



JUNOS

THE
FAST AND
THE FURIOUS:
PROPAGATION
DRIFT

A long time ago

in a job a year back



Anycast



**Asia
PoP**



**East Coast
PoP**



**West Coast
PoP**



**Server under
Theo de Raadt's house**



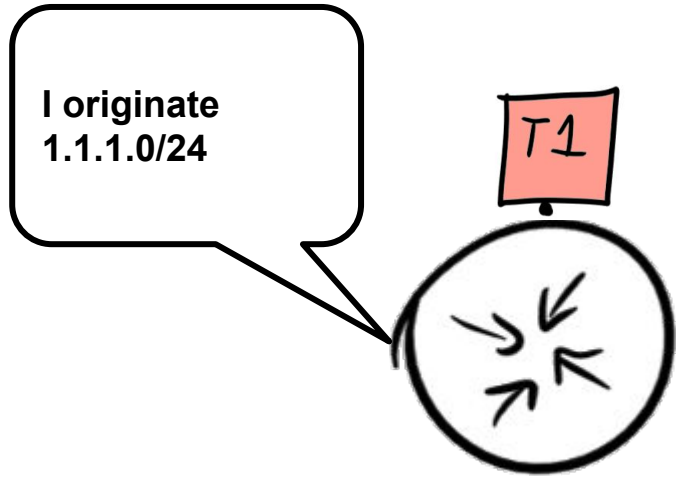





**I originate
1.1.1.0/24**

**I originate
1.1.1.0/24**

**I originate
1.1.1.0/24**

**I originate
1.1.1.0/24**



Most anycast
deployments



Most anycast
deployments



Better anycast deployments



EU:

AS_PATH: 13335 13335 179 1337

US:

AS_PATH: 13335 13335 179 1337

APAC:

AS_PATH: 13335 13335 179 1337

EU:

AS_PATH: 13335 13335 1299 1337

US:

AS_PATH: 13335 13335 1299 1337

APAC:

AS_PATH: 13335 13335 1299 1337

EU:


AS_PATH: 13335 13335 1299 1337


US:

AS_PATH: 13335 13335 1299 1337

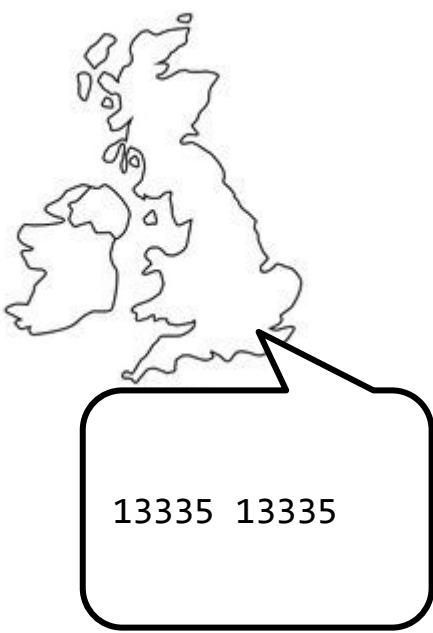
APAC:

