A study of RPKI deployment and discussion for improvement

RPKI is Coming of Age

Taejoong (Tijay) Chung
(https://tijay.github.io)

Assistant Professor
Rochester Institute of Technology

Outlines

- RPKI deployment and invalid route origins
  - RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins [IMC'19]
- Discussion (Follow-up works)

RPKI is Coming of Age

A Longitudinal Study of RPKI Deployment and Invalid Route Origins

Taejoong (Tijay) Chung§, Emile Aben†, Tim Bruijnzeels‡,
Balakrishnan Chandrasekaran△, David Choffnes*, Dave Levin+,
Bruce Maggs°◆, Alan Mislove*, Roland van Rijswijk-Deij‡±,
John Rula◆, Nick Sullivan※

§Rochester Institute of Technology, †RIPE NCC, ‡NLNetLabs,
△ Max Planck Institute for Informatics, *Northeastern University, +University of Maryland,
°Duke University, ±University of Twente, ◆Akamai Technologies, ※Cloudflare

RPKI is Coming of Age

A Longitudinal Study of RPKI Deployment and Invalid Route Origins

Resource PKI
(Public Key Infrastructure)

- Public Key Infrastructure framework designed to secure Internet's routing structure; specifically BGP (developed starting in 2008)

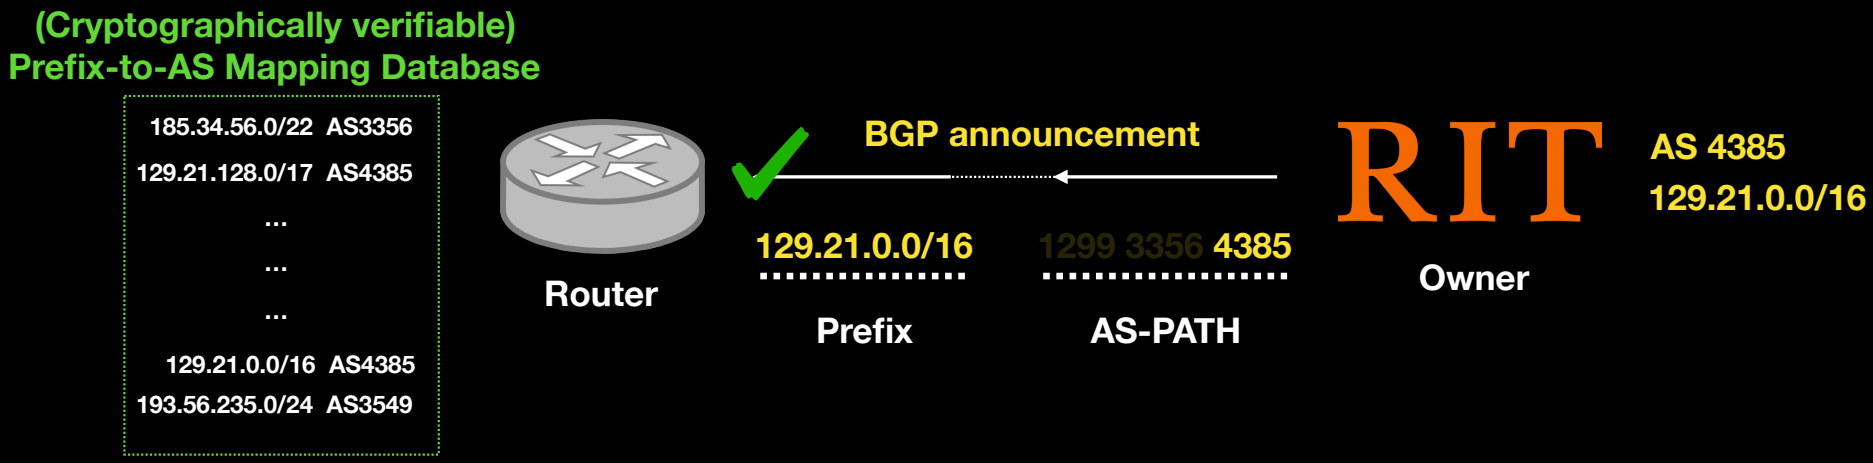

RPKI: How it works?

