Extending BGP Analysis with NetFlow

Doug Madory, dmadory@kentik.com

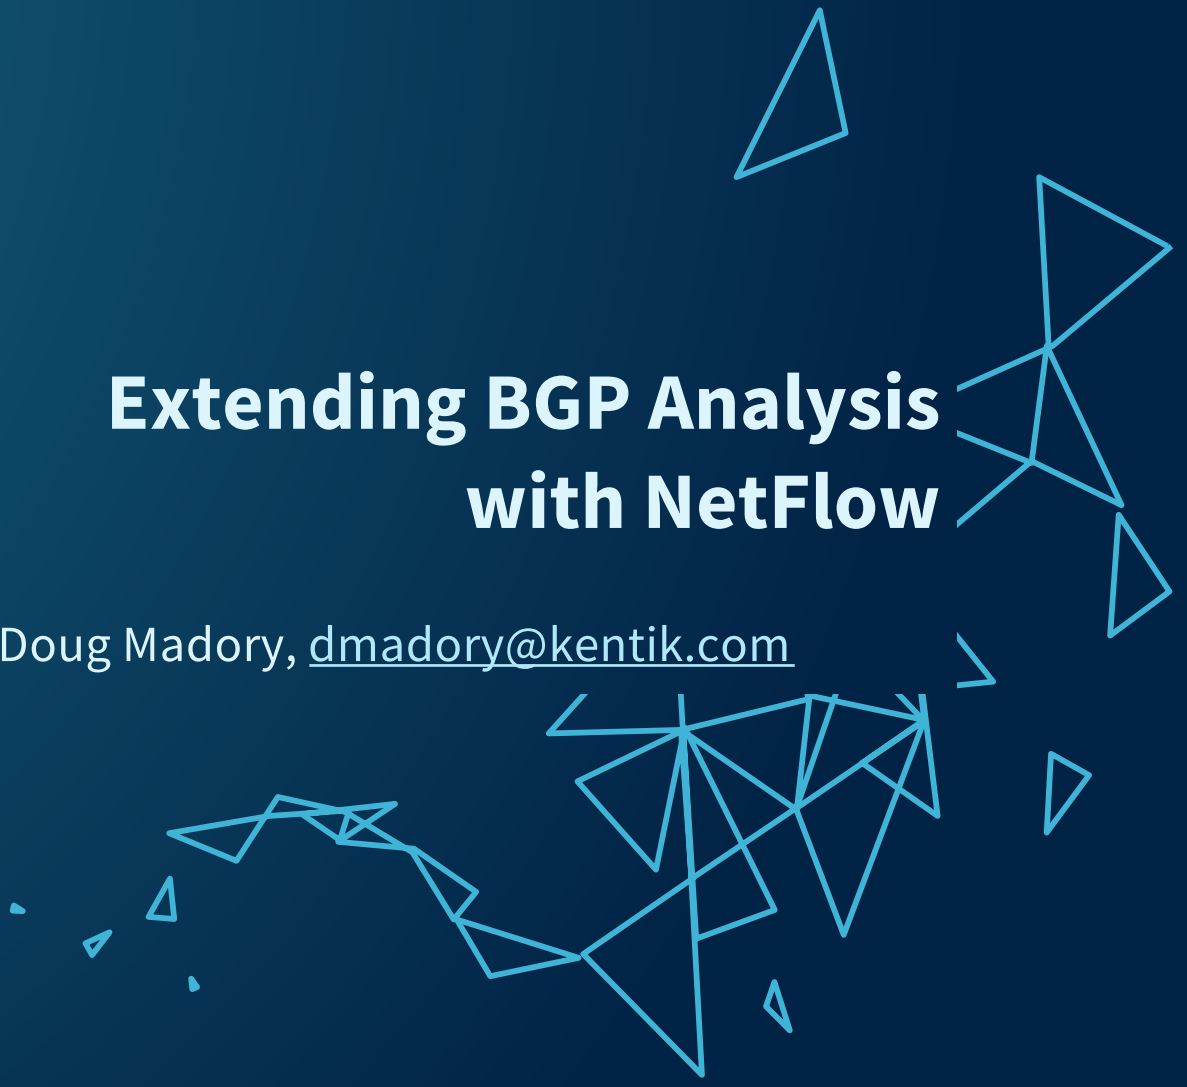

If a route falls in the forest and
no packet is around to hear it...

Zen Koan (updated)