AS_PATH: 13335 13335 1299 1337

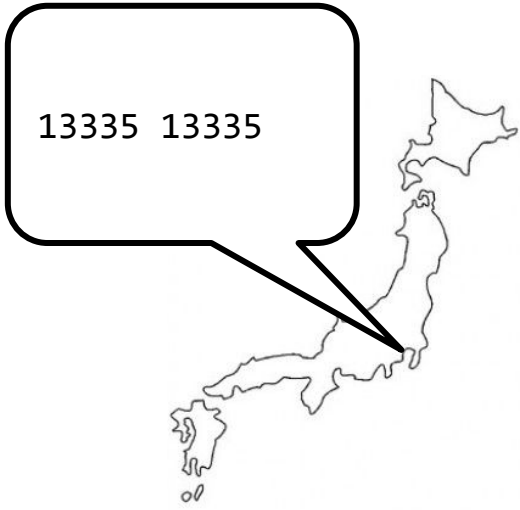




13335 13335




13335 13335

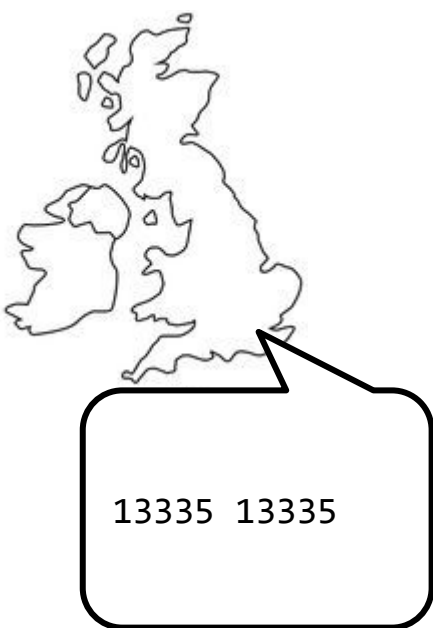


13335 13335

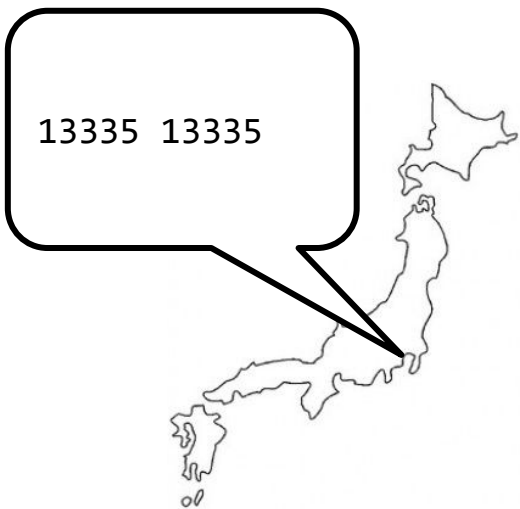
However, one day



13335 13335



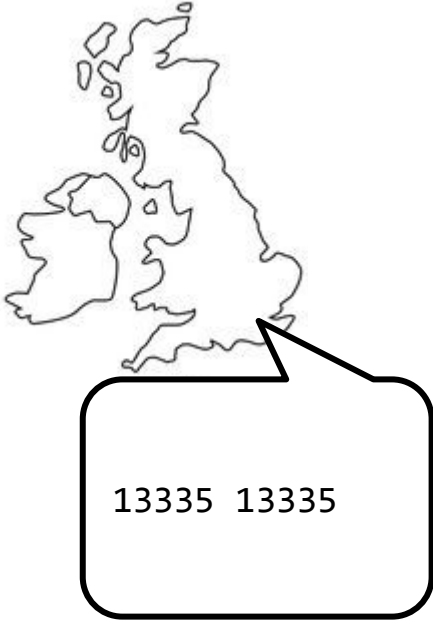
13335 13335



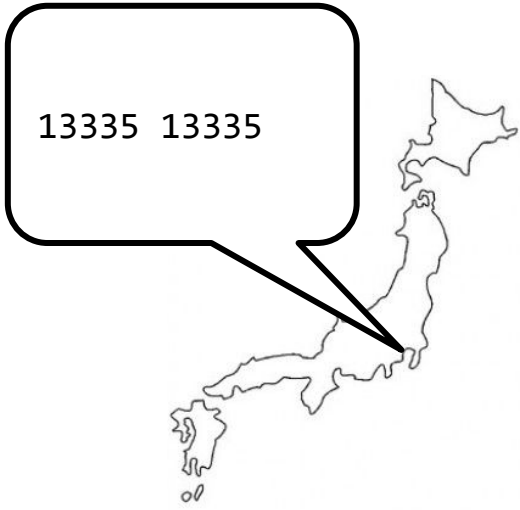
13335 13335





13335



13335 13335



13335 13335



This was a regrettable incident






Removing these preponds is now a
good idea

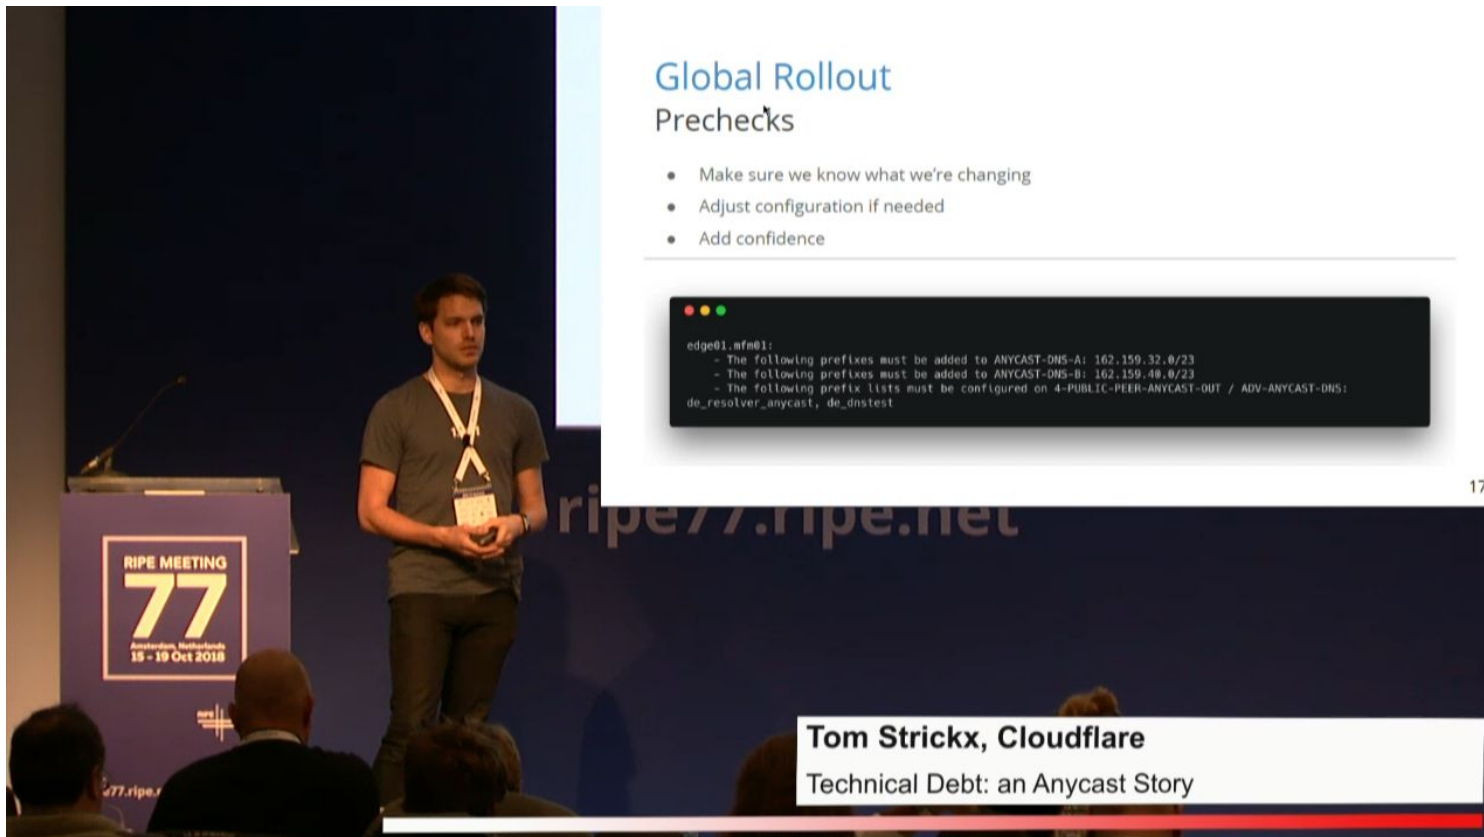


Removing these preponds is now a
good idea

But what if we do the same thing
again?