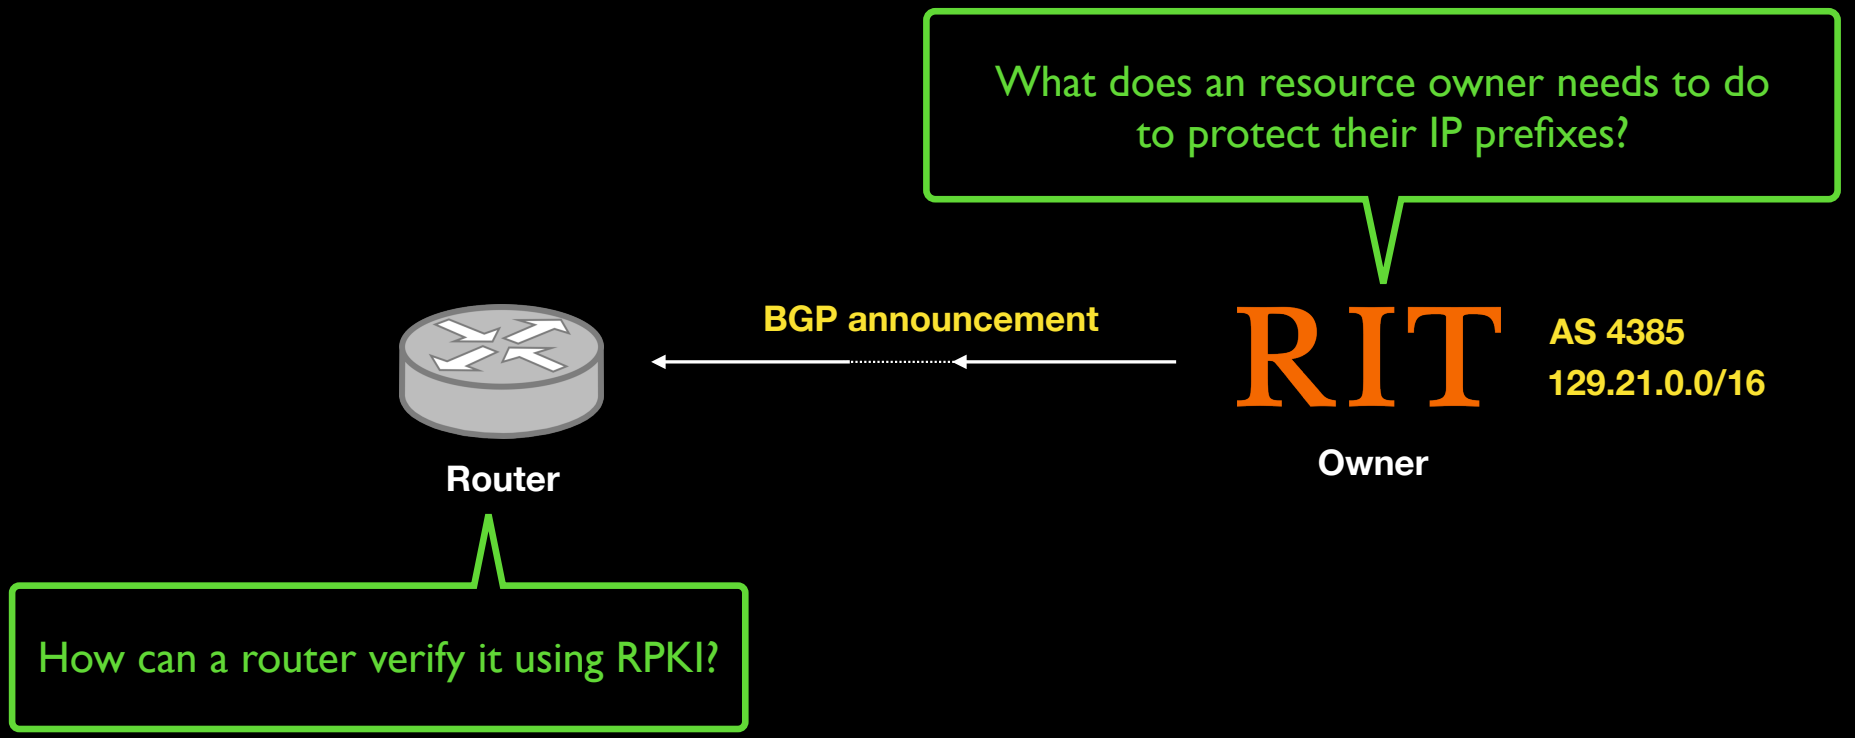

RPKI Structure

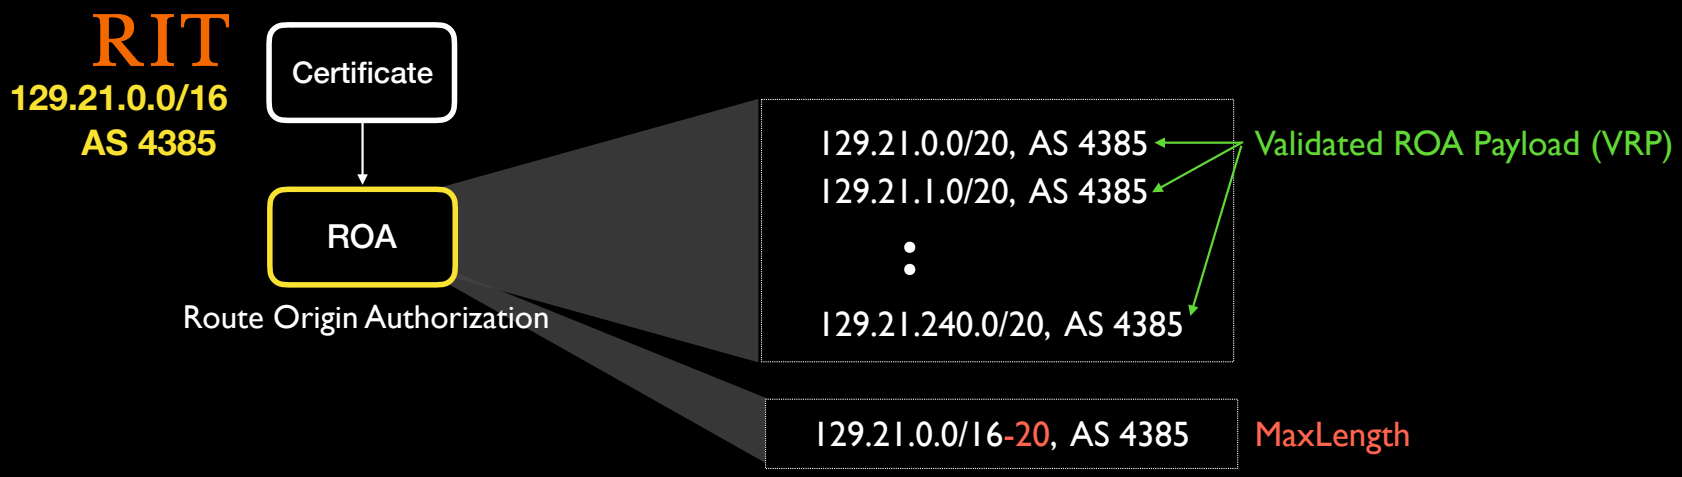

RPKI Structure

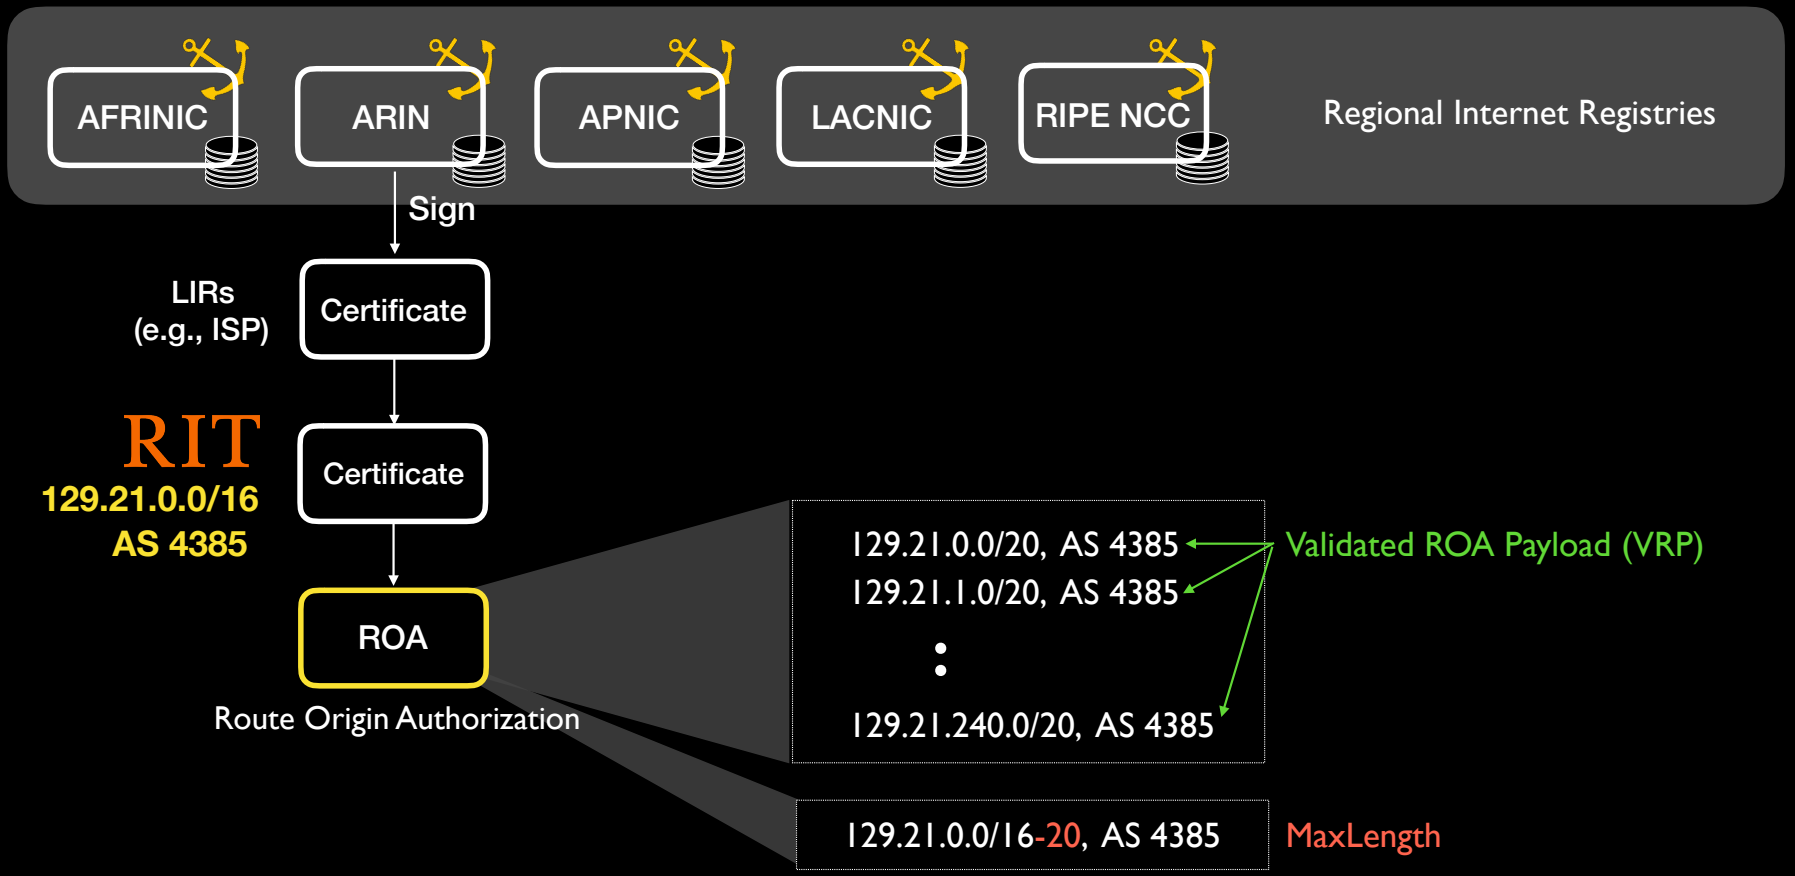

RPKI: How it works?

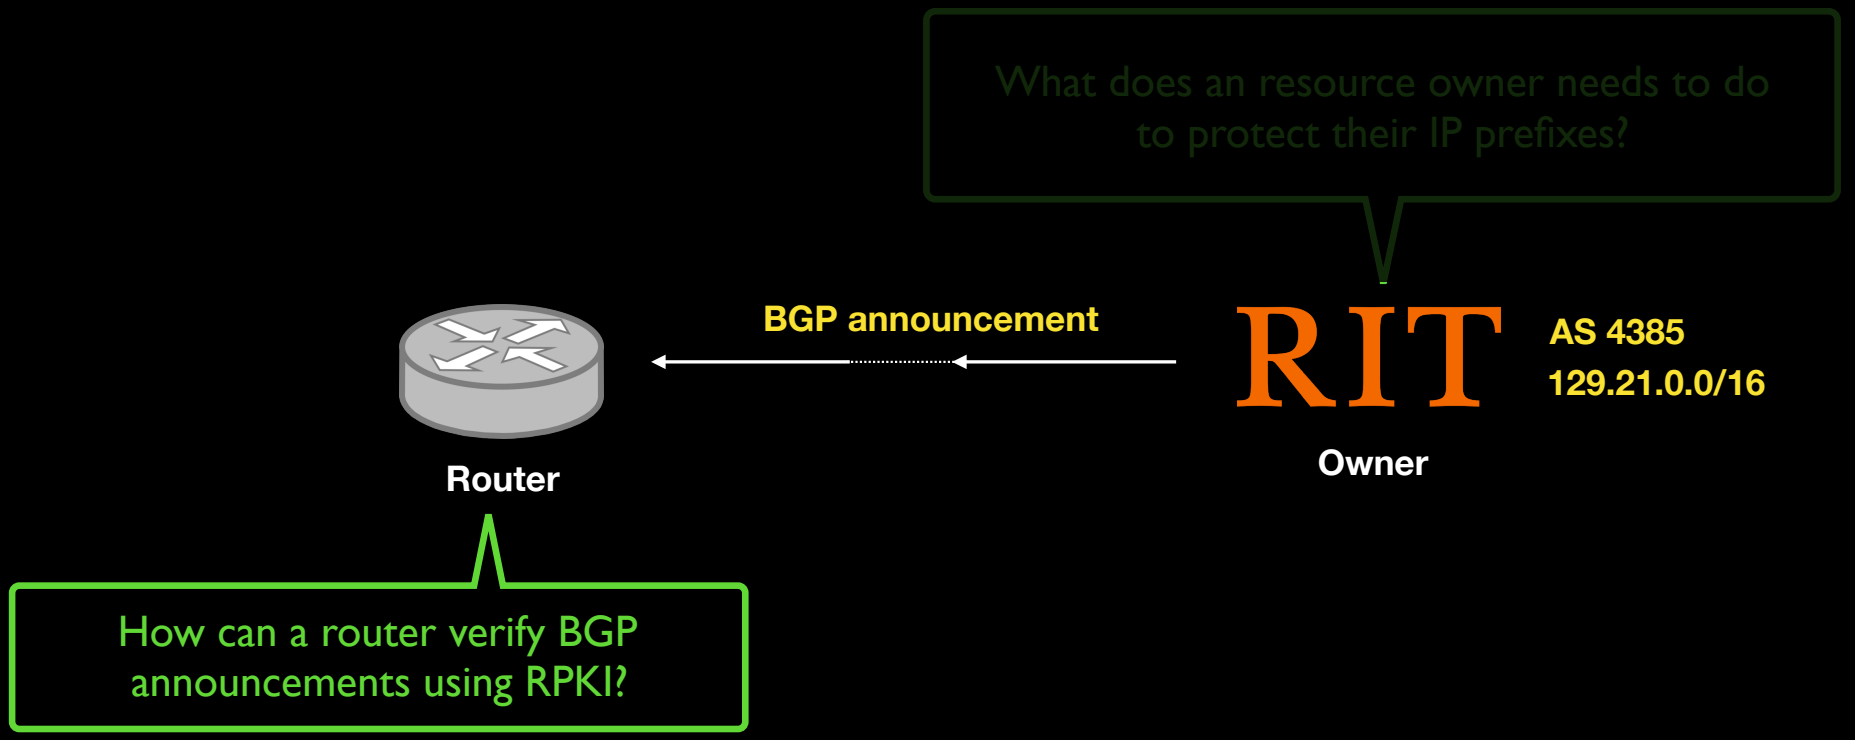

RPKI: How it works?
Validation process: Valid

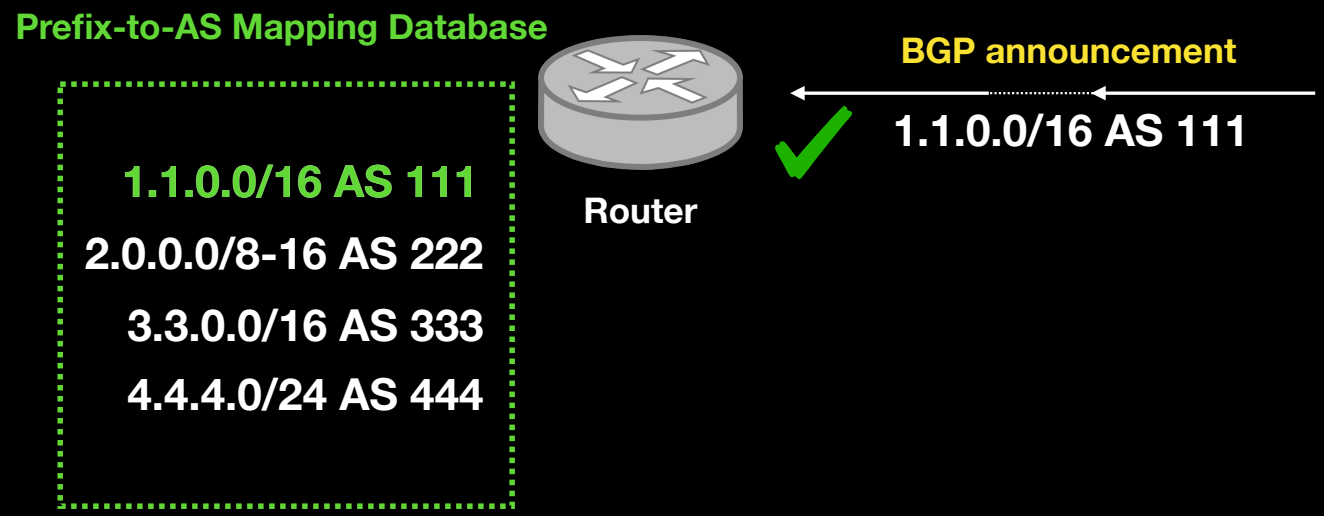

RPKI: How it works?
Validation process: Valid (w/ MaxLength)