A little about me…

BS in Computer Engineering

MS in Computer Engineering

Communications Engineer in US Air Force

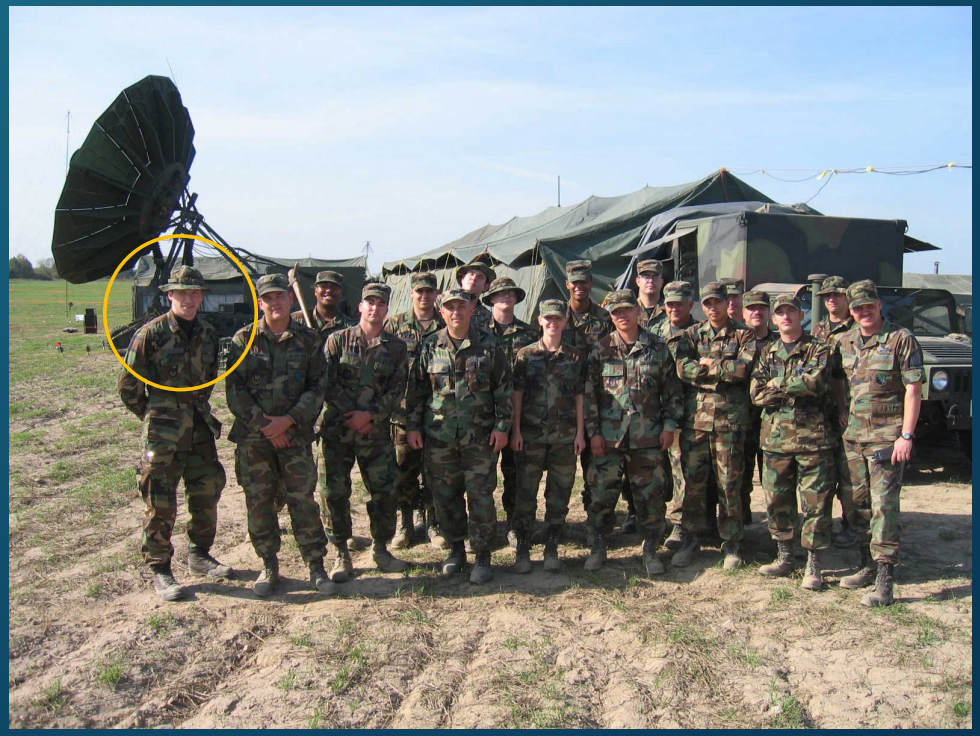

NATO Air Meet (Sep 2003)
Poznan, Poland

Al Udeid Air Base, Doha, Qatar
(following deployment to Iraq)


S.M.A.

Started work on BGP analysis in 2009

- **Renesys** (small startup working on internet measurement)
  - **Dyn Research** after acquisition by Dyn (2014)
  - **Oracle Internet Intelligence** after acquisition by Oracle (2017)
- Now **Kentik** (since 2020)

We published *many* BGP-related stories over the years

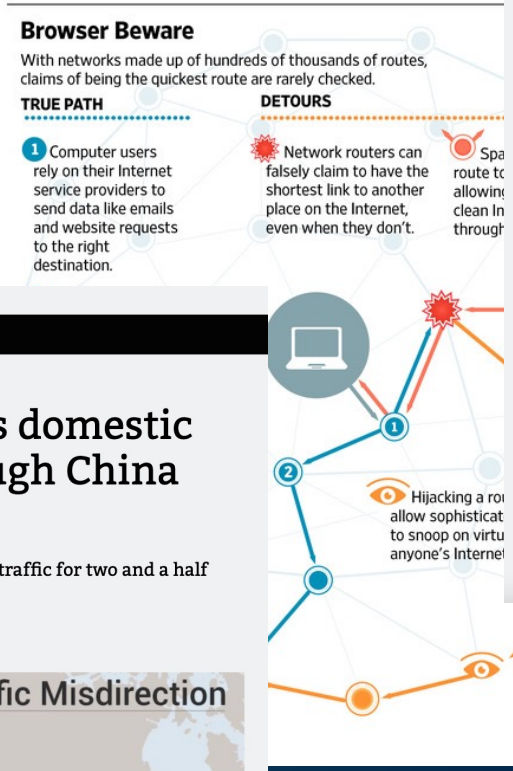

Analysis of those BGP incidents depended on traceroute

```
Traceroute from Helsinki to Ministry of Foreign
Affairs of Lithuania (May 23, 2013)
1   *
2   62.78.114.228   Helsinki, Finland     0.519
3   62.78.111.198   Helsinki, Finland     0.508
4   62.78.107.128   Tampere, Finland      8.669
5   62.78.107.135   Tampere, Finland     14.401
6   62.78.107.51    Tampere, Finland      8.694
7   194.68.123.212  Stockholm, Sweden    21.758
8   217.150.62.234  Moscow, Russia      156.642
9   217.150.62.233  Minsk, Belarus       44.710
10  84.15.6.213     Vilnius, Lithuania   66.443
11  213.226.128.18  Vilnius, Lithuania   66.613
12  195.22.173.222  Ministry of Foreign  68.120
                    Affairs of Lithuania
```

Traceroute during Belarus MITM BGP hijack

Legitimate route: … 13194 **24825**

Hijack route: … 20485 **6697** 56498

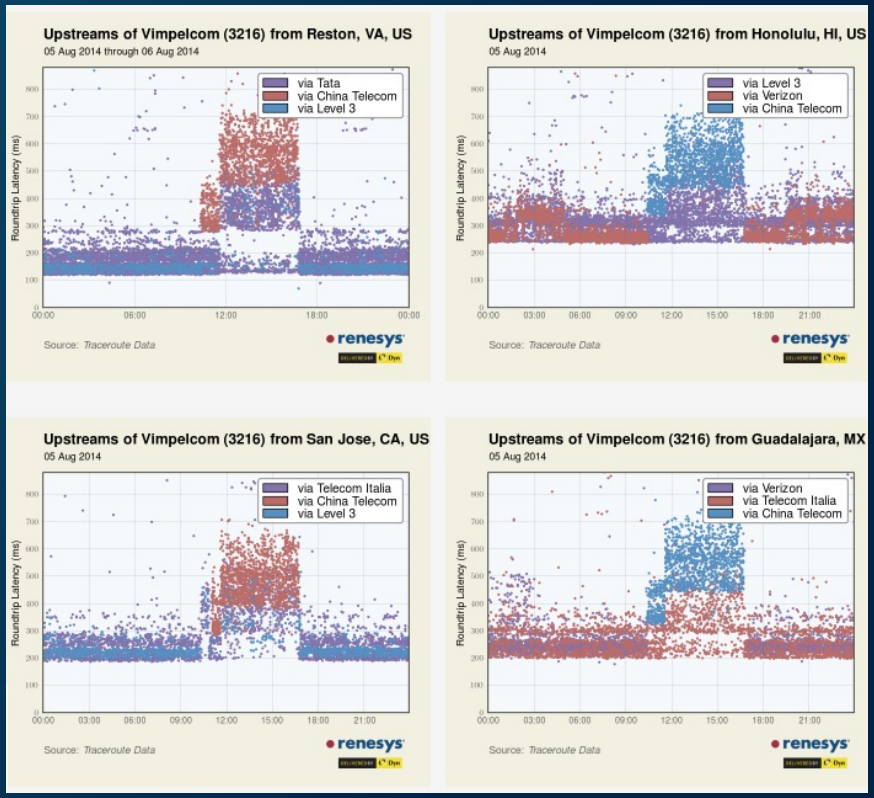

Latency path impact of China Telecom leaks impacting Vimpelcom of Russia.

