



Reviving BGP Zombies

Peering in the Routed Dead

Iliana Xytkou, Antonios Chariton, Fontas Dimitropoulos


What are BGP Zombies?



What are BGP Zombies?




What are BGP Zombies?



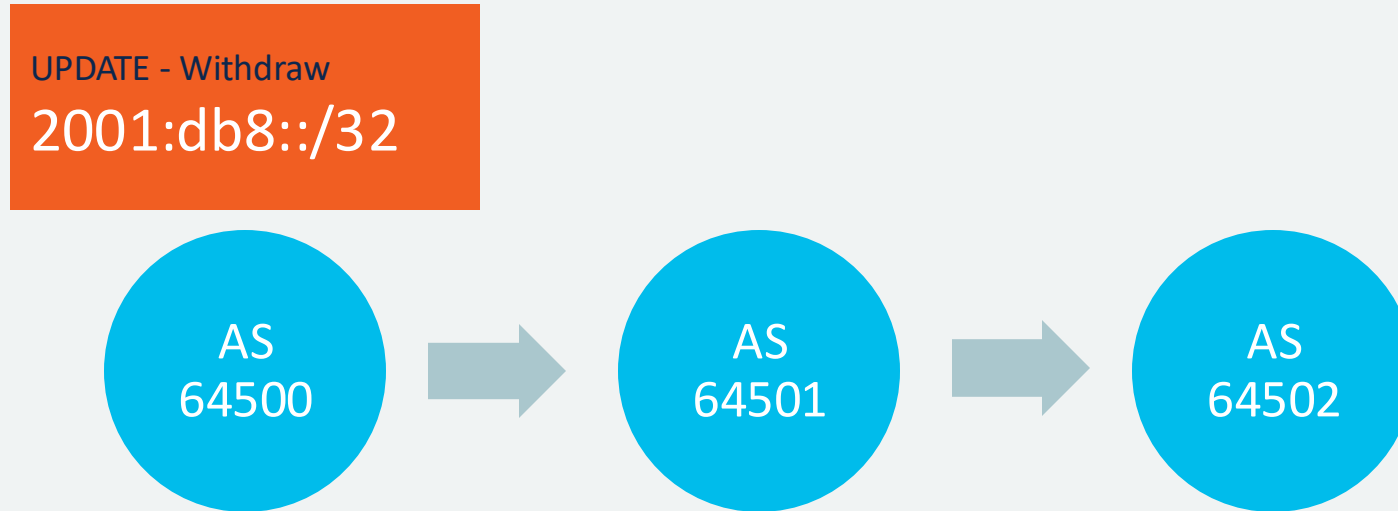
What are BGP Zombies?



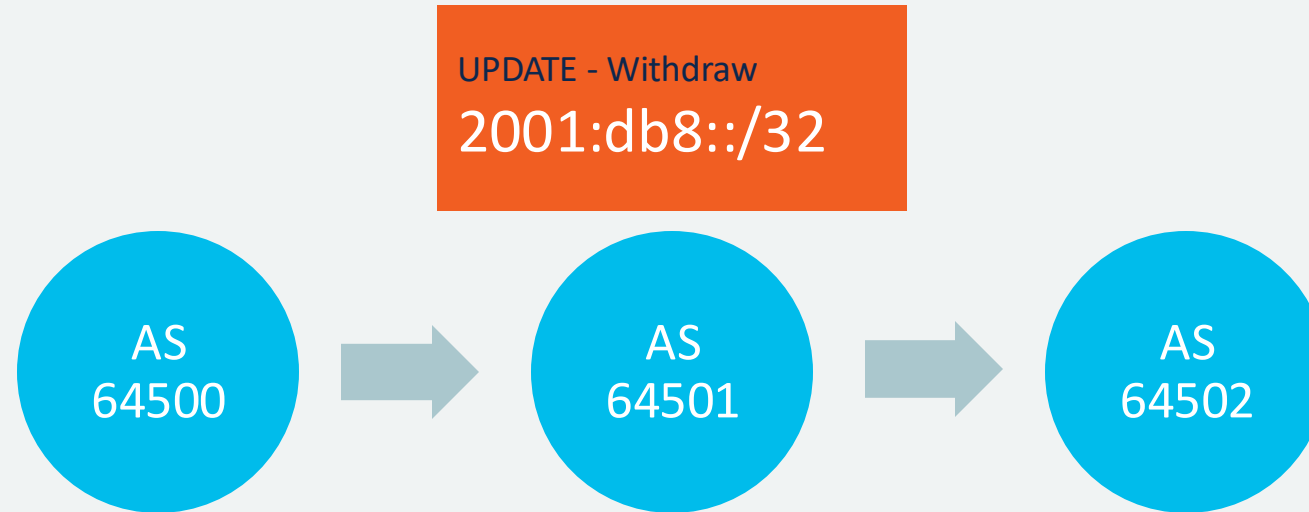
What are BGP Zombies?