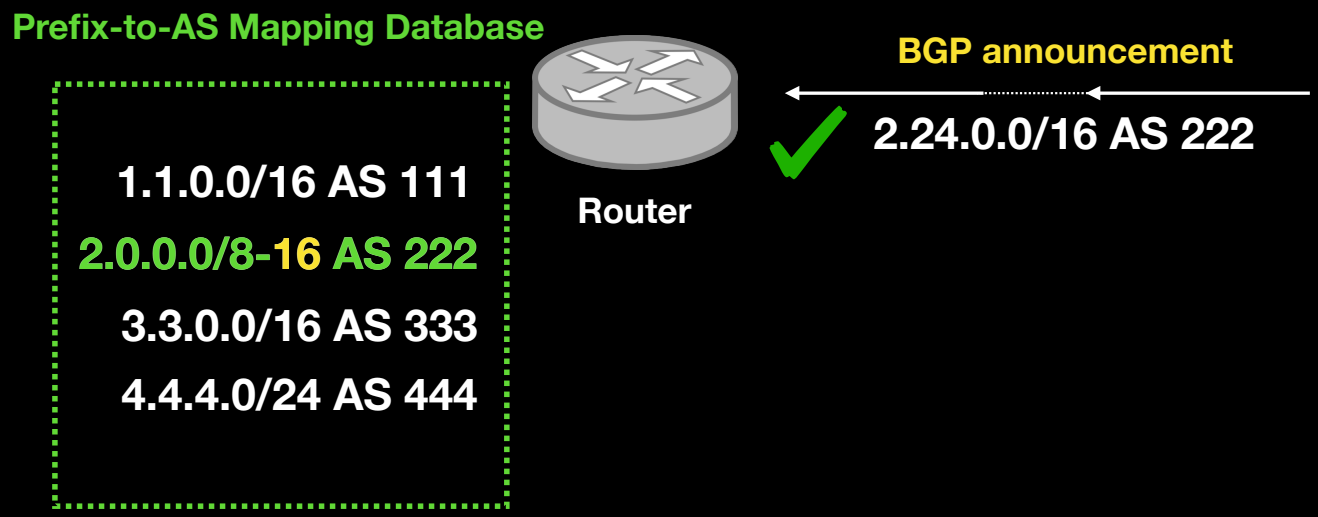

RPKI: How it works?
Validation process: Invalid (too-specific)

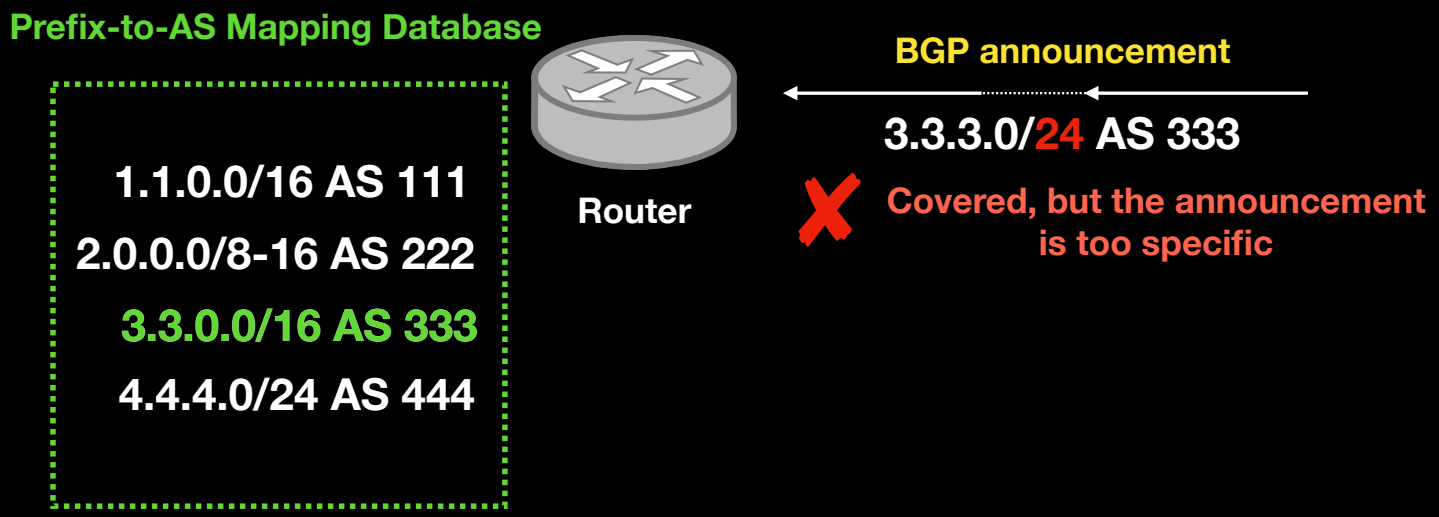

RPKI: How it works?
Validation process: Invalid (wrong ASN)

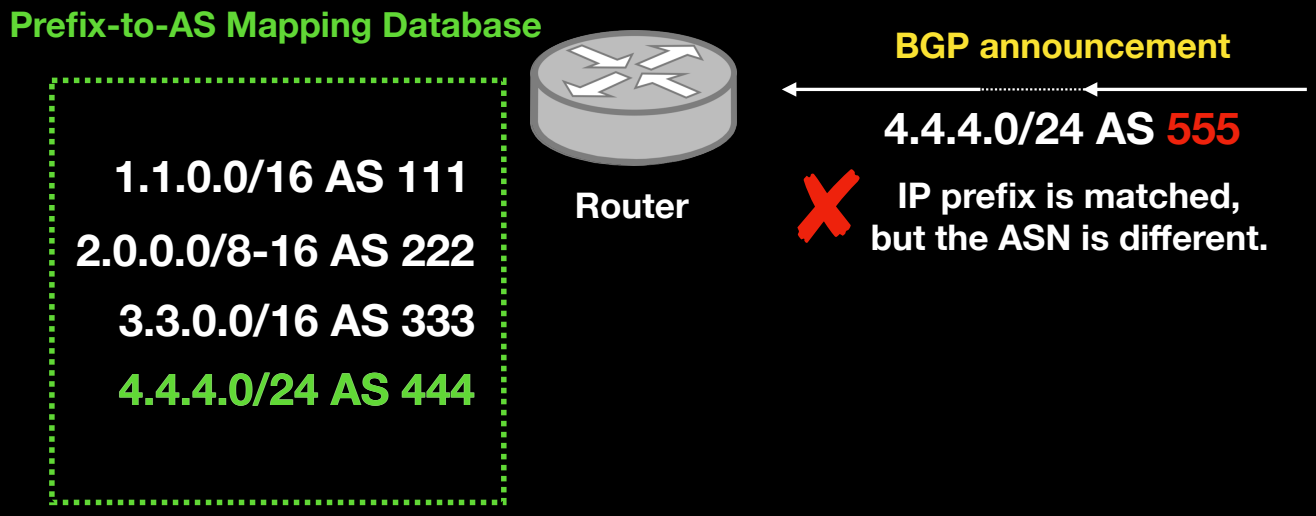

RPKI: How it works?
Validation process: Unknown (Uncovered)

