

ORACLE

Excessive BGP AS Path Prepending is a Self-Inflicted Vulnerability

Doug Madory

PTNOG 4
Lisbon, PT
5 December 2019

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 [1,2]


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

- As AS path lengths exceeded 255, Cisco routers c



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!



ars TECHNICA BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE STORE

BIZ & IT —
How China swallowed 15% of 'Net traffic for 18 minutes

In April 2010, 15 percent of all Internet traffic was suddenly diverted ...

NATE ANDERSON - 11/17/2010, 2:45 PM

In a **300+ page report** (PDF) today, the US-China Economic and Security Review Commission provided the US Congress with a detailed overview of what's been happening in China—including a curious incident in which 15 percent of the world's Internet traffic suddenly passed through Chinese servers on the way to its destination.

Here's how the Commission describes the incident, which took place earlier this year:

“

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~~

prepending-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 **prepending-to-all**.
- AS4671 leak of April 2014 (>320,000 prefixes)

... 2856 7862 7862 7862 7862 7862
146.23.208.0/21

^ Prepending-to-all




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

Impacts of Excessive Prepending During Leaks

- Much of the worst propagation of leaked routes during big leak events were due to routes being **prepending-to-all**.
- AS4788 leak of June 2015 (>260,000 prefixes)

... 174 12322 12322 12322 12322
82.224.0.0/12

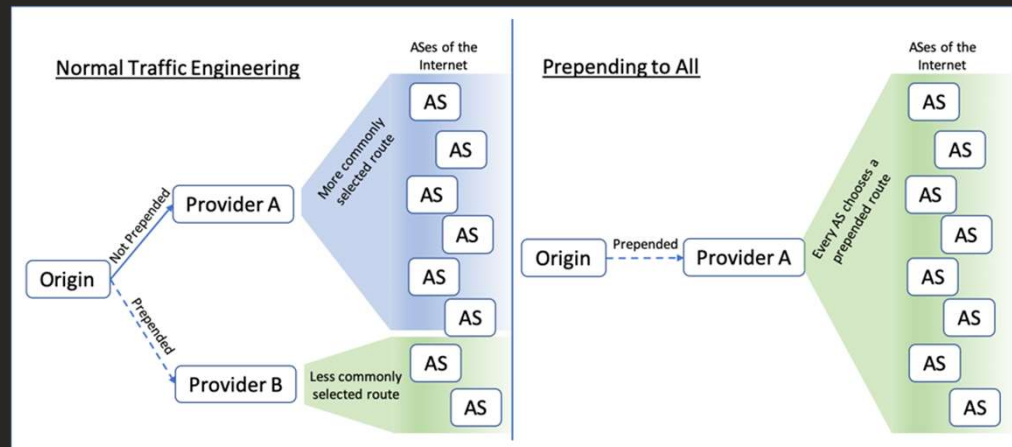
^ Prepending-to-all



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

Prepending to Everyone!

- Prepended-to-all prefixes are those seen as prependded 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 **prependded-to-all**? ...a lot!

Prepending in the Global Routing Tables

Prepending in the IPv4 Global Routing Table

Percentage of peers observing prepending by prefix, sorted in decreasing order




Source: BGP Data

INTERNET
INTELLIGENCE

ORACLE
Cloud Infrastructure

Prepending in the IPv6 Global Routing Table

Percentage of peers observing prepending by prefix, sorted in decreasing order




Source: BGP Data

INTERNET
INTELLIGENCE

ORACLE
Cloud Infrastructure

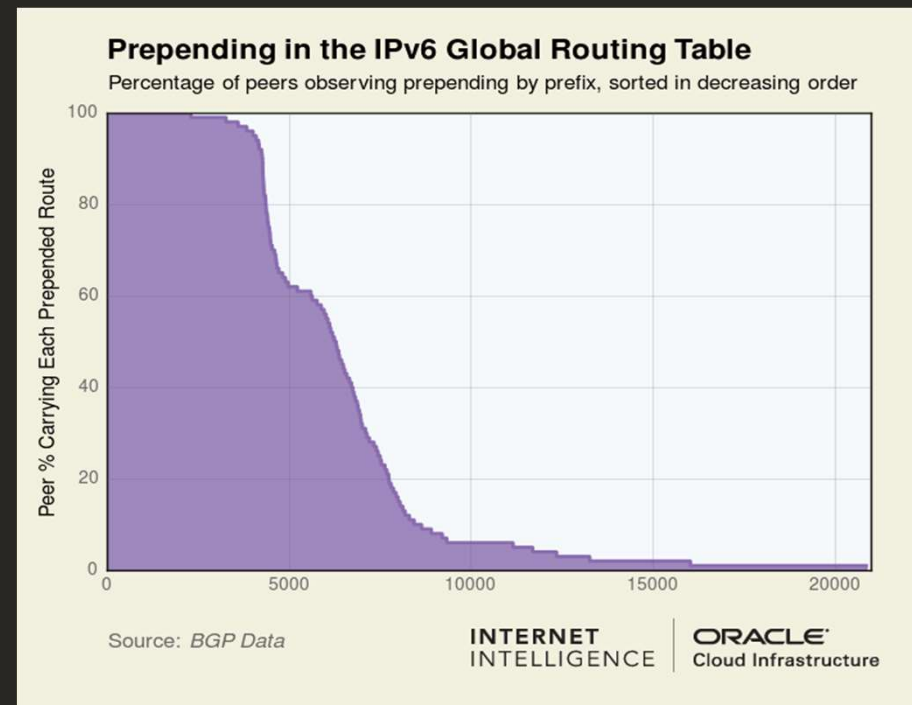
Prepending in the IPv4 Global Routing Table

- Prefixes prepended to >95% of ASes: >60k
 - 8% of IPv4 Global Routing Table (1/12)
 - Includes entities of every stripe: governments, financial institutions, even important parts of internet infrastructure.
- Prefixes prepended to >50% of ASes: >100k
 - 13.3% of IPv4 Global Routing Table.



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 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, 95.47.142.0/23 was “prepending-to-all” like so:

```
... 3255 197158 197158 197158 197158 197158 197158 197158 197158
197158 197158 197158 197158 197158 197158 197158 197158 197158 197158
197158 197158 197158 197158 197158
```

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

```
... ASXXX 3255 197158
```

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.

... 3255 197158 197158 197158 197158 197158 197158 197158 197158 197158
197158 197158 197158 197158 197158 197158 197158 197158 197158 197158
197158 197158 197158 197158

... ASXXX 3255 197158

Prepended vs non-prepended in the wild


Sometimes the impact of prepending isn't as straightforward.

This prefix is announced to the internet in two ways:

... 174 17356



... 6939 17356 17356 17356 17356 17356 17356 17356
17356 17356 17356 17356

- 58% of our peers chose prepending 6939 route
 - 42% chose non-prepend 174 routes
- Despite prepending, AS6939 is more popular due to extensive peering base of thousands of ASNs.



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 prepping when we reported👍 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.*

Prepending-to-All in Portugal



- Only ~50 / 740 PT prefixes are prepended-to-all (not bad!)
- AS13156 (Cabovisao, SA) announces 29 of them.
- Examples:
 - 84.91.20.0/22 and 217.129.236.0/22 are announced along:
... 6453 13156 13156
- 195.138.12.0/24 (Sociedade Interbancaria de Servicos, SA)
 - ... 8657 15525 6773 6773
- 104.244.8.0/22 (Aeroportos de Portugal)
 - ... 2860 201170 201170

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.

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.