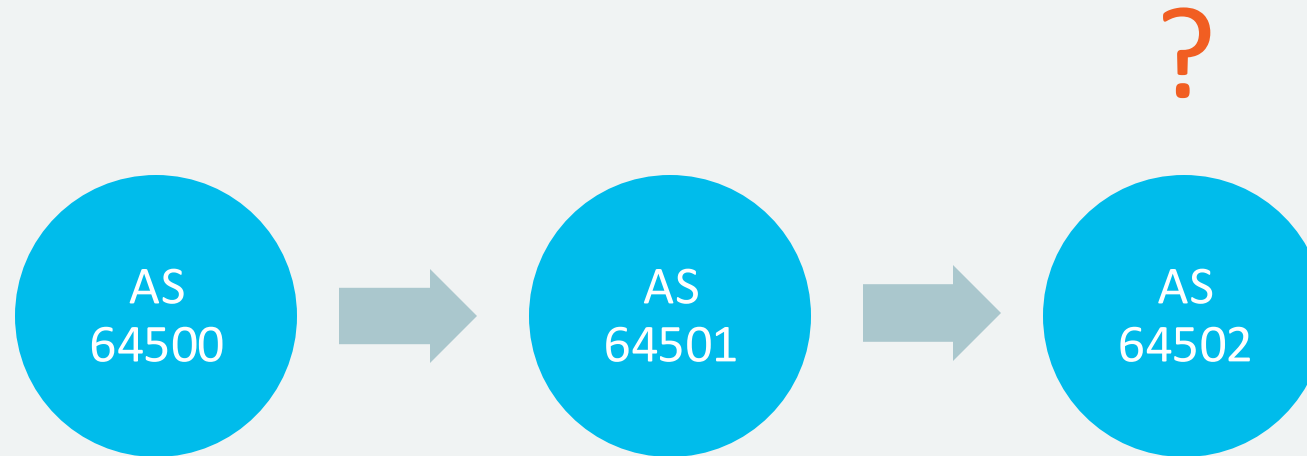
What are BGP Zombies?



