



Measuring RPKI ROV adoption with NetFlow

Doug Madory, dmadory@kentik.com, Kentik
Job Snijders, job@fastly.com, Fastly




Where are we with RPKI ROV adoption?

- Presently stands as the Internet's best defense against BGP hijacks due to typos or other BGP mishaps.
- Core challenge: broad deployment requires many individual actions.
 - *Why reject RPKI-invalids if no one is creating ROAs?*
 - *Why create ROAs if no one is rejecting RPKI-invalids?*



Where are we with RPKI adoption?

- 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>

Measuring RPKI deployment progress

- It takes two steps to reject an RPKI-Invalid BGP route.


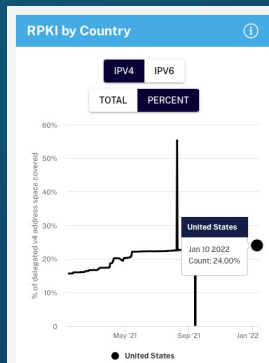
1 ROAs created to assert valid origin and prefix length.

2 Networks reject RPKI-invalids

How to evaluate progress?

Multiple resources (ex: NIST, RIPE)

Active area of research



Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering

Andreas Reuter
Freie Universität Berlin
andreas.reuter@fu-berlin.de

Randy Bush
IJ Research Lab / Dragon Research
randy@drag.com

Ílalo Cunha
Universidade Federal de Minas Gerais
cunha@dcc.ufmg.br

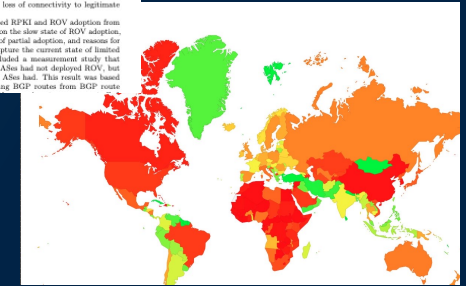
Ethan Katz-Bassett
USC / Columbia University
ethan.kb@usc.edu

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
m.waehlisch@fu-berlin.de

ABSTRACT
A proposal to improve routing security—Route Origin Authorization (ROA)—has been standardized. A ROA specifies which network is allowed to announce a set of Internet destinations. While some networks now specify ROAs, little is known about whether other networks check routes they receive against these ROAs, a process known as Route Origin Validation (ROV). Which networks blindly accept invalid routes? Which reject them outright? Which de-prioritize them if alternatives exist?
Recent analysis attempts to use uncontrolled experiments to characterize ROV adoption by comparing valid routes and invalid routes [5]. However, we argue that gaining a solid understanding of ROV adoption is impossible using currently available data sets and techniques. Our measurements suggest that, although some ISPs are not observed using invalid routes in uncontrolled experiments, they are actually using different routes for (non-security) traffic en-

can be used as part of the router's local BGP policy decisions, e.g., filtering routes that reflect invalid announcements or preferring valid ones. While the RPKI is fairly populated with ROAs and growing [9, 15, 23, 24], adoption of ROV and filtering has been negligible, according to operator gossip. A major reason for this is the lack of economic incentives. Since a significant share of invalid routes are due to misconfiguration [26], adopting ROV and filtering can even have adverse effects such as a loss of connectivity to legitimate network destinations.
A recent paper examined RPKI and ROV adoption from multiple angles, focusing on the slow state of ROV adoption, the security implications of partial adoption, and reasons for slow adoption [5]. To capture the current state of limited adoption, the paper included a measurement study that claimed that most large ASes had not deployed ROV, but that 9 of the 100 largest ASes had. This result was based on observations of existing RPKI routes from RPKI route



<https://stats.labs.apnic.net/roas>

Measuring RPKI deployment progress

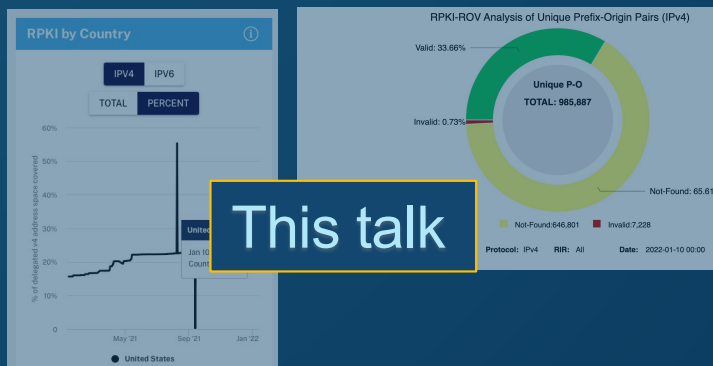
- It takes two steps to reject an RPKI-Invalid BGP route.

