

Analysis of Propagation of Regular, Extended, and Large BGP Communities

Lilia Hannachi and Kotikalapudi Sriram
lilia.hannachi@nist.gov ksriram@nist.gov

GROW WG Meeting
IETF 111
July 2021

Acknowledgements: The authors wish to thank Doug Montgomery, Oliver Borchert, Alexander Azimov, Jeff Haas, Jakob Heitz, Gyan Mishra, and Brian Dickson for comments/suggestions.

Abbreviations

RC = Regular Community

EC = Extended Community

LC = Large Community

Motivation for this Study

- There are applications in IETF drafts for which LC/EC transitivity is much needed (at least over a few hops), e.g.,
<https://datatracker.ietf.org/doc/html/draft-ietf-grow-route-leak-detection-mitigation>
- Can they get the transitivity of LC/EC that they need?
- There was an extensive discussion on the IDR/GROW WG lists about the propagation behaviors of LC/EC
<https://mailarchive.ietf.org/arch/browse/grow/?gbt=1&index=152EJ38HSvfapwg4F0SqFii69RM>
- This study focuses on measuring the propagation of EC and EC

Prior Work on Regular Communities

Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush, “BGP Communities: Even more Worms in the Routing Can,” Proceedings of IMC ’18. ACM, New York, NY, USA.

<https://dl.acm.org/doi/10.1145/3278532.3278557>

Vasileios Giotsas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger, “Inferring BGP Blackholing Activity in the Internet,” Proceedings of ACM Internet Measurements Conference, London, UK, November 2017 (IMC’17). <https://dl.acm.org/doi/pdf/10.1145/3131365.3131379>

Some Basics about the Measurement Methodology

- Using Routeviews and RIPE-RIS data
- Considering only unique {prefix, AS path, RC/LC/EC} combinations
- Conservative in the way # AS hops for propagation of RC/LC/EC is measured:
 - Example: AS Path: 202365 39533 3491 3356 21277 212573 208293 56484 205169, RC: 3491:3000
 - Assume RC 3491:3000 was added by AS 3941 but it may have been added by an earlier AS in the path
 - The UPDATE could have propagated faster than this AS Path and the collectors may not have gathered that data because of where the collectors' peers are
 - When the RC/LC/EC community ID (AS number) is not found in the associated AS path, the above measurement is skipped
- So, the estimates of RC/LC/EC propagation (#hops) are low estimates

Examples of UPDATES with RC and LC

- [TIME: 06/08/21 02:00:25, TYPE: TABLE_DUMP_V2/IPV4_UNICAST, PREFIX: 185.227.239.0/24, SEQUENCE: 654281, FROM: 194.50.19.4 AS202365, ORIGINATED: 06/03/21 08:49:23, ORIGIN: IGP, AS_PATH: 202365 39533 3491 3356 21277 212573 208293 56484 205169, NEXT_HOP: 194.50.19.4, COMMUNITY: 0:39533 3491:3000 3491:3011 3491:9002 39533:49666, LARGE_COMMUNITY: 3356:1:41 202365:2020:202409]

3356 is Level 3; 3491 is PCCW Global

- [TIME: 06/08/21 02:00:24, TYPE: TABLE_DUMP_V2/IPV4_UNICAST, PREFIX: 185.104.158.0/24, SEQUENCE: 638182, FROM: 208.94.118.10 AS40630, ORIGINATED: 06/05/21 00:04:44, ORIGIN: IGP, AS_PATH: 40630 2914 3356 21277 212573 208293 56484 200485, NEXT_HOP: 208.94.118.10, COMMUNITY: 2914:420 2914:1004 2914:2000 2914:3000 40630:100 40630:11701, LARGE_COMMUNITY: 3356:1:41]