But did a *single packet** get misdirected?

**non-measurement*

Now to add NetFlow into the Analysis

Bringing in an additional data source to understand operational impact

- Kentik has over 300 customers and almost half have opted-in to the use of their data as part of *aggregate analysis*.
  - Note: analysis is subject to biases of the customer set which includes (NSPs, CDNs and enterprises) and is skewed toward the US.
- Helps to answer questions that BGP and active measurement cannot.

BGP+NetFlow: Measurement the state of RPKI ROV

- Enormous progress in recent years as Tier-1 NSPs agreed to reject RPKI-Invalids.
  - NTT, GTT, Arelion (Telia), Cogent, Telstra, PCCW, Lumen, and more!
- According to NIST RPKI Monitor, the trend line is going in the right direction!

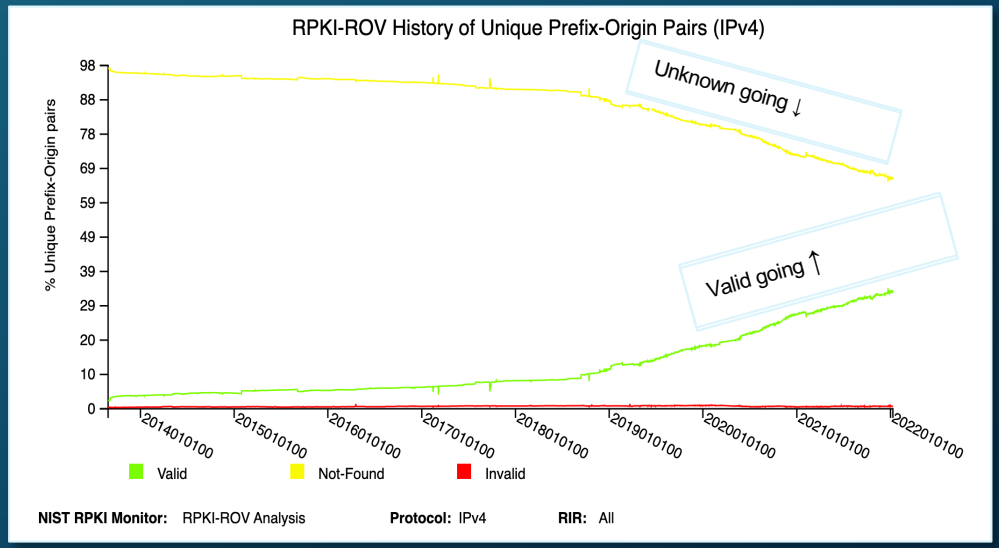

https://rpki-monitor.antd.nist.gov

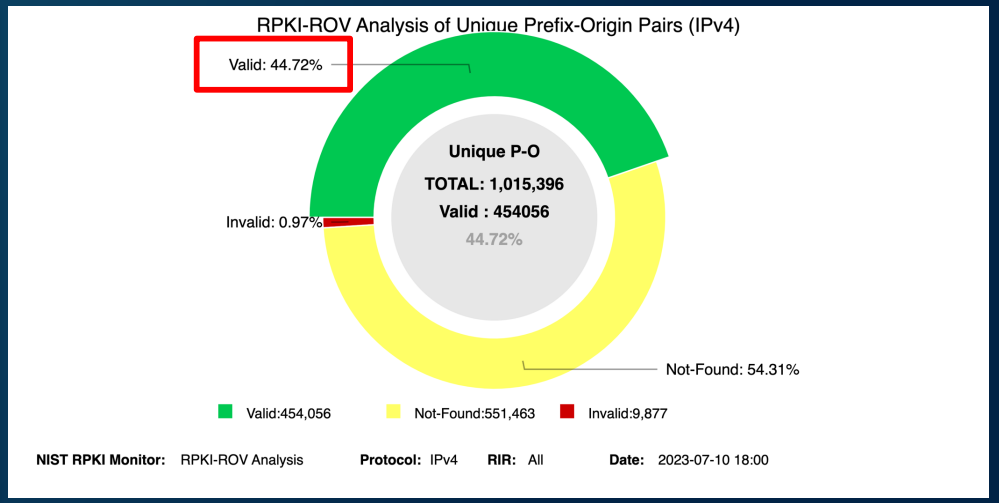

In February 2022, rate was only 34.1%.

What proportion of overall traffic is safeguarded* by that 44.72%?

* Eligible for the protection of RPKI

More than one might otherwise think

- 62.5% of traffic (bits/sec) going to “valid BGP routes”
  - Subject to biases in the data
- Much higher than count of BGP routes, IP space
- Due to major RPKI deployments at:
  - Content providers (Amazon, Google, Cloudflare, etc)
  - Eyeball networks (Comcast, Spectrum)
- *Some routes handle much more traffic than others*

