Freie Universität Berlin

Measuring RPKI-based Route Filtering

Andreas Reuter (andreas.reuter@fu-berlin.de)

Joint work with Matthias Wählisch, Randy Bush, Ethan Katz-Bassett, Italo Cunha, and Thomas C. Schmidt

Once upon a time ... someone is incorrectly announcing an IP prefix.

Incorrect origin AS

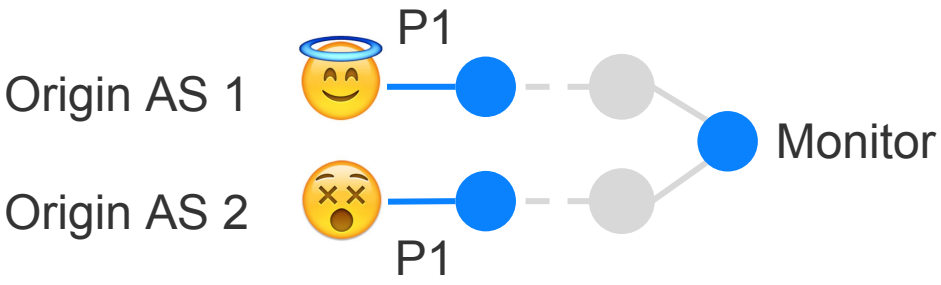

Incorrect prefix length

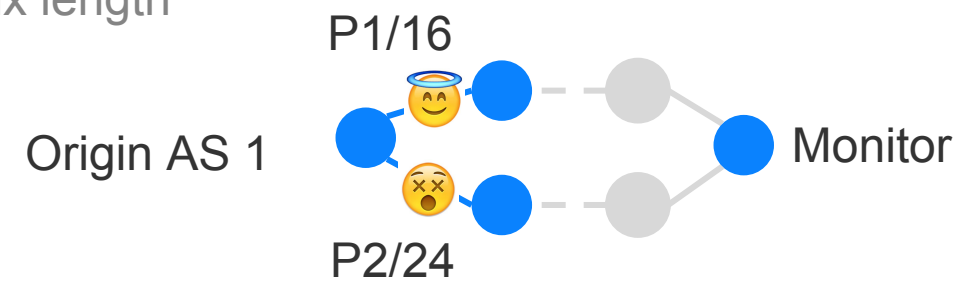

Once upon a time ... someone is incorrectly announcing an IP prefix.

For about 18 minutes on April 8, 2010, China Telecom advertised erroneous network traffic routes that instructed U.S. and other foreign Internet traffic to travel through Chinese servers.* Other servers around the world quickly adopted these paths, routing all traffic to about 15 percent of the Internet's destinations through servers located in China. This incident affected traffic to and from U.S. government (".gov") and military (".mil") sites, including those for the Senate, the army, the navy, the marine corps, the air force, the office of secretary of Defense, the National Aeronautics and Space Administration, the Department of Commerce, the National Oceanic and Atmospheric Administration, and many others. Certain commercial websites were also affected, such as those for Dell, Yahoo!, Microsoft, and IBM.[116]

Although the Commission has no way to determine what, if anything, Chinese telecommunications firms did to the hijacked data, incidents of this nature could have a number of serious implications. This level of access could enable surveillance of specific users or sites.† It could disrupt a data transaction and prevent a user from establishing a connection with a site. It could even allow a diversion of data to somewhere that the user did not intend (for example, to a "spoofed" site). Arbor Networks Chief Security Officer Danny McPherson has explained that the volume of affected data here could have been intended to conceal one targeted attack.[117] Perhaps most disconcertingly, as a result of the diffusion of Internet security certification authorities,‡ control over diverted data could possibly allow a telecommunications firm to compromise the integrity of supposedly secure encrypted sessions.§