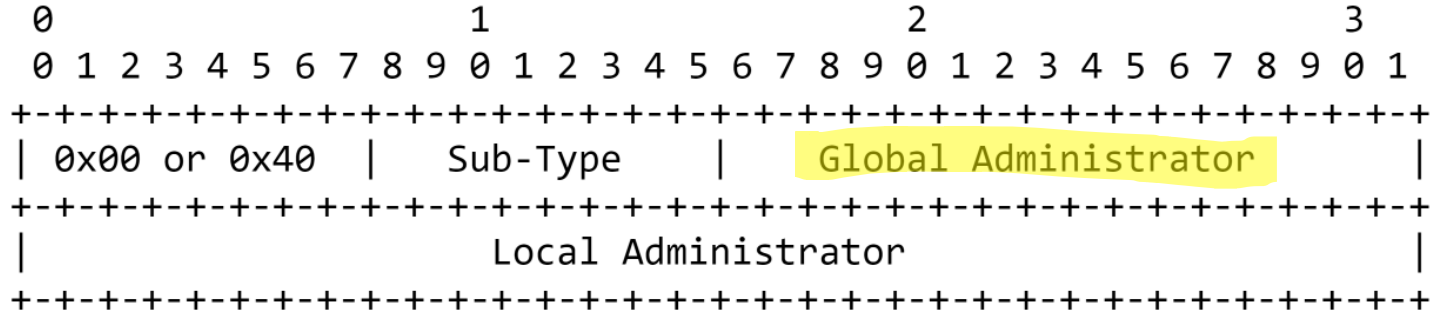
3356 is Level 3; 2914 is NTT

There is evidence Tier 1 Transit providers propagate RC and LC

Finding 2-Octet and 4-Octet AS in the Extended Community

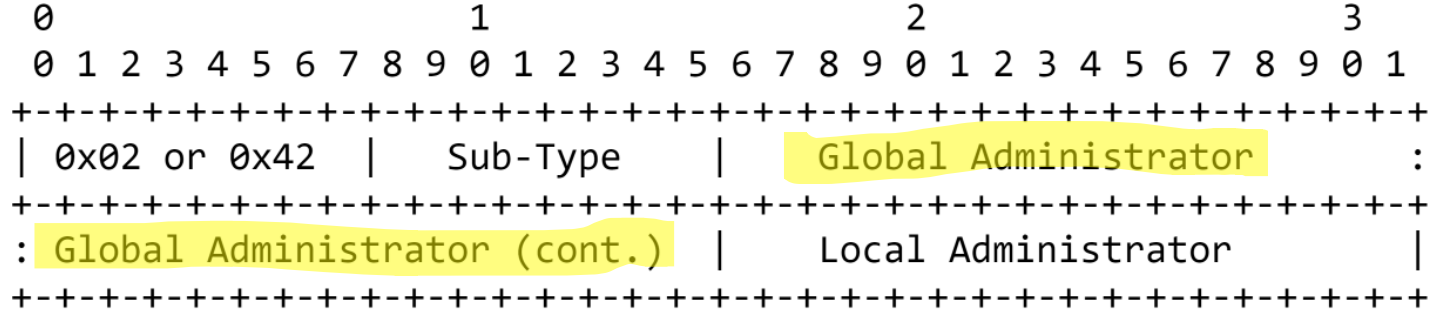
0x00	Transitive Two-Octet AS-Specific Extended Community (Sub-Types are defined in the "Transitive Two-octet AS-Specific Extended Community Sub-Types" registry)	[RFC7153]
0x02	Transitive Four-Octet AS-Specific Extended Community (Sub-Types are defined in the "Transitive Four-octet AS-Specific Extended Community Sub-Types" registry)	[RFC7153]
0x40	Non-Transitive Two-Octet AS-Specific Extended Community (Sub-Types are defined in the "Non-Transitive Two-octet AS-Specific Extended Community Sub-Types" registry)	[RFC7153]
0x42	Non-Transitive Four-Octet AS-Specific Extended Community (Sub-Types are defined in the "Non-Transitive Four-octet AS-Specific Extended Community Sub-Types" registry)	[RFC7153]

Location of the 2-octet AS number for EC Types 0x00 and 0x40



Example:
RFC 4360

Location of the 4-octet AS number for EC Types 0x02 and 0x42



Example:
RFC 5668

Example of an UPDATE with EC

[TIME: 07/15/21 01:30:00, TYPE: BGP4MP/MESSAGE/Update, FROM: 2001:7f8:13::a520:5206:1 AS205206, TO: 2001:7f8:13::a501:2654:1 AS12654, ORIGIN: IGP, **ASPATH**: 205206 33891 **59947** 212705, **UNKNOWN_ATTR(192, 16, 8)**: 00 03 **ea 2b** 00 00 ea 2b, MP_REACH_NLRI(IPv6 Unicast), NEXT_HOP: 2001:7f8:13::a520:5206:1, NEXT_HOP: fe80::8a90:900:772:7446, COMMUNITY: 33891:33893 33891:33897 33891:40001 59947:400 59947:59947, ANNOUNCE, 2602:feda:b70::/48]

EC = **UNKNOWN_ATTR(192, 16, 8)**: 00 03 **ea 2b** 00 00 ea 2b

Global Administrator ID = **ea 2b** = **ASN 59947**

ASPATH: 205206 33891 **59947** 212705

The EC has propagated at least two hops

Some Notes about Regular Communities

Community name	Hex	ASN:X notation
Internet	0x00000000	0:0
BLACKHOLE	0xFFFF029A	65535:666
NO_EXPORT	0xFFFFFFFF01	65535:65281
NO_ADVERTISE	0xFFFFFFFF02	65535:65282
NO_EXPORT_SUBCONFED	0xFFFFFFFF03	65535:65283
NOPEER	0xFFFFFFFF04	65535:65284
Wide range for private use	0x00010000 to 0xFFFEFFFF	1:0 to 65534:65535

Not seen in the measurements

Did not look exhaustively for these

- Zero or 65535 in the first 2 octets does not mean it is an AS number; it is a WKC
- 0:0 is a WKC called “Internet” and the Update should propagate Internet-wide
- Blackhole Community may be encoded in two ways: Example 2914:666 (using the administrator’s AS number) or 65535:666 (using the WKC)

Blackhole Community: AS Distance and Propagation

- AS Distance is the distance in # hops the Blackholed prefix (client) is from the RTBH provider
- AS distance 0 means that the prefix is originated from the same ISP/IXP that is the RTBH provider
- Example: AS path = 25160 3356 12956 6147 and RC = 3356:666
 - This means that the client is at AS 6147 (origin AS) and AS 3356 is the RTBH provider
 - AS Distance to RTBH provider = 2
 - Propagation (#hops): The Blackhole Community propagated 3 hops in this case (AS 6147 to AS 25160)