1 ROAs created to assert valid origin and prefix length.

2 Networks reject RPKI-invalids

How to evaluate progress?

Multiple resources (ex: NIST, RIPE)



This talk

Active area of research

Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering

Andreas Reuter
Freie Universität Berlin
andreas.reuter@fu-berlin.de

Randy Bush
IJ Research Lab / Dragon Research
randy@drag.com

Ílalo Cunha
Universidade Federal de Minas Gerais
cunha@dcc.ufmg.br

Ethan Katz-Bassett
USC / Columbia University
ethan.kb@usc.edu

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
m.waehlisch@fu-berlin.de

ABSTRACT

A proposal to improve routing security—Route Origin Authorization (ROA)—has been standardized. A ROA specifies which network is allowed to announce a set of Internet destinations. While some networks now specify ROAs, little is known about whether other networks check routes they receive against these ROAs, a process known as Route Origin Validation (ROV). Which networks blindly accept invalid routes? Which reject them outright? Which de-prioritize them if alternatives exist?

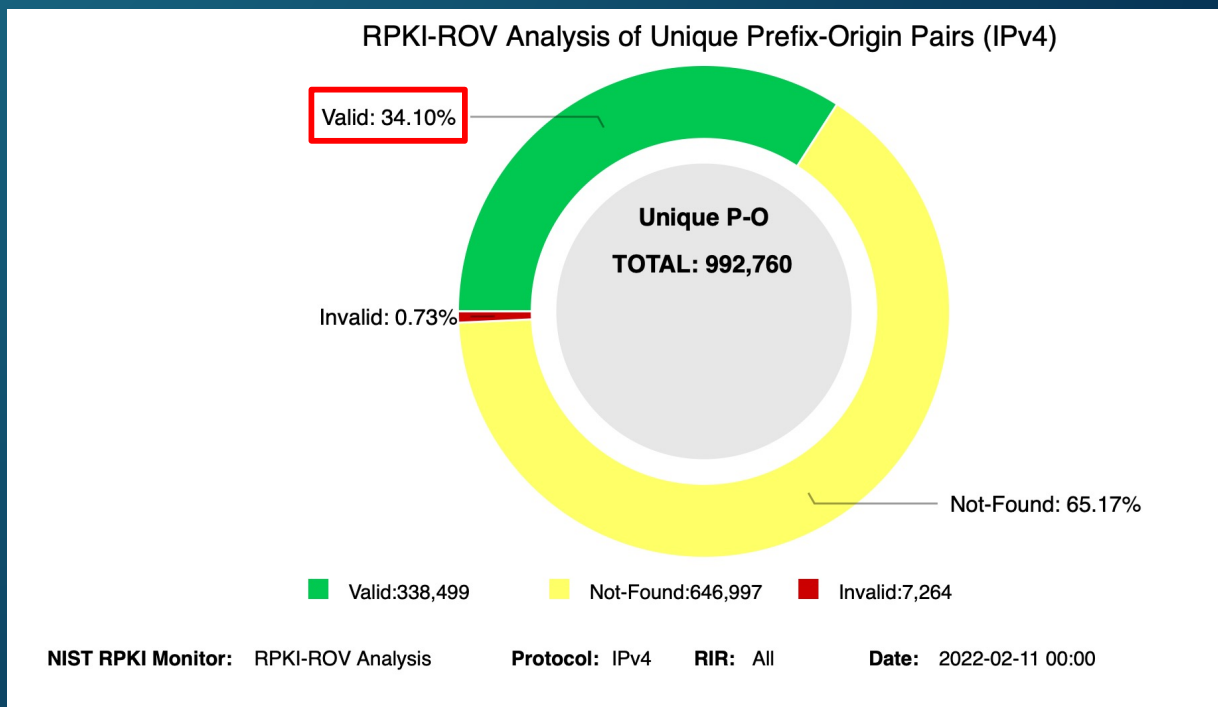
Recent analyses attempt to use uncontrolled experiments to characterize ROV adoption by comparing valid routes and invalid routes [5]. However, we argue that gaining a solid understanding of ROV adoption is impossible using currently available data sets and techniques. Our measurements suggest that, although some ISPs are not observed using invalid routes in uncontrolled experiments, they are actually using different routes for (non-security) traffic and can be used as part of the router's local BGP policy decisions, e.g., filtering routes that reflect invalid announcements or preferring valid ones. While the RPKI is fairly populated with ROAs and growing [1, 15, 23, 24], adoption of ROV and filtering has been negligible, according to operator gossip. A major reason for this is the lack of economic incentives. Since a significant share of invalid routes are due to misconfiguration [25], adopting ROV and filtering can even have adverse effects such as a loss of connectivity to legitimate network destinations.

