


Visualizing Routing Incidents in 3D

Doug Madory
Director of Internet
Analysis

PTNOG 4
Lisbon, PT
5 December 2019



Scourge of route leaks continue

Serverless | M* | CLL | Events | White

The Register
Biting the hand that feeds IT

WARE | SECURITY | DEVOPS | BUSINESS | PERSONAL TECH | SCIENCE | EMERGENT TECH | BOOTNOT

Data Centre ► **Networks**

BGP super-blunder: How Verizon today sparked a 'cascading catastrophic failure' that knackered Cloudflare, Amazon, etc

'Normally you'd filter it out if some small provider said they own the internet'

By Kieren McCarthy in San Francisco 24 Jun 2019 at 19:01 61 SHARE ▼

Updated Verizon sent a big chunk of the internet down a black hole this morning – and caused outages at Cloudflare, Facebook, Amazon, and others – after it wrongly accepted a network misconfiguration from a small ISP in Pennsylvania, USA.

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

THANKS, BGP. —

BGP event sends European mobile traffic through China Telecom for 2 hours

Improper leak to Chinese-government-owned telecom lasts up to two hours.

DAN GOODIN · 6/8/2019, 6:05 PM

Traffic Path during

Impact often measured simply by prefix count



Data Centre ► Networks

BGP super-blunder: How Verizon

"It all started when new internet routes for more than **20,000 IP address prefixes** – roughly two per cent of the internet – were wrongly announced..."

'Normally you'd filter it out if some small provider said they own the internet'

By Kieren McCarthy in San Francisco 24 Jun 2019 at 19:01 61 SHARE ▼

Updated Verizon sent a big chunk of the internet down a black hole this morning – and caused outages at Cloudflare, Facebook, Amazon, and others – after it wrongly accepted a network misconfiguration from a small ISP in Pennsylvania, USA.



THANKS, BGP. —

BGP event sends European mobile traffic through China Telecom for 2 hours

"...Safe Host improperly updated its routers to advertise it was the proper path to reach what eventually would become more than **70,000 Internet routes**..."



Prefix count is one-dimensional and lacks nuance

“more than 20,000 IP address prefixes”

“more than 70,000 Internet routes”

Weaknesses of a one-dimensional measure of a leak


- Not every leaked route is accepted by the same number of ASes
- Not every leaked route is in circulation for the same amount of time
- There is often a long tail of prefixes that didn't propagate far or for very long, but are included in the “prefix count” metric.

“There has to be a better way!”


“more than 20,000 IP address prefixes”

“more than 70,000 Internet routes”

- Need to include propagation and duration to improve our understanding
- Resulting in a 3-dimensional view of an incident:
 - prefixes (x-axis), duration (y-axis), propagation (z-axis)