Blackhole Community Contributes Differently to RC Transitivity Counts for ASes

- For a Blackhole Community, the propagation happens across the whole AS path
 - Example: AS path = 25160 3356 12956 6147 and RC = 3356:666
 - The RC propagation count for each of the ASes shown in the path increments by 1 except the left-most (i.e., AS 25160)
 - The Community ID (C-ID) value in the RC may be 65535 (WKC for BLACKHOLE)
- When it is some other Community type, we compare the ASN in the C-ID with ASNs in the path
 - Example: AS path = 207910 204092 30781 3356 20940 and RC = 3356:903
 - The RC counts for only ASes 204092 30781 3356 increment by 1

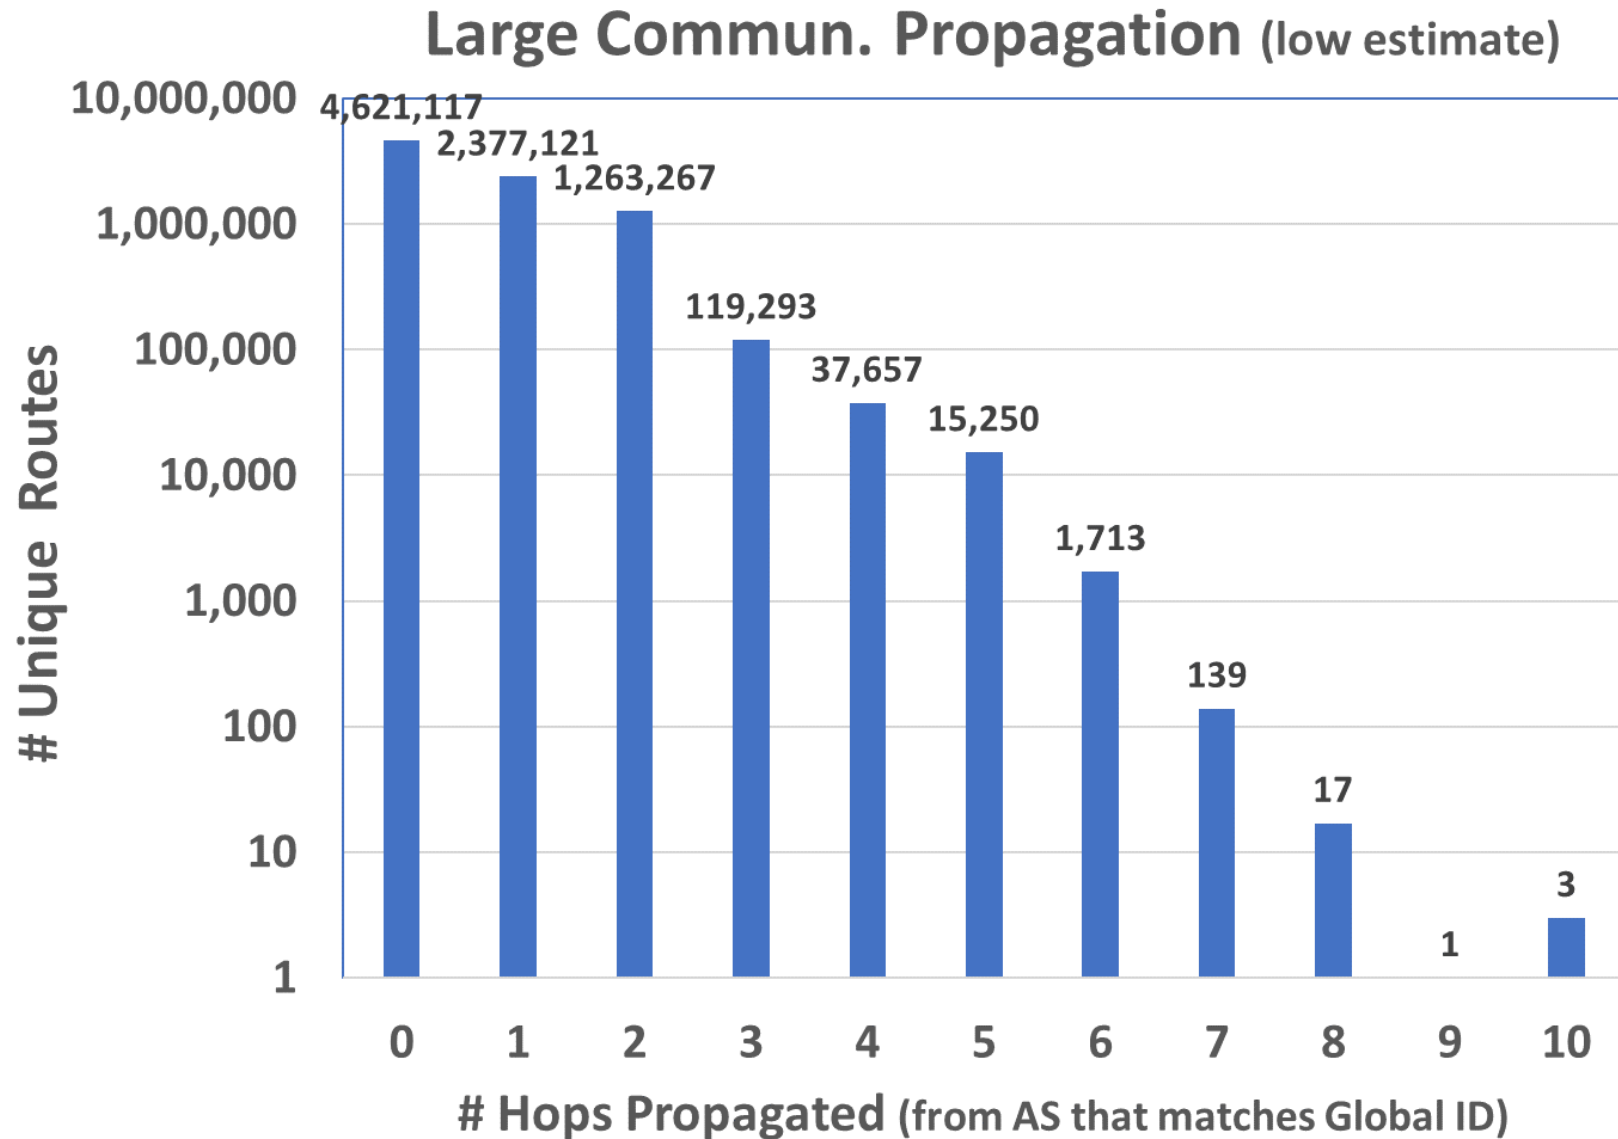
Note: The second bullet also applies to LC/EC Transitivity Counts for ASes