A recent paper examined RPKI and ROV adoption from multiple angles, focusing on the slow state of ROV adoption, the security implications of partial adoption, and reasons for slow adoption [5]. To capture the current state of limited adoption, the paper included a measurement study that claimed that most large ASes had not deployed ROV, but that 9 of the 100 largest ASes had. This result was based on observations of existing BGP routes from BGP route

<https://stats.labs.apnic.net/roas>

Where are we with ROA creation?

- NIST RPKI Monitor reports that *only 34.1% of IPv4 BGP routes* are presently signed. *



Two RPKI unknown routes for each RPKI valid one.

Question:

What proportion of overall traffic is safeguarded by that 34.1%?

*32.6% of IPv6 routes are RPKI-Valid

Back in 2019 NTT/pmacct introduced NetFlow + RPKI

- Offered as a capability to evaluate impact of rejecting RPKI-Invalids on traffic levels
- Kentik was challenged, heeded the challenge!
- The rest of this talk focuses on what Kentik learned so far from its aggregate data

Analysing traffic in context of rejecting RPKI invalids using pmacct

Job Snijders [job at ntt.net](mailto:job@ntt.net)

Tue Feb 12 18:15:54 UTC 2019

- Previous message (by thread): [Clueful Contact at IPVolume.net ?](#)
- Next message (by thread): [Route Filtering Update](#)
- Messages sorted by: [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#)

Dear all,

Whether to deploy RPKI Origin Validation with an "invalid == reject" policy really is a business decision. One has to weigh the pros and cons: what are the direct and indirect costs of accepting misconfigurations or hijacks for my company? what is the cost of deploying RPKI? What is the cost of honoring misconfigured RPKI ROAs? There are a few thousand misconfigured ROAs, what does this mean for me?

To answer these questions, Paolo Lucente and myself worked to extend

Kind regards,


Job

ps. Dear Kentik & Deepfield, please copy+paste this feature! We'll happily share development notes with you, you can even look at pmacct's source code for inspiration. :-)

<https://mailman.nanog.org/pipermail/nanog/2019-February/099522.html>

Job's presentation at DKNOG-9 (March 2019)

- Earlier datapoint. NTT traffic based on RPKI status (DKNOG 9, March 2019)



- *We've come a long way since then!*

Also from Job's presentation at DKNOG-9

- Job's prediction: Given the consolidation of the Internet industry, only a few major companies needed to deploy RPKI before we saw large benefits.

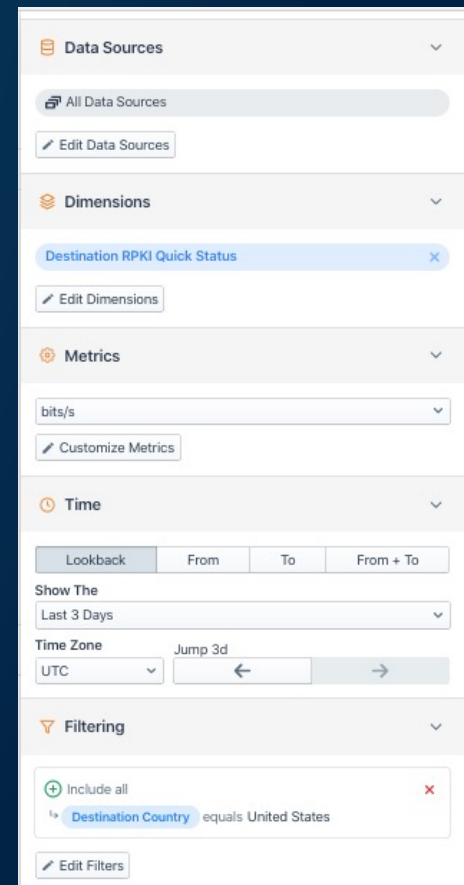
Not everyone needs to do RPKI

- Because of the centralization of the web, if a select few companies deploy RPKI Origin Validation - millions of people benefit
- (google, cloudflare, amazon, pch/quad9, facebook, akamai, fastly, liberty global, comcast, etc...)
- I think only 20 companies or so need to do Origin Validation for there to be big benefits...
- <https://dyn.com/blog/bgp-dns-hijacks-target-payment-systems/>

dknog

Kentik's perspective can deepen understanding of RPKI

- 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.
- Kentik's NetFlow analytics platform annotates flow records with an RPKI evaluation of route of destination IP upon intake.
 - Originally built to understand how much traffic would be lost by dropping invalids.
 - Can also be used to understand RPKI from a traffic-volume perspective.