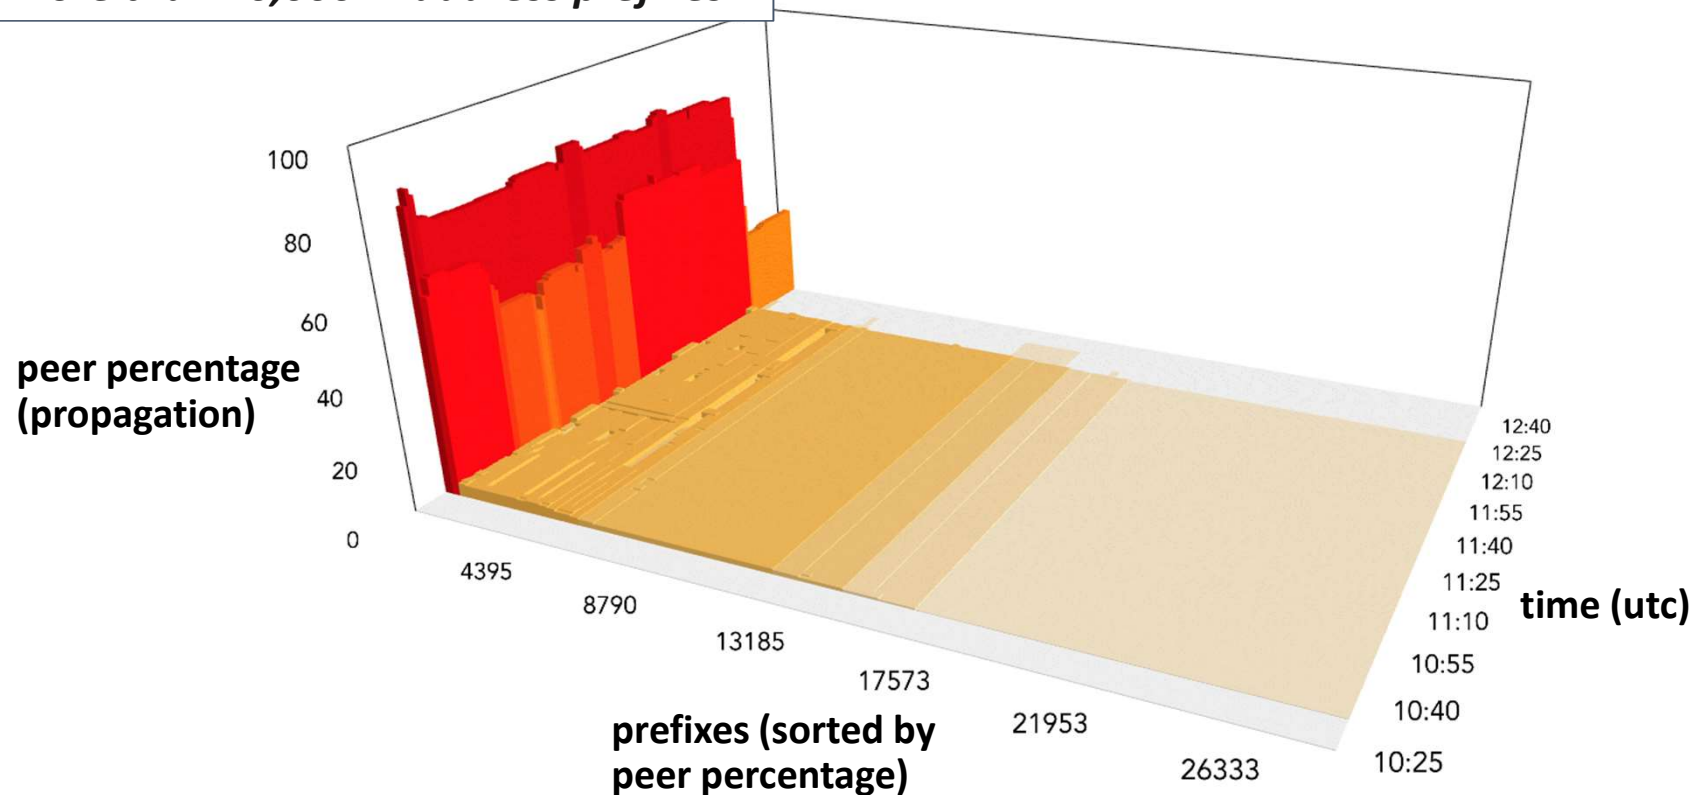


Global propagation of all routes for duration of leak would be a solid box:



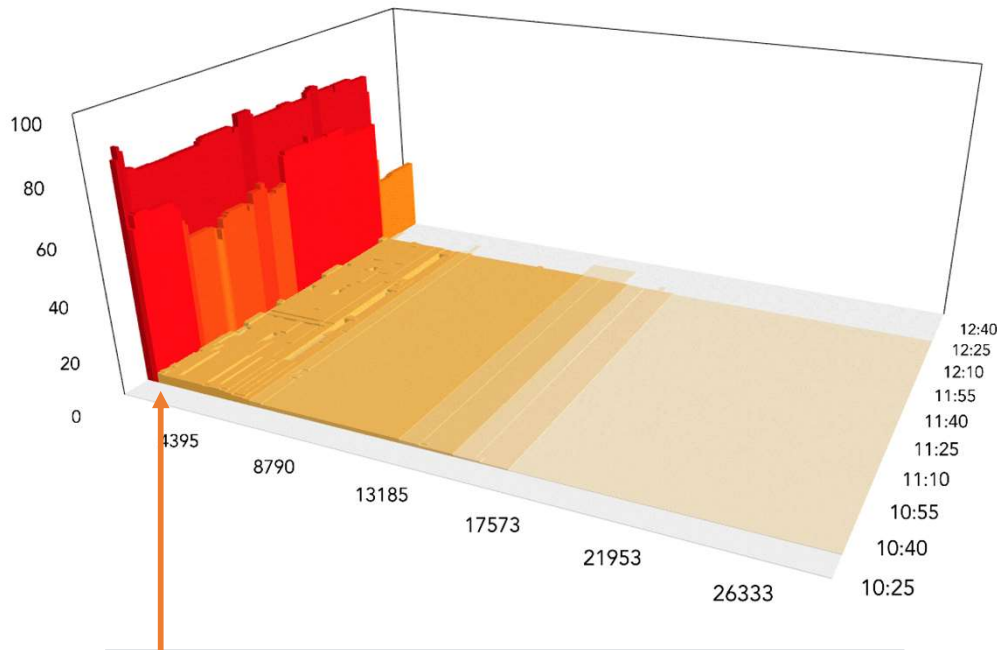
3-dimensional view of routing leak

“more than 20,000 IP address prefixes”



Analysis of potential RPKI filtering

"more than 20,000 IP address prefixes"



Optimizer generated ~263 more-specifics that were widely circulated.

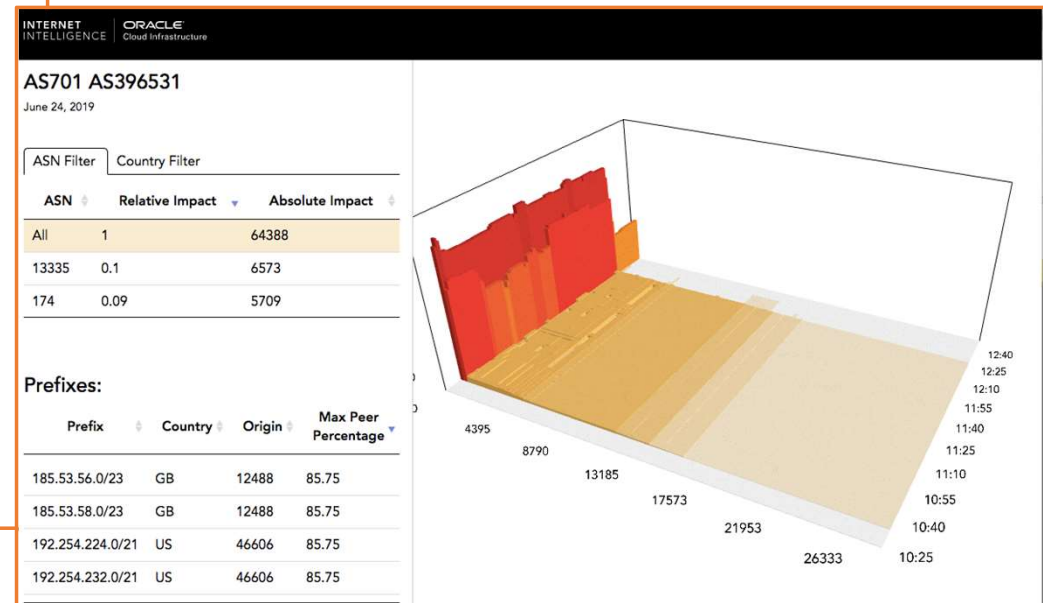
- Had RPKI invalids been dropped during the leak, here's how the 29k leaked routes would have fared:
26873 RPKI:UNKNOWN
2145 RPKI:VALID
130 RPKI:INVALID_LENGTH
28 RPKI:INVALID_ASN
- RPKI would have only filtered 158 leaked routes (0.5%)
 - 66 of 80 Cloudflare prefixes
- A lot of work remains to be done to reduce the incidences of RPKI:UNKNOWN, but there were 13x more RPKI:VALID than RPKI:INVALID

This analysis can be automated!!

- New website will be available at: {URL TBD}
- Will publish interactive autopsies of significant routing leaks soon after they occur.*
- In addition, a history of previous incidents will be available for comparison and research.