Propagation Analysis Using the RIPE-RIS UPDATE Data

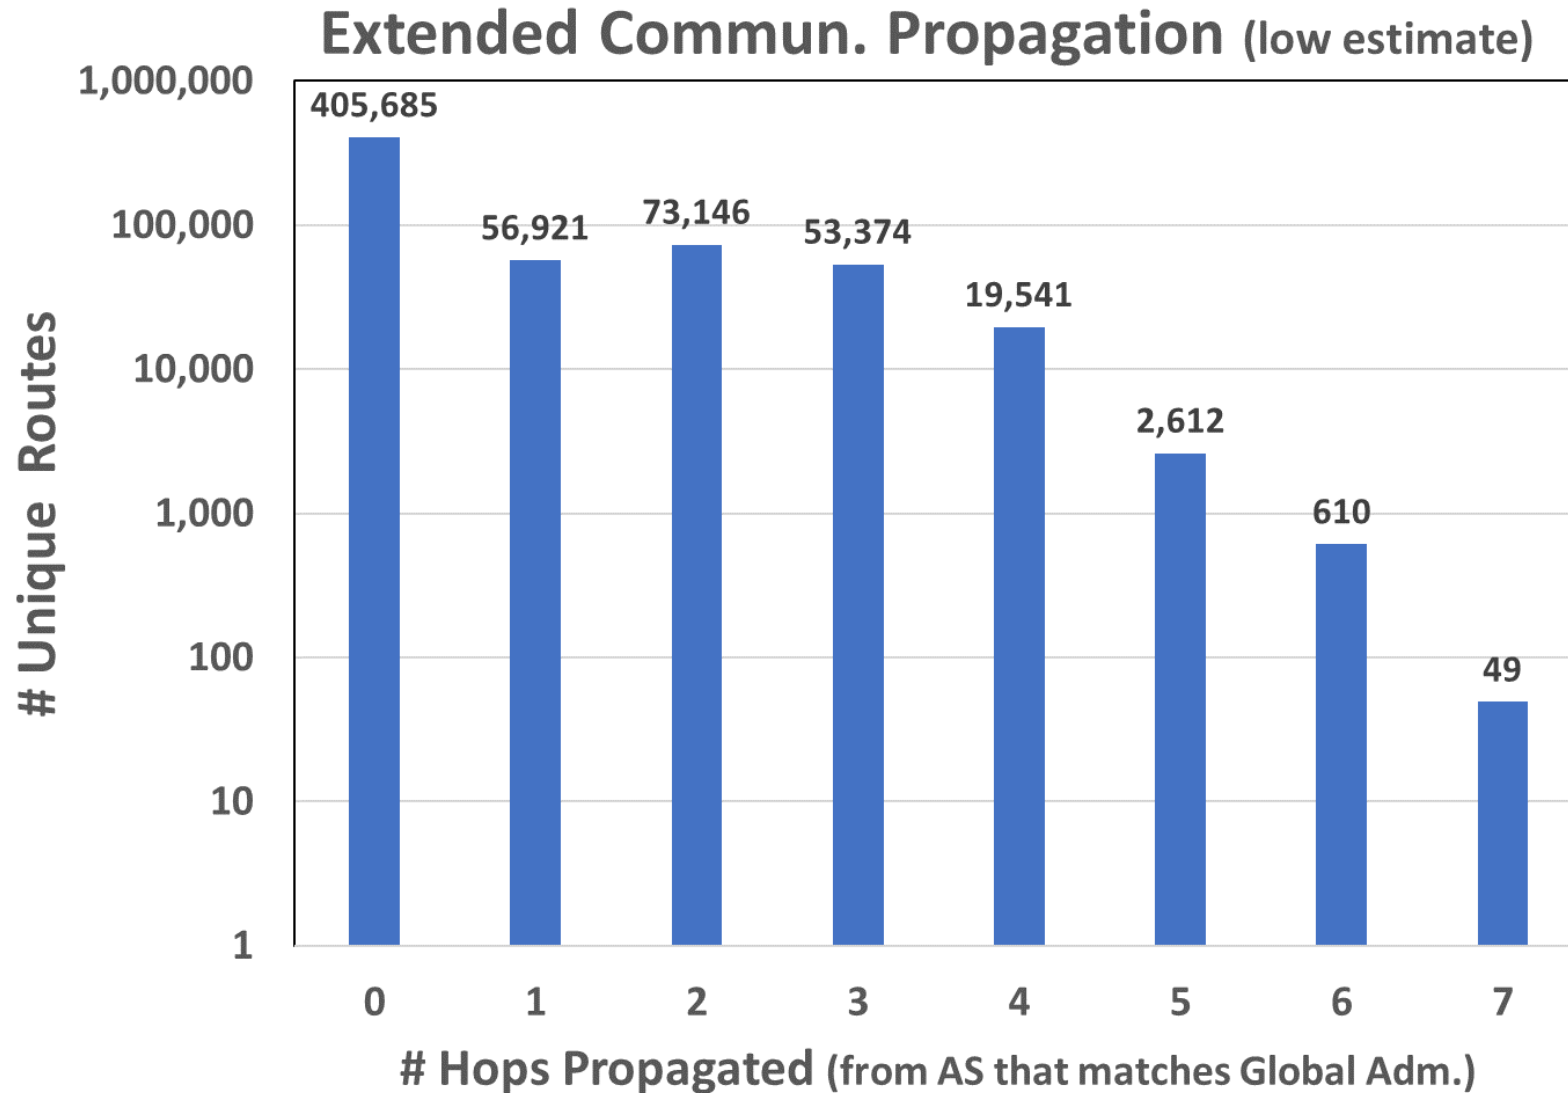
Collector : rrc03 From 2021-07-15 00:00 To 2021-07-15 23:55 (one whole day)