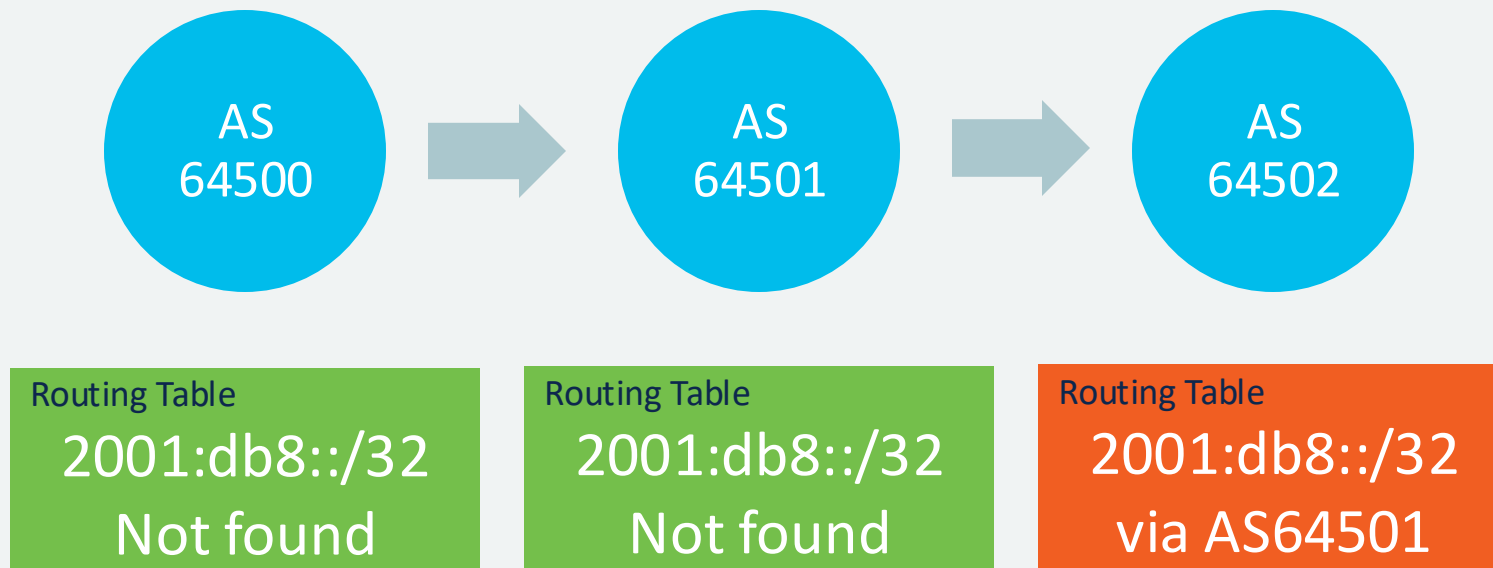
What are BGP Zombies?




What are BGP Zombies?



What are BGP Zombies?



What are BGP Zombies?



Why is this a problem?



Why is this a problem?

Effects of Zombies

Non-exhaustive

- Deaggregated prefixes for Traffic Engineering / DDoS Mitigation may remain in some locations
- Paths that are no longer there may appear and be used
 - Routing loops
 - Dropped traffic
- Sold or revoked IP space may cause small percentage of traffic being redirected
- False BGP Hijack Alerts / RPKI Invalids
- The number of prefixes on the Internet will keep going up



This is a problem both if **your** prefixes are stuck **and** if **others'** prefixes are stuck in **yours** or **your upstreams'** routers.



Why is this a problem?

UPDATE - Withdraw
2001:db8::/32



Why is this a problem?

UPDATE - Announce
2001:db8::/32



Is this a big problem?



Is this a big problem?

BGP Zombies: an Analysis of Beacons Stuck Routes

Romain Fontugne, Esteban Bautista, Colin Petrie, Yutaro Nomura, Patrice
Abry, Paulo Gonçalves, Kensuke Fukuda, Emile Aben



Is this a big problem?

