ORACLE

Excessive BGP AS Path Prepending is a Self-Inflicted Vulnerability

Doug Madory

Virtual LACNOG
October 2020

What is AS_PATH Prepending?

- A technique used to de-prioritize a route by artificially increasing AS_PATH length.
- “Prepending” is repeating an ASN in AS_PATH – typically to a subset of adjacent ASes.

… 3356 4192 4192 7160 208.72.91.0/24

- Assuming all other criterion are equal, BGP route selection prefers the shorter AS path length (i.e. non-prepended route).

But prepending can also be problematic

Rarely the direct cause of problems, with one notable exception:

- Feb 2009: Internet-wide outages caused by a single errant routing announcement. In this incident, AS47868 announced its one prefix with an extremely long AS path. [1,2]

- Big difference in MikroTik vs Cisco config
  - Admin entered ASN instead of prepend count
  - *47868 modulo 256 = 252 prepends*

- As AS path lengths exceeded 255, Cisco routers crashed

https://dyn.com/blog/the-flap-heard-around-the-world/

https://dyn.com/blog/longer-is-not-better/

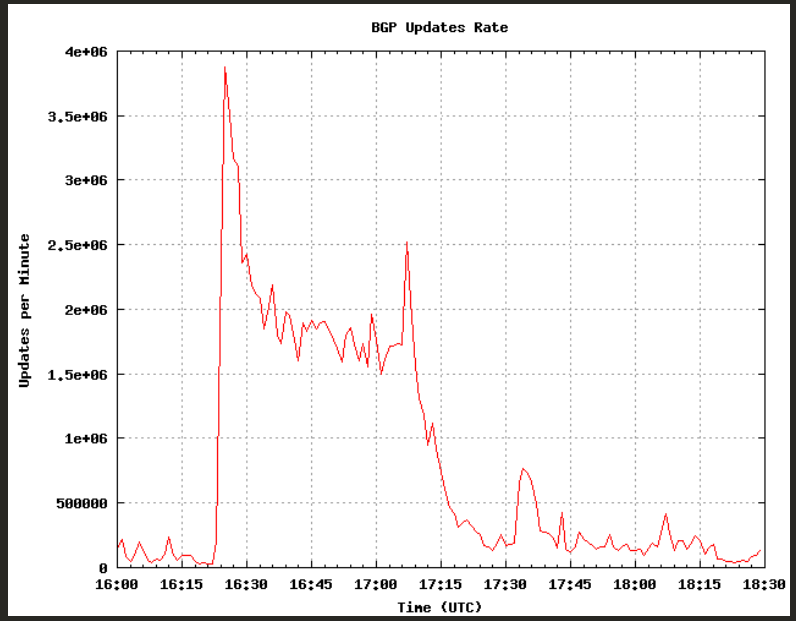

China **did not** hijack 15% of all internet traffic

- Most impact was constrained to Chinese routes.
- However, two of the top five most-propagated leaked routes were US routes!