Global Rollout Prechecks

- Make sure we know what we're changing
- Adjust configuration if needed
- Add confidence

```
edge01.mfm01:  
- The following prefixes must be added to ANYCAST-DNS-A: 162.159.32.0/23  
- The following prefixes must be added to ANYCAST-DNS-B: 162.159.48.0/23  
- The following prefix lists must be configured on 4-PUBLIC-PEER-ANYCAST-OUT / ADV-ANYCAST-DNS:  
de_resolver_anycast, de_dnstest
```

17

Tom Strickx, Cloudflare
Technical Debt: an Anycast Story

<https://ripe77.ripe.net/archives/video/2222/>



Change done without any major
issues



It was observed that not all carriers came into sync at the same time

[side note]



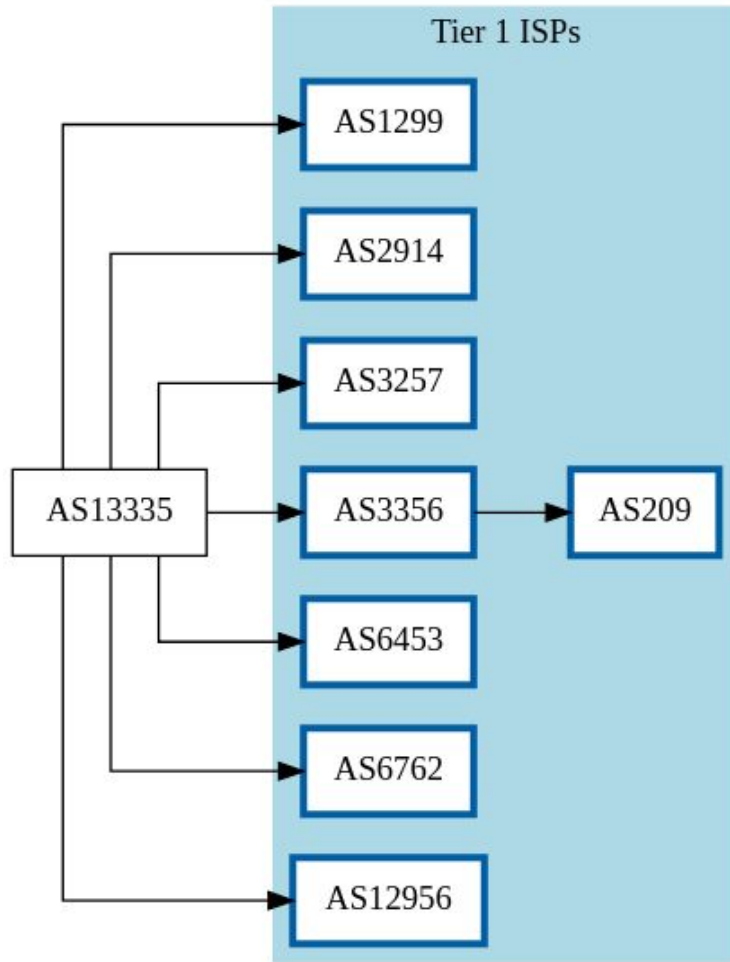
How long does it take for carriers to send route updates globally?



How long does it take for carriers to send route updates globally?



Who is the fastest/slowest?





The Stance 
@YourStancePolls

Follow



Snail Race

Vote & Retweet

13%



26%



47%



14%



39,689 votes • Final results



The Stance 
@YourStancePolls

Follow



Snail Race

Vote & Retweet




39,689 votes • Final results


Pro tip:


When presented with 4 options of the same, people overwhelmingly pick the 3rd one