RIPE RIS Beacons

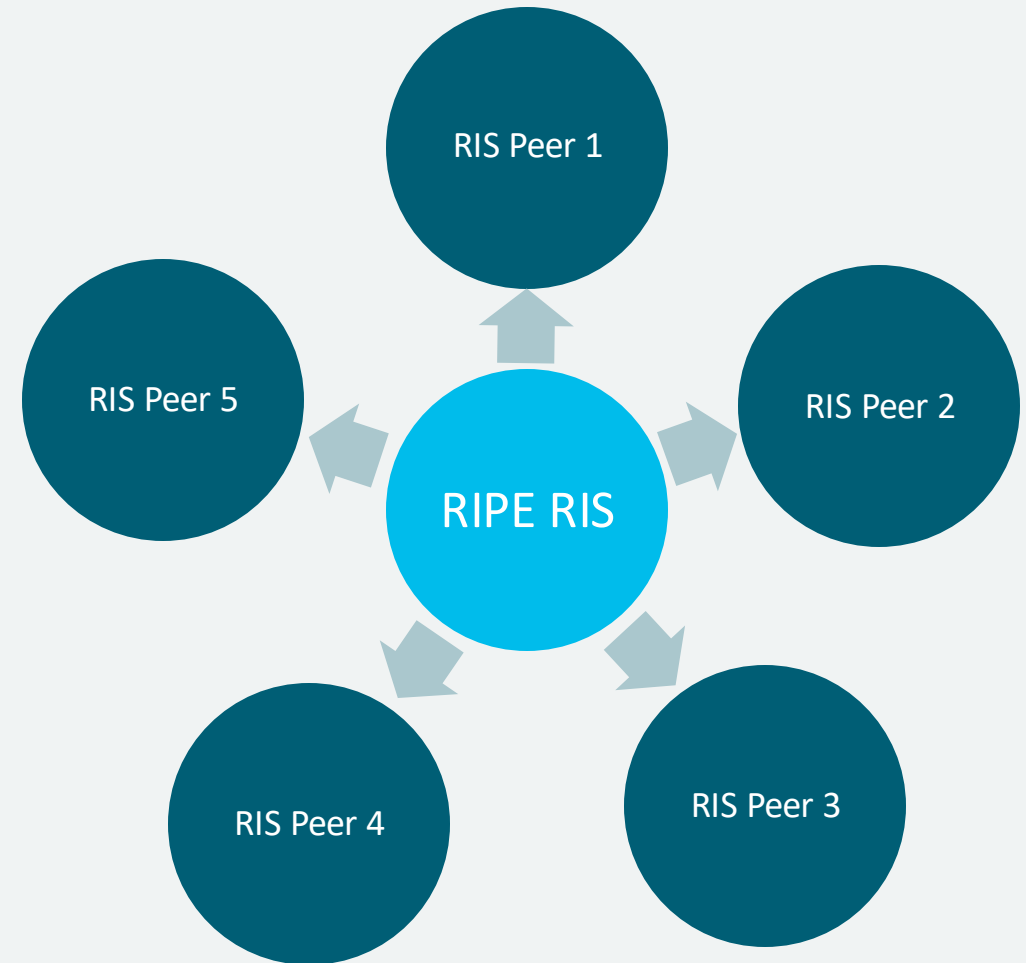
- Announce its prefix every 4 hours (00:00, 04:00, ...)
- Withdraw the prefix 2 hours later (02:00, 06:00, ...)

UPDATE - Announce

```
2001:7fb:fe01::/48  
84.205.65.0/24
```


UPDATE - Withdraw

```
2001:7fb:fe01::/48  
84.205.65.0/24
```




Is this a big problem?

RIPE RIS Beacons



Is this a big problem?

RIPE RIS Beacons



Is this a big problem?

Zombie Outbreaks

Fontugne et al, 2019

Start	End	IPv6 Outbreaks	IPv4 Outbreaks
2017-03-01	2017-04-28	591	1'732
2017-10-01	2018-12-28	1'202	384
2018-07-19	2018-08-31	686	520



Is this a big problem?

Zombie Outbreaks

Fontugne et al, 2019

Start	End	IPv6 Outbreaks	IPv4 Outbreaks
2017-03-01	2017-04-28	591	1'732
2017-10-01	2017-12-28	1'202	384
2018-07-19	2018-08-31	686	520



Is this a big problem?