Upstream AS		Leak AS		Average Start	Average Duration	Prefix Count (All)	Prefix Count (Significant)
+	Zayo (AS6461)	APEXn Pty Ltd (AS38195)		2019-07-11 21:41	00:05:20	20870	11411
+	Global Cloud Xchange (AS15412)	AboveNet Taiwan (AS17408)		2019-06-29 08:39	00:09:10	4722	4716
+	Kazakhtelecom (AS9198)	KVANT Telekom (AS43727)		2019-06-25 20:43	00:25:49	1766	1766
+	Verizon Business (AS701)	Allegheny Technologies Incorporated (AS396531)		2019-06-24 10:35	00:44:17	29253	14610
+	China Telecom (AS4134)	Safe Host SA (AS21217)		2019-06-06 10:25	00:18:15	78252	15373
+	Republican Unitary Enterprise National Traffic Exchange Center (AS60280)	Beltelecom (AS6697)		2019-05-15 21:56	00:08:38	9718	8632
+	Lanka Bell Limited (AS45224)	Sri Lanka Telecom, Internet Service Provider, IX (AS45489)		2019-04-21 15:32	00:03:17	527	526

Showing 1 to 28 of 28 entries



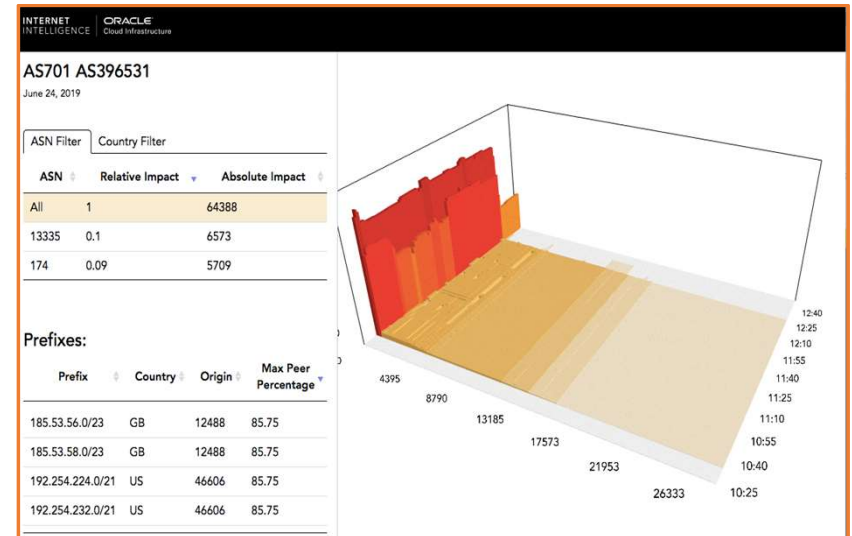
*Significant = More than 100 prefixes and seen by at least 10% of our peer set

*Soon = As soon as we can verify the analysis.



Explore a routing incident using filters

- Interface includes filters by origin & country-level geo.
- Lists most affected prefixes by max peer percentage for any selected origin or country.
- List of most impacted origins and countries by impact:
 - Impact = sum(area under curve for selected filter)
- Absolute impacts from different incidents can be directly compared.



The Ultimate Routing Leak Myth: China Telecom (April 2010)



[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:


“

For about 18 minutes on April 8, 2010, China Telecom advertised erroneous network traffic routes that instructed US and other foreign Internet traffic to travel through Chinese servers. Other servers around the world quickly adopted these paths, routing all

The Ultimate Routing Leak Myth: China Telecom (April 2010)

“15% of internet traffic for 18 minutes”

- Obviously, biggest problem: **routes != traffic**
- But also, not all of the routes were widely circulated
- For argument's sake, let's we assume routes = traffic
 - If 15% of all traffic was redirected, each route would need to be propagated to 100% of the internet. Like this →
- It was isn't even close.



The Ultimate Routing Leak Myth: China Telecom (April 2010)