Total # Updates (raw)	52,765,351
# Unique {Prefix, Path, RC}	69,643,850
# Unique {Prefix, Path, LC}	12,802,541
# Unique {Prefix, Path, EC}	1,143,809
# Unique {Prefix, Path, Transitive EC}	925,082
# Unique {Prefix, Path, Transitive 2-Octet AS EC}	826,575
# Unique {Prefix, Path, Transitive 4-Octet AS EC}	98,507
# Unique {Prefix, Path, RC = Any:666}	264,557
# Unique {Prefix, Path, RC = 65535:666}	21
# Unique {Prefix, Path, RC = 0:Any}	10,293,308
# Unique {Prefix, Path, RC = Any:65535}	195,462
# Unique {Prefix, Path, RC = 0:0}	18,583
# routes that have LC with Global ID = 0	12

Propagation of Large Communities




Propagation of Extended Communities



- A few peers of the collector send large numbers of routes with EC that propagated zero hops. The same is true for RC and LC also.
- E.g., the 405,685 routes at 0 hops are all from one AS

Extended Community Frequency by Type



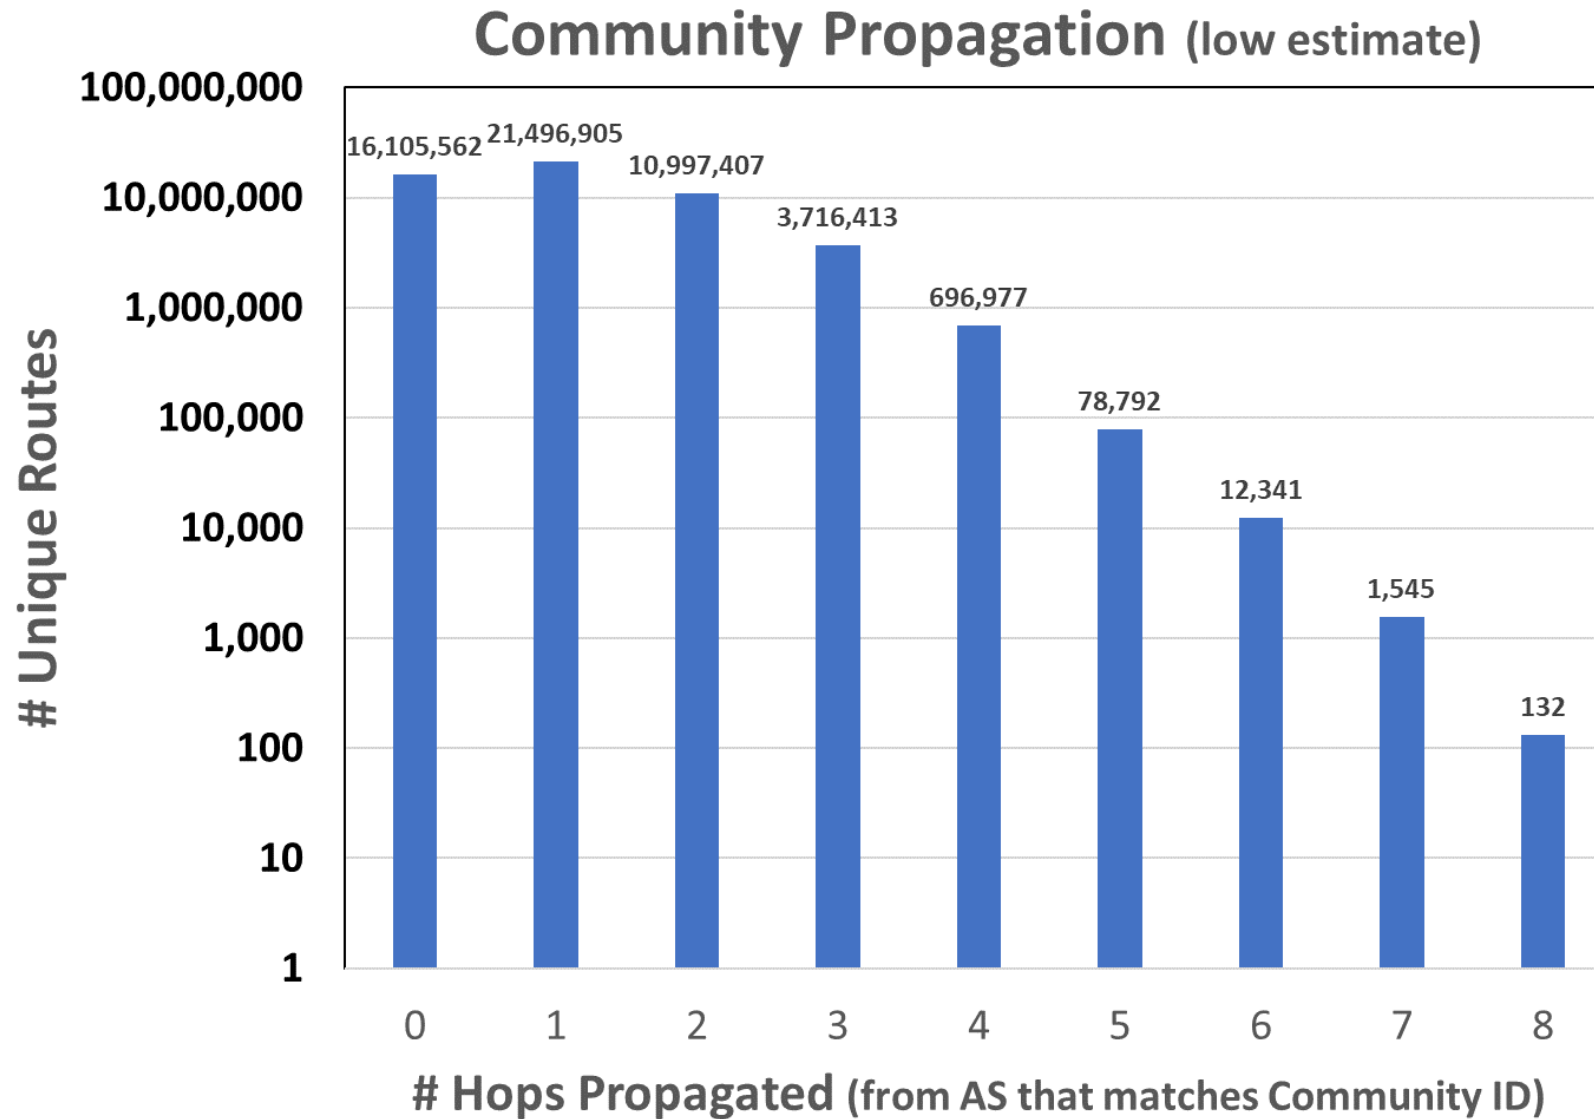
- Only the Transitive ECs are seen propagating on the Internet.
- Non-transitive ECs are not seen 😊

Count of ASes that Propagate RC, LC, EC


# ASes that Propagate RC	3,298	
# ASes that Propagate LC	262	
# ASes that Propagate EC	527	
# ASes that Propagate RC, LC, or EC	3,872	
Percentage of Transit ASes that Propagate RC, LC, or EC	32%	