Zombie Outbreaks

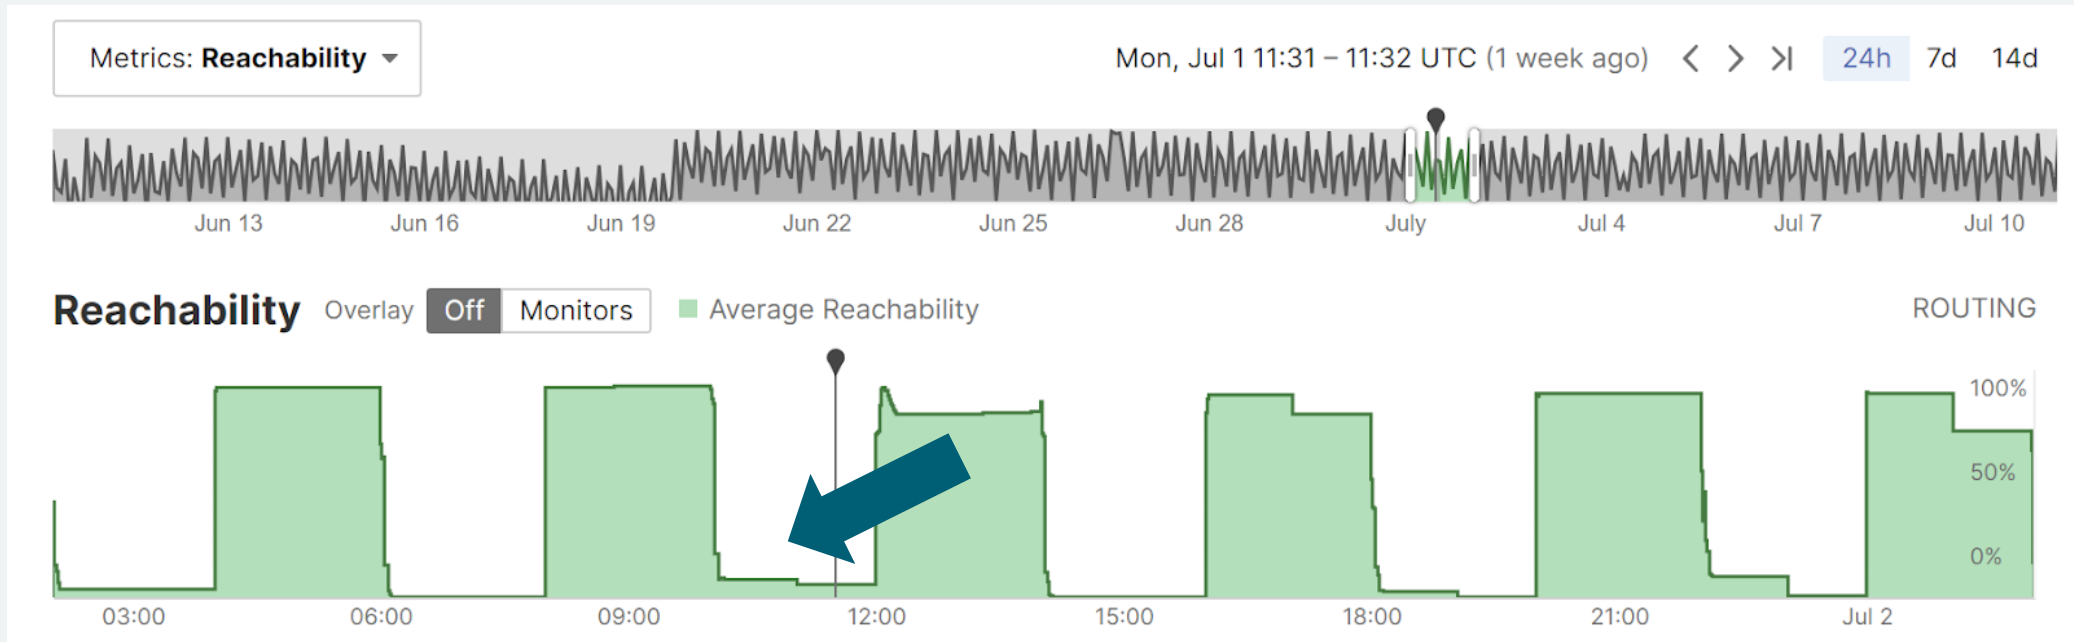
Start	End	Study IPv6 Outbreaks	Study IPv4 Outbreaks	Our IPv6 Outbreaks	Our IPv4 Outbreaks
2017-03-01	2017-04-28	591	1'732	610	1'781
2017-10-01	2017-12-28	1'202	384	1'378	705
2018-07-19	2018-08-31	686	520	745	536

We are using RIB Dumps, UPDATEs, and STATE
The study used the LG and filtered the results with UPDATEs




Is this a big problem?