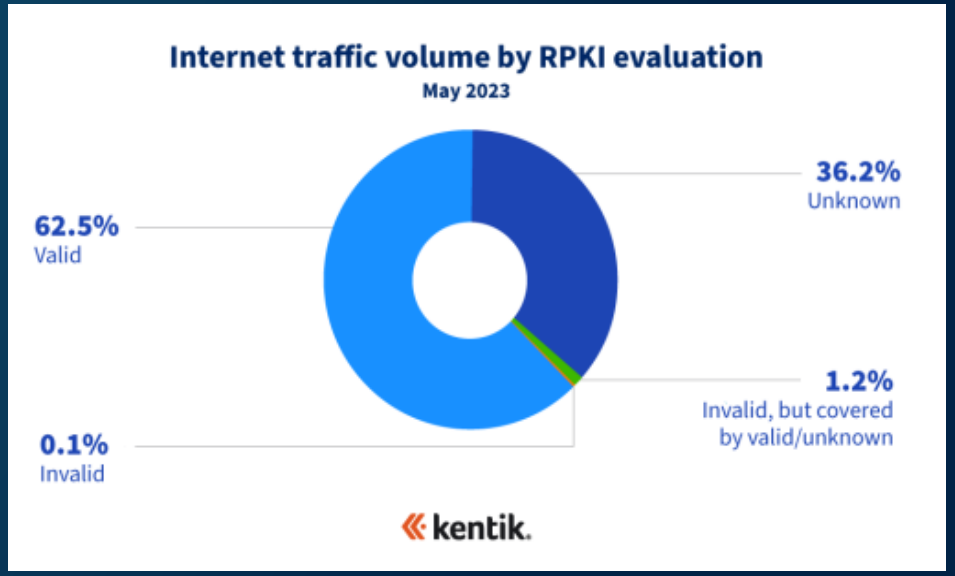

https://www.kentik.com/blog/exploring-the-latest-rpki-rov-adoption-numbers/

What about rejection of RPKI-invalid routes?

- ROAs alone are useless if only a few networks are rejecting invalid routes.
- Recent analysis shows propagation of RPKI-invalid routes is half or less than other types.

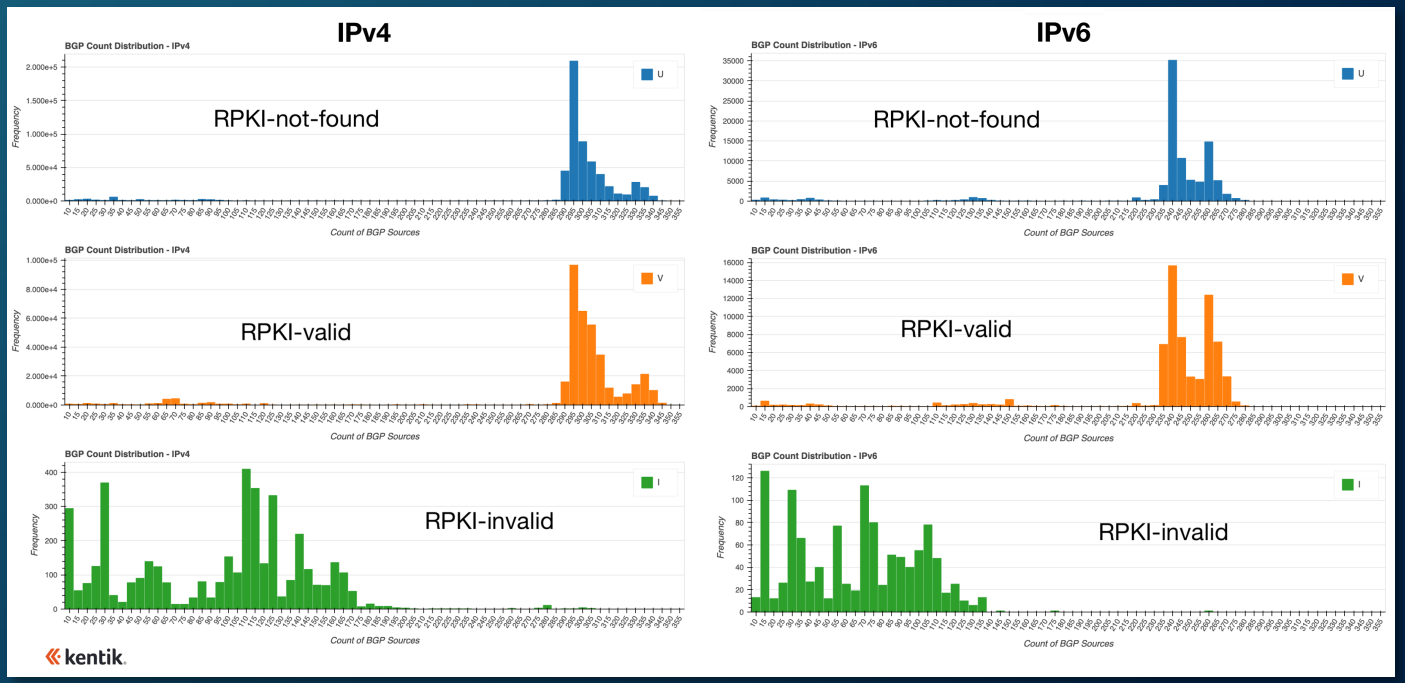

https://www.kentik.com/blog/how-much-does-rpki-rov-reduce-the-propagation-of-invalid-routes/