When asked for a random number between 0-9, people overwhelmingly pick 7

#1







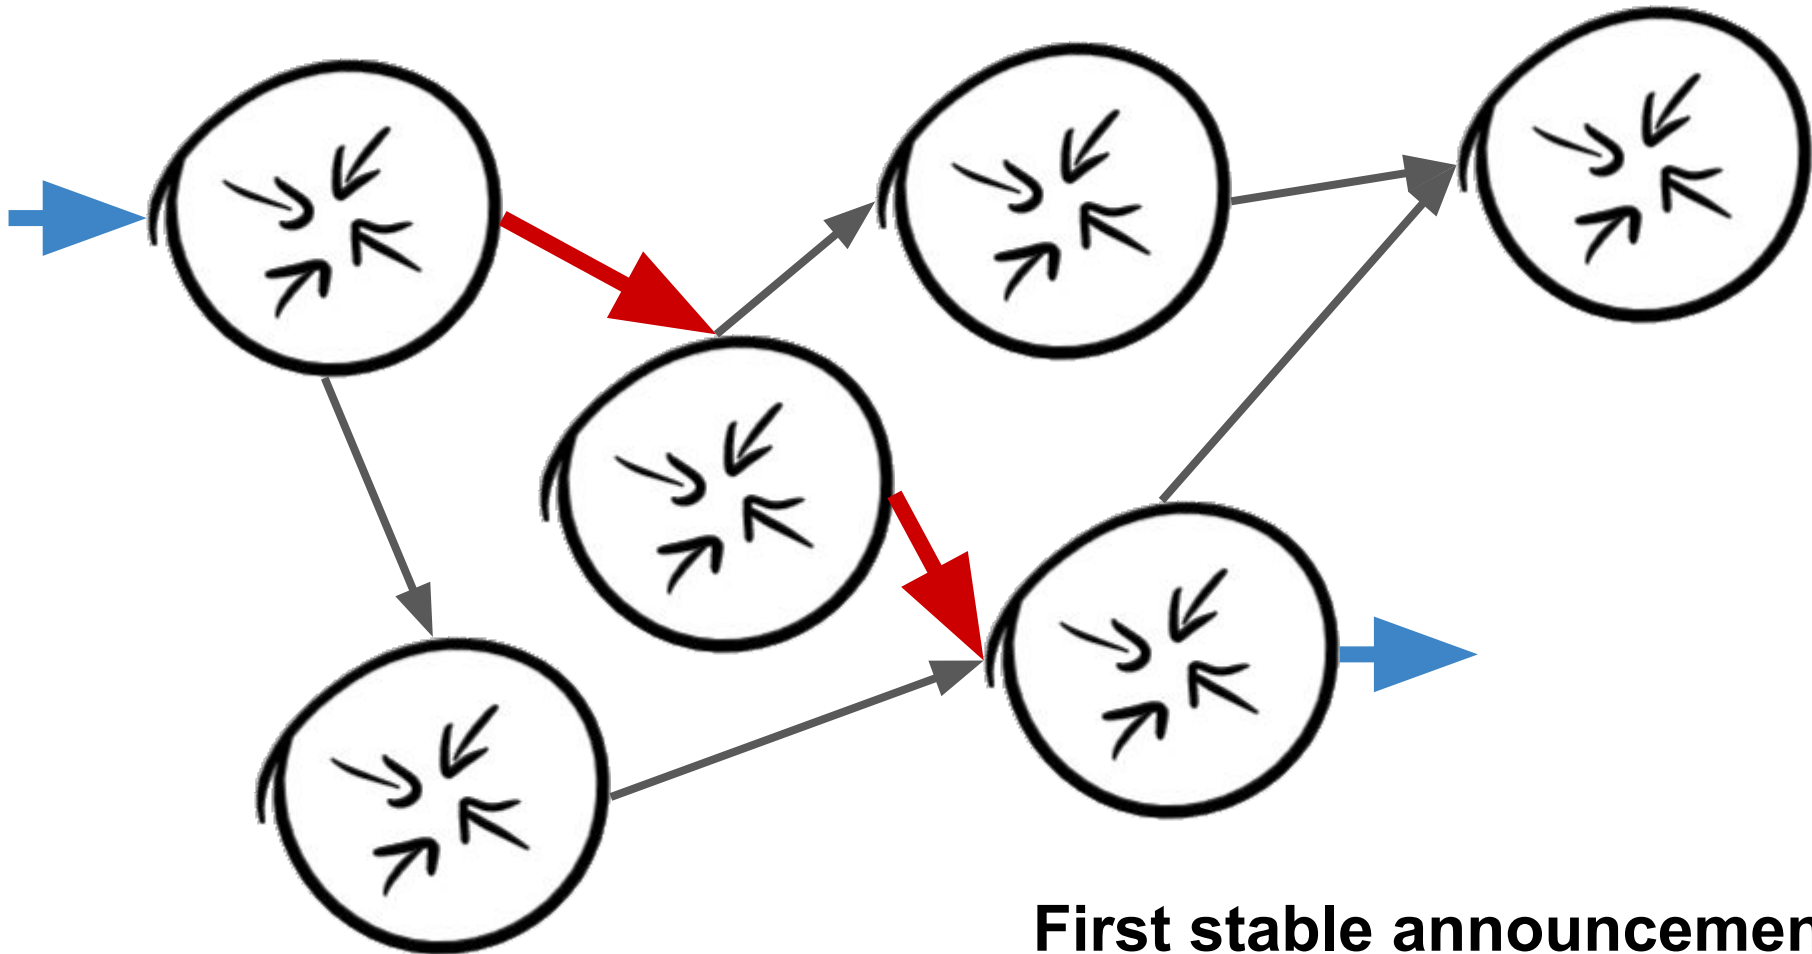


First announcement wins

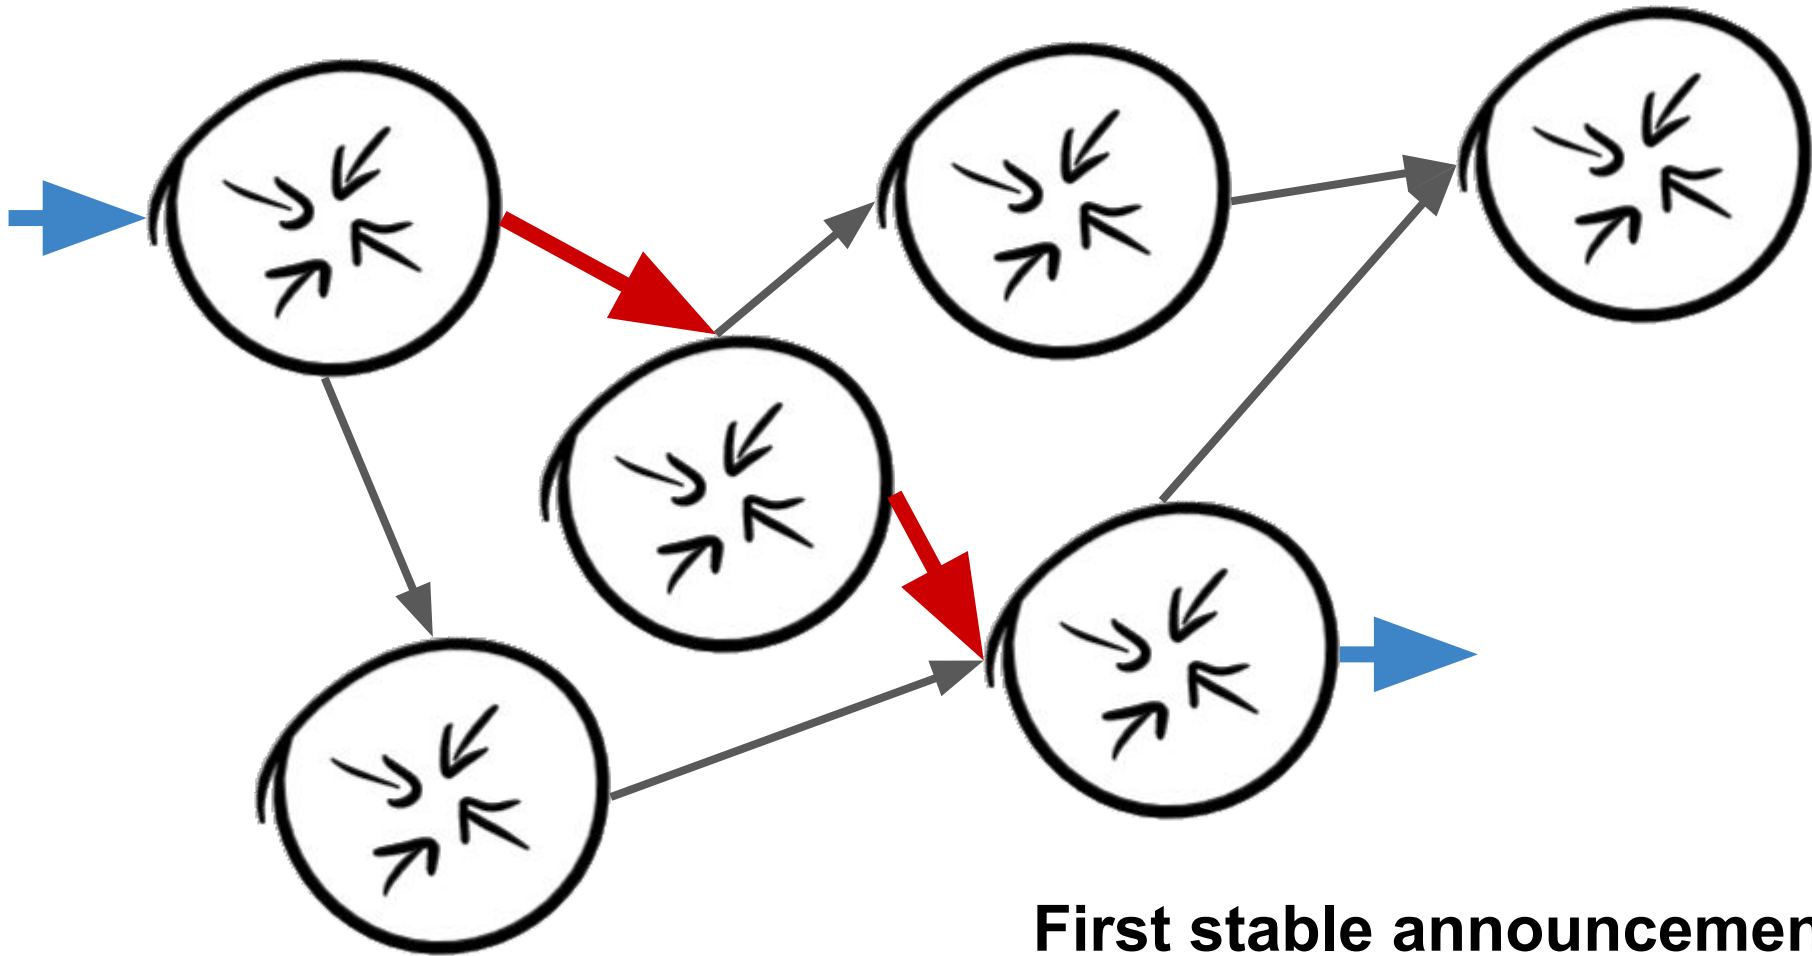
#2









**First stable announcement
wins**



**First stable announcement
wins**





First announcement wins

Race setup

- Announce a /24 to one T1 at a time

Race setup

- Announce a /24 to one T1 at a time
- 16 PoP's with these carriers are tcpdump'ing their bgp feeds

Race setup



- Announce a /24 to one T1 at a time
- 16 PoP's with these carriers are tcpdump'ing their bgp feeds
- pcap's are collected post experiment

Race setup


- Announce a /24 to one T1 at a time
- 16 PoP's with these carriers are tcpdump'ing their bgp feeds
- pcap's are collected post experiment
- Program parses the packet captures BGP traffic and using the high precision timestamps, calculates latency.

Race setup

- Announce a /24 to one T1 at a time
- 16 PoP's with these carriers are tcpdump'ing their bgp feeds
- pcap's are collected post experiment
- Program parses the packet captures BGP traffic and using the **high precision timestamps**, calculates latency.