# Tier 1 / Major ISPs that Propagate RC	22	out of 25
# Tier 1 / Major ISPs that Propagate LC	3	*
# Tier 1 / Major ISPs that Propagate EC	2	*
# Tier 1 / Major ISPs that Propagate RC, LC, or EC	22	out of 25

* Hit or miss depending on the date of the data collection

Propagation of Regular Communities



Blackhole Communities - ASN:666



- Includes 65535:666 (WKC)

Top 25 ASes that Propagate RC/LC/EC

Units are # Unique {prefix, AS path, and RC or LC or EC}

Ordered according to RC

	Top 25 ASNs	Tier 1 (Y/N)	RC	LC	EC	Total (RC+LC+EC)
1	34927	N	3,668,623	0	0	3,668,623
2	2914	Y	3,656,266	0	0	3,656,266
3	3356	Y	3,206,724	0	0	3,206,724
4	1299	Y	2,845,117	0	0	2,845,117
5	3491	Y	2,746,308	0	0	2,746,308
6	174	Y	2,072,735	0	0	2,072,735
7	6762	Y	1,781,164	0	0	1,781,164
8	3257	Y	1,592,597	0	179	1,592,776
9	6453	Y	1,515,968	0	0	1,515,968
10	8943	N	920,784	0	0	920,784
11	6830	Y	867,892	0	0	867,892
12	33891	N	764,024	0	0	764,024
13	20495	N	684,331	693,953	0	1,378,284
14	39533	N	628,885	0	0	628,885
15	57866	N	550,573	0	0	550,573
16	9002	Y	403,464	0	0	403,464
17	13030	N	384,532	0	0	384,532
18	58299	N	372,127	0	0	372,127
19	6461	Y	360,749	0	0	360,749
20	3320	Y	339,234	0	0	339,234
21	12956	Y	315,732	0	0	315,732
22	9498	N	305,893	0	0	305,893
23	2603	N	281,690	0	5,904	287,594
24	6939	Y	187,961	0	0	187,961
25	37468	N	177,652	0	8,829	186,481