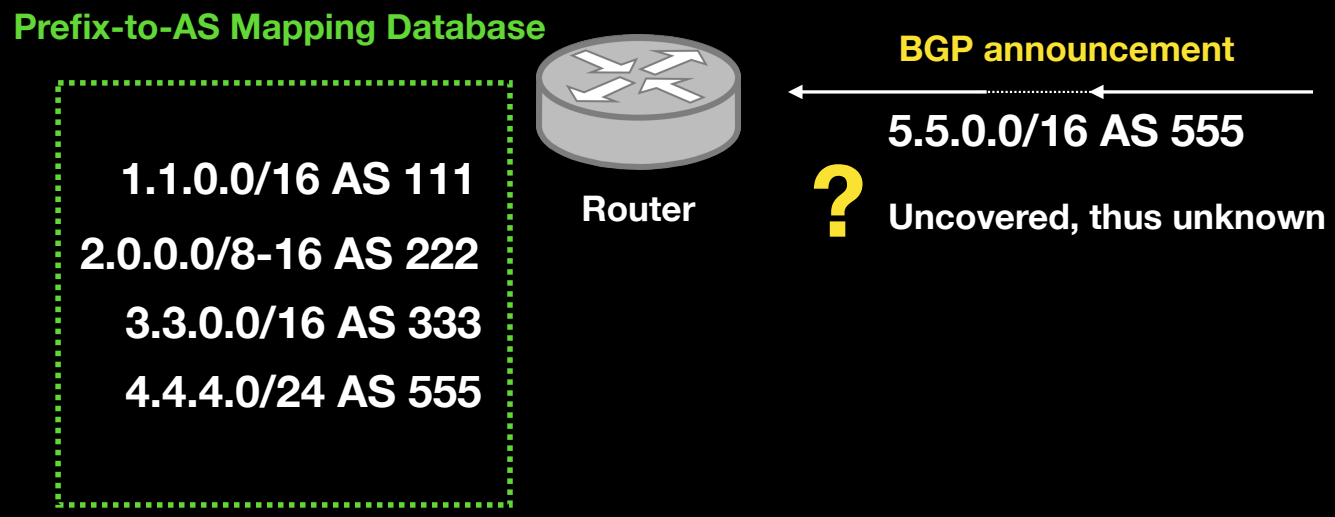

RPKI: How it works?
Validation Process

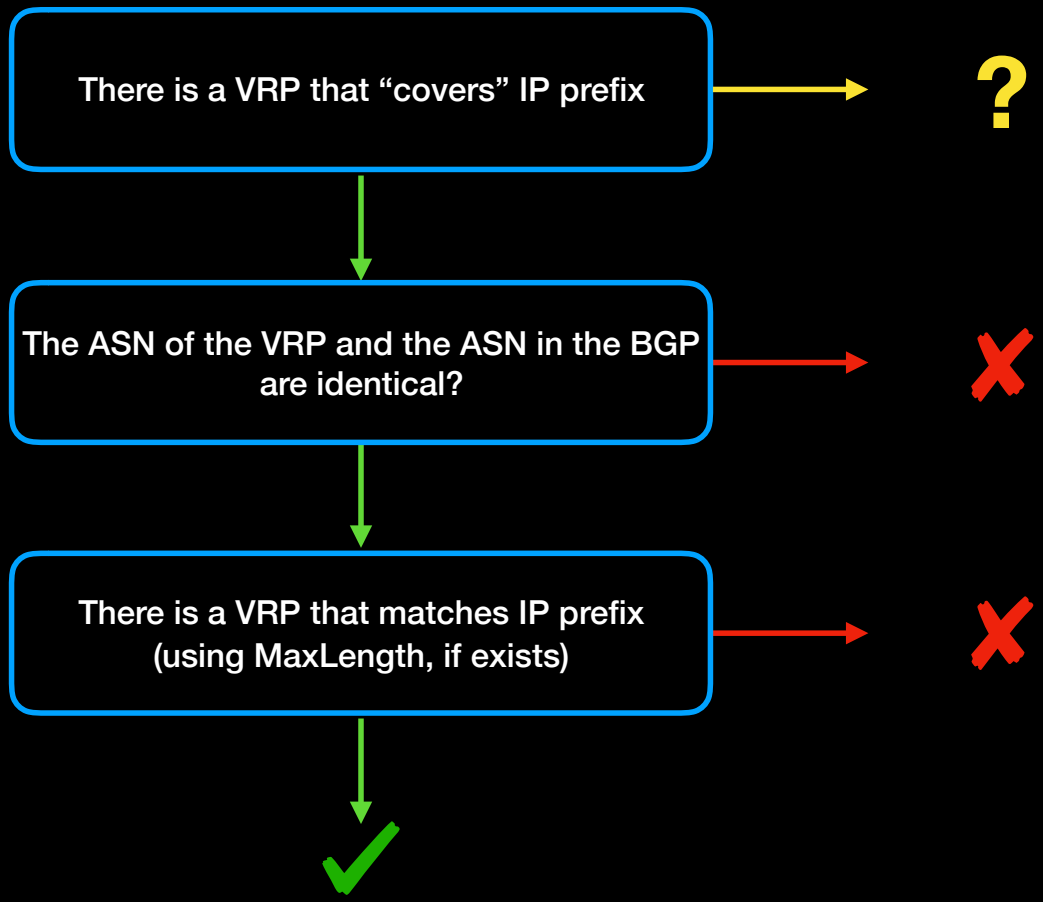

Datasets (1)
RPKI Objects

	Measurement Period*	VRPs (from the latest snapshot)	
		Number	Percent of ASes
APNIC	2011-01 ~ 2019-02	14,025	8.14%
LACNIC	2011-01 ~ 2019-02	4,510	9.33%
RIPENCC	2011-01 ~ 2019-02	40,830	16.04%
ARIN	2012-09 ~ 2019-02	4,575	1.47%
AFRINIC	2011-01 ~ 2019-02	176	3.30%

*https://ftp.ripe.net/rpki

Deployment: VRPs

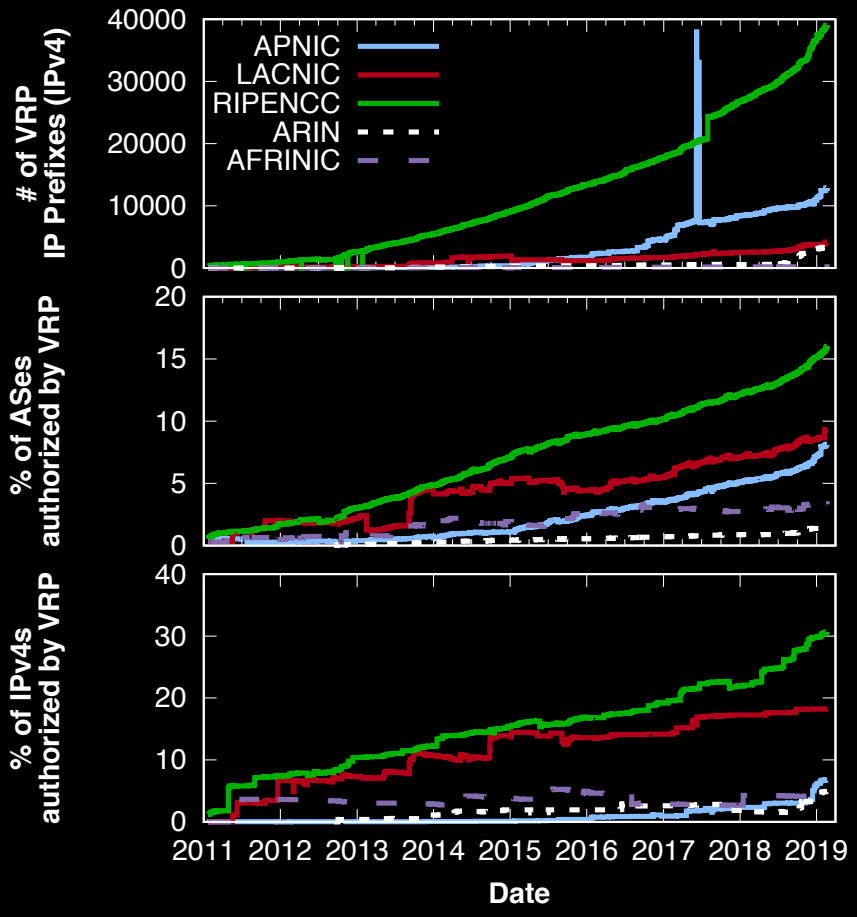

A general increasing trend in adoption of RPKI!

It varies significantly between RIRs:
1.38% (ARIN) ~ 15.11% (RIPENCC) of ASes and
2.7% (AFRINIC) ~ 30.6% (RIPENCC) of IPv4 addressesare authorized by VRPs

* AS4775, AS10091, AS9299

Datasets (2)
BGP Announcements

	Measurement Period	# of	
		VPs	Prefixes
RIPE-RIS	2011-01 ~ 2018-12	24	905K
RouteViews	2011-01 ~ 2018-12	23	958K
Akamai	2017-01 ~ 2018-12	3,300	1.94M