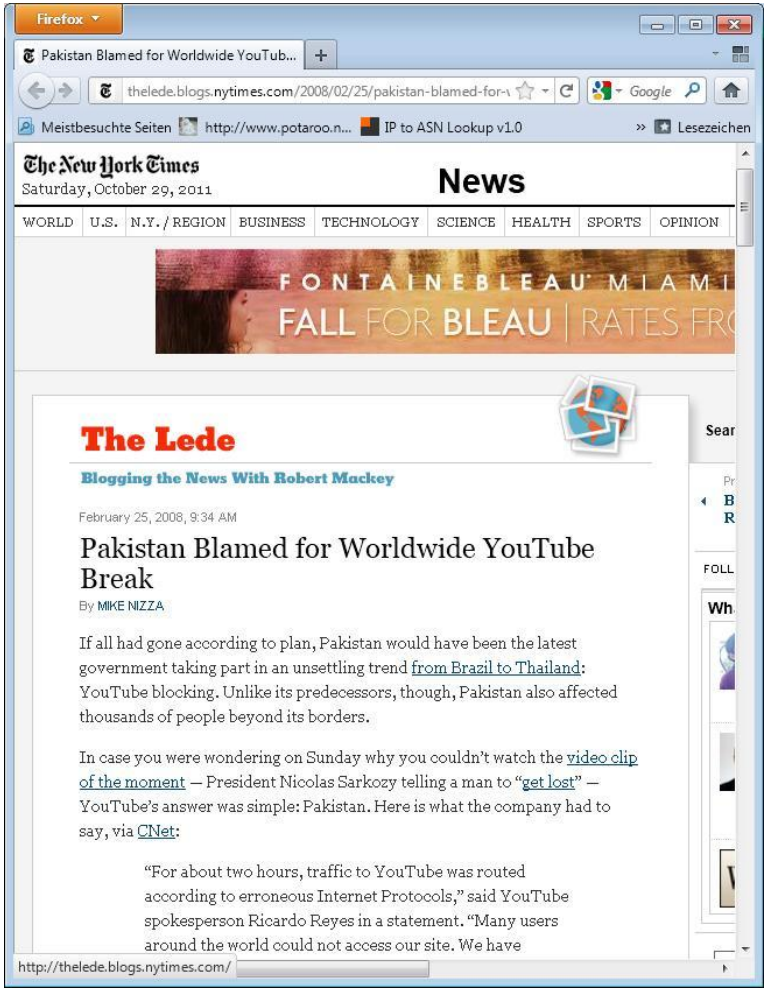

Background: RPKI-based Filtering

Prefix hijacking prevention using Resource Public Key Infrastructure

ROA Data + **Route Origin Validation** + **Local Policy**

Attestation object which AS is valid to announce IP prefix

Router operation to verify BGP Updates based on ROA data

Decide handling of invalid BGP routes (drop?)

Problem Statement & Challenge

Prefix hijacking prevention using Resource Public Key Infrastructure

ROA Data + **Route Origin Validation** + **Local Policy**

Attestation object which AS is valid to announce IP prefix

Router operation to verify BGP Updates based on ROA data

Decide handling of invalid BGP routes (drop?)

Goal: Measure the adoption of RPKI-based filter policies.