Using NetFlow to explore a BGP incident

BGP Hijacks Targeting Cryptocurrency Services

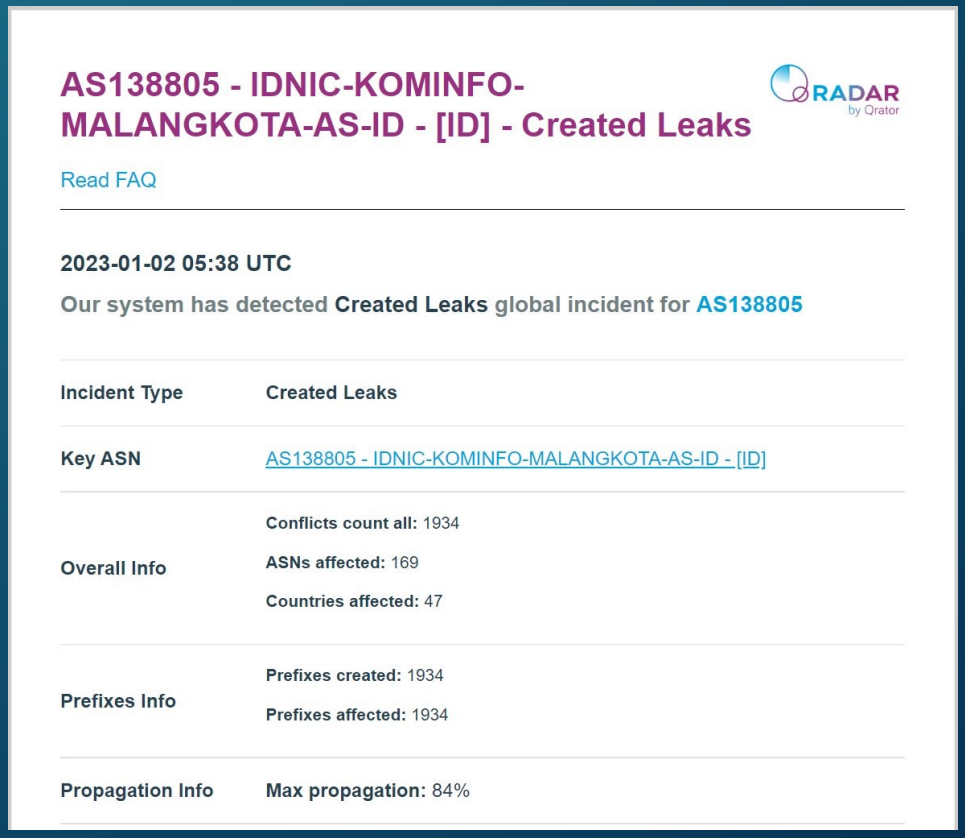

AS138805 - IDNIC-KOMINFO-MALANGKOTA-AS-ID - [ID] - Created Leaks

RADAR by Qrator

Read FAQ

2023-01-02 05:38 UTC

Our system has detected **Created Leaks** global incident for **AS138805**

Incident Type	**Created Leaks**
Key ASN	AS138805 - IDNIC-KOMINFO-MALANGKOTA-AS-ID - [ID]
Overall Info	**Conflicts count all:** 1934 **ASNs affected:** 169 **Countries affected:** 47
Prefixes Info	**Prefixes created:** 1934 **Prefixes affected:** 1934
Propagation Info	**Max propagation:** 84%

- AS138805 of Indonesia leaked several thousand routes learned from one transit provider (TELIN, AS7713) to another transit provider (Lintasarta, AS4800)
  - 05:37 UTC on 2 January 2023
- Biggest impact was to “regional routes”
  - BGP routes with intentionally limited propagation
  - When a leak occurs, there is nothing for the leaked route to compete against.

BGP impact of this leak

- The upper stacked time series is a measure of route propagation over time.
- Shows how our BGP sources reach this prefix by each upstream of the origin.

- Normally about 20% of our sources see this route at all (17.4% via Hurricane Electric).
- During the leak it jumps up to about 70% (with 51.7% suddenly seeing it via the Toronto Internet exchange TorIX, AS11670).

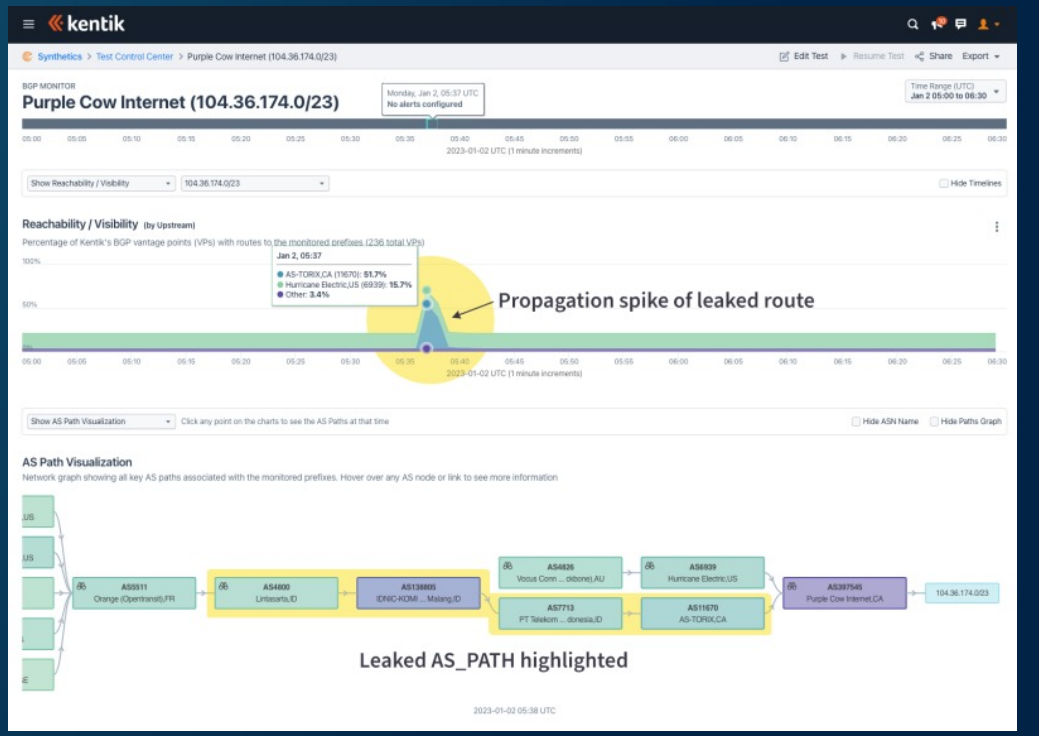

What was the operational impact?

Aggregate NetFlow view of the leak

- NetFlow is annotated with AS_Path of src and dst IPs, from perspective of router
- Enables user to query for NetFlow records matching the AS_Path leak subsequence.

- Then we can discover *a portion of* misdirected internet traffic by country, dest ASN, among other dimensions.