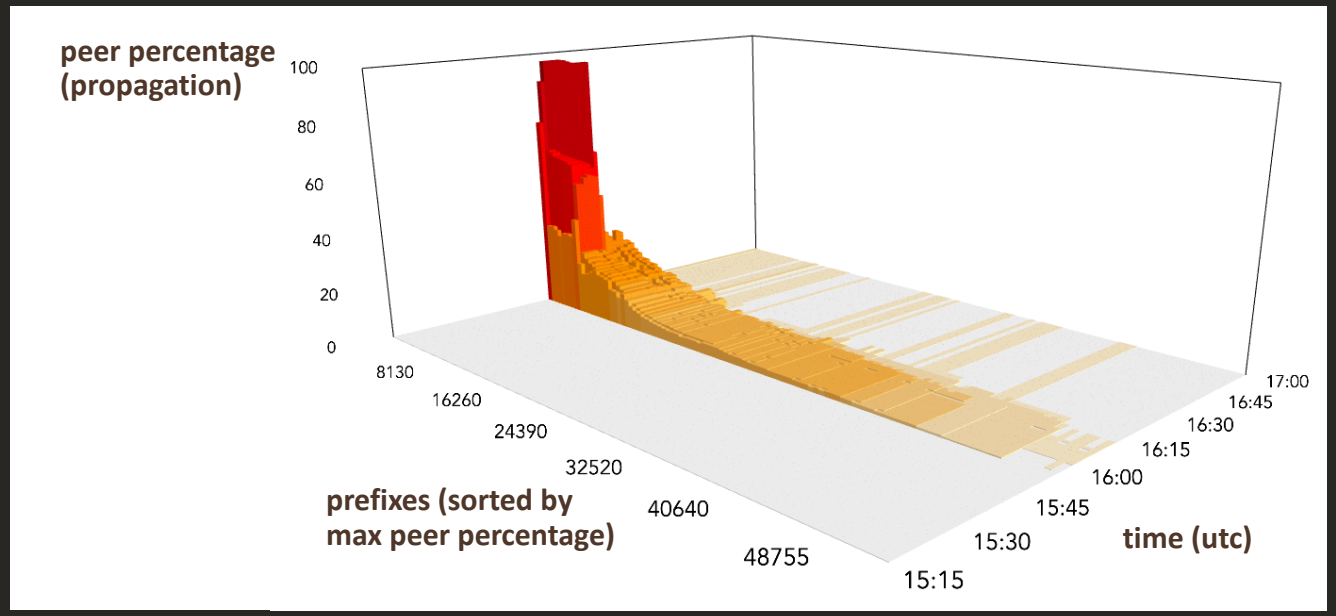

China **did not** hijack 15% of all internet traffic

- Why were two of the most-propagated leaked routes from the US?

  12.5.48.0/21 and 12.4.196.0/22 were announced to the internet along following excessively prepended AS path:

  … 3257 7795 12163 12163 12163 12163 12163 12163

- We termed this:

  ~~hijack me please~~
  ~~I hate myself~~
  prepended-to-all

Prefix	Country	Origin	Max Peer Percentage
218.30.222.0/24	CN	4134	95.58
59.42.0.0/16	CN	4134	87.91
12.4.196.0/22	US	12163	87.61
12.5.48.0/21	US	12163	87.61
59.52.0.0/14	CN	4134	87.61

Impacts of Excessive Prepending During Leaks

- Much of the worst propagation of leaked routes during big leak events were due to routes being **prepended-to-all**.

AS4671 leak of April 2014
(>320,000 prefixes)

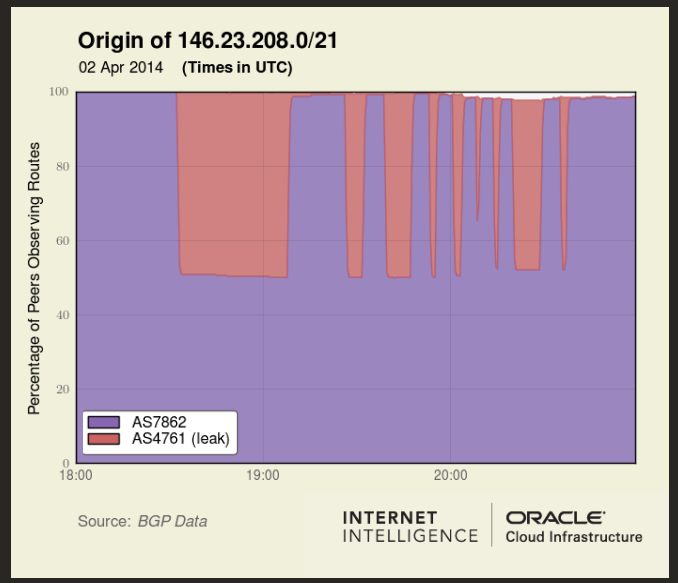

AS4788 leak of June 2015
(>260,000 prefixes)

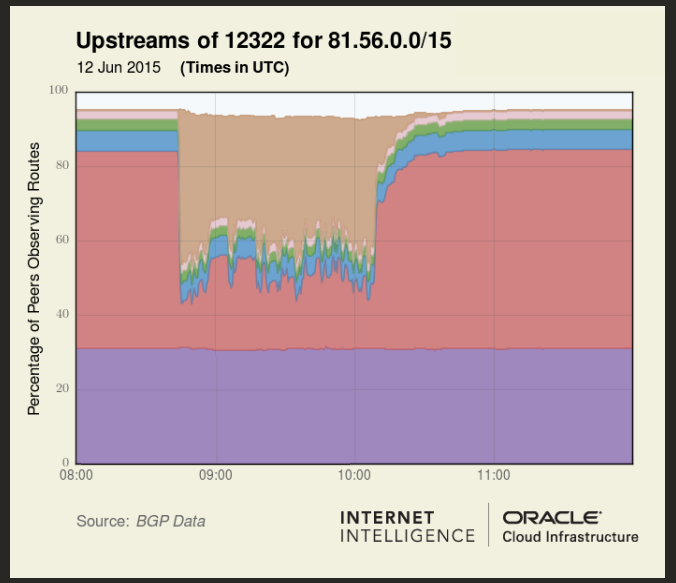

Both 146.23.208.0/21 and 82.224.0.0/12 were prepended-to-all and highlighted in blog posts.

https://dyn.com/blog/indonesia-hijacks-world/

https://dyn.com/blog/global-collateral-damage-of-tmnet-leak/

Prepending to Everyone!