What proportion of traffic goes to signed routes?

- Kentik tracks four cases of RPKI outcome.
 1. Valid
 2. Unknown
 3. Invalid
 4. Invalid – but covered by valid/unknown

Note #4 only exists in the analysis-plane and is not part of IETF/BGP/Routing!

Example of #4:

Origin AS	Announcement		Description
AS6128	24.38.0.0/17		Cablevision Systems Corp.
AS33759	24.38.10.0/24	 	Regeneron (C03272042)

Address has 0 hosts associated with it.

Only ~1/3 of BGP routes have ROAs - but how much traffic?

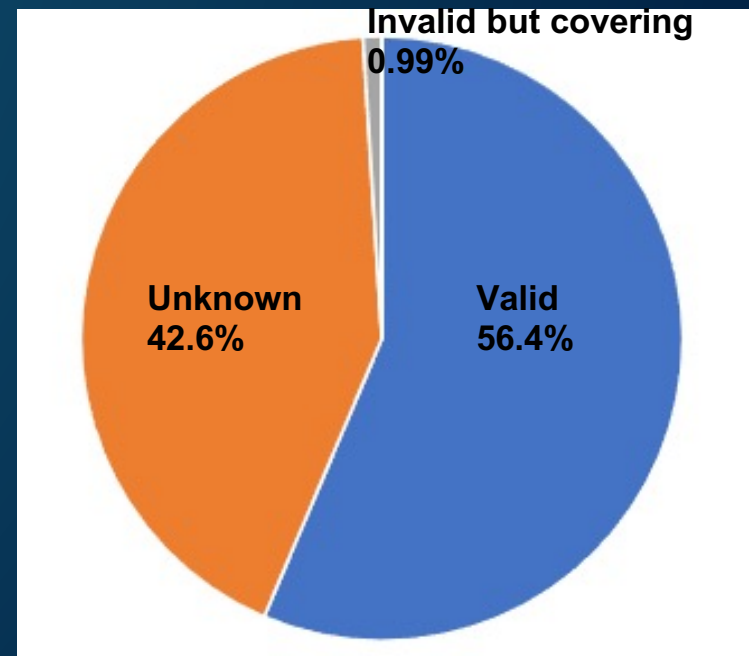
Period of analysis: 29 Jan 2022 00:00 UTC to 5 Feb 2022 00:00 UTC (7 days)

Main Observations*

- 0.99% of traffic volume is 'Invalid but covering'
- 42.6% is Unknown
- 56.4% is Valid
- 0.01% is Invalid

Traffic to invalid routes is infinitesimal.