Problem Statement & Challenge

Prefix hijacking prevention using Resource Public Key Infrastructure

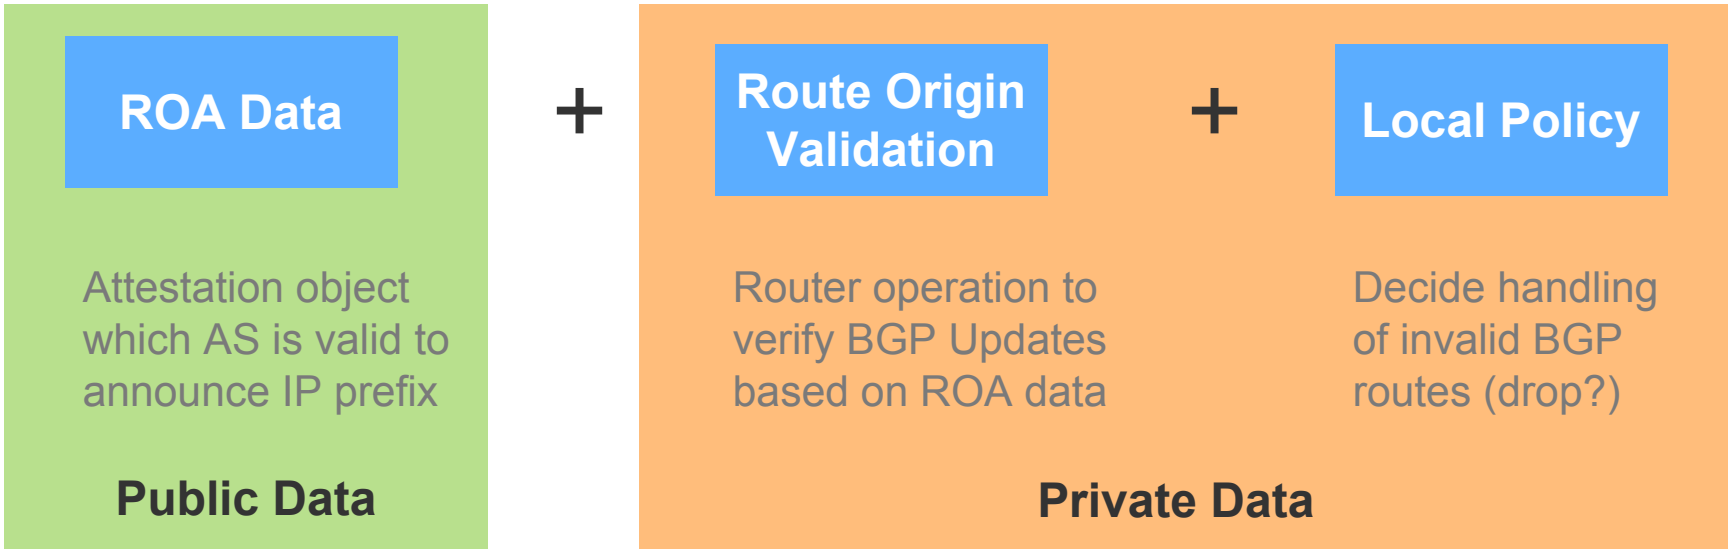

Goal: Measure the adoption of RPKI-based filter policies.

Challenge: Private data must be inferred from measurements.