- Prepended-to-all prefixes are those seen as prepended by all (or nearly all) of the ASes of the internet.
- In this configuration, prepending is no longer shaping route propagation.
- It is simply incentivizing ASes to choose *another origin* if one were to suddenly appear whether by mistake or otherwise.

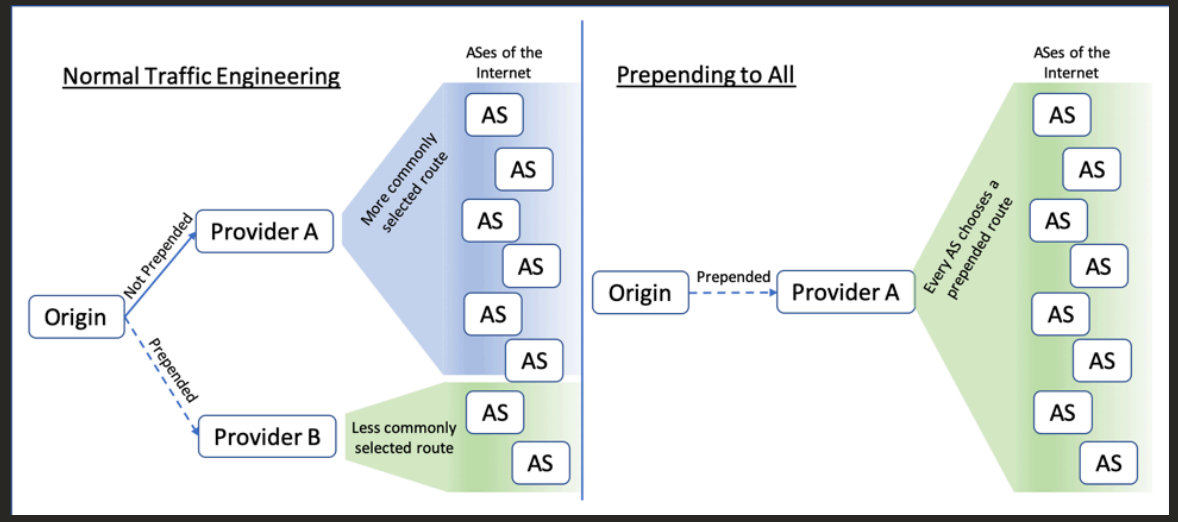

- How many prefixes are **prepended-to-all**? …a lot!