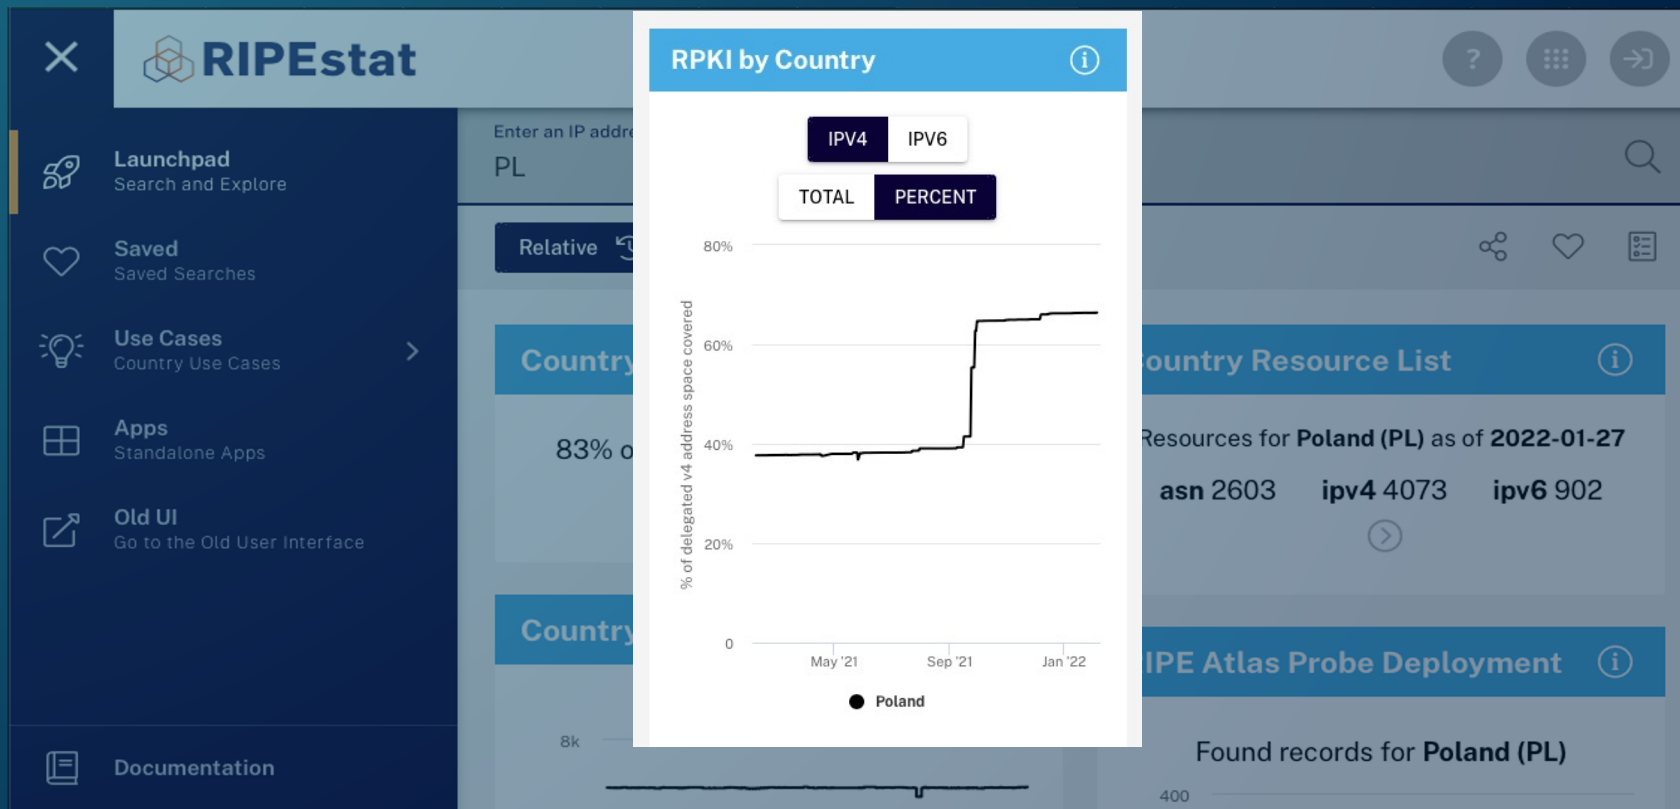
- Not a reason to not drop invalids.



*Combined IPv4 + IPv6

Comparing metrics for ROA creation by country

- RIPEstat reports % of IP address space <https://stat.ripe.net/app/launchpad/>



For example, how is the US doing with ROA creation?

United States



58.5% of bits/sec (NetFlow)*

24.2% of IPv4 space (RIPEstat)

20.1% of IPv6 space

Why?

Major RPKI deployments

- Eyeball networks

- Comcast (AS7922)

99.7%

- Spectrum (AS20115)

99.9%

- Content providers

- Amazon (AS16509)

100%

- Google (AS15169)

100%

- Cloudflare (AS13335)

93.3%

Valid %*

Maybe not a majority of BGP routes, but these companies account for a lot of US traffic!