Two principle approaches

Uncontrolled experiments

No coupling between triggering reason and observed event

Use existing BGP dumps for observations

Controlled experiments

Trigger events by actively changing BGP updates or ROAs

Use existing BGP dumps for observations, being clear on potential visibility

Vantage Point Visibility Matters

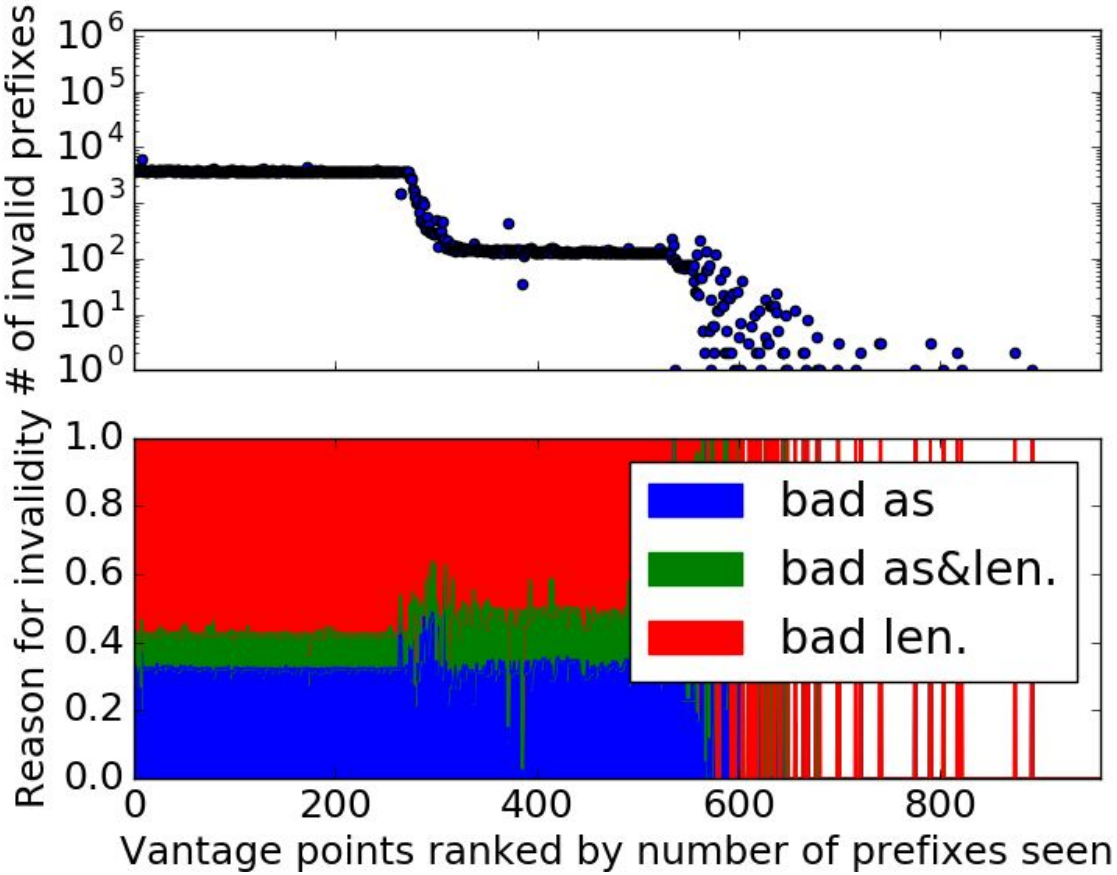