Prepending in the Global Routing Tables

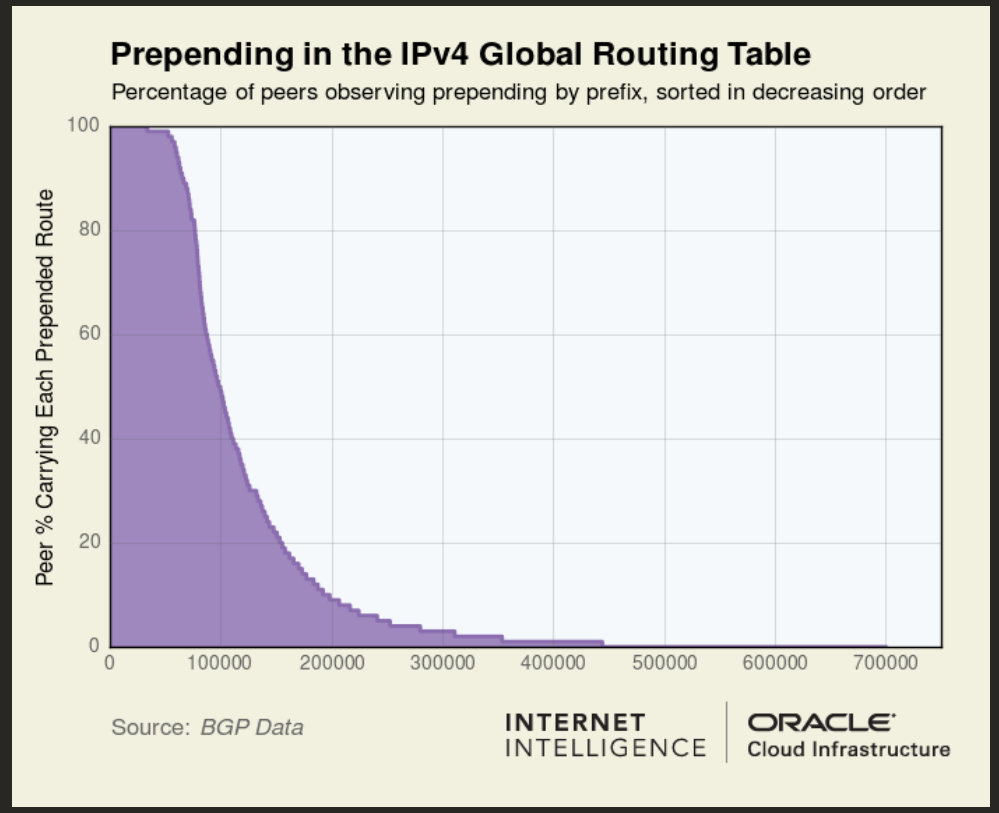

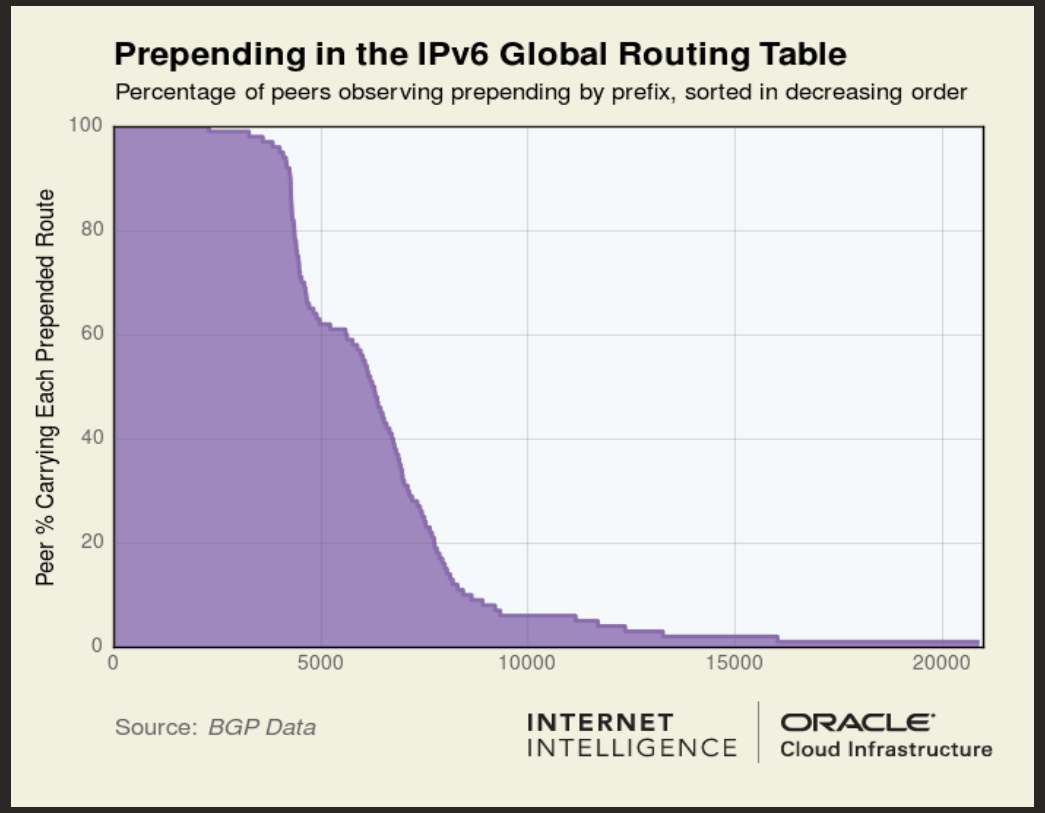

Prepending in the IPv4 Global Routing Table

- Prefixes prepended to 95%+ of ASes: >60k
  - 8% of IPv4 Global Routing Table (1/12)
  - Includes entities of every stripe: govts, banks, internet infrastructure, etc.

- Prefixes prepended to 50%+ of ASes: >100k
  - 13.3% of IPv4 Global Routing Table.

