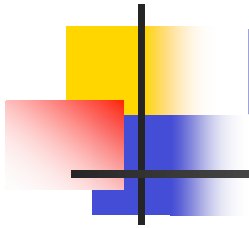

Looking at Deaggregation

- 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

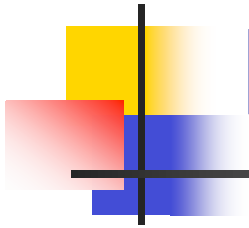

Deaggregation Factor

- 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"

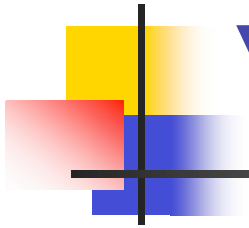

“Original Internet” — 2007/02

Total Prefixes

- Global BGP Table
  - 213k prefixes
- North America
  - 104k prefixes
- Europe & Middle East
  - 44k prefixes

Deaggregation Factor

- Global Average
  - 1.86
- North America
  - 1.70
- Europe & Middle East
  - 1.53

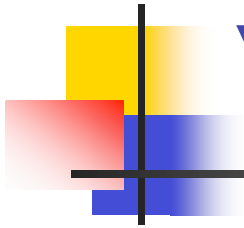

“Newer Internet” — 2007/02

Total Prefixes

- Global BGP Table
  - 210k prefixes
- Asia & Pacific
  - 48k prefixes
- Africa
  - 3k prefixes
- Latin America & Caribbean
  - 14k prefixes

Deaggregation Factor

- Global Average
  - 1.85
- Asia & Pacific
  - 2.49
- Africa
  - 2.67
- Latin America & Caribbean
  - 3.47