➔ Vantage points have limited visibility

➔ Observations might be misattributed to RPKI-based filtering

Vantage Point Visibility Matters

Per-Origin Prefix Visibility

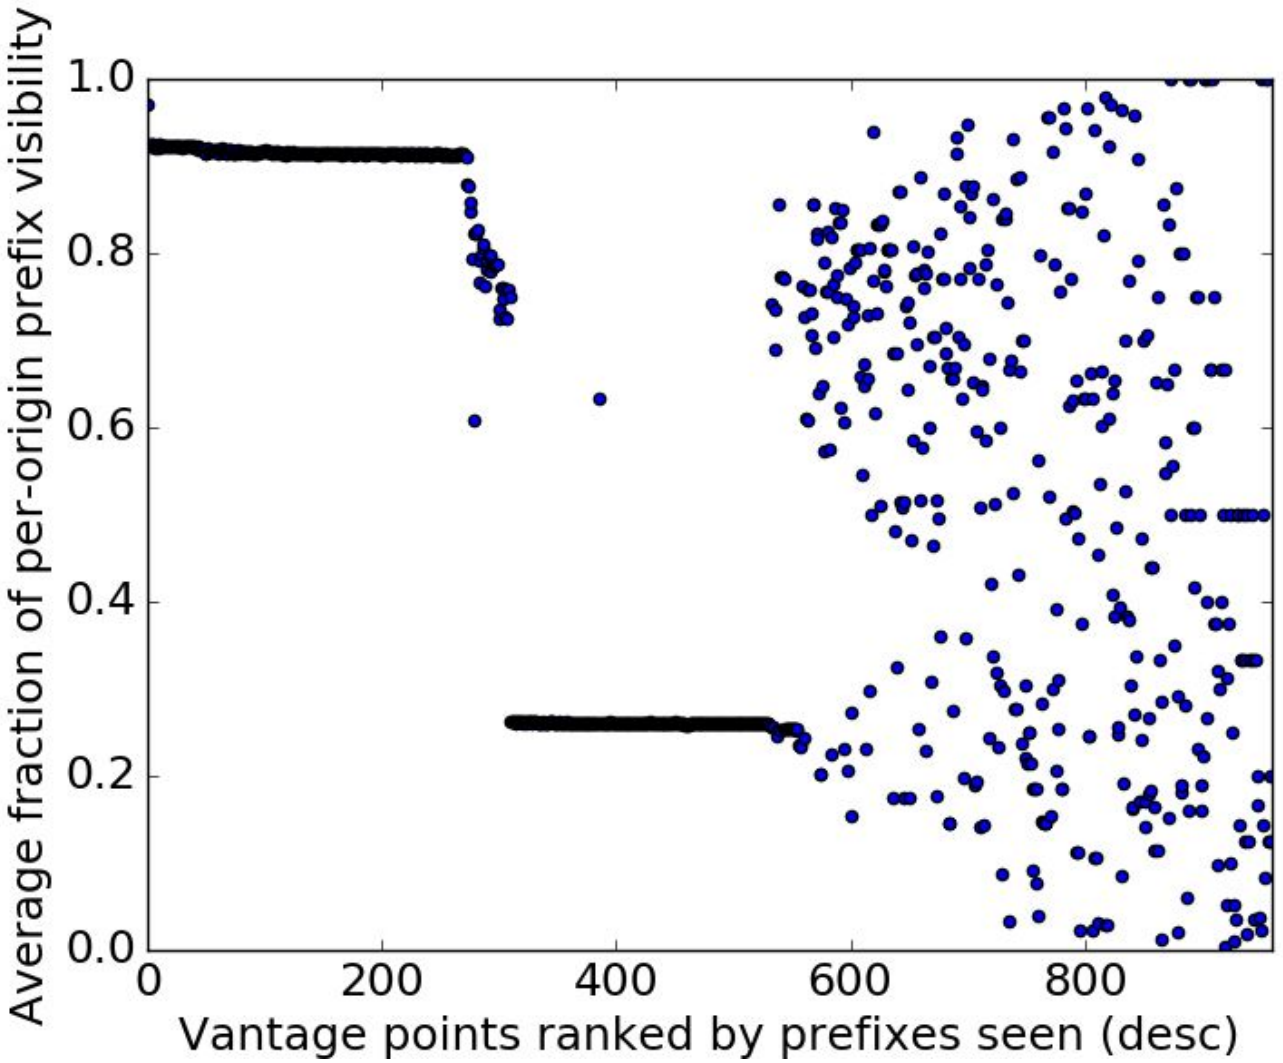

Invalid Announcements: Path Diversity

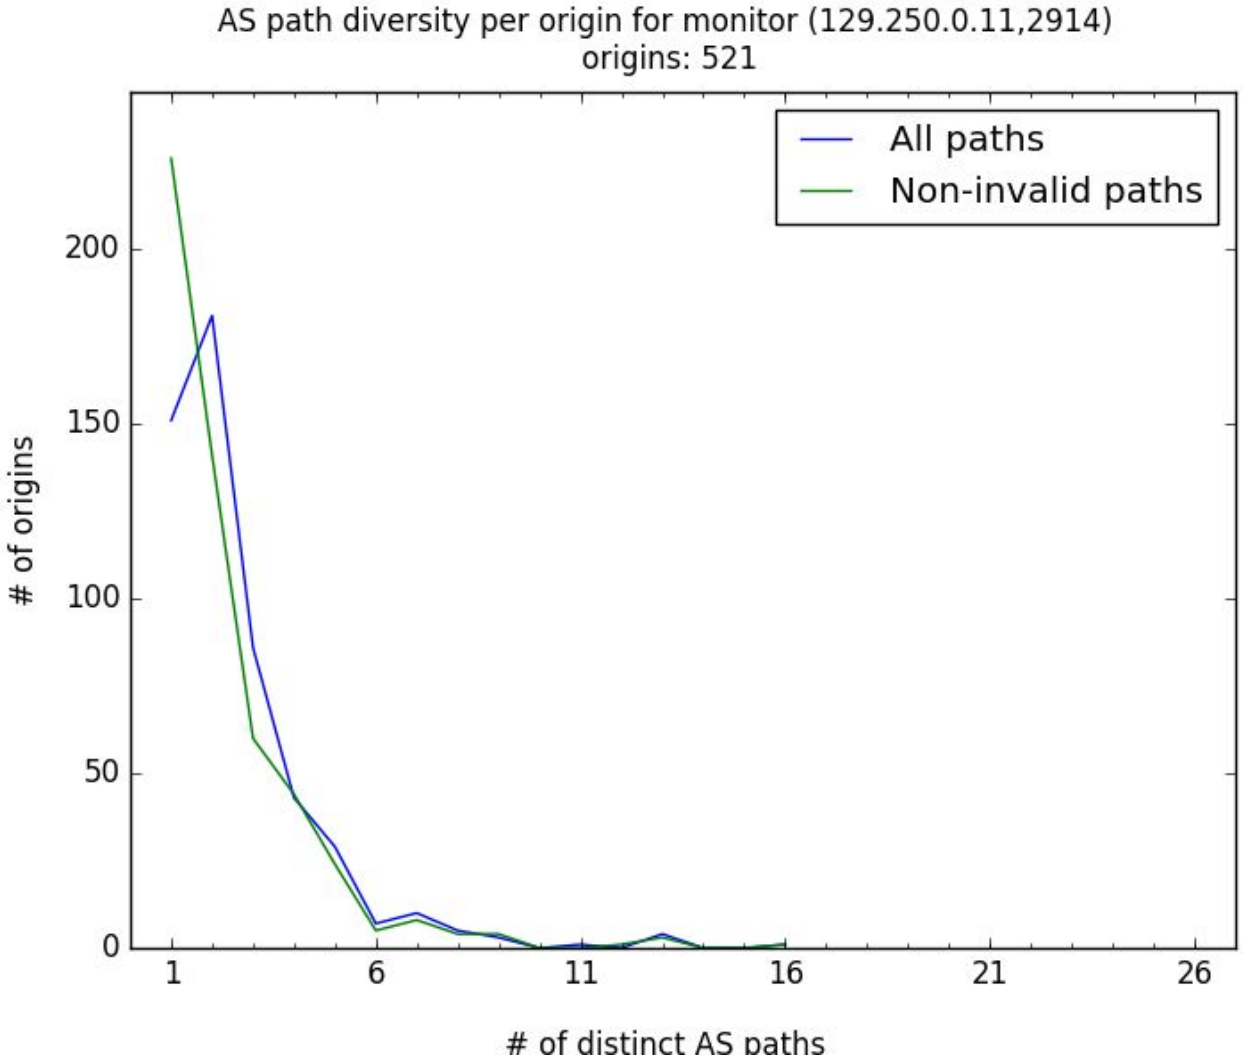

Invalid Announcements: Path Diversity

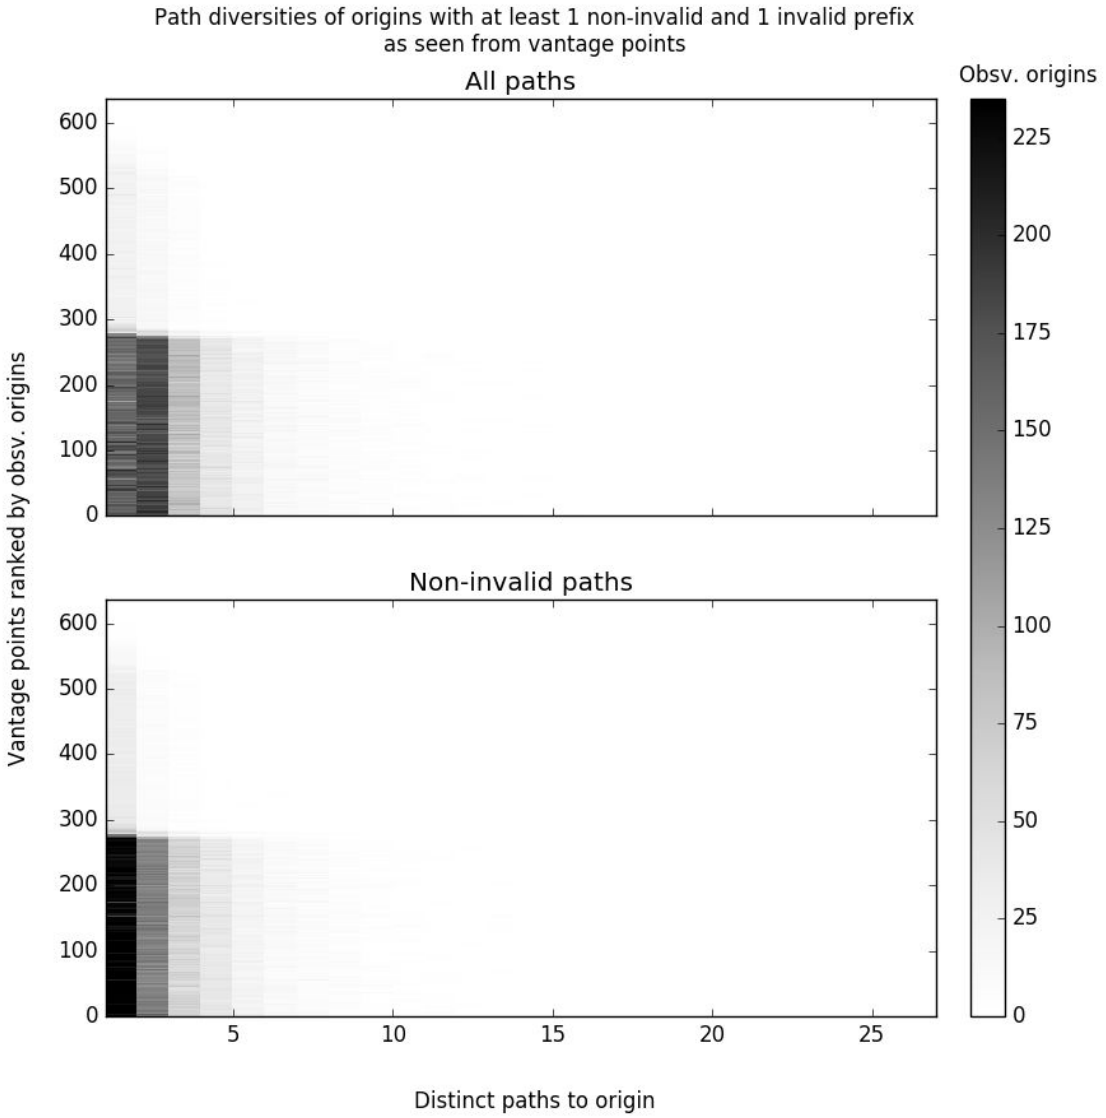

We need controlled experiments

Uncontrolled experiments can lead to incorrect inference.

Can we compare 2 paths (valid/invalid) to infer route origin validation?

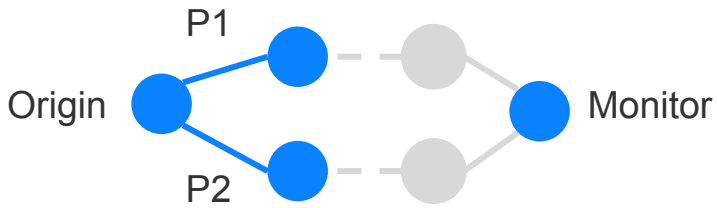

Observation	Different paths
Interpretation	ROV on path
Alternate explanation	

We need controlled experiments

Uncontrolled experiments can lead to incorrect inference.