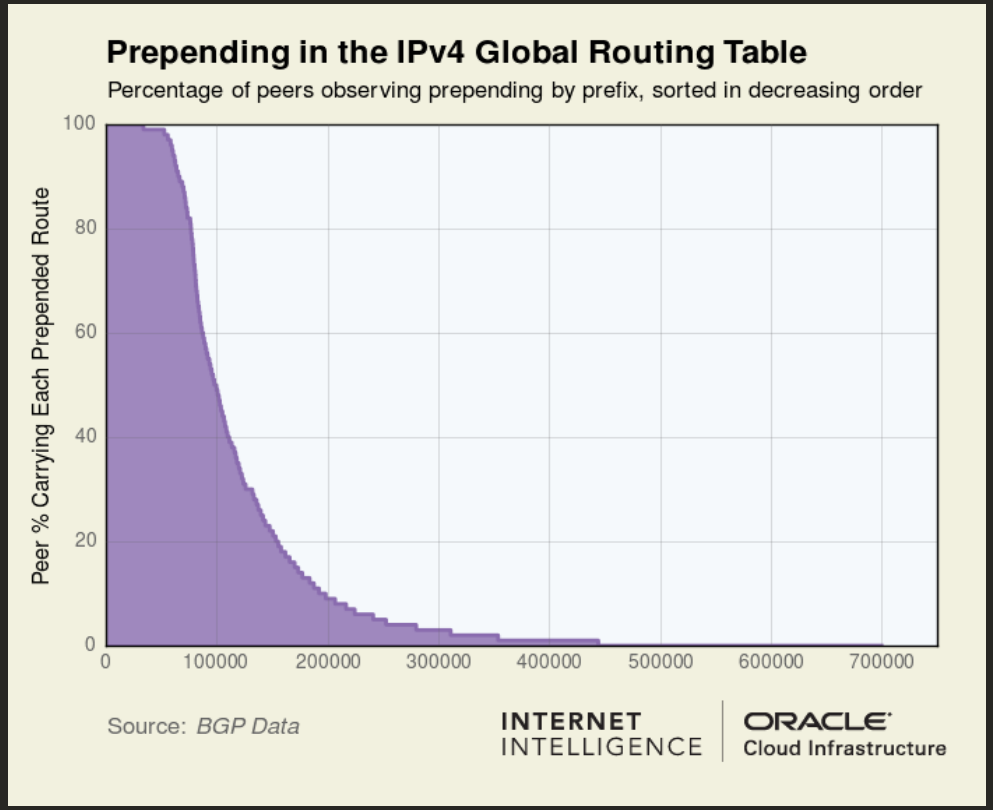

Prepending in the IPv4 Global Routing Table
In LACNIC Region

- Top 10 LACNIC countries (by overall v4 prefix count)
  - Brazil has highest prepend-to-all count
  - AR, CO, EC have highest percentages
    - AS14080 & AS27738 are biggest sources
  - Mexico, Peru have lowest percentages

- *13 of 15 prefixes from Guyana prepended-to-all by AS19863*

country	total prefixes	prepended-to-all	percentage	largest source	% from source
Brazil	58951	**6816**	11.6%	AS53237	3.2%
Mexico	18713	765	**4.1%**	AS17072	38.2%
Argetina	11908	3143	26.4%	AS20207	6.6%
Colombia	8139	1922	23.6%	AS14080	78.7%
Chile	4434	436	9.8%	AS20015	14.9%
Costa Rica	4243	696	16.4%	AS11830	63.1%
Ecuador	3209	808	25.2%	AS27738	73.4%
Peru	2202	83	**3.8%**	AS6147	53.0%
Venezuela	1791	158	8.8%	AS22313	39.9%
Panama	1312	276	21.0%	AS21599	37.7%

Prepending in the IPv6 Global Routing Table

- Prefixes prepended to 95%+ ASes: >3k
  - 5.6% of IPv6 Global Routing Table
- Prefixes prepended to 50%+ ASes: >6k
  - 8.6% of IPv6 Global Routing Table