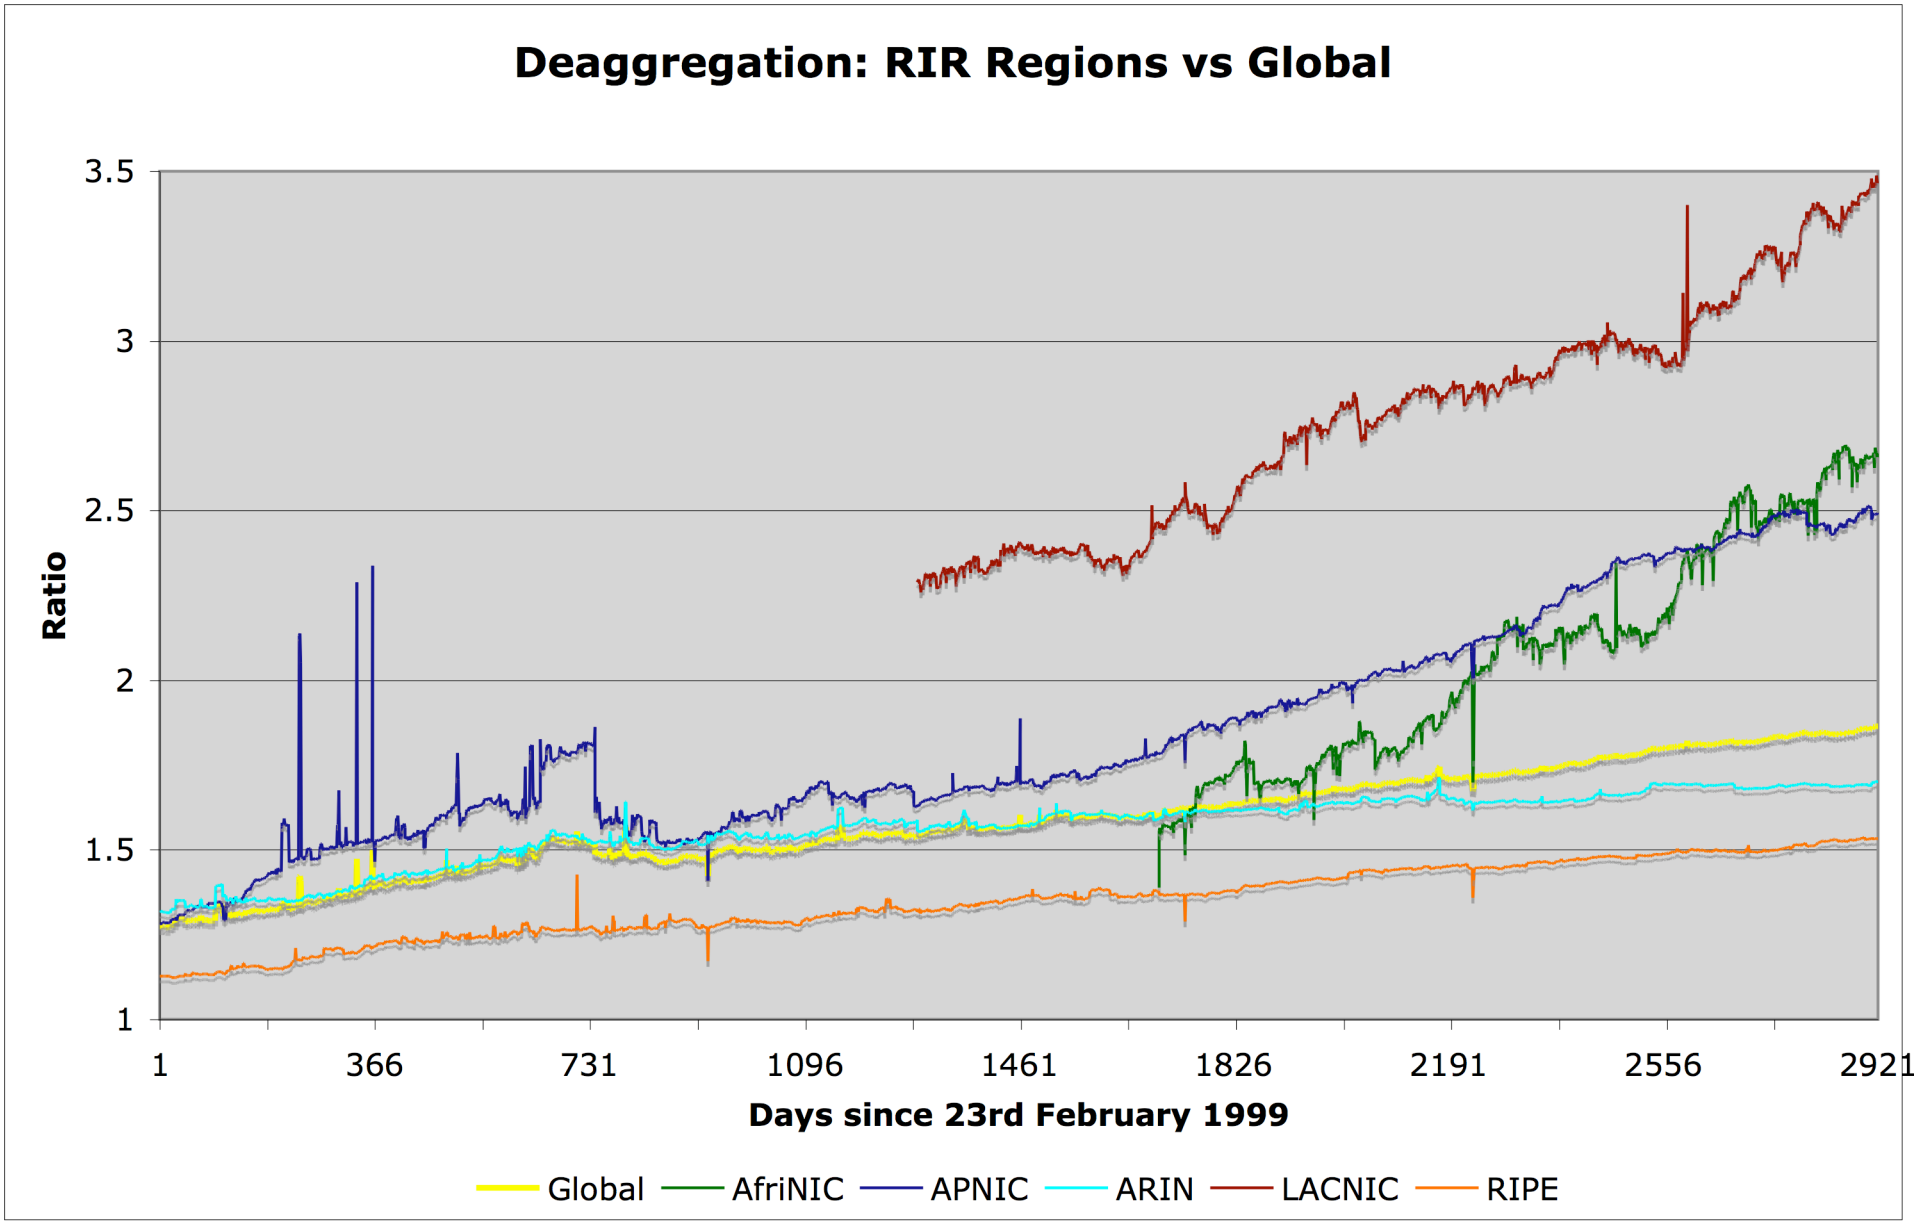
Deaggregation: RIR Regions vs Global
Ratio
3.5
3
2.5
2
1.5
1
1
366
731
1096
1461
1826
2191
2556
2921
Days since 23rd February 1999
Global
AfriNIC
APNIC
ARIN
LACNIC
RIPE