INTERNET
INTELLIGENCE | ORACLE

AS4134 AS23724

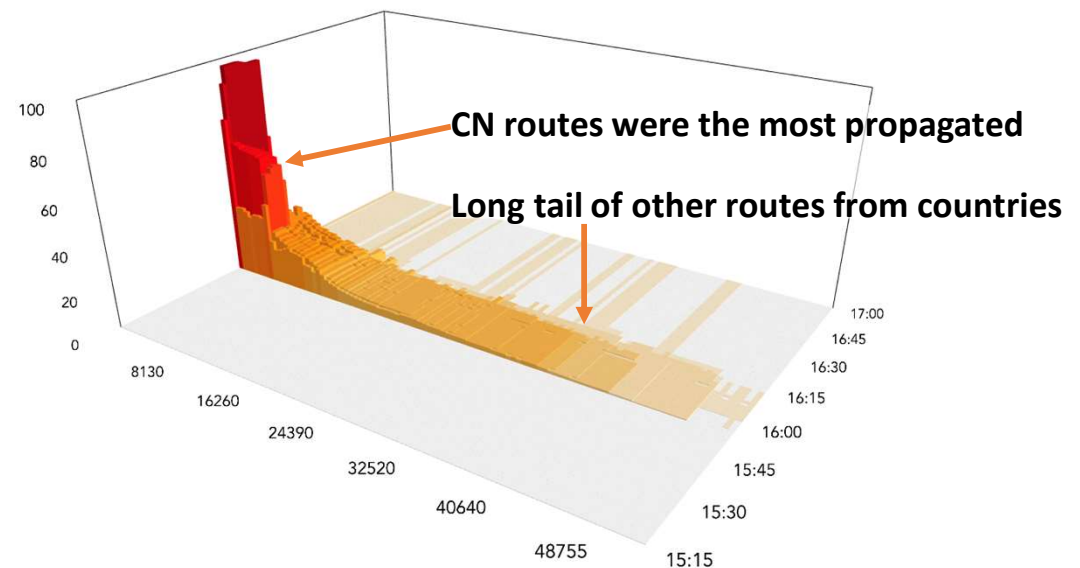
April 8, 2010

ASN Filter Country Filter

ASN	Prefixes	Relative Impact	Absolute Impact
All	54165	1	3742097
4134	10384	0.7	2619090
4538	562	0.03	114953

Prefixes:

Prefix	Country	Origin	Max Peer Percentage
202.100.192.0/19	CN	4134	97.05
202.100.224.0/19	CN	4134	96.76



The Ultimate Routing Leak Myth: China Telecom (April 2010)

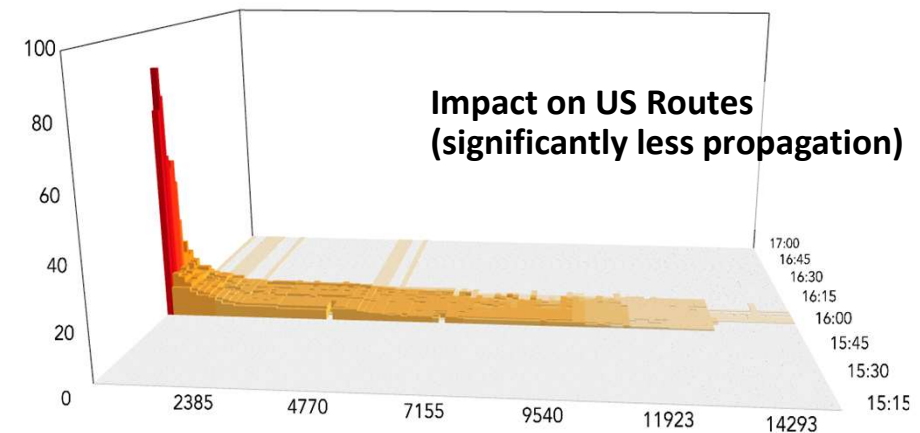
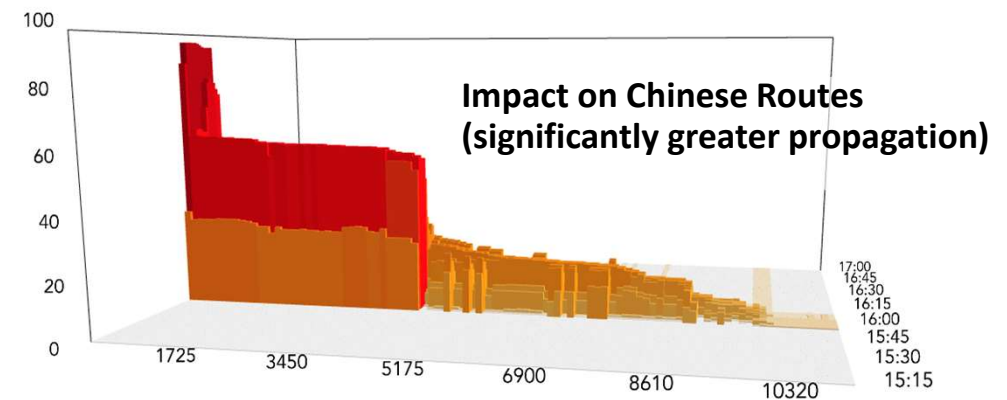
- Better than simply counting prefixes, we can measure “impact” by aggregate propagation:

$\text{pfx_count} * \text{duration} * \text{peer_percentage}$

April 8, 2010			
ASN Filter		Country Filter	
ASN	Prefixes	Relative Impact	Absolute Impact
All	54165	1	3742097
CN	11460	0.74	2756164
US	15873	0.08	290987

- 74% (CN) vs 8% (US)
- Impact was only 4.6% of theoretical max

“15% **0.07%** of internet traffic **route propagation** for 18 minutes”



* Widely propagated US prefixes due to prepending



Revisiting big leaks from the past: Indosat, April 2014

INTERNET
INTELLIGENCE | ORACLE

AS4761

April 2, 2014

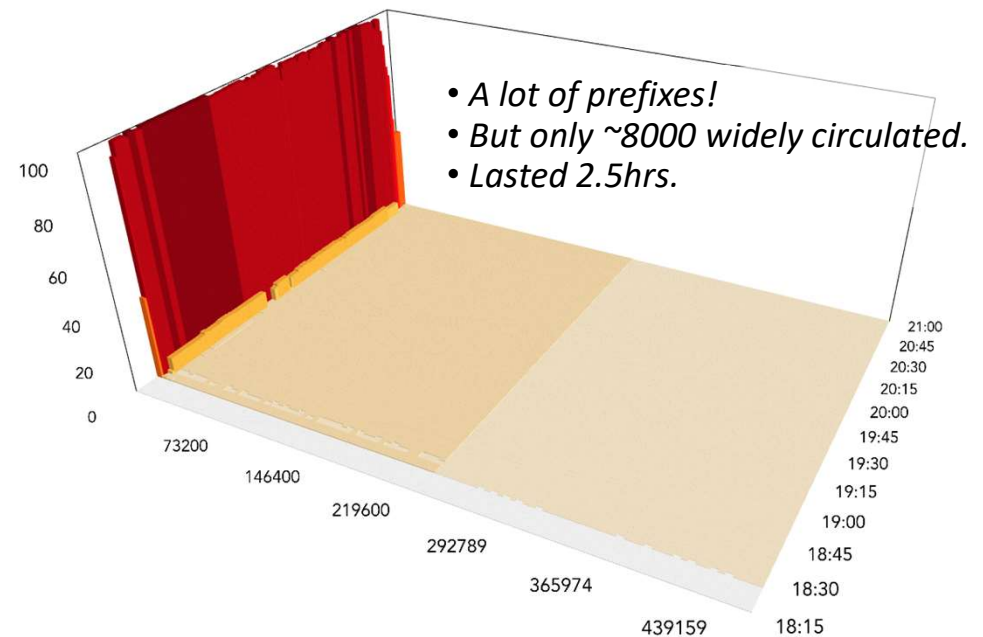
ASN Filter

Country Filter

ASN	Prefixes	Relative Impact	Absolute Impact
All	487949	1	22684033
ID	8387	0.5	11311281
US	144519	0.14	3121307

Prefixes:

Prefix	Country	Origin	Max Peer Percentage
121.54.24.0/21	PH	10139	100
121.54.32.0/21	PH	10139	100



Revisiting big leaks from the past: TMnet, June 2015

INTERNET
INTELLIGENCE

ORACLE

AS3549 AS4788

June 12, 2015


ASN Filter

Country Filter

ASN	Prefixes	Relative Impact	Absolute Impact
All	264636	1	135725355
CN	13808	0.21	28454289
AU	11494	0.15	20520048

Prefixes:

Prefix	Country	Origin	Max Peer Percentage
58.26.216.0/22	MY	4788	98.17
1.32.104.0/22	MY	4788	97.65