More than 46 Billion BGP announcements

Deployment:
BGP announcements w/ RPKI

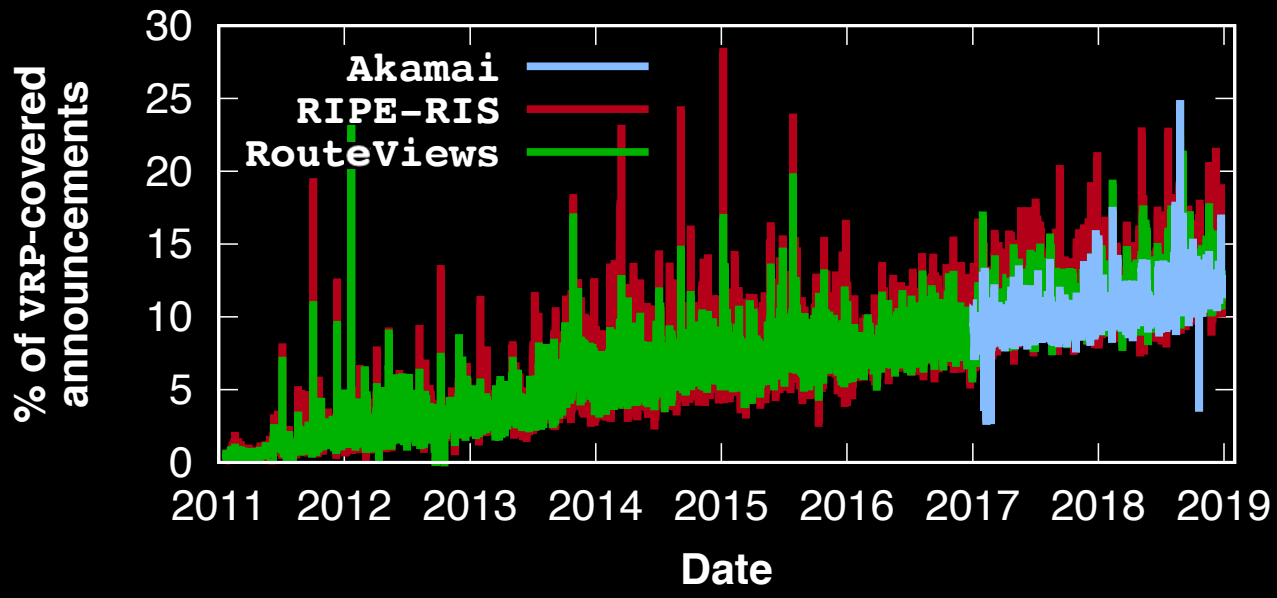

Deployment

RPKI-enabled BGP announcements are consistently increasing

RPKI validation over BGP announcements

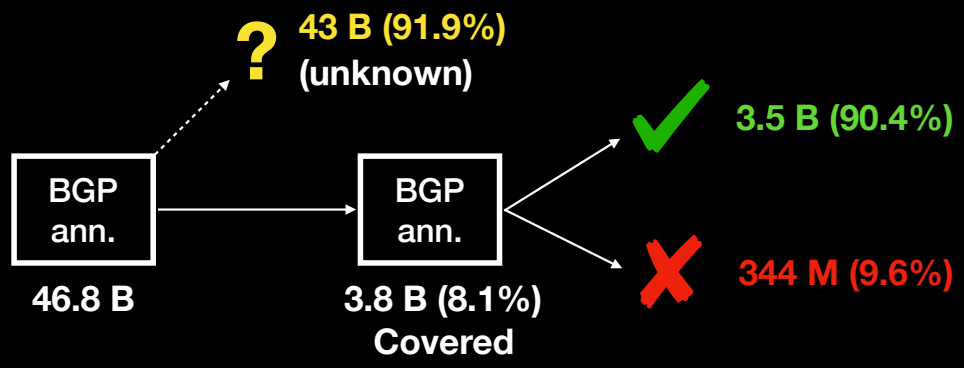

RPKI validation over BGP announcements

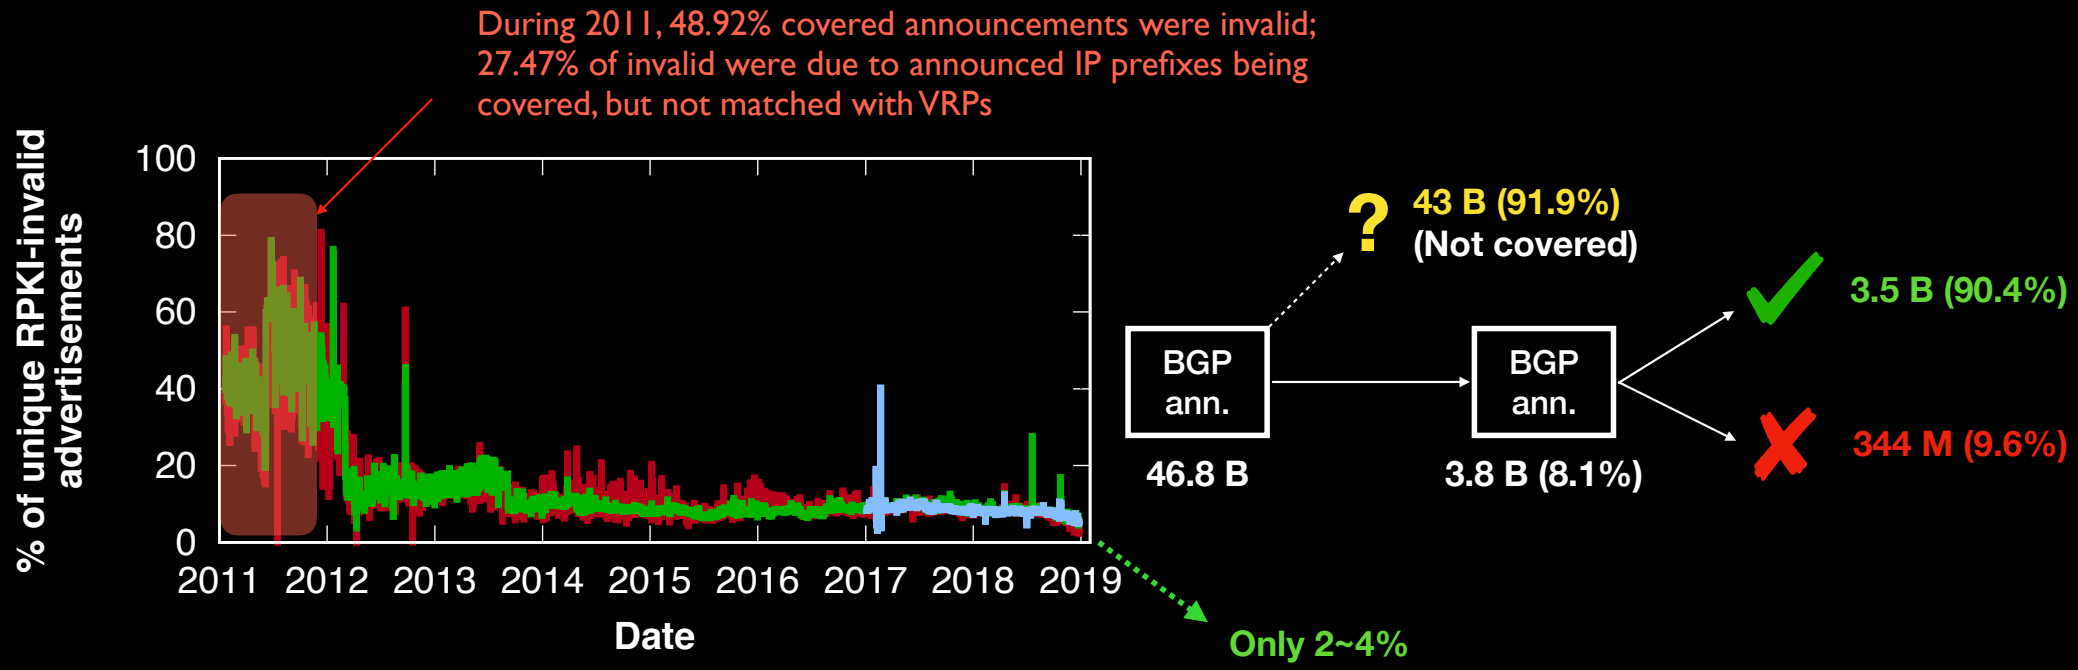

21

Then, why are they invalid?

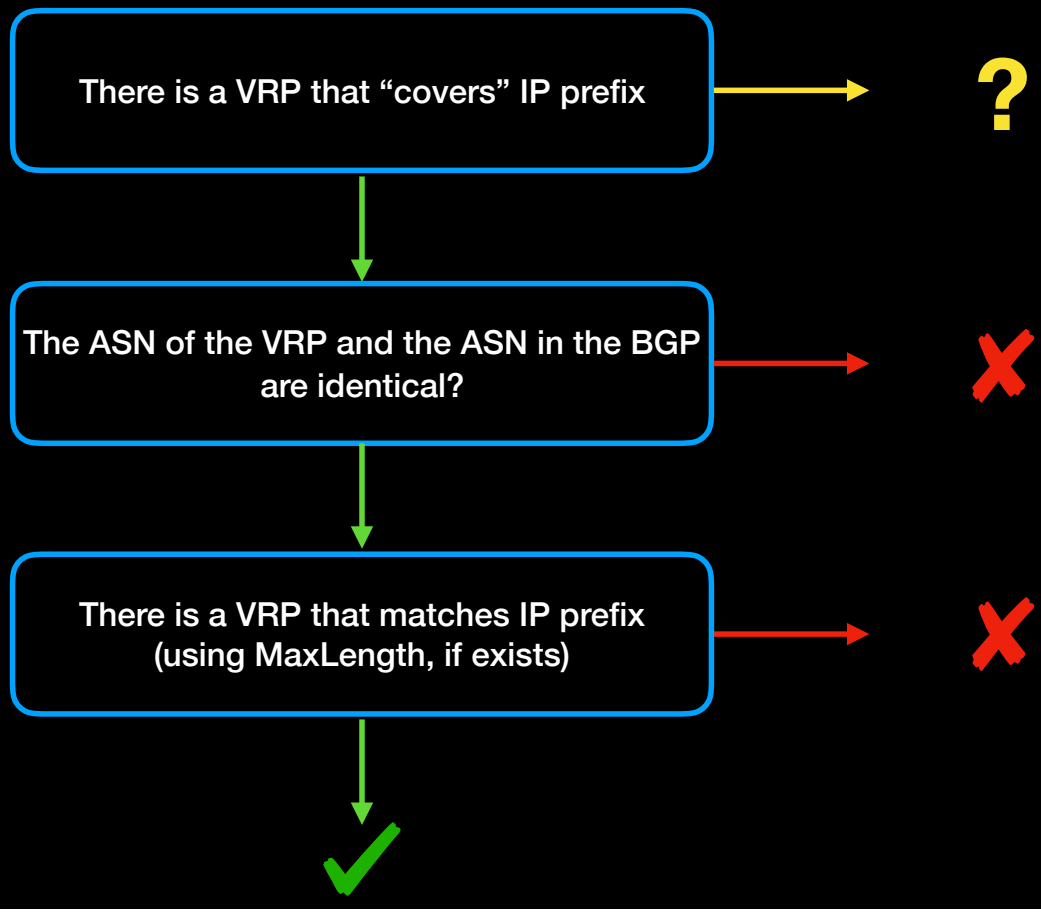

Then, why are they invalid?

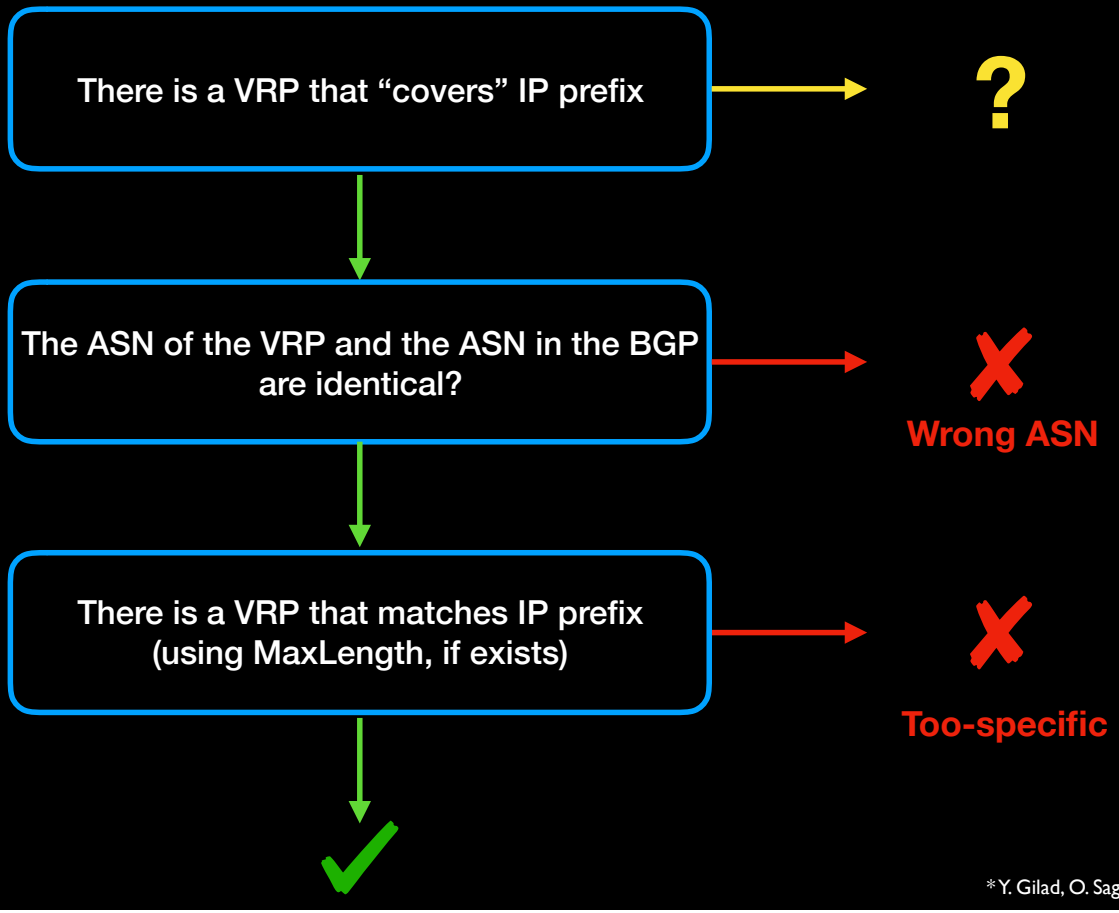

Potential Reasons:

- Malicious hijacking attacks?

Potential Reasons:

- Misunderstanding of ROAs (VRPs) of network operators*
- Stale ROAs
- …

*Y. Gilad, O. Sagga, and S. Goldberg. MaxLength Considered Harmful to the RPKI. CoNEXT, 2017.