JNP10003-RE1-LT



10

g100093

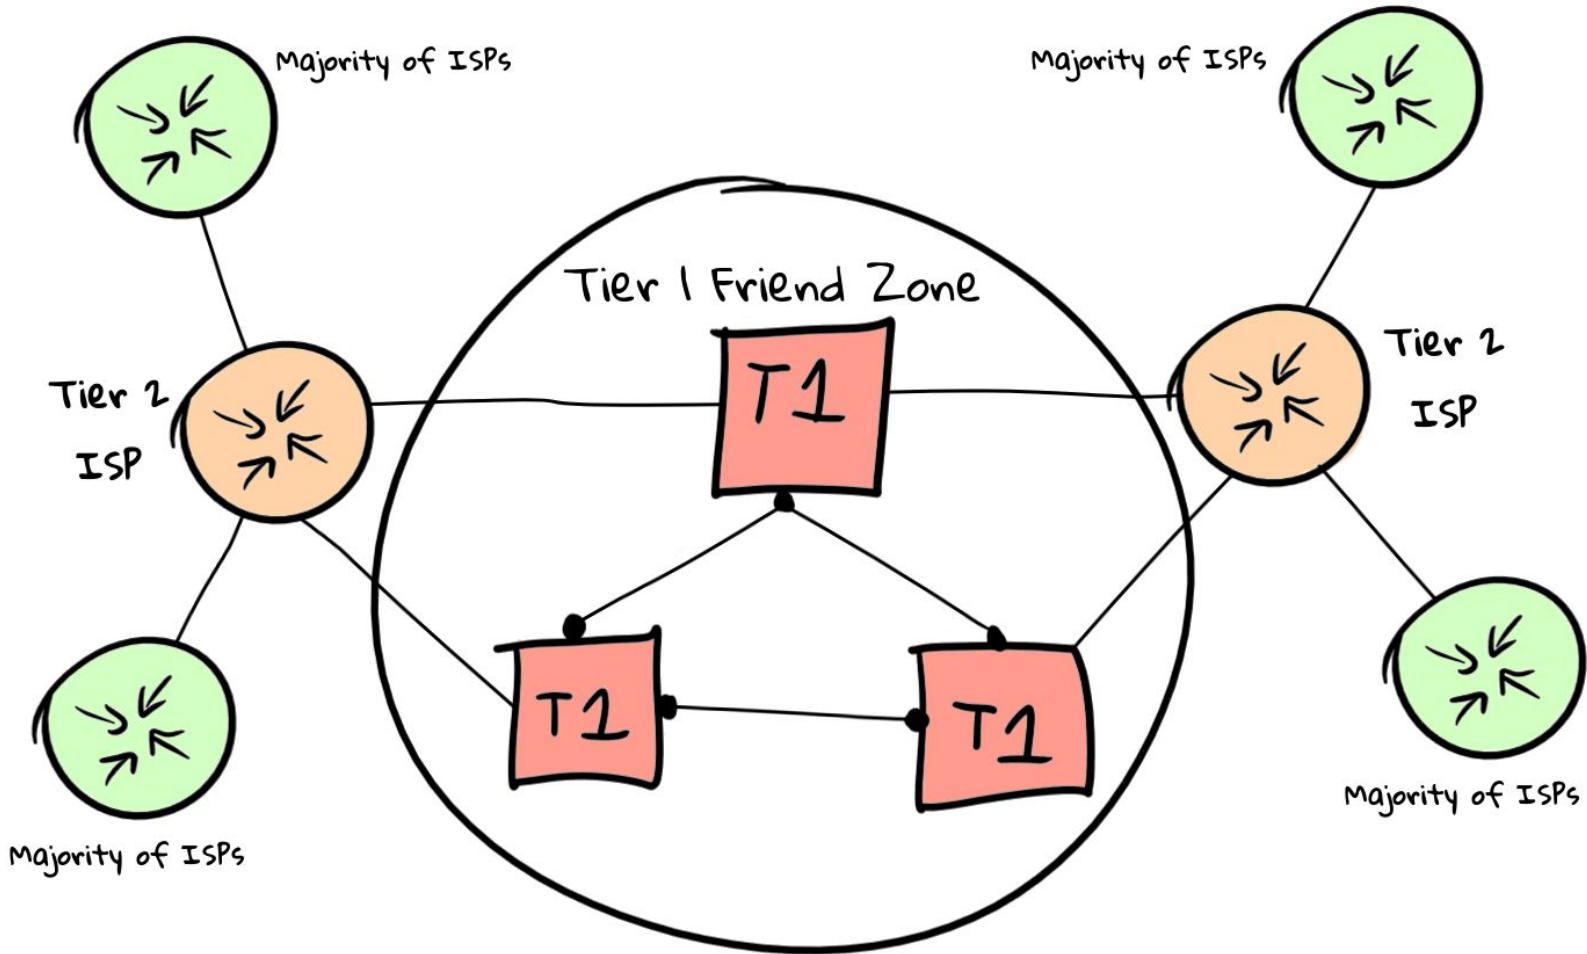

```
[16:48:37] ben@metropolis:~/new-server-race$ sort bgp-race.tsv
```

```
1553612072316966000 2019-03-26 14:54:32.316966 +0000 GMT ewr 64515,65534,20473,1299,62240,62240,206924
1553612073531538000 2019-03-26 14:54:33.531538 +0000 GMT wildfire 206924
1553612074054742000 2019-03-26 14:54:34.054742 +0000 GMT sea 64515,65534,20473,3356,6453,6453,62240,62240,206924
1553612074117278000 2019-03-26 14:54:34.117278 +0000 GMT cdg 64515,65534,20473,1299,6453,6453,62240,62240,206924
1553612074124066000 2019-03-26 14:54:34.124066 +0000 GMT fra 64515,65534,20473,1299,6453,6453,62240,62240,206924
1553612074149076000 2019-03-26 14:54:34.149076 +0000 GMT lax 64515,65534,20473,3257,6453,6453,62240,62240,206924
1553612074256535000 2019-03-26 14:54:34.256535 +0000 GMT yyz 64515,65534,20473,3356,6453,6453,62240,62240,206924
1553612074349282000 2019-03-26 14:54:34.349282 +0000 GMT nrt 64515,65534,20473,2914,6453,6453,62240,62240,206924
1553612074362145000 2019-03-26 14:54:34.362145 +0000 GMT sjc 64515,65534,20473,3257,6453,6453,62240,62240,206924
1553612074371090000 2019-03-26 14:54:34.37109 +0000 GMT mia 64515,65534,20473,3257,6453,6453,62240,62240,206924
1553612074376447000 2019-03-26 14:54:34.376447 +0000 GMT ams 64515,65534,20473,1299,6453,6453,62240,62240,206924
1553612074434826000 2019-03-26 14:54:34.434826 +0000 GMT lhr 64515,65534,20473,1299,6453,6453,62240,62240,206924
1553612074457718000 2019-03-26 14:54:34.457718 +0000 GMT dfw 64515,65534,20473,2914,6453,6453,62240,62240,206924
1553612074526157000 2019-03-26 14:54:34.526157 +0000 GMT ord 64515,65534,20473,2914,6453,6453,62240,62240,206924
1553612074552698000 2019-03-26 14:54:34.552698 +0000 GMT fra 64515,65534,20473,3257,6453,6453,62240,62240,206924
1553612074597688000 2019-03-26 14:54:34.597688 +0000 GMT sin 64515,65534,20473,2914,6453,6453,62240,62240,206924
1553612075053011000 2019-03-26 14:54:35.053011 +0000 GMT lhr 64515,65534,20473,2914,6453,6453,62240,62240,206924
1553612075244176000 2019-03-26 14:54:35.244176 +0000 GMT syd 64515,65534,20473,2914,6453,6453,62240,62240,206924
1553612095008991000 2019-03-26 14:54:55.008991 +0000 GMT dfw 64515,65534,20473,3356,6453,6453,62240,62240,206924
1553612266956741000 2019-03-26 14:57:46.956741 +0000 GMT ewr 64515,65534,20473,1299,6453,6453,62240,62240,206924
1553612391552044000 2019-03-26 14:59:51.552044 +0000 GMT wildfire 206924
```


Time between last (ms)	Wall time (s)	Location	Raw Path	Upstreaming AS	Rx is di
0	0	wildfire	206924	-	
523.209984ms	0.523209984	sea	64515,65534,20473,3356,6453,6453,62240,62240,206924	6453	No
62.530048ms	0.585740032	cdg	64515,65534,20473,1299,6453,6453,62240,62240,206924	6453	No
6.789888ms	0.59252992	fra	64515,65534,20473,1299,6453,6453,62240,62240,206924	6453	No
25.010176ms	0.617540096	lax	64515,65534,20473,3257,6453,6453,62240,62240,206924	6453	No
107.45984ms	0.724999936	yyz	64515,65534,20473,3356,6453,6453,62240,62240,206924	6453	No
92.75008ms	0.817750016	nrt	64515,65534,20473,2914,6453,6453,62240,62240,206924	6453	No
12.859904ms	0.83060992	sjc	64515,65534,20473,3257,6453,6453,62240,62240,206924	6453	No
8.950016ms	0.839559936	mia	64515,65534,20473,3257,6453,6453,62240,62240,206924	6453	No
5.350144ms	0.84491008	ams	64515,65534,20473,1299,6453,6453,62240,62240,206924	6453	No
58.380032ms	0.903290112	lhr	64515,65534,20473,1299,6453,6453,62240,62240,206924	6453	No
22.889984ms	0.926180096	dfw	64515,65534,20473,2914,6453,6453,62240,62240,206924	6453	No
68.439808ms	0.994619904	ord	64515,65534,20473,2914,6453,6453,62240,62240,206924	6453	No
26.540032ms	1.021159936	fra	64515,65534,20473,3257,6453,6453,62240,62240,206924	6453	No
44.990208ms	1.066150144	sin	64515,65534,20473,2914,6453,6453,62240,62240,206924	6453	No
455.329792ms	1.521479936	lhr	64515,65534,20473,2914,6453,6453,62240,62240,206924	6453	No
191.160064ms	1.71264	syd	64515,65534,20473,2914,6453,6453,62240,62240,206924	6453	No
s	21.47745997	dfw	64515,65534,20473,3356,6453,6453,62240,62240,206924	6453	No
ns	193.4252101	ewr	64515,65534,20473,1299,6453,6453,62240,62240,206924	6453	No