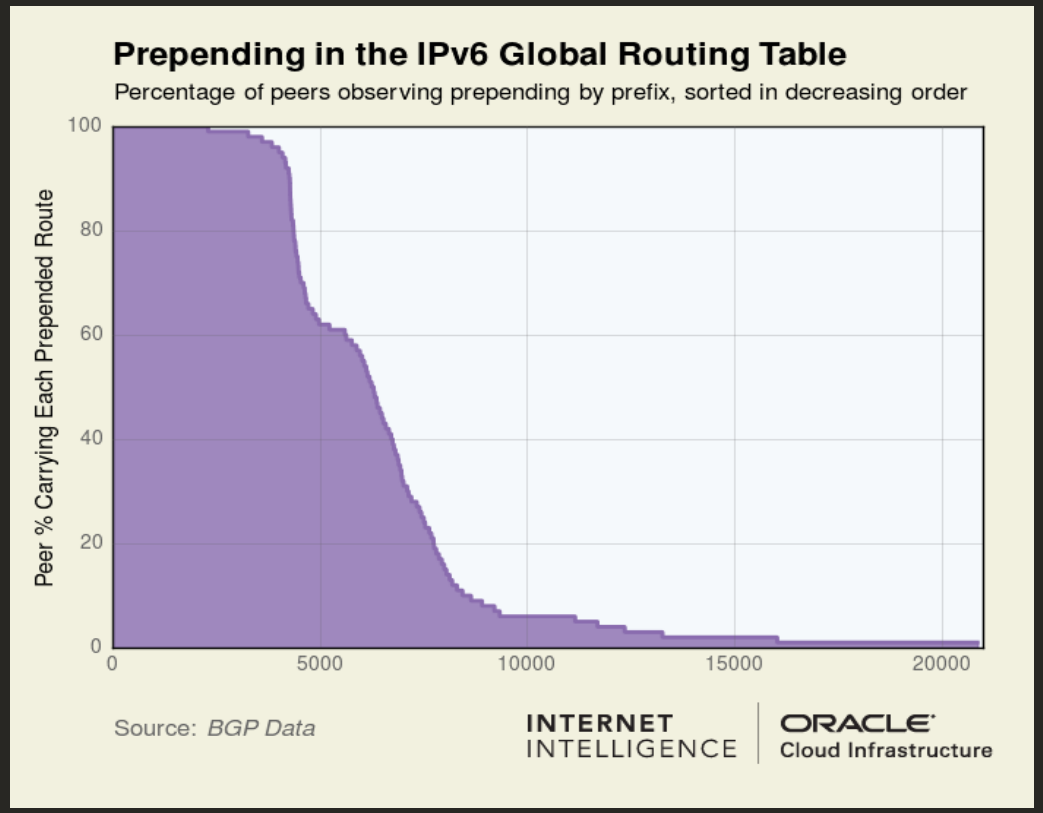

Prepending in the IPv6 Global Routing Table
In LACNIC Region

- Top 10 LACNIC countries (by overall v6 prefix count)
  - Brazil has highest prepend-to-all count
  - CO, EC have highest percentages
    - AS14080 & AS27738: primary sources again
  - Peru has lowest percentage

- *AS17072 is the source of nearly all prepending-to-all in MX*
- *AS52468 announces prepend-to-all in CR, PA, GT*

country	total prefixes	prepended-to-all	percentage	largest source	% from source
Brazil	12408	**741**	6.0%	AS263124	3.5%
Mexico	3441	155	4.5%	AS17072	85.8%
Argentina	671	104	15.5%	AS28073	26.9%
Colombia	516	107	20.7%	AS14080	53.3%
Ecuador	502	112	22.3%	AS22724	53.6%
Peru	431	4	0.9%	AS6147	50.0%
Chile	411	26	6.3%	AS264814	26.9%
Costa Rica	132	6	4.5%	AS52468	50.0%
Panama	117	13	11.1%	AS52468	15.4%
Guatemala	93	4	4.3%	AS52468	100.0%

Prepending is frequently employed in an excessive manner such that it renders routes vulnerable to disruption or misdirection – accidental or otherwise

What's the Risk?

On a recent day, 190.56.128.0/18 was "prepended-to-all" like so:

… 1299 14754

An attacker might announce the same prefix with a fabricated AS path like the following:

… **ASXXX** 1299 14754 14754

Would redirect a portion of traffic to this prefix via ASXXX

What's the Risk?

- The length of prepending gives the attacker room to craft an AS path that would appear plausible, comply with origin validation, and not be detected by off-the-shelf route monitoring.

… 1299 14754

… **ASXXX** 1299 14754 14754

Is Prepending-To-All a growing problem?

What happens when we run these stats over time? Is there a trend?

