CSCI-351
Data communication and Networks

Lecture 17: BGP + Security (aka RPKI)

Warning: This may be hard to understand. Do not lose yourself during the class and keep asking questions

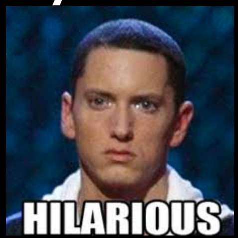

2

THE POWER OF FALSE ADVERTISING —

How an Indonesian ISP took down the mighty Google for 30 minutes

Internet's web of trust let a company you never heard of block your Gmail.

SEAN GALLAGHER - 11/6/2012, 11:07 AM

Google's services went offline for many users for nearly a half-hour on the evening of November 5, thanks to an erroneous routing message broadcast by Moratel, an Indonesian telecommunications company. The outage might have lasted even longer if it hadn't been spotted by a network engineer at CloudFlare who had a friend in a position to fix the problem.

The root cause of the outage was a configuration change to routers by Moratel, apparently intended to block access to Google's services from within Indonesia. The changes used the Border Gateway Protocol to "advertise" fake routes to Google servers, shunting traffic off to nowhere. But because of a misconfiguration, the BGP advertisements "leaked" through a peering connection in Singapore and spread to the wider Internet through Moratel's connection to the network of Hong Kong-based backbone provider PCCW. Google was interrupted in a similar way in 2008, when Pakistan Telecom moved to block access to YouTube in Pakistan because of an order from the Pakistani government.

Tom Paseka, a networking engineer at the content distribution network and Web security provider Cloudflare, spotted the source of the outage. "When I figured out the problem," Paseka wrote in CloudFlare's blog this morning, "I contacted a colleague at Moratel to let him know what was going on. He was able to fix the problem at around 2:50 UTC / 6:50pm PST. Around 3 minutes later, routing returned to normal and Google's services came back online."

By Marie Huillet

APR 24, 2018

MyEtherWallet Warns That A "Couple" Of Its DNS Servers Have Been Hacked

Update: Data from EtherScan shows that over $150k worth of ETH has been stolen in the DNS hack. Starting from 07:17 this morning, 179 inbound transactions totaling 216.06 ETH were sent to ETH address 0x1d50588C0aa11959A5c28831ce3DC5F1D3120d29. At 10:15, the attacker sent 215 ETH to 0x68ca85dbf8eba69fb70ecdb78e0895f7cd94da83.

And more..

BGP attacks hijack Telegram traffic in Iran

With so many users in Iran, it's unsurprising that potentially state-sponsored groups would want an access point into the banned app.

By Charlie Osborne for Zero Day | November 6, 2018 -- 11:44 GMT (03:44 PST) | Topic: Security

Mutually Agreed Norms for Routing Security (MANRS) 27 April 2018 EN ES

What Happened? The Amazon Route 53 BGP Hijack to Take Over Ethereum Cryptocurrency Wallets

By Aftab Siddiqui

Technical Engagement Manager for Asia-Pacific

For two hours, a large chunk of European mobile traffic was rerouted through China

It was China Telecom, again. The same ISP accused last year of "hijacking the vital internet backbone of western countries."

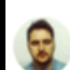

By Catalin Cimpanu for Zero Day | June 7, 2019 -- 19:41 GMT (12:41 PDT) | Topic: Security

BGP Hijacking: how it works (high-level view)