Ordered according to LC

	Top 25 ASNs	Tier 1 (Y/N)	RC	LC	EC	Total (RC+LC+EC)
	57463	N	27,403	2,685,415	0	2,712,818
	20495	N	684,331	693,953	0	1,378,284
	204092	N	0	186,756	0	186,756
	20473	N	78,996	130,121	0	209,117
	34872	N	0	59,781	0	59,781
	24961	N	108,090	7,350	0	115,440
	31898	N	2,736	5,720	0	8,456
	140731	N	0	4,811	0	4,811
	44684	N	0	4,675	0	4,675
	198949	N	0	4,467	0	4,467
	49832	N	1,273	3,220	0	4,493
	25682	N	0	2,272	0	2,272
	206499	N	0	2,250	0	2,250
	48646	N	360	1,195	0	1,555
	200365	N	0	1,148	0	1,148
	54825	N	599	1,018	0	1,617
	131284	N	0	920	0	920
	202562	N	0	817	0	817
	8315	N	50	767	7	824
	131675	N	0	743	0	743
	1930	N	0	741	0	741
	13238	N	1,407	731	0	2,138
	35280	N	50,127	673	0	50,800
	213045	N	0	582	0	582
	140938	N	0	516	0	516

Ordered according to EC

	Top 25 ASNs	Tier 1 (Y/N)	RC	LC	EC	Total (RC+LC+EC)
1	264556	N	0	0	41,032	41,032
2	61568	N	123,055	0	29,247	152,302
3	43531	N	68,727	0	17,637	86,364
4	11164	N	36,069	0	12,664	48,733
5	37468	N	177,652	0	8,829	186,481
6	2603	N	281,690	0	5,904	287,594
7	8866	N	9,965	0	5,866	15,831
8	12430	N	2,938	0	5,349	8,287
9	265126	N	0	0	3,806	3,806
10	267420	N	0	0	3,775	3,775
11	267613	N	0	0	3,738	3,738
12	21277	N	872	0	2,859	3,731
13	5713	N	4,980	0	2,638	7,618
14	212708	N	0	0	2,580	2,580
15	51375	N	2,825	0	2,570	5,395
16	50710	N	362	0	2,440	2,802
17	62240	N	308	0	2,033	2,341
18	36992	N	10,377	0	1,935	12,312
19	61599	N	8,151	0	1,776	9,927
20	8717	N	75	0	1,747	1,822
21	19551	N	12,632	0	1,654	14,286
22	138840	N	0	0	1,648	1,648
23	8452	N	923	0	1,578	2,501
24	34224	N	2,042	0	1,452	3,494
25	196844	N	0	0	1,401	1,401

Conclusions, Questions, and Future Work

- The analyses show evidence that LC and EC propagate multiple hops (in addition to RC).
- Considering this, what are WG thoughts on support for use of LC/EC in applications (drafts) where transitivity is needed?
- Would the WG make any new recommendations for the operators concerning handling of LC/EC?
- Future work: We (authors) are open to suggestions for additional data analysis.