RIPE RIS Beacons




Is this a big problem?

RIPE RIS Beacons



Is this a big problem?

Aggregator Address



10.XX.YY.ZZ

└──────────────────────────┘
24-bit seconds since
beginning of month UTC

01/11/24 11:00 CET (~Now)
10.[36'000] = 10.0.140.160



Is this a big problem?

Zombie Outbreaks

Start	End	Study IPv6 Outbreaks	Study IPv4 Outbreaks	Our IPv6 Outbreaks	Our IPv4 Outbreaks	IPv6 Outbreaks	IPv4 Outbreaks
2017-03-01	2017-04-28	591	1'732	610	1'781	610	1'319
2017-10-01	2017-12-28	1'202	384	1'378	705	1'370	478
2018-07-19	2018-08-31	686	520	745	536	514	226



Is this a big problem?

Research Findings

- ~19% of RIPE RIS <PeerAS, BeaconPfx> were not affected by BGP Zombies
- 50% of pairs have ~0.26% probability of falling for them
- On average, ~1.6% of IPv6 and ~0.5% of IPv4 probability for a pair to see Zombies
- Over 90-95% of the time, the Zombie path was NOT the best path




Can we improve our understanding?



Can we improve our understanding?

RIPE RIS Beacons



BGP Clock



Can we improve our understanding?

2a0d:3dc1:HHMM::/48

Every 15'



Can we improve our understanding?

2a0d:3dc1:(HH)(MM+dd%15)::/48

Every 15'



Can we improve our understanding?



RPKI ROA

2a0d:3dc1::/32-48 AS210312



Can we improve our understanding?

BGP Clock

- Prefixes recycled every 24h / 15d, not every 4 hours
 - Allows us to see beyond the 1h30m – 2h mark into the unknown
 - Many more prefixes – 4 / Hour -> More data to study
 - Originated from AS210312 to over 1'700 direct adjacencies
-
- Ben Cartwright-Cox called this the Route Cycler, as I found out later :)



Can we improve our understanding?