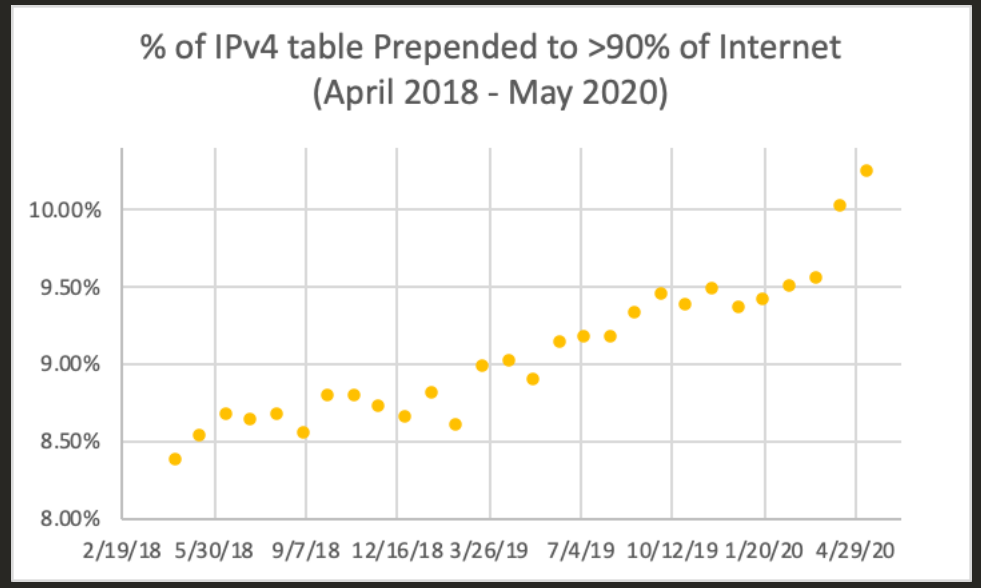

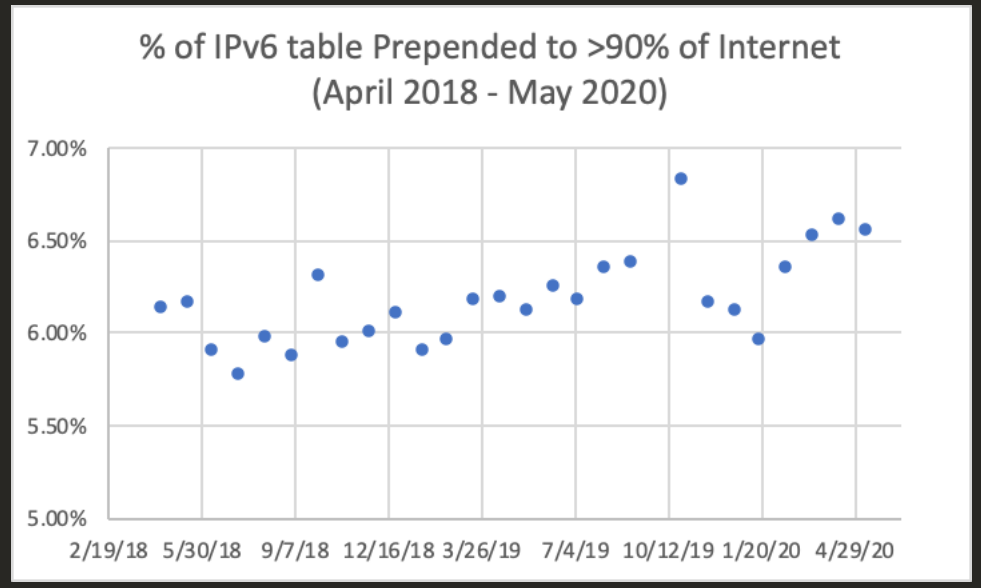

Yes! % of IPv4 table that is prepended-to-all is growing at 0.5%/year

- IPv6 table is growing slower: 0.2%/year

An inadvertent origin leak could also disrupt traffic to these routes. Accidents happen, so why deliberately put your routes at risk?

Why does prepending-to-all happen?

We wanted to know, so we asked some folks doing this. Is it intentional?

... 3356 19256 7955 30321 30321 30321 162.212.148.0/23

We asked Burning Man NetOps about their excessive prepending.

They immediately fixed it. 👍

Why does prepending-to-all happen?

We wanted to know, so we asked some folks doing this.

- Cloudflare, Google also removed the excessive prepending when we reported it to them. 👍
- Most either didn’t respond or claimed it was an “operational issue” and it remains.

Why does prepending-to-all happen?