* the following findings are of my own and I encourage you to cross check this work, also I don't want to get sued by TATA



Time between last (ms)	Wall time (s)	Location	Raw Path	Upstreaming AS	Rx is d
0	0	wildfire		206924	-
2425.149952ms	2.425149952	yyz	64515,65534,20473,3257,174,174,174,174,62240,62240,622	174	No
2202.050048ms	4.6272	yyz	64515,65534,20473,1299,174,174,174,174,62240,62240,622	174	No
3328.009984ms	7.955209984	yyz	64515,65534,20473,3356,174,174,174,174,62240,62240,622	174	No
1141.260032ms	9.096470016	sea	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
85.090048ms	9.181560064	lhr	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
158.010112ms	9.339570176	syd	64515,65534,20473,17819,4826,174,62240,62240,62240,62	174	No
230.729984ms	9.57030016	dfw	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
56.239872ms	9.626540032	ewr	64515,65534,20473,1299,174,174,174,174,62240,62240,622	174	No
65.819904ms	9.692359936	nrt	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
107.650048ms	9.800009984	fra	64515,65534,20473,3257,174,174,174,174,62240,62240,622	174	No
18.660096ms	9.81867008	lax	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
49.769984ms	9.868440064	mia	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
33.319936ms	9.90176	cdg	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
18.200064ms	9.919960064	ams	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
74.170112ms	9.994130176	sjc	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
4ms	9.998130176	sin	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
16898.519808ms	26.89664998	ord	64515,65534,20473,2914,174,174,174,174,62240,62240,622	174	No
531.63008ms	27.42828006	dfw	64515,65534,20473,3356,174,174,174,174,62240,62240,622	174	No



* the following findings are of my own and I encourage you to cross check this work, also I don't want to get sued by Cogent, <3 the fellows in the back

Time between last (ms)	Wall time (s)	Location	Raw Path	Upstreaming AS	Rx is d
0	0	wildfire		206924	-
721.139968ms	0.721139968	yyz	64515,65534,20473,1299,3257,62240,62240,62240,206924	3257	No
64.240128ms	0.785380096	lax	64515,65534,20473,3257,62240,62240,206924	3257	Yes
9.420032ms	0.794800128	dfw	64515,65534,20473,3257,62240,62240,206924	3257	Yes
52.909824ms	0.847709952	sea	64515,65534,20473,3257,62240,62240,206924	3257	Yes
21.350144ms	0.869060096	sjc	64515,65534,20473,3257,62240,62240,206924	3257	Yes
44.879872ms	0.913939968	mia	64515,65534,20473,3257,62240,62240,206924	3257	Yes
15.830016ms	0.929769984	nrt	64515,65534,20473,2914,3257,62240,62240,62240,206924	3257	No
82.909952ms	1.012679936	fra	64515,65534,20473,1299,3257,62240,62240,62240,206924	3257	No
20.080128ms	1.032760064	cdg	64515,65534,20473,3257,62240,62240,206924	3257	Yes
36.649984ms	1.069410048	sin	64515,65534,20473,2914,3257,62240,62240,62240,206924	3257	No
19.660032ms	1.08907008	ams	64515,65534,20473,1299,3257,62240,62240,62240,206924	3257	No
19.129856ms	1.108199936	lhr	64515,65534,20473,1299,3257,62240,62240,62240,206924	3257	No
256.130048ms	1.364329984	syd	64515,65534,20473,2914,3257,62240,62240,62240,206924	3257	No
19.130112ms	1.383460096	ord	64515,65534,20473,3257,62240,62240,206924	3257	Yes
208.519936ms	1.591980032	lhr	64515,65534,20473,2914,3257,62240,62240,62240,206924	3257	No
281.209856ms	1.873189888	fra	64515,65534,20473,3257,62240,62240,206924	3257	Yes
88.730112ms	1.96192	nrt	64515,65534,20473,2914,3257,62240,62240,62240,206924	3257	No
114745.520128ms	116.7074401	ewr	64515,65534,20473,1299,3257,62240,62240,62240,206924	3257	No



* the following findings are of my own and I encourage you to cross check this work, also I don't want to get sued by GTT

Time between last (ms)	Wall time (s)	Location	Raw Path	Upstreaming AS	Rx is
0	0	wildfire	206924	-	
632.25984ms	0.63225984	cdg	64515,65534,20473,1299,62240,62240,206924	1299	Yes
11.300096ms	0.643559936	fra	64515,65534,20473,1299,62240,62240,206924	1299	Yes
107.099904ms	0.75065984	sea	64515,65534,20473,3356,1299,62240,62240,206924	1299	No
76.75008ms	0.82740992	ams	64515,65534,20473,1299,62240,62240,206924	1299	Yes
17.64992ms	0.84505984	lhr	64515,65534,20473,2914,1299,62240,62240,206924	1299	No
40.519936ms	0.885579776	yyz	64515,65534,20473,3257,1299,62240,62240,206924	1299	No
27.640064ms	0.91321984	mia	64515,65534,20473,3257,1299,62240,62240,206924	1299	No
59.910144ms	0.973129984	sjc	64515,65534,20473,2914,1299,62240,62240,206924	1299	No
8.199936ms	0.98132992	dfw	64515,65534,20473,1299,62240,62240,206924	1299	Yes
9.37984ms	0.99070976	lax	64515,65534,20473,3257,1299,62240,62240,206924	1299	No
10.920192ms	1.001629952	ewr	64515,65534,20473,1299,62240,62240,206924	1299	Yes
5.879808ms	1.00750976	yyz	64515,65534,20473,1299,62240,62240,206924	1299	Yes
61.510144ms	1.069019904	lhr	64515,65534,20473,1299,62240,62240,206924	1299	Yes
137.139968ms	1.206159872	ord	64515,65534,20473,3257,1299,62240,62240,206924	1299	No
211.780096ms	1.417939968	nrt	64515,65534,20473,2914,1299,62240,62240,206924	1299	No
12.279808ms	1.430219776	sin	64515,65534,20473,2914,1299,62240,62240,206924	1299	No
499.030016ms	1.929249792	nrt	64515,65534,20473,2914,1299,62240,62240,206924	1299	No
198.569984ms	2.127819776	syd	64515,65534,20473,2914,1299,62240,62240,206924	1299	No
21135.510016ms	23.26332979	mia	64515,65534,20473,1299,62240,62240,206924	1299	Yes
3196.370176ms	26.45969997	ord	64515,65534,20473,1299,62240,62240,206924	1299	Yes