2a0d:3dc1:2327::/48

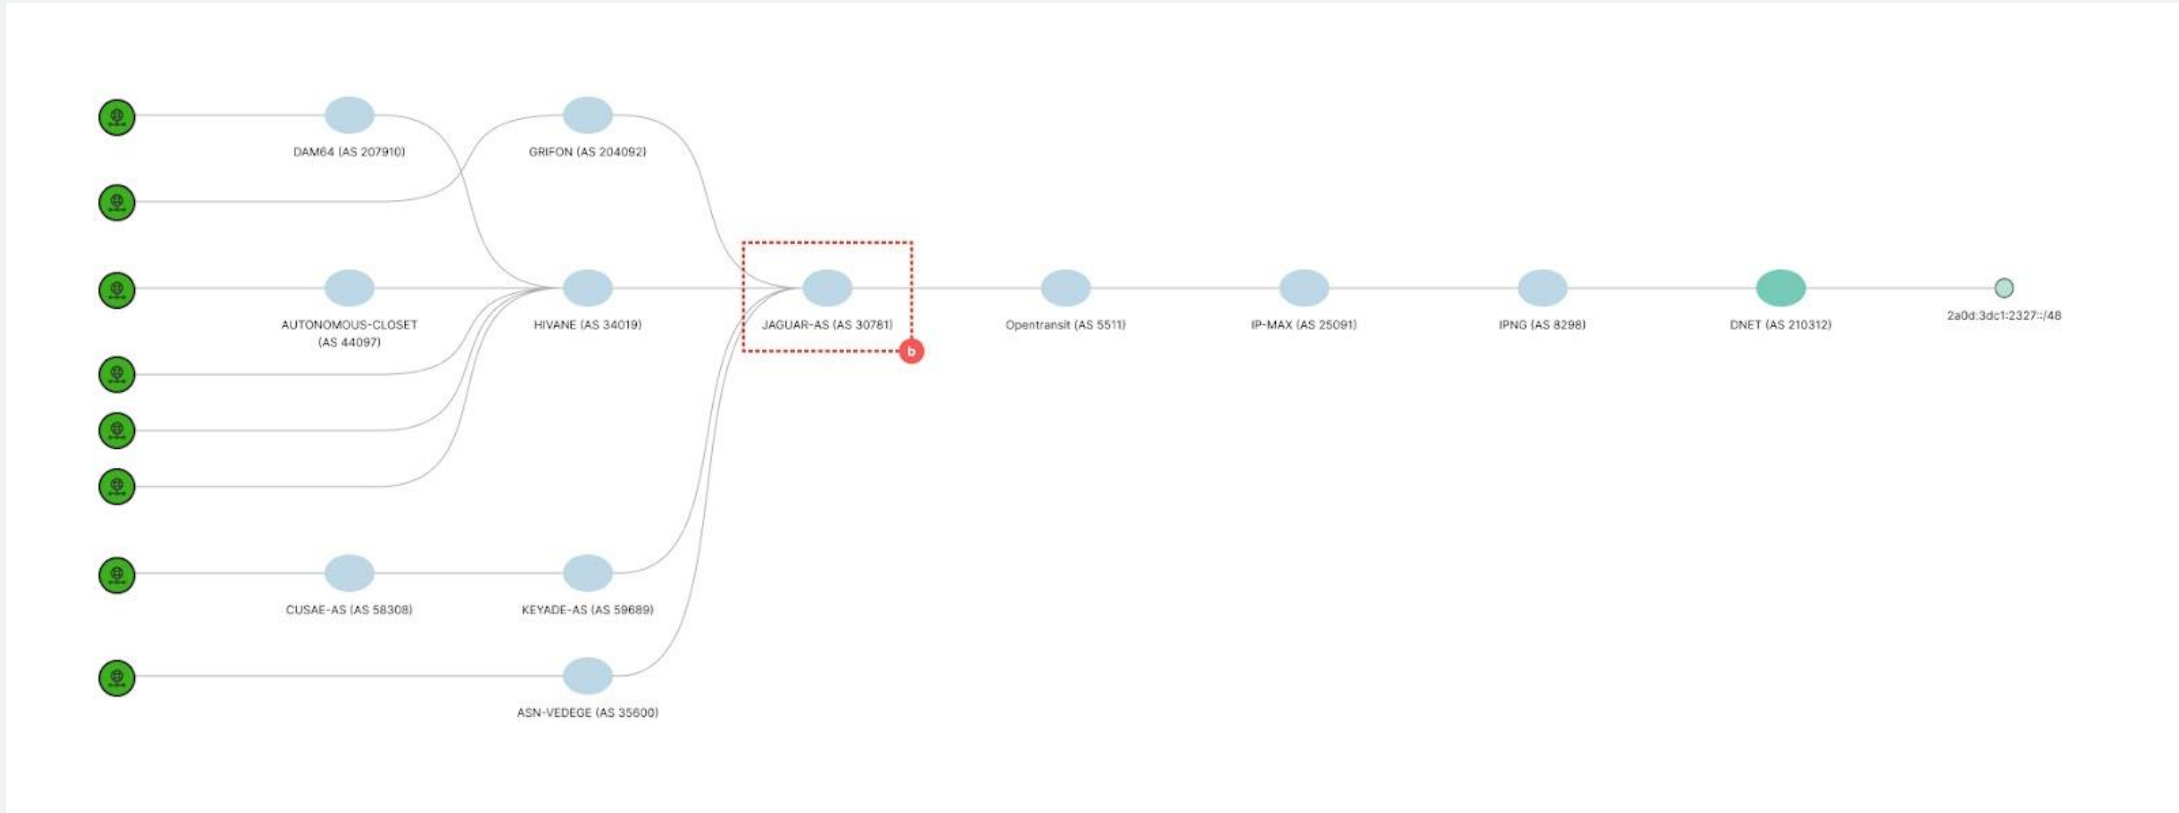
Findings

- Stuck in 8 RIPE RIS Peers (6 Unique ASNs)
- Common Subpath: 30781 5511 25091 8298 210312
- Stuck in Free Pro SAS in France (>200 ASNs in Cone)
 - Probably all 200+ ASes in the Cone were “infected”, but only 6 had RIS Peers



Can we improve our understanding?