Backup slides

Propagation Analysis Using the Routeviews RIB Data

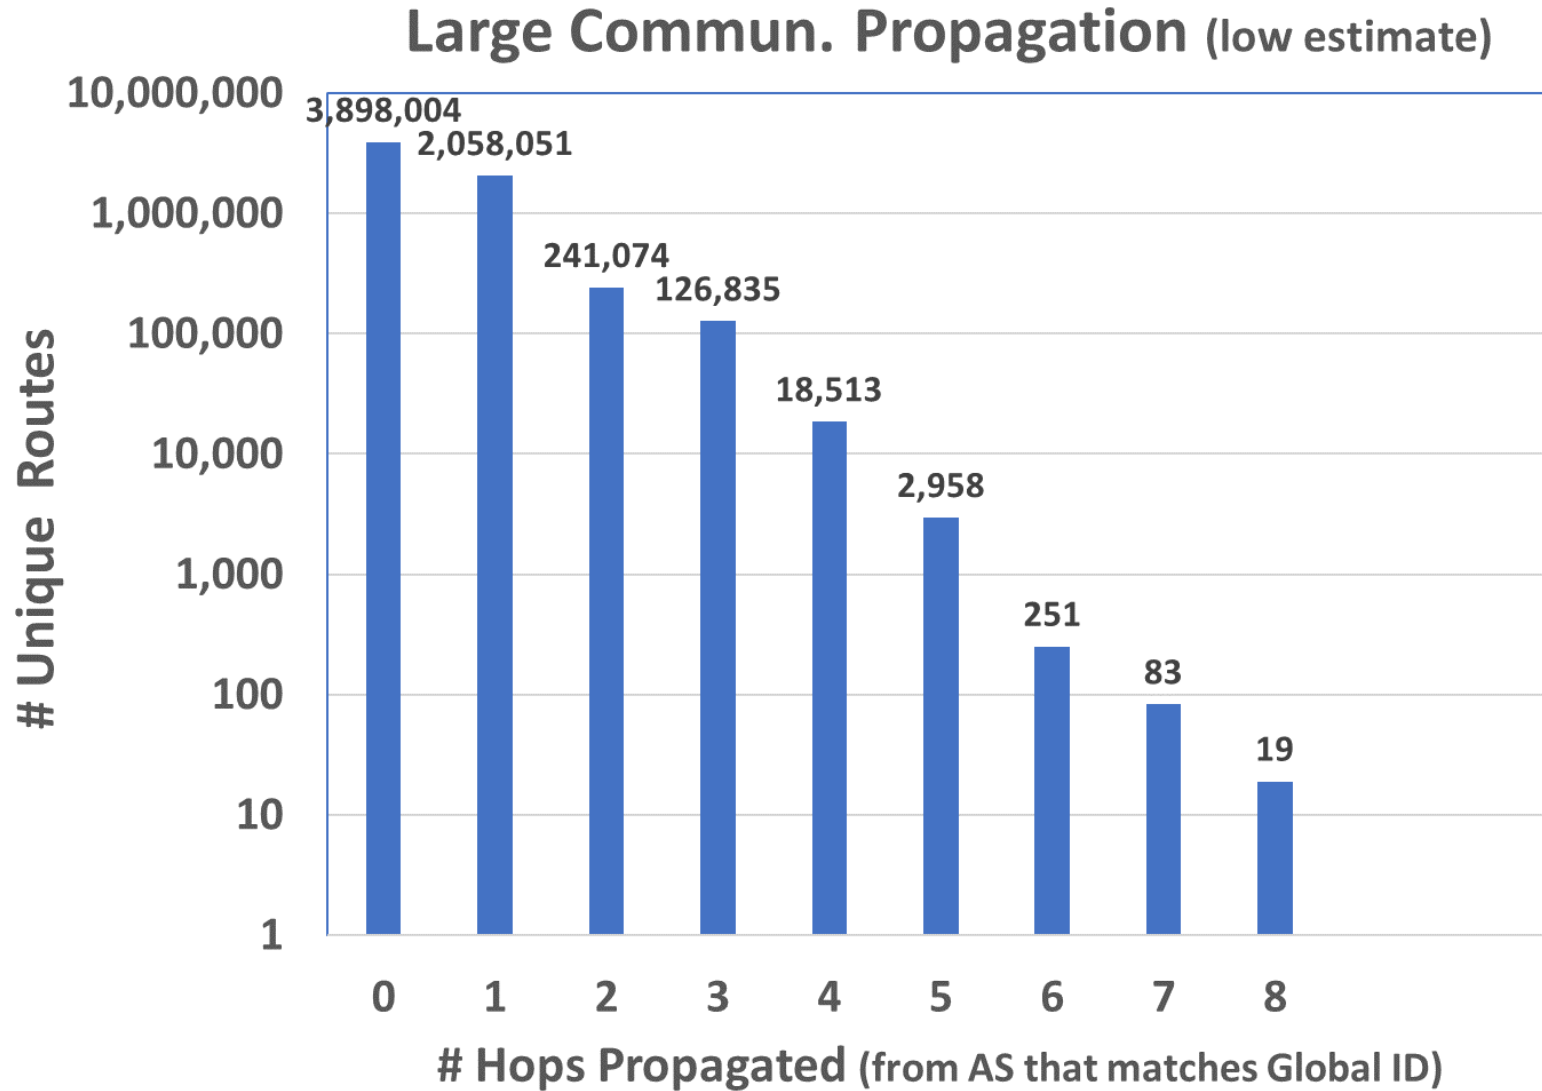
Collector : RouteViews3 From 2021-07-15 00:00 To 2021-07-15 00:00)

Total # Updates (raw)	35,590,333
# Unique {Prefix, Path, RC}	128,706,731
# Unique {Prefix, Path, LC}	12,677,829
# Unique {Prefix, Path, RC = Any:666}	1,868,810
# Unique {Prefix, Path, RC = 65535:666}	223
# Unique {Prefix, Path, RC = 0:Any}	10,025,210
# Unique {Prefix, Path, RC = Any:65535}	41,287
# Unique {Prefix, Path, RC = 0:0}	25,482
# routes that have LC with Global ID = 0	14

Propagation of Regular Communities



Propagation of Large Communities




Count of ASes that Propagate RC, LC

# ASes that Propagate RC	3,443	
# ASes that Propagate LC	270	
# ASes that Propagate RC or LC	3,661	
# Tier 1 / Major ISPs that Propagate RC	22	out of 25
# Tier 1 / Major ISPs that Propagate LC	2	*
# Tier-1 ASes that Propagate RC or LC	22	out of 25

* Hit or miss depending on the date of the data collection

Blackhole Communities - ASN:666



- Includes 65535:666 (WKC)