*Combined IPv4 + IPv6

Many countries are doing better than earlier stats suggest

United States




58.5% of bits/sec*


24.2% of IPv4 space

20.1% of IPv6 space


*Combined IPv4 + IPv6




ROA Measurement (%) Western Hemisphere



ROA Measurement (%) Europe

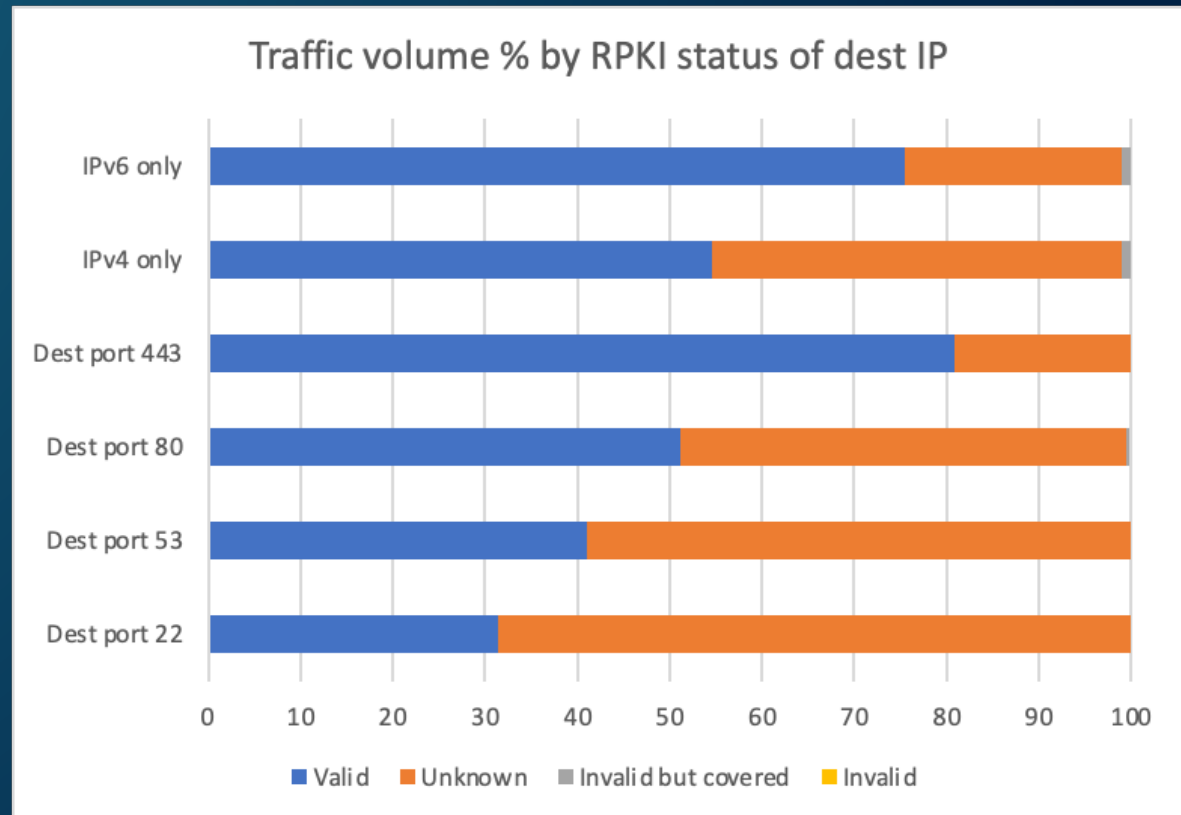


ROA Measurement (%) Asia & Australia

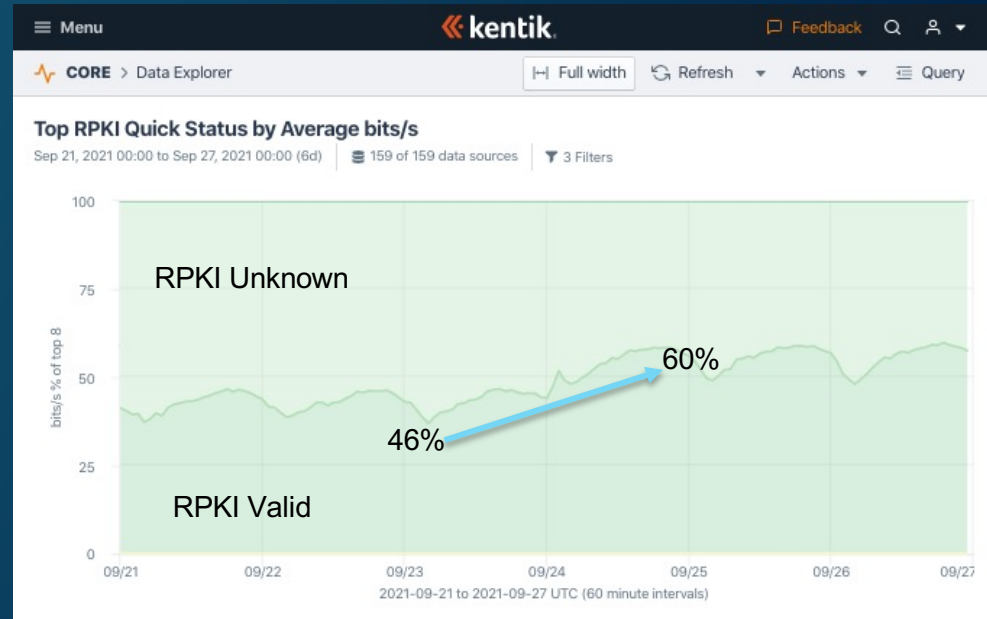
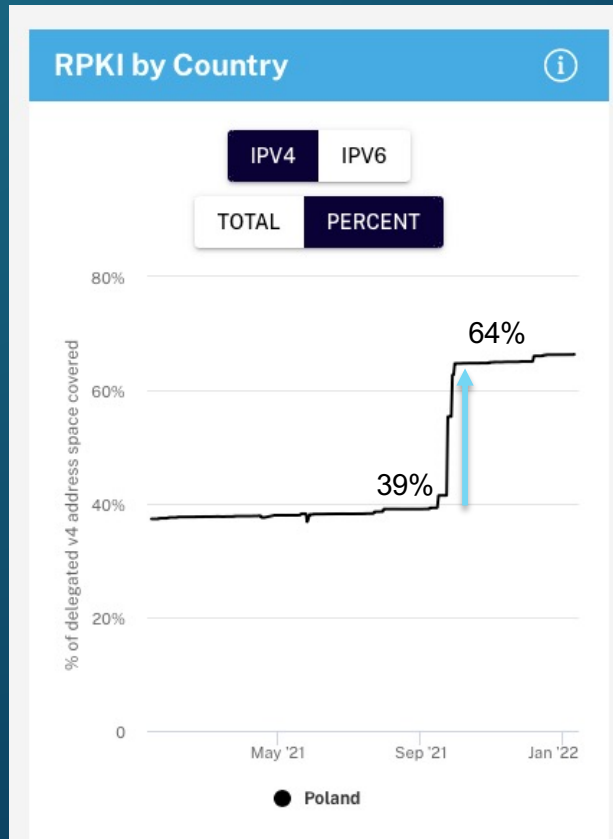