Too specific vs. Wrong ASNs

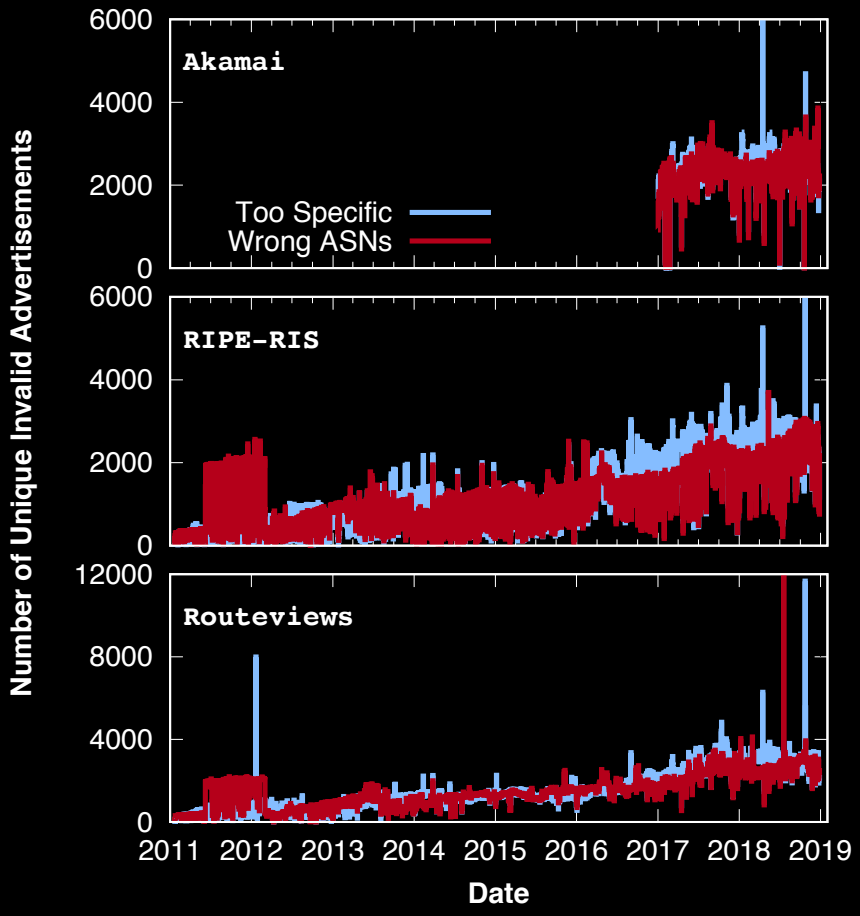

Too specific vs. Wrong ASNs

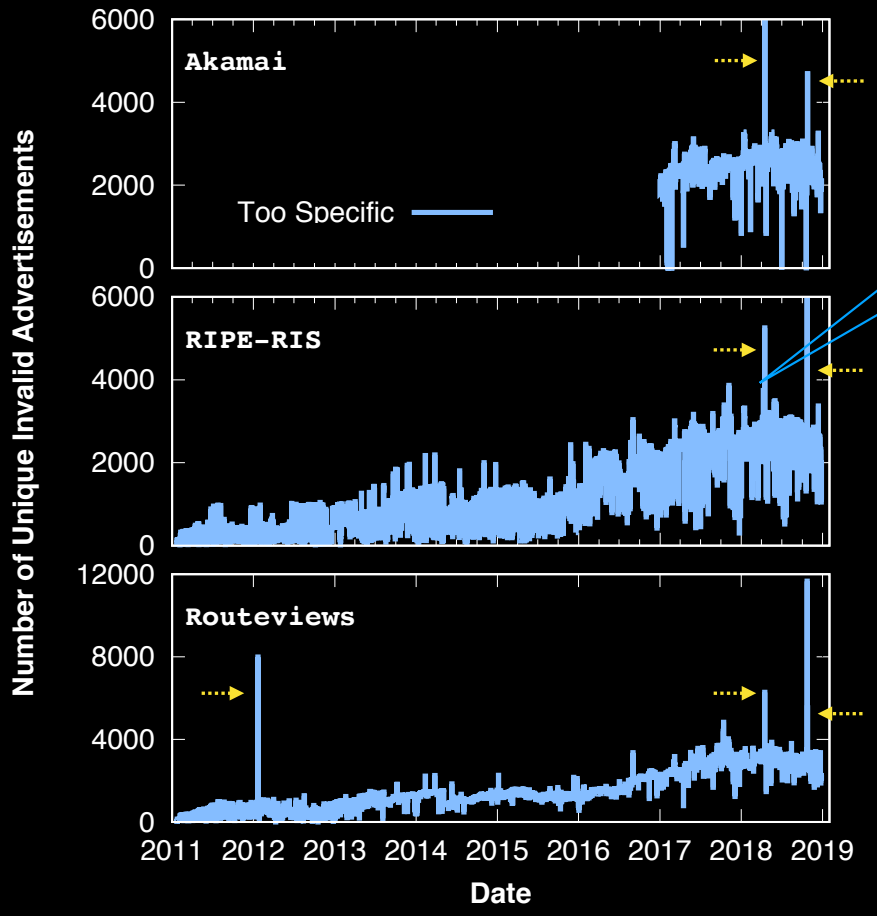

AS 5089 (Virgin Media Limited)

On April 16, 2018,
3,200 IP prefixes are more specific than the VRPs; none of them specified MaxLength

AS12322 (Free SAS)

6 ROAs for 7,671 (96.0%) IP prefixes are more specific than the VRPs (w/o MaxLength)

January 21, 2012

January 22, 2012

Added the MaxLength to include more specific IP prefixes

8,800 IP prefixes went invalid failing to specify a proper value for MaxLength

October 23, 2018

Wrong ASN

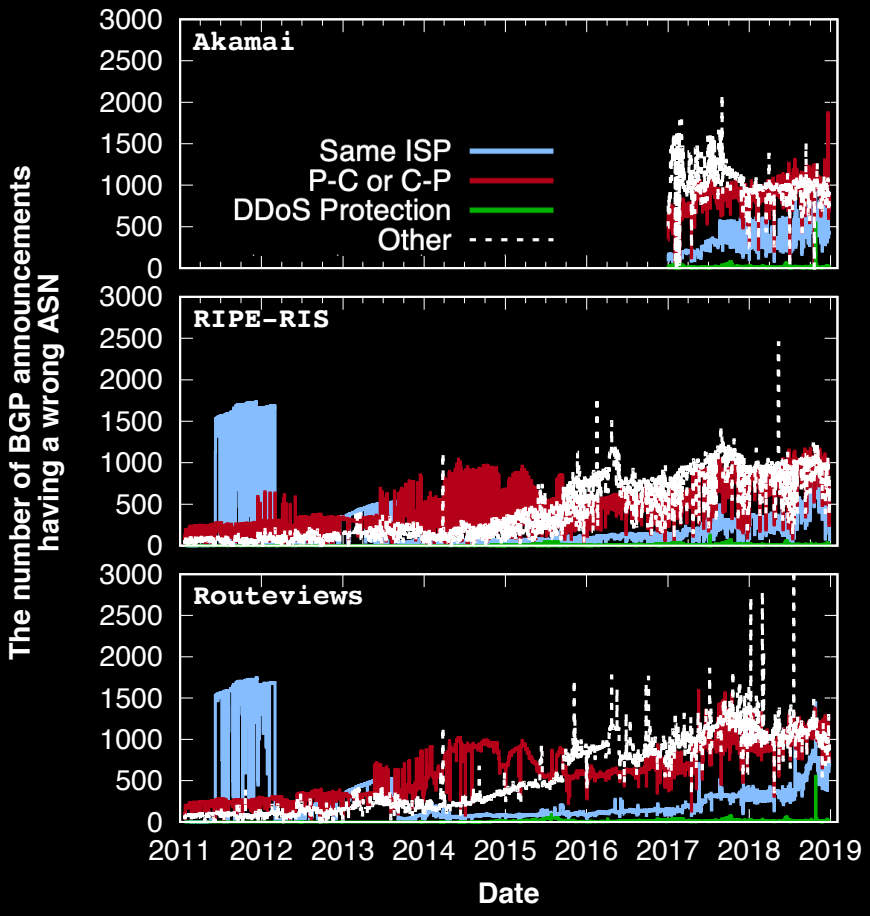

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource “scrubbing” of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don’t know, but it could be malicious (e.g., hijacking)

Wrong ASN: Same ISP

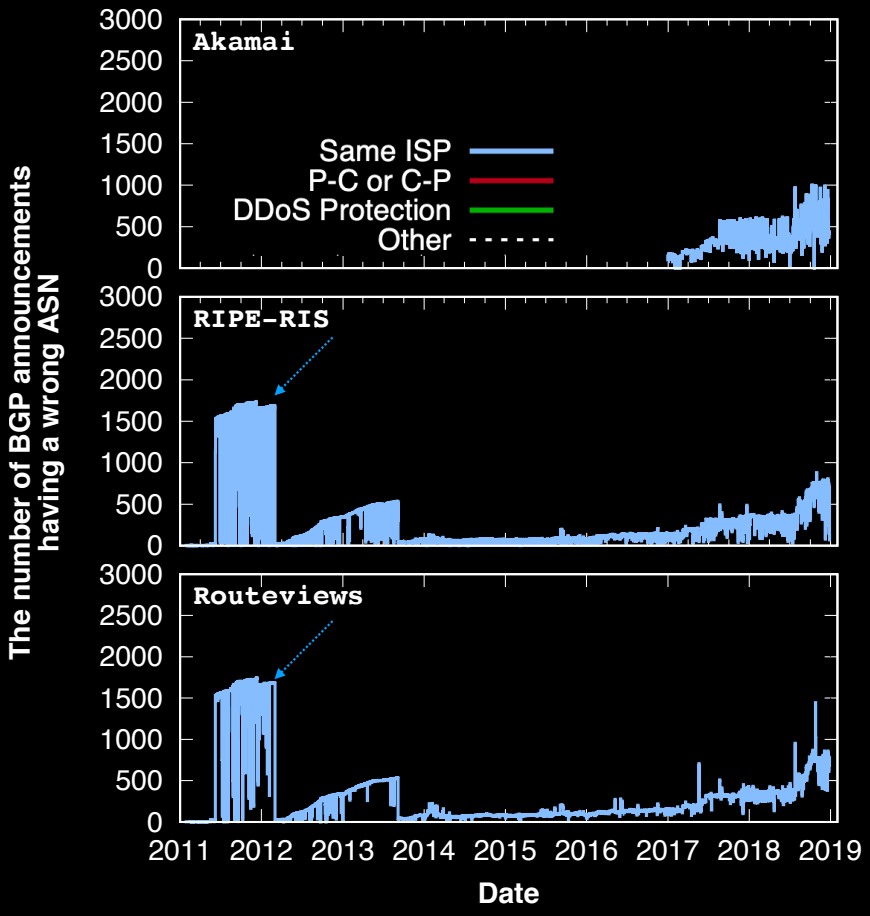

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource "scrubbing" of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don't know, but it could be malicious (e.g., hijacking)

Telmex Columbia S.A. manages two ASes (AS 10620, 14080)
AS 10620 announced 1,500 prefixes supposed to be from AS 14080 for 9 months