* the following findings are of my own and I encourage you to cross check this work, also I don't want to get sued by Telia

Time between last (ms)	Wall time (s)	Location	Raw Path	Upstreaming AS	Rx is
0	0	wildfire	206924	-	
18508.079872ms	18.50807987	sea	64515,65534,20473,3356,62240,62240,206924	3356	Yes
365.55008ms	18.87362995	yyz	64515,65534,20473,3356,62240,62240,206924	3356	Yes
241.890048ms	19.11552	lhr	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
251.180032ms	19.36670003	cdg	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
174.459904ms	19.54115994	mia	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
87.279872ms	19.62843981	fra	64515,65534,20473,1299,3356,3356,3356,3356,62240	3356	No
6.09024ms	19.63453005	sin	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
6.909952ms	19.64144	ewr	64515,65534,20473,1299,3356,3356,3356,3356,62240	3356	No
53.309952ms	19.69474995	ams	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
212.290048ms	19.90704	dfw	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
72.140032ms	19.97918003	lax	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
0.909824ms	19.98008986	nrt	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
187.480064ms	20.16756992	sjc	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
194.109952ms	20.36167987	syd	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
10094.810112ms	30.45648998	ewr	64515,65534,20473,3356,62240,62240,206924	3356	Yes
3963.629824ms	34.42011981	mia	64515,65534,20473,3356,62240,62240,206924	3356	Yes
1207.540224ms	35.62766003	ord	64515,65534,20473,2914,3356,3356,3356,3356,62240	3356	No
1684.079872ms	37.3117399	sjc	64515,65534,20473,3356,62240,62240,206924	3356	Yes
476.240128ms	37.78798003	fra	64515,65534,20473,3356,62240,62240,206924	3356	Yes
434.519808ms	38.22249984	dfw	64515,65534,20473,3356,62240,62240,206924	3356	Yes
1264.510208ms	39.48701005	ord	64515,65534,20473,23352,3356,62240,62240,206924	3356	No
5106.689792ms	44.59369984	ams	64515,65534,20473,3356,62240,62240,206924	3356	Yes
1426.160128ms	46.01985997	ord	64515,65534,20473,3356,62240,62240,206924	3356	Yes
695.45984ms	46.71531981	lhr	64515,65534,20473,3356,62240,62240,206924	3356	Yes
3801.630208ms	50.51695002	cdg	64515,65534,20473,3356,62240,62240,206924	3356	Yes



* the following findings are of my own and I encourage you to cross check this work, also I don't want to get sued by CenturyLink or Level 3 or whoever owns you seriously it's getting confusing now

Time between last (ms)	Wall time (s)	Location	Raw Path	Upstreaming AS	Rx is d
0	0	wildfire	206924	-	
814.89024ms	0.81489024	ewr	64515,65534,20473,2914,62240,62240,206924	2914	Yes
90.649856ms	0.905540096	fra	64515,65534,20473,2914,62240,62240,206924	2914	Yes
290.150144ms	1.19569024	lax	64515,65534,20473,2914,62240,62240,206924	2914	Yes
0.57984ms	1.19627008	cdg	64515,65534,20473,2914,62240,62240,206924	2914	Yes
121.449984ms	1.317720064	yyz	64515,65534,20473,1299,2914,62240,62240,206924	2914	No
5.100032ms	1.322820096	lhr	64515,65534,20473,2914,62240,62240,206924	2914	Yes
24.04992ms	1.346870016	nrt	64515,65534,20473,2914,62240,62240,206924	2914	Yes
33.970176ms	1.380840192	ams	64515,65534,20473,1299,2914,62240,62240,206924	2914	No
1.12ms	1.381960192	dfw	64515,65534,20473,3257,2914,62240,62240,206924	2914	No
29.569792ms	1.411529984	sea	64515,65534,20473,3356,2914,62240,62240,206924	2914	No
81.39008ms	1.492920064	sjc	64515,65534,20473,3257,2914,62240,62240,206924	2914	No
68.790016ms	1.56171008	sea	64515,65534,20473,2914,62240,62240,206924	2914	Yes
63.350016ms	1.625060096	ord	64515,65534,20473,3257,2914,62240,62240,206924	2914	No
57.07008ms	1.682130176	mia	64515,65534,20473,3257,2914,62240,62240,206924	2914	No
75.670016ms	1.757800192	ams	64515,65534,20473,2914,62240,62240,206924	2914	Yes
108.919808ms	1.86672	sjc	64515,65534,20473,2914,62240,62240,206924	2914	Yes
196.820224ms	2.063540224	mia	64515,65534,20473,2914,62240,62240,206924	2914	Yes
75.819776ms	2.13936	syd	64515,65534,20473,2914,62240,62240,206924	2914	Yes
54.889984ms	2.194249984	dfw	64515,65534,20473,2914,62240,62240,206924	2914	Yes
0.020224ms	2.194270208	ord	64515,65534,20473,2914,62240,62240,206924	2914	Yes
17.499904ms	2.211770112	sin	64515,65534,20473,2914,62240,62240,206924	2914	Yes



* the following findings are of my own and I encourage you to cross check this work, also I don't want to get sued by NTT and I still hope we can be friends Job S



The dying breaths





JUNOS

THE
FAST AND

THE **FURIOUS:**
PROPAGATION
DRIFT



JUNOS

THE
FAST AND
THE FURIOUS:
PROPAGATION
DRIFT



Questions?


THE
FAST AND
THE FURIOUS:
PROPAGATION
DRIFT

 @Benjojo12 /  AS206924

Questions?

ben@benjojo.co.uk

THE
FAST AND
THE FURIOUS:
PROPAGATION
DRIFT



bgp.tools

Using bgp.he.net?

Using bgpview.io?

Using radar.qrator.net?

**I've been building a
(opinionated but neutral)
routing table utility in my
spare time!**

**Check it out! I really
want feedback**



THE
FAST AND
THE FURIOUS:
PROPAGATION
DRIIFT