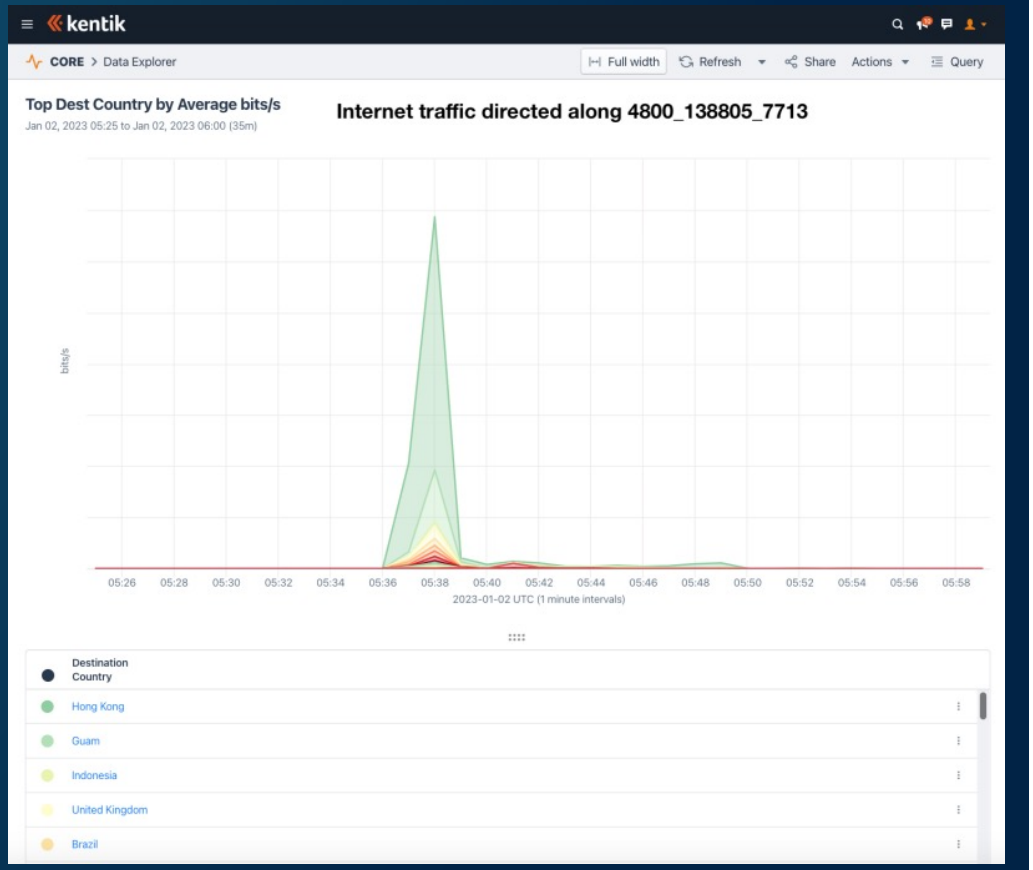

Aggregate NetFlow view of the leak

- Query for NetFlow records destined for the IP space in the most propagated leaked routes.
- Reveals two discrete impacts:
  1. A drop in traffic due to packet loss
  2. Separately, the portion of the traffic that followed “4800 138805 7713”
- No hard-and-fast rule about these two types of impacts. Varies by incident.

More examples here: https://www.kentik.com/blog/new-year-new-bgp-leaks/

NetFlow helps us to better understand:

1. RPKI deployment
2. Operational impact of leaks/hijacks

BGP Hijacks Targeting Cryptocurrency Services

- Attack against Celer Bridge (August 2022)
- Previous attacks against cryptocurrency services
  - Etherwallet (Apr 2018)
  - Klayswap (Feb 2022)
- What can be done to prevent these attacks?

https://www.kentik.com/blog/bgp-hijacks-targeting-cryptocurrency-services/

Attack against Celer Bridge (August 2022)

- Celer Bridge is a service which allows users to convert between cryptocurrencies.
- Attacker used a BGP hijack to gain control of a portion of Amazon's IP address space hosting Celer Bridge infrastructure.
- Hijack allowed attacker to impersonate part of the Celer Bridge infrastructure.
- Attacker issued malicious smart contracts, redirecting digital assets to attacker's wallet.

Why was this BGP hijack successful?

See: https://www.coinbase.com/blog/celer-bridge-incident-analysis

Attack against Celer Bridge (August 2022)

- Attacker needed to ensure malicious BGP announcements wouldn't get filtered
  1. Inserted bogus route objects for QuickhostUK in AltDB (free RIR alternative)

```
irrd.log-20220817.gz:31106270-ADD 96126
irrd.log-20220817.gz:31106280-
irrd.log-20220817.gz:31106281-as-set:     AS-SET209243
irrd.log-20220817.gz:31106306-descr:      quickhost set
irrd.log-20220817.gz:31106332-members:    AS209243, AS16509
irrd.log-20220817.gz:31106362:mnt-by:     MAINT-QUICKHOSTUK
irrd.log-20220817.gz:31106392-changed:    crussell()quickhostuk net 2022081
irrd.log-20220817.gz:31106438-source:     ALTDB
<br />
irrd.log-20220817.gz:31147549-ADD 96127
irrd.log-20220817.gz:31147559-
irrd.log-20220817.gz:31147560-route:      44.235.216.0/24
irrd.log-20220817.gz:31147588-descr:      route
irrd.log-20220817.gz:31147606-origin:     AS16509
irrd.log-20220817.gz:31147626:mnt-by:     MAINT-QUICKHOSTUK
irrd.log-20220817.gz:31147656-changed:    crussell()quickhostuk net 2022081
irrd.log-20220817.gz:31147702-source:     ALTDB
```

Credit: Siyuan Miao of Misaka on NANOG list

Attack against Celer Bridge (August 2022)

- Attacker needed to ensure malicious BGP announcements wouldn't get filtered
  1. Inserted bogus route objects for QuickhostUK in AltDB (free RIR alternative)
  2. Attacker altered the AS_PATH to appear to be originated by an Amazon ASN.