Wrong ASN: Provider — Customer Relationship

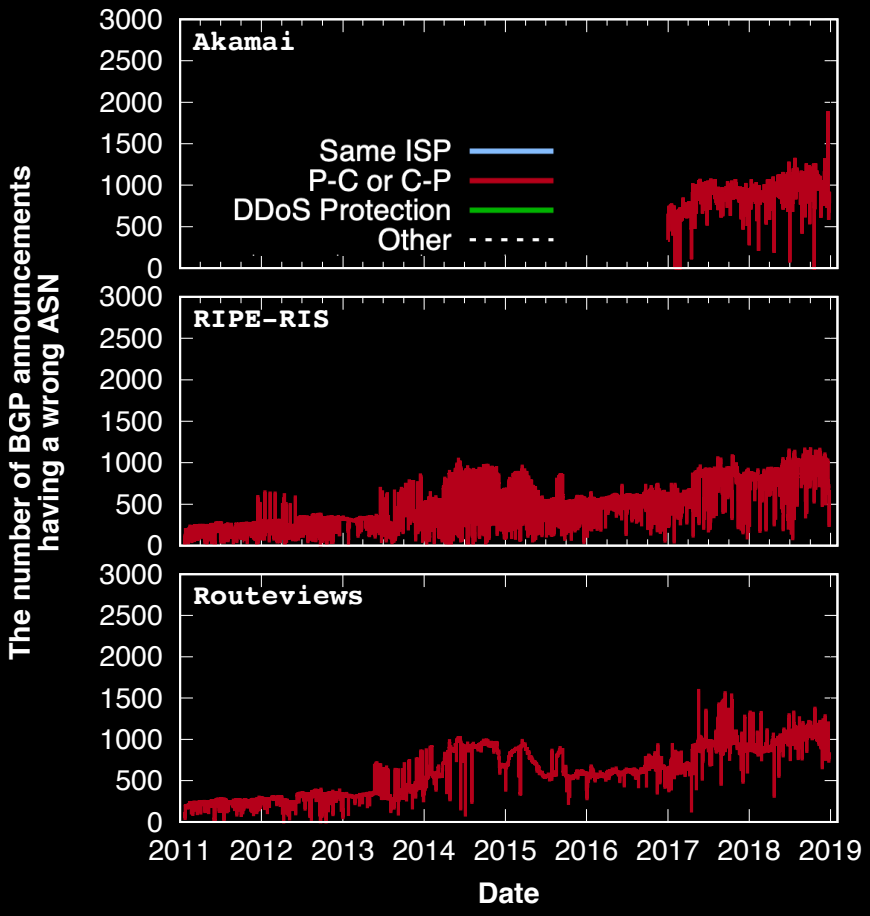

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource "scrubbing" of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don't know, but it could be malicious (e.g., hijacking)

P-C and C-P are quite prevalent; mainly due to providers that have not updated after leasing to the IP prefixes customers (up to 89.45%) such as AS 6128 (CableVision Systems) allocating to 9 different ASes

Wrong ASN: DDoS Protection

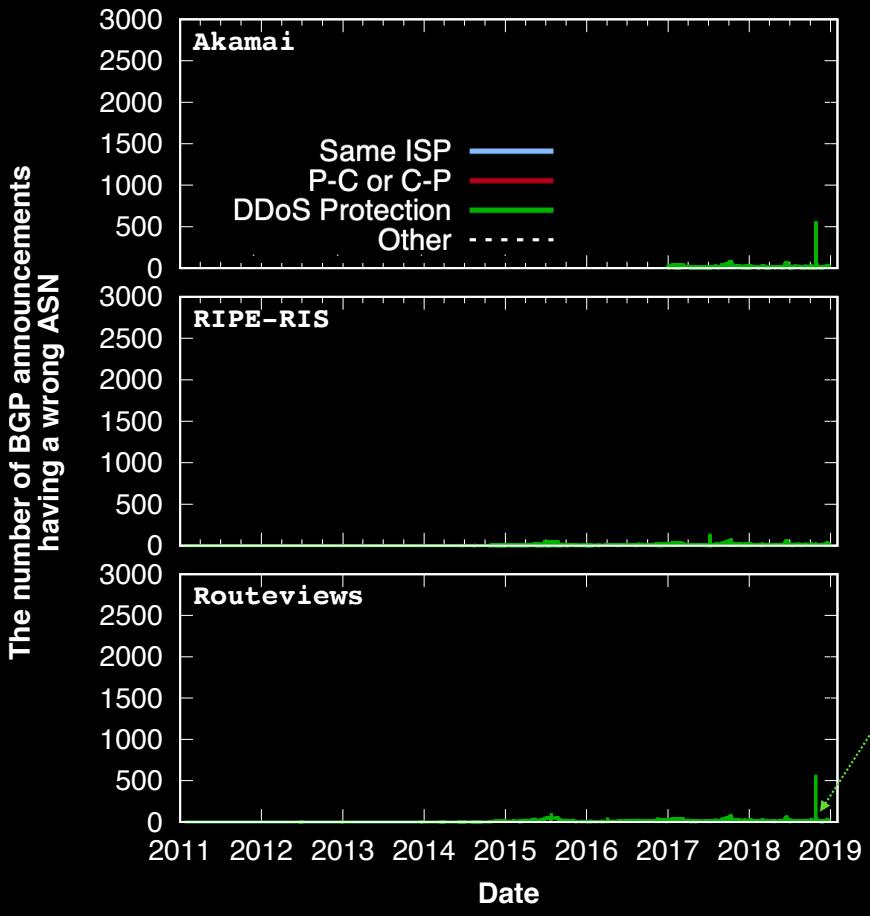

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource "scrubbing" of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don't know, but it could be malicious (e.g., hijacking)

We rarely see announcements from DDoS protection services
AS 26415 (Verisign) announced 6 IP prefixes of AS 13285 (TalkTalk)
AS 19905 (Neustar) announced 1 IP prefix of AS 21599

Wrong ASNs: The others (possibly suspicious)

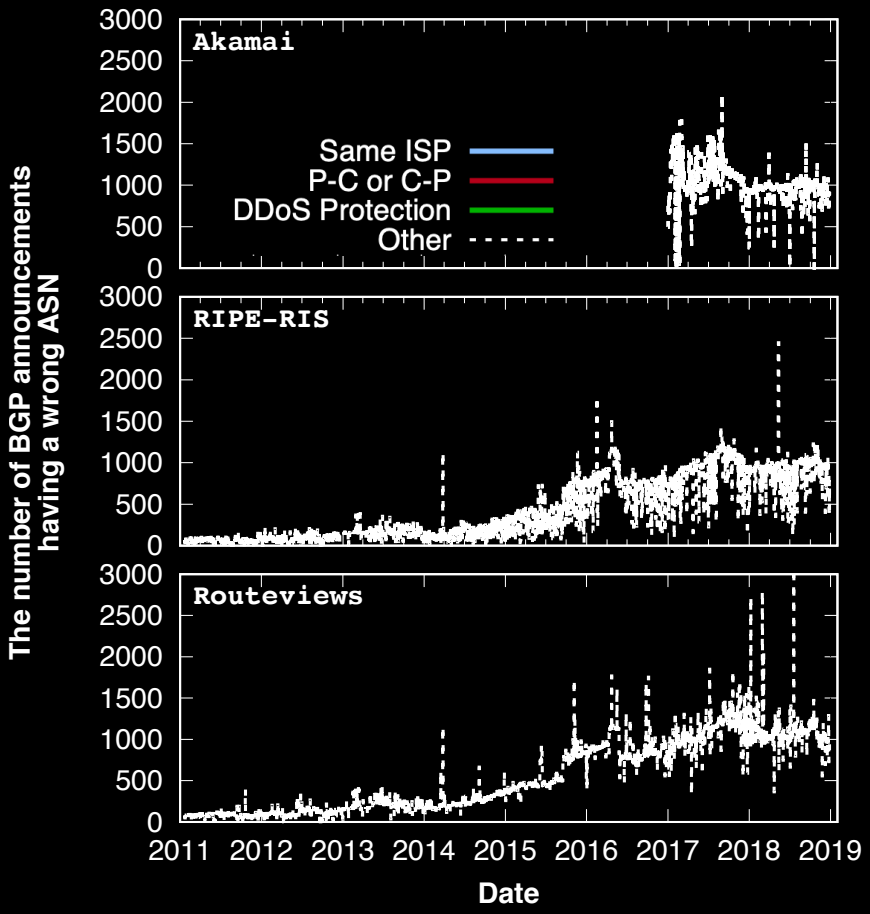

(1) AS 37468 (Angola Cables) announced more than 2,500 IP prefixes owned by 82 ASes on May 11, 2018 and 15,000 IP prefixes owned by 1,554 ASes on July 19, 2018

(2) Targeted attack: AS 55649 (a private ISP in Hong Kong) announced 1,091 IP prefixes owned by 12 ASes, 10 of which are in China on February 28, 2018

(3) Targeted attack: 401 IP prefixes owned by AS 27738 (Ecuadortelecom S.A.) are announced by 743 ASes on January 7, 2018?

Traffic from "the others" category

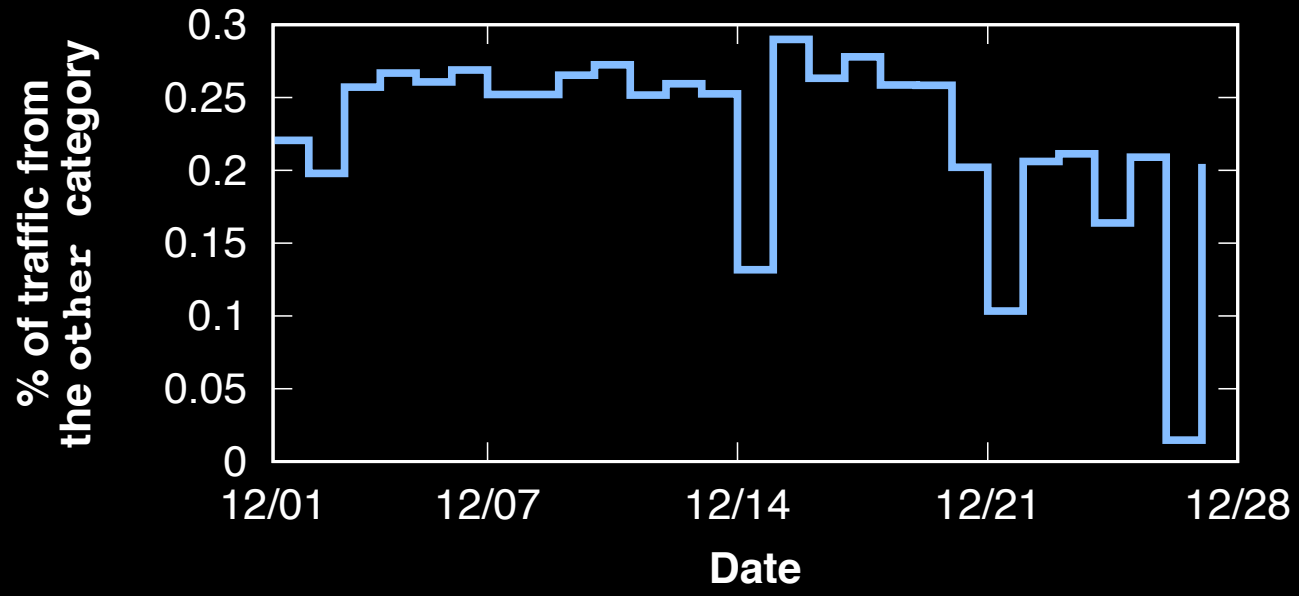

Amount of Traffic

The portion of all HTTP/S traffic coming from the other category is very small (less than 0.3%)