Visualization



Can we improve our understanding?

2a0d:3dc1:2233::/48

Findings

- Stuck in 24 RIPE RIS Peers (21 Unique ASNs)
- Common Subpath: 33891 25091 8298 210312
- Stuck in Core Backbone GmbH in Germany (>2'000 ASNs in Cone)
 - Probably all ASes in the Cone were “infected”, but only 21 had RIS Peers



Can we improve our understanding?

2a0d:3dc1:1737::/48

Findings

- Stuck in 7 RIPE RIS Peers
- Common Subpath: 24961 210312
- Stuck in WIIT AG / myLoc in Germany (>200 ASNs in Cone)
 - Probably all 200+ ASes in the Cone were “infected”, but only 7 had RIS Peers



Can we improve our understanding?




RPKI ROA

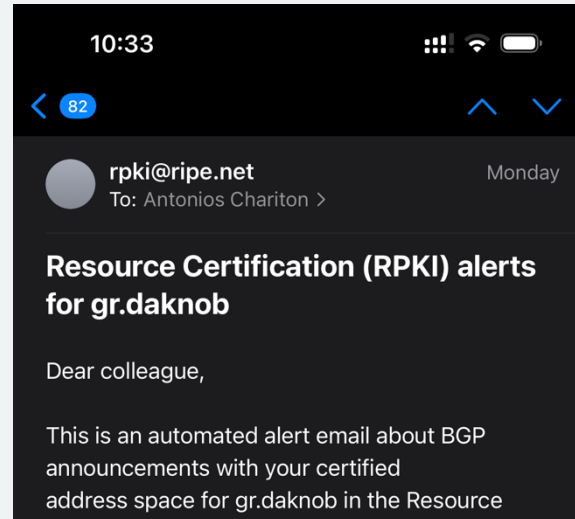
2a0d:3dc1::/32-48 AS210312




Can we improve our understanding?



Can we improve our understanding?



3.5 Months Later!



Can we improve our understanding?

Routes still stuck...

Over 4 months later, we can see:


- ~~1~~ in RIPE RIS (Lasted 2.5 mo)
- 3 in bgp.tools
- ~~4~~ 2 in bgp.he.net (Half life of 3 months!)
- RPKI Invalid for > 3-4 months `~_(\ツ)_/~`




What did we learn?



Thresholds Matter



Thresholds Matter



What did we learn?


Noisy Peers

Peer IP	ASN	1h30m Threshold	3h Threshold
176.119.234.201	211509 Rudakov Ihor	9.91%	9.06%
2001:678:3f4:5::1	211509 Rudakov Ihor	9.91%	9.06%
2a0c:9a40:1031::504	211380 Simulhost Limited	7%	6.88%



What did we learn?

Thresholds Matter



What did we learn?

Stuck routes up over time?

What happened

- We observed zombies increasing at ~160'
- Prefixes that withdrew 10' earlier, are coming back
- There's a new Announcement!
- Common subpath: 4637 1299 25091 8298 210312
- Telstra Global, with >5'000 ASNs in Cone
- Session reset? Filter update?
- Reinfections can happen!



Can we do something?



Can we do something?

Yes!

draft-ietf-idr-bgp-sendholdtimer:

- Practically an RFC already
- Adds a SendHoldTimer in addition to the HoldTimer
- Tears down sessions if messages can't be sent (not just received)
- Addresses XX% of stuck route causes
- **Ask for support from your vendor!**
- Has to be added, included in stable releases, operators have to upgrade, ???, profit!



Learn more