AS_Path: … 1299 209243 14618

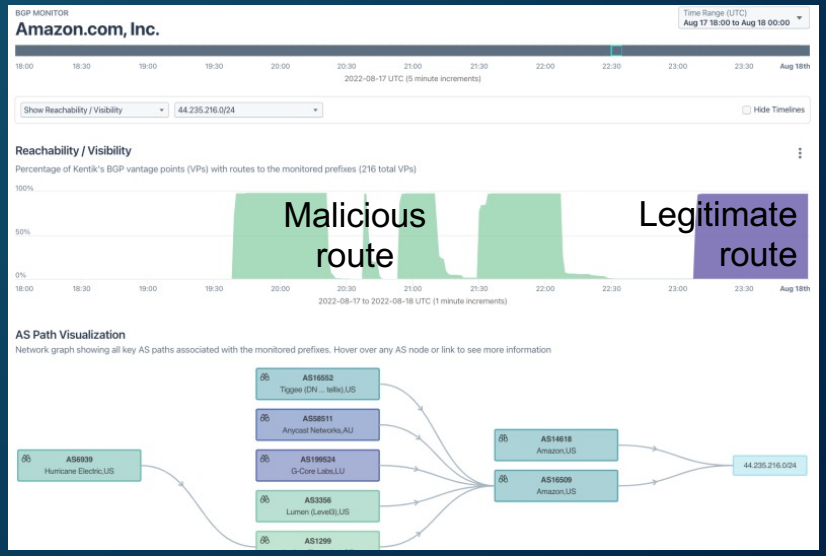

https://twitter.com/DougMadory/status/1562089866321698819

Attack against Celer Bridge (August 2022)

- Attacker needed to ensure malicious BGP announcements wouldn't get filtered
  1. Inserted bogus route objects for QuickhostUK in AltDB (free RIR alternative)
  2. Attacker altered the AS_PATH to appear to be originated by an Amazon ASN.

- Amazon didn't begin announcing this identical /24 until 23:07 UTC (in purple), an hour after the last hijack was finished.
- According to Coinbase's timeline, victims had cryptocurrency stolen in separate events between 19:51 and 21:49 UTC.

Previous attacks against cryptocurrency services

- **Apr 2018** Amazon's authoritative DNS service was hijacked in order to redirect certain DNS queries to an imposter DNS service, as is illustrated below.
  - Imposter auth DNS server returned bogus responses for myetherwallet.com, misdirecting users to an imposter version of MyEtherWallet's website.

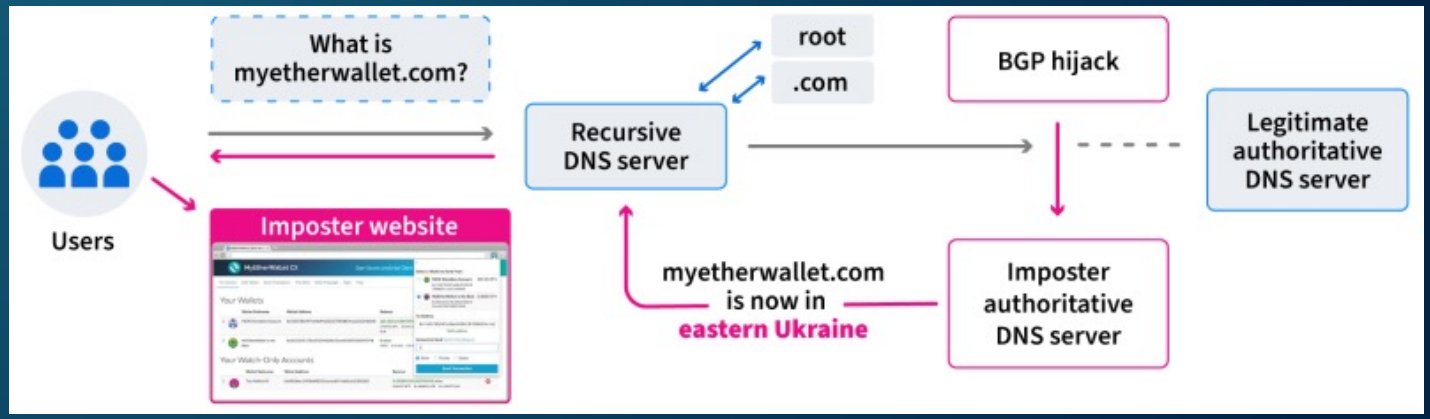

Previous attacks against cryptocurrency services

- **Apr 2018** Amazon's authoritative DNS service was hijacked in order to redirect certain DNS queries to an imposter DNS service, as is illustrated below.
- **Feb 2022** Attackers went after the users of the KLAYswap cryptocurrency exchange by performing a BGP hijack of the IP space of a South Korean hosting provider (Kakao).
  - Kakao was hosting a javascript library that was loaded when users were on KLAYswap.
  - BGP hijack enabled the attackers to impersonate Kakao and return a malicious code.