Can we compare 2 paths (valid/invalid) to infer route origin validation?

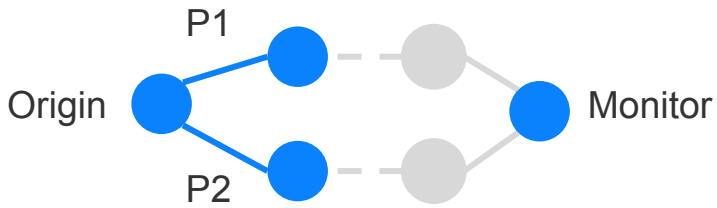

Observation	Different paths
Interpretation	ROV on path
Alternate explanation	Traffic engineering by origin, ROAs not updated.

We need controlled experiments

Uncontrolled experiments can lead to incorrect inference.

Can we compare 2 paths (valid/invalid) to infer route origin validation?

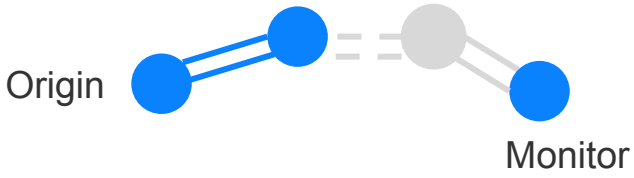

Observation	Same paths
Interpretation	No ROV
Alternate explanation	ROV policy to prefer valid, but no valid covering exist for the invalid prefix

We need controlled experiments

Uncontrolled experiments can lead to incorrect inference.

Can we compare 2 paths (valid/invalid) to infer route origin validation?

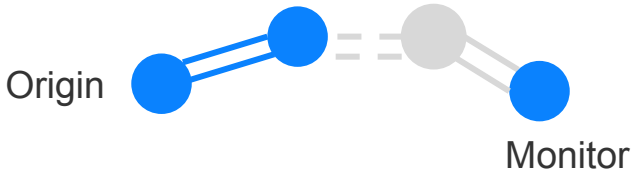

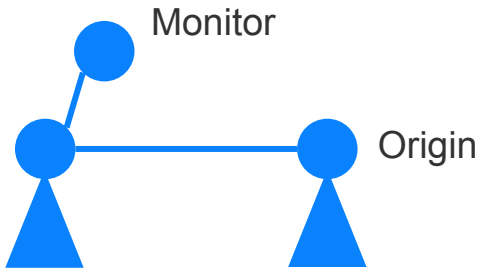

Observation	Same paths	No path to invalid
Interpretation	No ROV	ROV on path
Alternate explanation	ROV policy to prefer valid, but no valid covering exist for the invalid prefix	Limited visibility: Peer route not exported to monitor.

Can we compare two paths to infer route origin validation based on uncontrolled experiments?

No! We need controlled experiments!

Controlled experiments: Hand-crafted ROAs *and* BGP Updates

+ You know your peers
+ Reproduce observations
+ Detailed analysis of subtle filter policies
+ ...
+ You know your policies
+ Independent of external events
+ Iterative approach: results can inform later interpretation

Controlled Experiments: Setup

Using the **PEERING** testbed infrastructure, announce prefixes P_A and P_E:

➔ Prefixes are both /24, from same /16 block
➔ Both have same route object in RIR DB (that of the /16)
➔ ROA exists for both prefixes, making our announcement VALID