Other interesting observations in RPKI classified traffic

- IPv6 (75.4% valid)
- IPv4 (54.5% valid)
- Port 443 (80.8% valid)
- Port 80 (51.5% valid)
- Port 53 (41.1% valid)
- Port 22 (31.4% valid)



Increases in ROAs leads to increases in valid traffic (Poland)



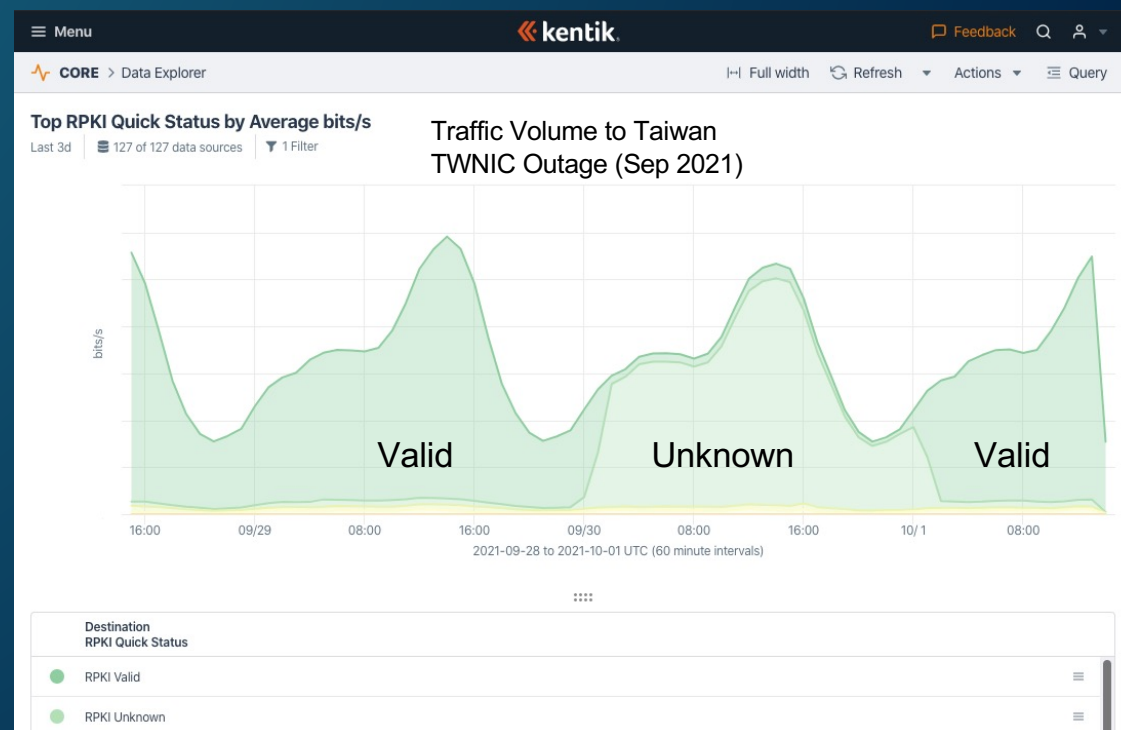
Combined IPv4 + IPv6

Not 1-to-1 but movements are correlated.