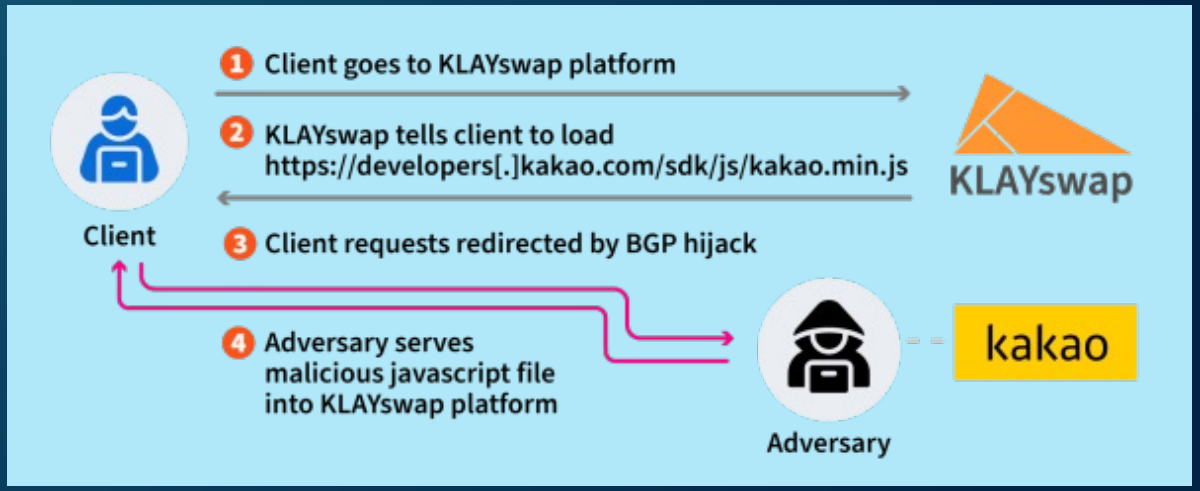

How do we prevent these attacks?

- *While these incidents involved cryptocurrency services, the underlying issues are universal and can affect any organization that uses internet-based services.*
- Monitoring
  - DNS – fire off alert if agent elicits a response doesn't match expected results
  - BGP – unexpected upstream of AS209243 for Amazon should have been suspicious
- RPKI ROV
  - Amazon had an ROA for the prefix that was hijacked, so why didn't RPKI ROV help?

How do we prevent these attacks?

- RPKI ROV
  - Amazon had an ROA for the prefix that was hijacked, so why didn't RPKI ROV help?
  - Very liberal ROAs: 3 different Amazon ASNs can all announce parts of this address space with prefixes ranging in size from a /10 all the way down to a /24

See **RFC 9319** : The Use of maxLength in the Resource Public Key Infrastructure (RPKI)

Leaving the maxLength field blank in a ROA has the same effect as setting the maxLength field to match the prefix.

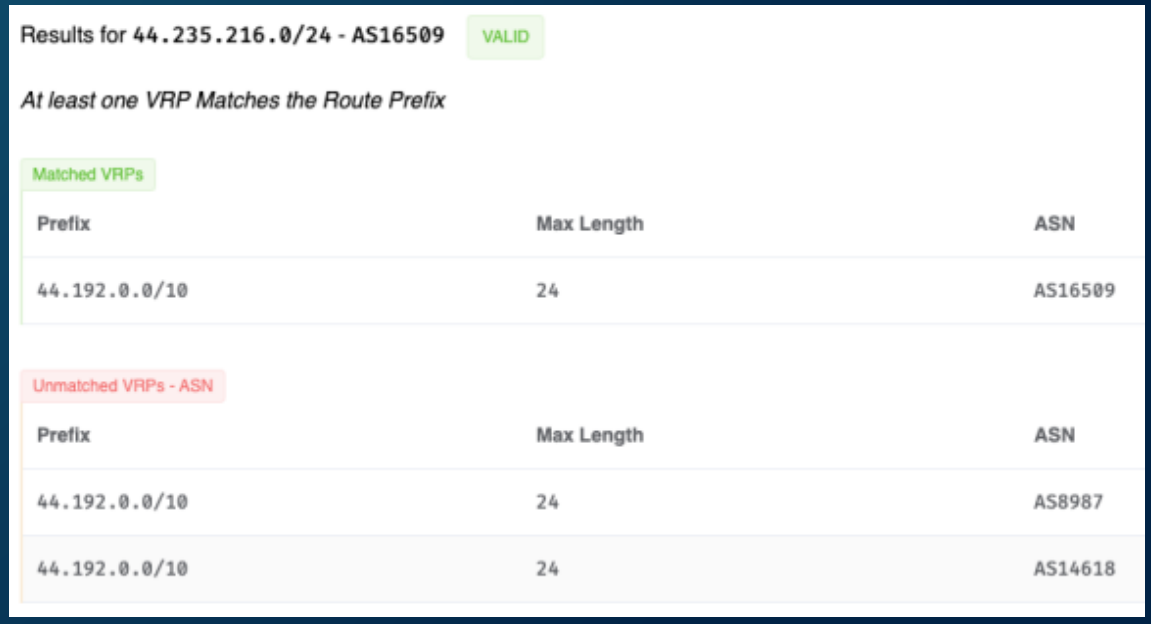

Results for `44.235.216.0/24` - `AS16509` VALID

At least one VRP Matches the Route Prefix

Matched VRPs

Prefix	Max Length	ASN
44.192.0.0/10	24	AS16509

Unmatched VRPs - ASN

Prefix	Max Length	ASN
44.192.0.0/10	24	AS8987
44.192.0.0/10	24	AS14618

How do we prevent these attacks?

- RPKI ROV
  - Amazon had an ROA for the prefix that was hijacked, so why didn't RPKI ROV help?
  - Very liberal ROAs: 3 different Amazon ASNs can all announce parts of this address space with prefixes ranging in size from a /10 all the way down to a /24
- Need BGPSEC to eliminate impersonation of ASes.
  - Protection only extends via contiguous BGPSEC-aware ASes
  - Adoption by major cloud providers and network service providers alone could severely limit the efficacy of AS impersonations
  - Partial deployment does offer benefits.

Facts I'd like to become common knowledge in networking:

1. The majority of internet traffic is directed to RPKI-valid routes,

2. Route propagation is cut in half when evaluated as RPKI-invalid.

We must keep marching forward

Expect to hear from a certain routing security evangelist soon!

- Peerlock ✅
- Using RPKI to cleanup IRR ✅
- RPKI ✅
  - To reduce impacts of fat-fingers.
- BGPSEC
  - To eliminate origin impersonation.
- ASPA (IEFT draft)
  - To reduce impacts of adjacency leaks.


kentik®
Thank you
dmadory@Kentik.com
@dougmadory