Theory 1: Poor Housekeeping - The AS forgets to remove the prepending for one of its transit providers when it is no longer needed.

Theory 2: Return Path Influence – AS attempting to de-prioritize traffic from transit providers over settlement-free peers.

Why does this happen?

Theory 3: Mistakes Abound - There are simply a lot of errors in BGP routing. Consider the prepended AS path of 181.191.170.0/24 below:

… 52981 267429 267429 267492 267492 267429 267429 267492 267492 267429 267429 267492 267492 267429

*In case your eyes didn't catch it, the prepending here involves a mix of two distinct ASNs (2674**29** and 2674**92**) with the last two digits transposed.*

Conclusions

- Long AS paths (whether due to prepending or not) incur risk of disruption
  - In the event another AS originates the same prefix with a shorter AS path
- Network operators should ensure prepending is absolutely necessary
  - *Many of your networks have excessive prepending (ask me for examples)*
- With 8% of IPv4 and 5.6% of IPv6 global routing tables presently prepended to *everyone*, this traffic engineering technique is significantly overused.

Further Reading…

- AS-Path Prepending…
  - Pedro Marcos (et al)
  - Paper at ACM IMC 2020

https://conferences.sigcomm.org/imc/2020/accepted/

AS-Path Prepending: there is no rose without a thorn

Pedro Marcos
FURG
pbmarcos@furg.br

Lars Prehn
MPI for Informatics
lprehn@mpi-inf.mpg.de

Lucas Leal
UFRGS
lsleal@inf.ufrgs.br

Alberto Dainotti
CAIDA/UCSD
alberto@caida.org

Marinho Barcellos
University of Waikato
marinho.barcellos@waikato.ac.nz

Anja Feldmann
MPI for Informatics
anja@mpi-inf.mpg.de

ABSTRACT

Inbound traffic engineering (ITE) is an essential task for Autonomous Systems (ASes). It is the process of announcing routes to, e.g., maximize revenue or minimize congestion. AS Path Prepending (ASPP) is an easy to use well-known ITE technique that routing manuals show as one of the first

Prepending (ASPP) [15, 22, 76], selective or more-specific prefix announcements [27], BGP communities[23, 63], or Multi Exit Discriminator (MED) values [25, 43].

In this paper, we focus on ASPP and the controversy it generates—as we explain next. ASPP is a straightforward, easy-to-use technique that is often mentioned among the first ITE techniques by router vendors [19, 21, 26, 36, 45]

- AS Path Prepending
  - Draft IETF Best Current Practice

https://datatracker.ietf.org/doc/draft-ietf-grow-as-path-prepending/

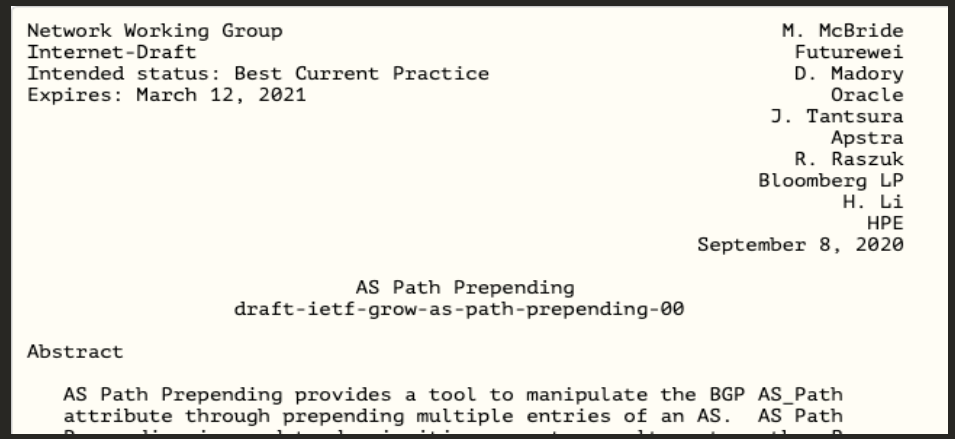

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: March 12, 2021

M. McBride
Futurewei
D. Madory
Oracle
J. Tantsura
Apstra
R. Raszuk
Bloomberg LP
H. Li
HPE

September 8, 2020

AS Path Prepending
draft-ietf-grow-as-path-prepending-00

Abstract

AS Path Prepending provides a tool to manipulate the BGP AS_Path attribute through prepending multiple entries of an AS. AS Path

Thank you

Doug Madory
@InternetIntel
Oracle Internet Intel

Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.

The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.