Africa Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
8452	244	238	TEDATA
15475	155	151	Nile Online
6713	144	133	Itissalat Al-MAGHRIB
5536	123	108	Internet Egypt Network
33783	111	106	EEPAD TISP TELECOM & INTERNET
24835	86	80	RAYA Telecom - Egypt
2561	64	62	Egyptian Universities Network
3741	290	58	The Internet Solution
15706	55	51	Sudatel Internet Exchange Aut
23889	50	37	MAURITIUS TELECOM
33766	33	32	Nyala Communications Pty Ltd
2905	121	27	The Internetworking Company o
2018	138	25	Tertiary Education Network
8524	30	24	AUCEGYPT Autonomous System
12455	27	24	Jambonet Autonomous system
21280	27	23	Swift Global Kenya Ltd.Is an
33776	27	22	Starcomms Nigeria Limited
33774	42	20	AS Number for Telecom Algeria
21491	22	20	UTL On-line is RF broadband I
15804	18	17	AS of The Way Out Internet So

Asia & Pacific Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
4755	1123	1046	Videsh Sanchar Nigam Ltd. Aut
4134	1275	998	CHINANET-BACKBONE
9498	970	904	BHARTI BT INTERNET LTD.
9583	1076	708	Sify Limited
17488	598	574	Hathway IP Over Cable Interne
7545	587	512	TPG Internet Pty Ltd
18101	528	505	Reliance Infocom Ltd Internet
4766	766	457	Korea Telecom (KIX)
4668	461	451	LG-EDS Systems Inc.
17676	503	438	Softbank BB Corp.
4812	500	427	China Telecom (Shanghai)
9443	444	369	Primus Telecommunications
17974	338	324	PT TELEKOMUNIKASI INDONESIA
9829	315	300	BSNL National Backbone
10139	300	293	Meridian Telekoms
17849	370	276	Telecommunications Technology
9929	287	250	China Netcom Corp.
4780	271	242	Digital United Inc.
9394	248	241	CHINA RAILWAY Internet(CRNET)
9800	278	225	CHINA UNICOM

North America Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
18566	992	983	Covad Communications
4323	1333	982	Time Warner Telecom
11492	981	965	Cable One
6478	1124	930	AT&T Worldnet Services
22773	732	691	Cox Communications, Inc.
19262	757	579	Verizon Global Networks
5668	571	555	CenturyTel Internet Holdings,
6197	1025	525	BellSouth Network Solutions,
19916	568	514	OLM LLC
855	566	496	Canadian Research Network
15270	507	473	PaeTec.net -a division of Pae
7029	511	434	Alltel Information Services,
3602	522	422	Sprint Canada, Inc.
33588	432	406	Bresnan Communications, LLC.
20115	801	380	Charter Communications
6517	414	378	Yipes Communications, Inc.
2386	1112	375	AT&T Data Communications Serv
721	633	359	DLA Systems Automation Center
11139	353	327	Cable & Wireless Dominica
7011	799	324	Citizens Utilities

Latin America Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
8151	1042	831	UniNet S.A. de C.V.
11830	482	463	Instituto Costarricense de El
16814	329	321	NSS, S.A.
22047	307	296	VTR PUNTO NET S.A.
11172	348	263	Servicios Alestra S.A de C.V
11556	237	233	Cable-Wireless Panama
14117	239	226	Telefonica del Sur S.A.
6471	253	223	ENTEL CHILE S.A.
7303	251	221	Telecom Argentina Stet-France
6147	228	204	Telefonica Del Peru
10481	181	173	Prima S.A.
20299	174	149	NEWCOM AMERICAS
18822	145	135	TELEFONICA MANQUEHUE
14522	140	131	SatNet S.A.
7910	152	120	ANDINET ON LINE
10620	134	119	TVCABLE BOGOTA
19169	143	118	Telconet
6503	199	116	AVANTEL, S.A.
23216	153	111	RAMtelecom Telecomunicaciones
19429	146	109	E.T.B.

EU & Middle East Aggregation Savings Summary

ASN	No of Nets	Poss Savings	Description
24863	328	306	LINKdotNET AS number
5416	219	204	BATELCO-BH
20858	197	194	EgyNet
8551	204	176	Bezeq International
12479	177	171	Uni2 Autonomous System
3352	201	168	Ibernet, Internet Access Netw
3215	239	147	France Telecom Transpac
9121	172	147	TTnet Autonomous System
3269	213	142	TELECOM ITALIA
5486	158	140	Euronet Digital Communication
9116	161	137	Goldenlines main autonomous s
6830	169	130	UPC Distribution Services
702	540	122	UUNET - Commercial IP service
12715	140	113	Jazz Telecom S.A.
30890	207	110	SC Kappa Invexim SRL
8866	127	107	Bulgarian Telecommunication C
9051	146	96	INCONET Autonomous System
3300	182	93	AUCS Communications Services
12302	122	89	MobiFon S.A.
25233	105	88	Awalnet

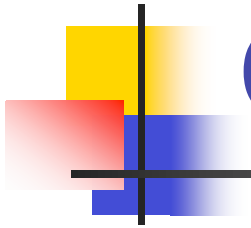

Observations

- 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

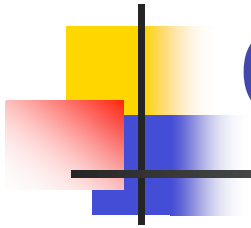

Conclusion

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