P_A serves as anchor and stays VALID

Announcement of P_E becomes INVALID periodically by changing ROA

Controlled Experiments: First Results

State of P_E

Case 1

valid

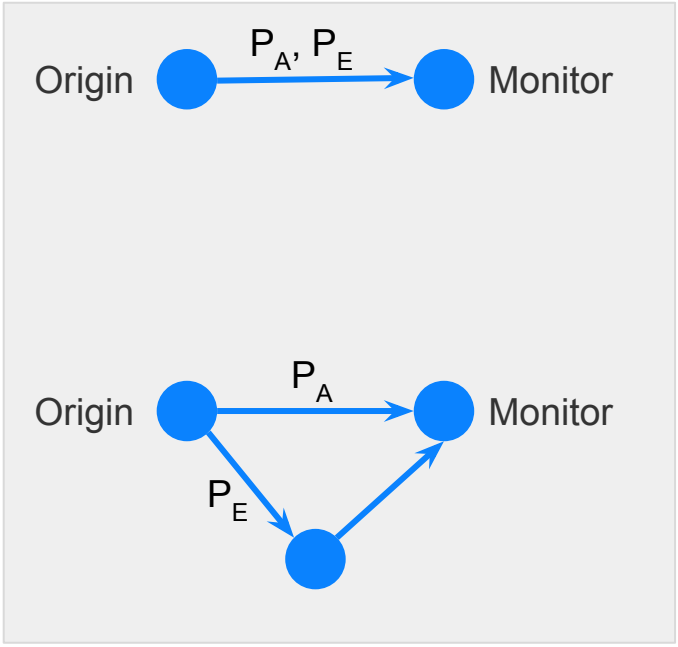

invalid

Monitor filters direct routes,
not routes via provider

Controlled Experiments: First Results

State of P_E

	Case 1	Case 2
valid		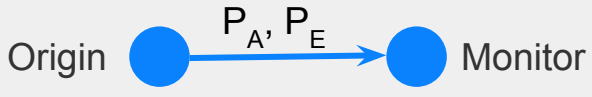
invalid	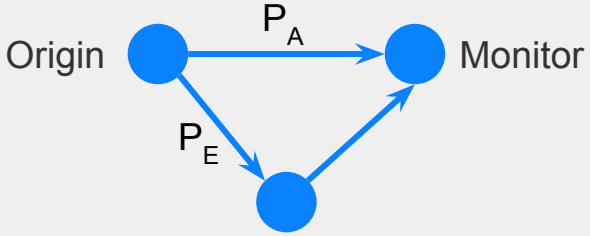	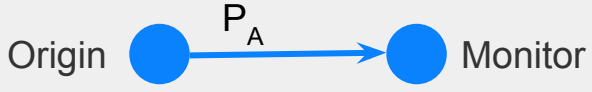
	Monitor filters direct routes, not routes via provider	Monitor has no path for invalid P_E. Either filtered by monitor or some AS on path

Controlled Experiments: Oddities

State of P_E

valid

invalid

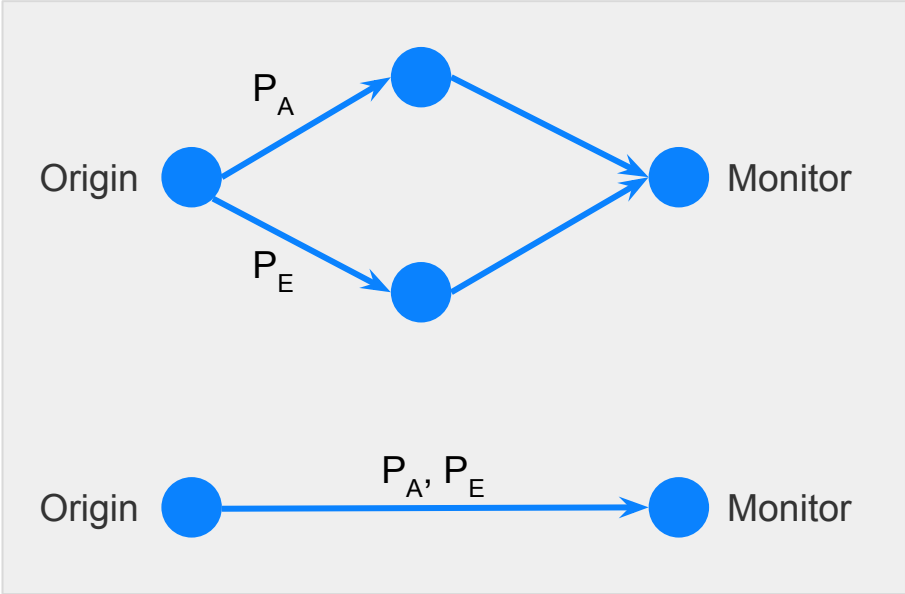

Monitor chooses separate routes when P_E is valid, same routes when P_E is invalid

Possible explanation: Monitor has long refresh interval for ROA data

Conclusion

- ➔ Uncontrolled experiments are not sufficient to infer RPKI-filtering policy
- ➔ Controlled experiments show that RPKI-based filtering is virtually non-existent
  - ◆ 2 AS found and confirmed, none of them in top 100 AS (ranked by customer cone size)
- ➔ Some oddities still unexplained. Work in progress.

Next Steps

➔ Refinement of measurement methodology

➔ Establish a live monitoring system with public access

We need your help to improve measurement coverage!

➔ Establish direct peering with PEERING testbed

➔ Peer with public route collectors