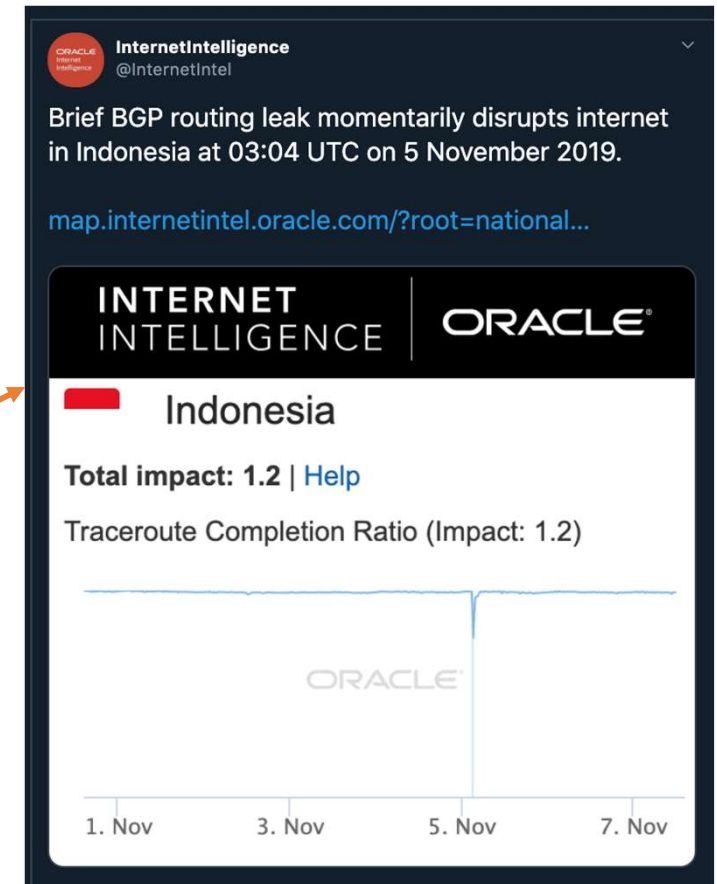
Biggest impacts of all time!

- Using the same formula for impact, we can compare different events through time.

Top 5

Leaker	Impact	Date
AS4788	135,725,355	Jun 12, 2015
AS4761	22,684,033	Apr 2, 2014
AS41095	22,272,707	Oct 10, 2019
AS3303	10,959,010	Feb 19, 2019
AS58944	8,279,144	Nov 5, 2019

- Skewed towards more recent events due to growth of global routing table.




Conclusion

- We need to include the dimensions of propagation and duration.
- It's time we had a better metric than simply prefix count.
 - Suggestion: Count of leaked prefixes seen by >1% of peers.
 - More esoteric suggestion: Impact as measured by aggregate propagation
- RPKI can help contain leaks but needs greater participation
 - More signed routes & more dropping of invalids
- We hope that these interactive routing leak autopsies will help inform discussion around routing leaks.

Stop saying China Telecom hijacked 15% of internet! 😊

-
- **Doug Madory**
 - **@InternetIntel**
 - **Oracle Internet Intel**

Copyright © 2019, Oracle and/or its affiliates. All rights reserved.






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.

Don't we already have BGP leak analyzers?

- Jared Mauch's leakinfo.cgi and BGPstream take similar approaches of looking for three "BIG" networks in the AS path of a BGP message
- This message-by-message approach gets dominated by ephemeral "leaks" which exist only momentarily during convergence from one routing state to another.
- Most often ephemeral leaks occur when a prefix is withdrawn and ASes frantically exchange routing info to exchange a viable route.
- Ephemeral leaks help identify where filtering is lacking, but generally have little operational impact due to their brevity.



[Search this data](#)

Source	Time	Prefix	AS_PATH	Contact?	Score	blame_asn
rrc00-45.12.69.254-1566148361.43-8023399	2019-08-18 13:12:41.945465	170.244.168.0/22	58057 34549 3320 701 5511 52468 267858	3320	3	3320
rrc00-80.77.16.114-1566148361.52-5260667	2019-08-18 13:12:41.661801	170.244.168.0/22	34549 3320 701 5511 52468 267858	3320	3	3320
rrc00-80.77.16.114-1566148361.52-5260661	2019-08-18 13:12:41.619739	170.244.168.0/22	34549 3320 701 5511 52468 267858	3320	3	3320
http://archive.routeviews.org/bgpdata/2019.08/UPDATES/updates.20190818.1645.bz2	2019-08-18 13:02:54.8381	149.106.96.0/19	3561 3910 3356 2914 13536 13536 19999	3561	3	3910
http://archive.routeviews.org/bgpdata/2019.08/UPDATES/updates.20190818.1645.bz2	2019-08-18 13:02:54.835821	149.106.0.0/19	3561 3910 3356 2914 13536 13536 19999	3561	3	3910
http://archive.routeviews.org/bgpdata/2019.08/UPDATES/updates.20190818.1645.bz2	2019-08-18 13:02:54.833508	149.106.64.0/19	3561 3910 3356 2914 13536 13536 19999	3561	3	3910