TWNIC Outage in September 2021



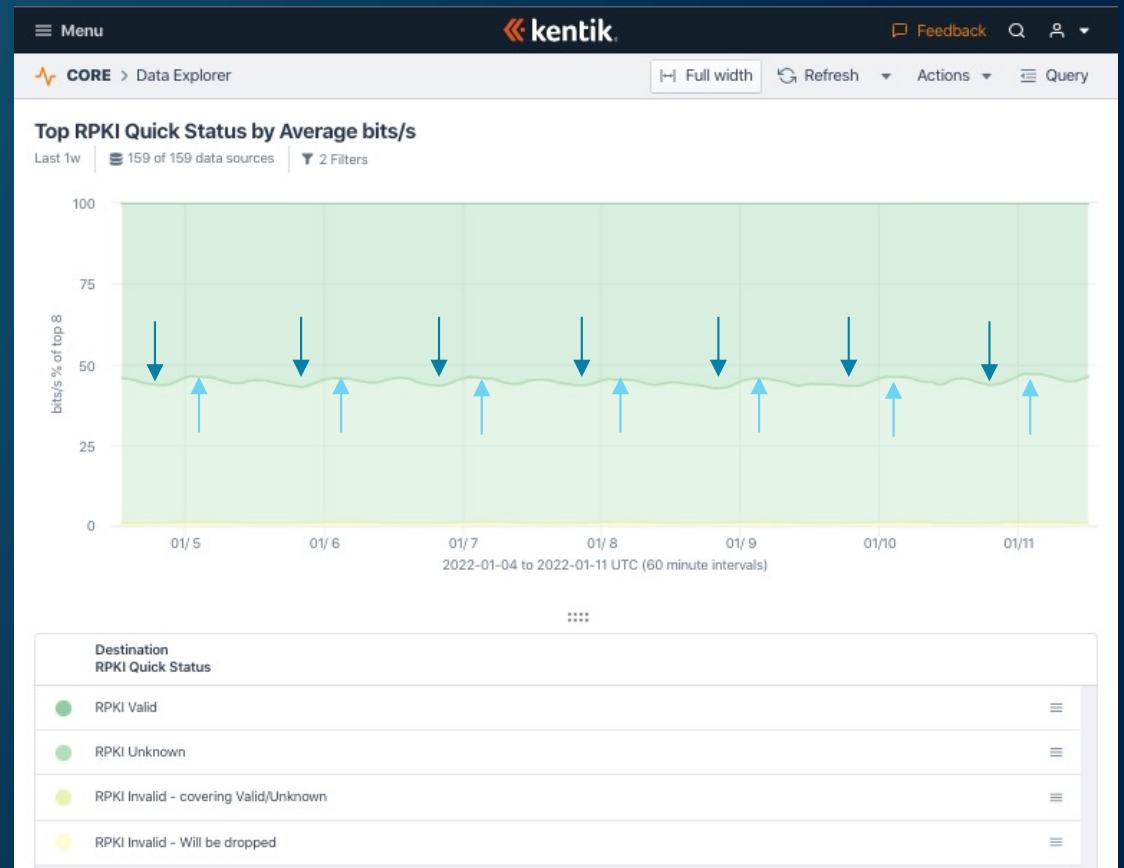
- Valid traffic was briefly Unknown
- No disruptions



Combined IPv4 + IPv6

Weird phenomena! valid:unknown ratio fluctuates over time

- Valid traffic
 - High: 20:00 UTC (57%)
 - Low: 2:00 UTC (54%)
- May be linked to shifts in user behavior when connecting via mobile vs fixed-line Internet.



Combined IPv4 + IPv6

Best Current Practice – Reject RPKI-Invalid BGP routes!

Rejecting RPKI-Invalid routes on EBGP sessions...

1. Protects a majority of your outbound traffic from BGP hijacks due to typos, BGP mishaps.
2. Not a risk to legitimate traffic.

Other BCPs include:

1. Do NOT modify LOCAL_PREF based on validation states
2. Do NOT set / remove BGP communities based on validation states

Security issues like CVE-2021-41531 / CVE-2021-3761 are examples of how not following the above BCP could result in massive BGP churn!

https://bgpfilterguide.nlnog.net/guides/reject_invalids/



Thanks for your attention!

If you have ANY questions regarding RPKI, please reach out to the BGP A-Team:

dmadory@Kentik.com



job@fastly.com