Case-study: BGPStream

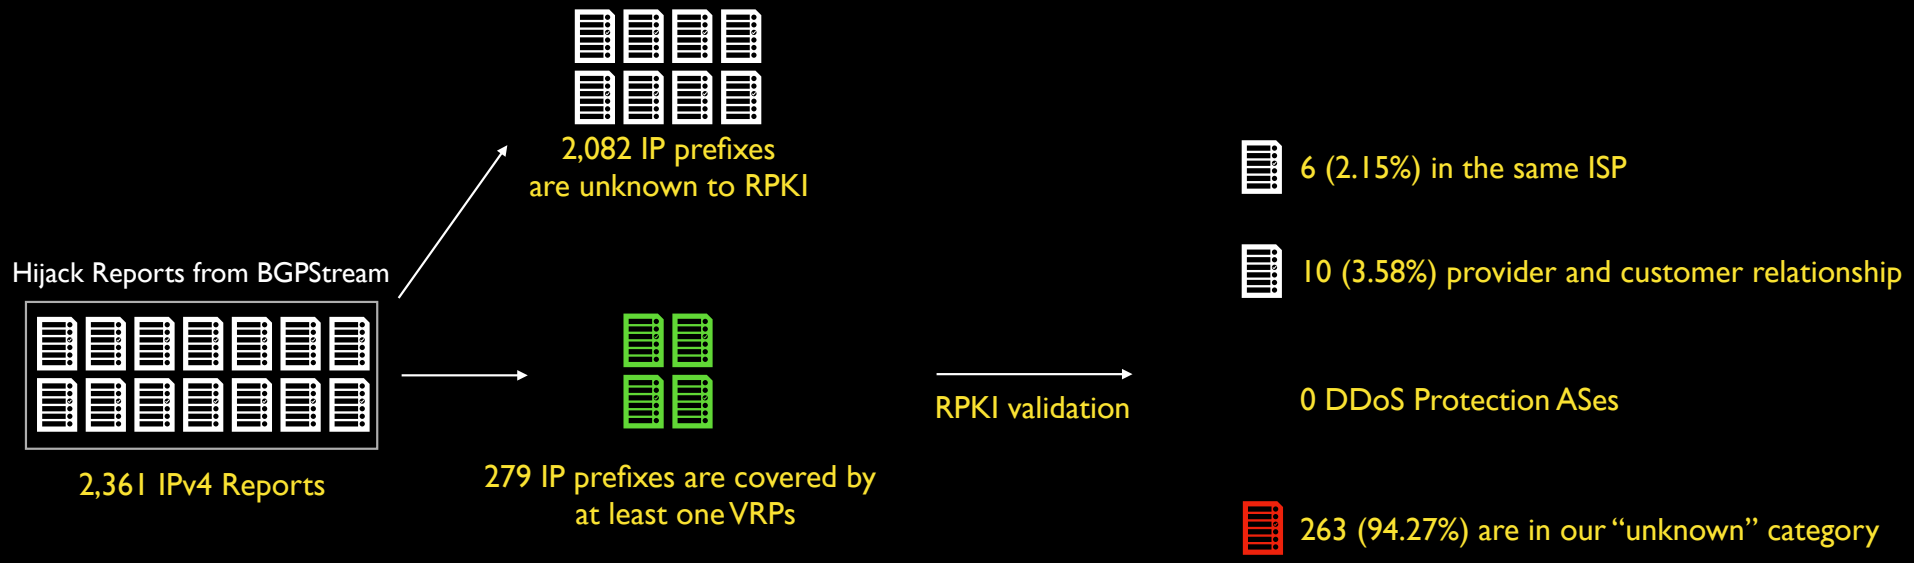

Conclusion and Discussion

- RPKI has been widely deployed
  - RPKI Objects: 2.7% (AFRINIC) ~ 30.6% (RIPENCC) of the total IPv4 space is covered
  - BGP announcements: 8.1% of BGP announcements are covered
- 2~4 % of (verifiable) BGP announcements are invalid!
  - Too specific announcements
  - Wrong ASNs

Datasets

- All the datasets and source codes are available here:
  - https://rpki-study.github.io

Discussion

D1: Identifying hijacking attempt

- Hijacking detection was never the goal of RPKI; the goal was to be able to filter out BGP updates with unauthorized announcements; however, as RPKI coverage expands and data quality keeps improving, invalid announcements detected by RPKI may become a valuable source of evidence of malicious intent.

- How can we identify hijacking attempt with high confidence?

D2: IRR vs. RPKI

- Internet Routing Registry (IRR) is a database managed by RIRs other entities containing ASNs and IP prefixes
  - Often criticized that nobody has a complete list; downloadable using ftp (sometimes without any authentication mechanism)
  - Many network operators rely on IRRs to filter or verify the BGP announcements
  - How many of them actually verifiable using RPKI? — currently communicating with RIPE NCC to fetch historical IRR datasets

D3: IPv4 vs. IPv6 (BGP Quality)

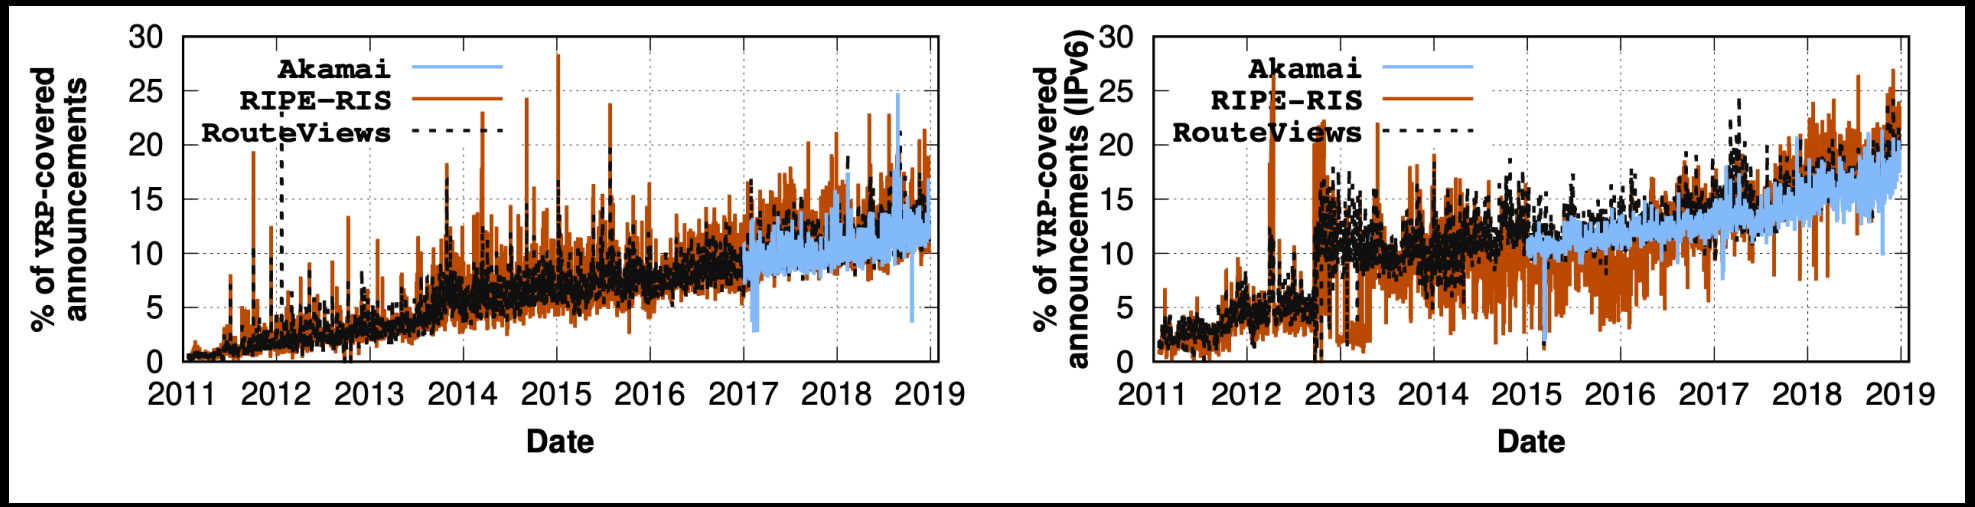

- Coverages are not that different; however, the % of IPv6 invalid announcements is 3x more than that of IPv4
  - Don't know why yet; still analyzing..

D4. Identifying RPKI-validating ASes

- Passive approach
  - Analyzing AS_PATH; if invalid IP prefixes are advertised, all ASes on the AS_PATH are not validating (but the opposite doesn't hold)
- Active approach
  - (Ben Cox and Job Snijders) Pinging two destinations; one is covered by valid ROA, and the other one is invalid (on purpose)
- Others?

D5. MaxLength

- MaxLength:
  - pros: it is efficient and gives flexibility for network operators
  - cons: if some sub prefixes are not actually advertised, those are vulnerable to forged-origin sub-prefix hijack:
    - Announcing sub-prefix that are not advertised by the owner.
    - "MaxLength Considered Harmful to the RPKI" [CoNext'17]
- Minimal ROAs:
  - The IP prefixes being advertised == The IP prefixes specified on ROAs (w/ MaxLength)
  - How many ROAs with the MaxLength enabled are actually minimal ROAs

QNA

Backup

Too-specific and MaxLength attribute

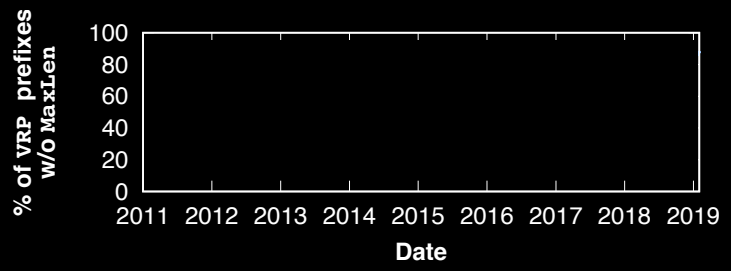

The use of MaxLength has been decreasing

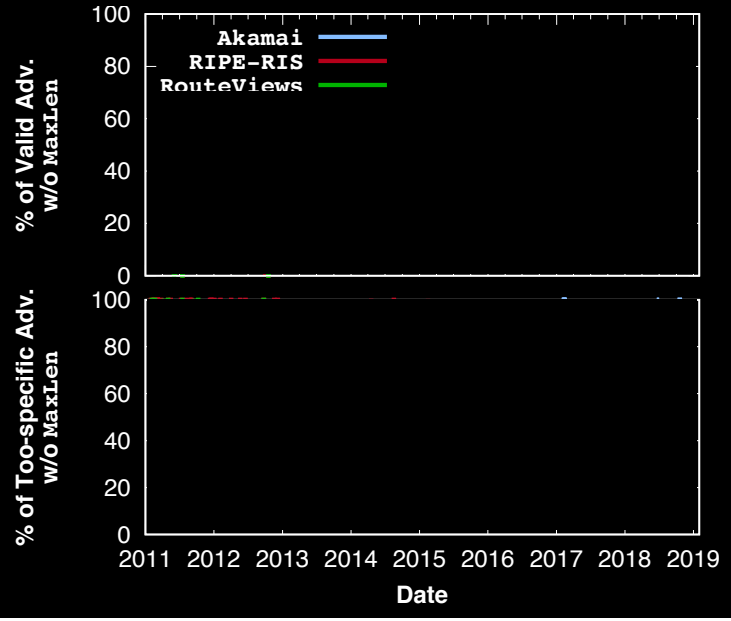

52.3% of the valid IP prefixes are validated through VRPs with the MaxLength attribute

92% of too-specific announcements are due to VRPs that do not have the MaxLength attribute