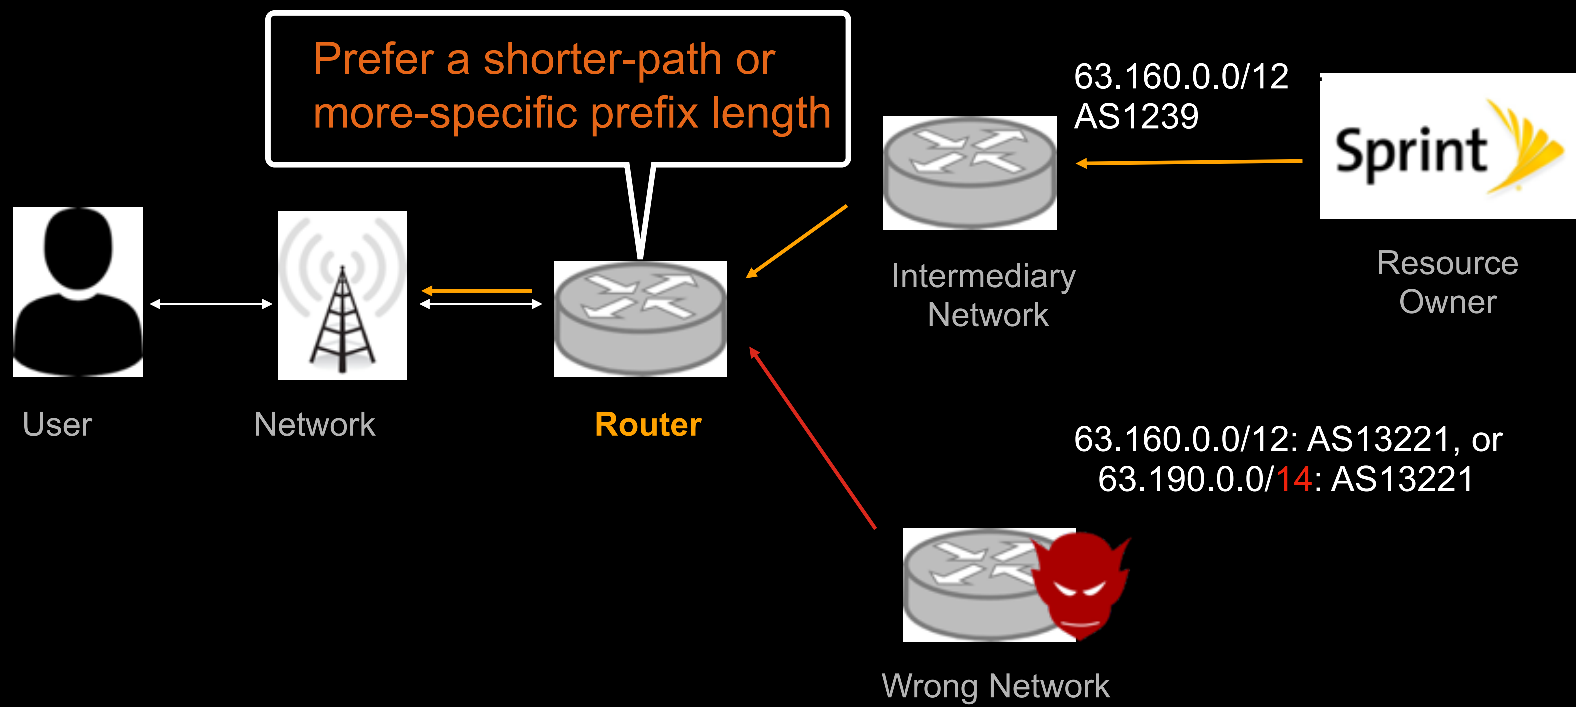

Resource PKI
(Public Key Infrastructure)

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

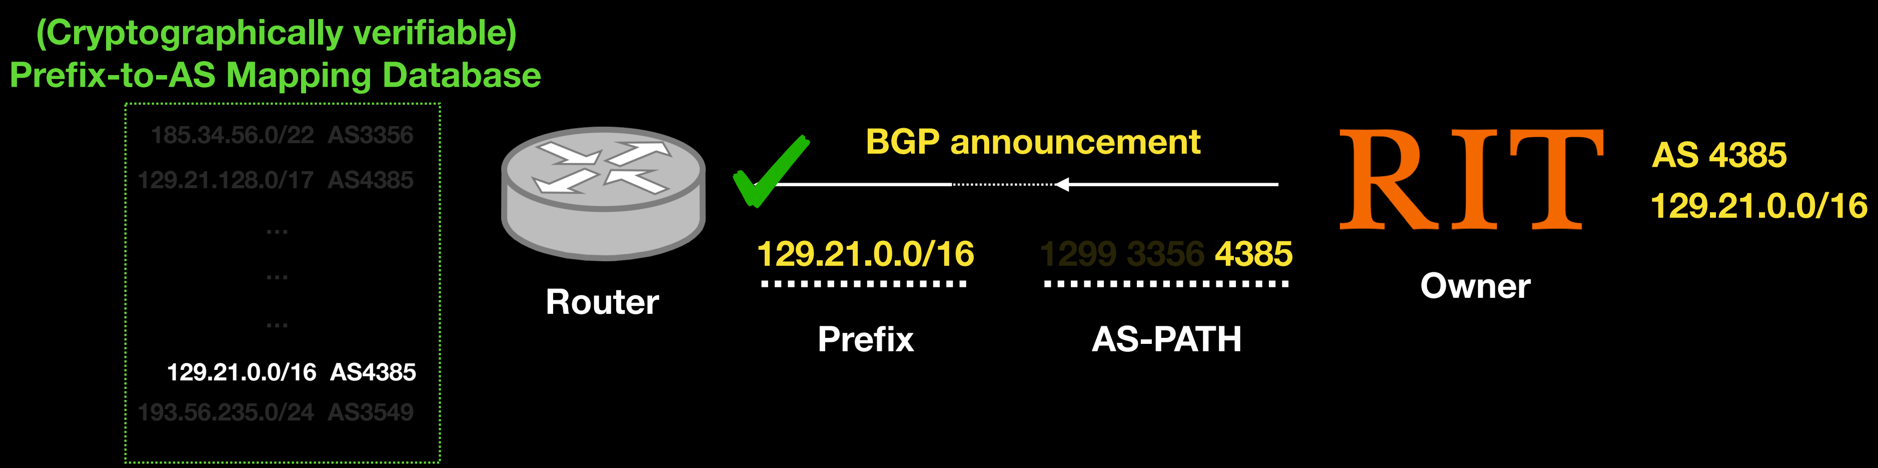

RPKI: How it works?

What does an resource owner needs to do to protect their IP prefixes?

BGP announcement

Router

RIT

Owner

AS 4385
129.21.0.0/16

How can a router verify it using RPKI?

RPKI Structure

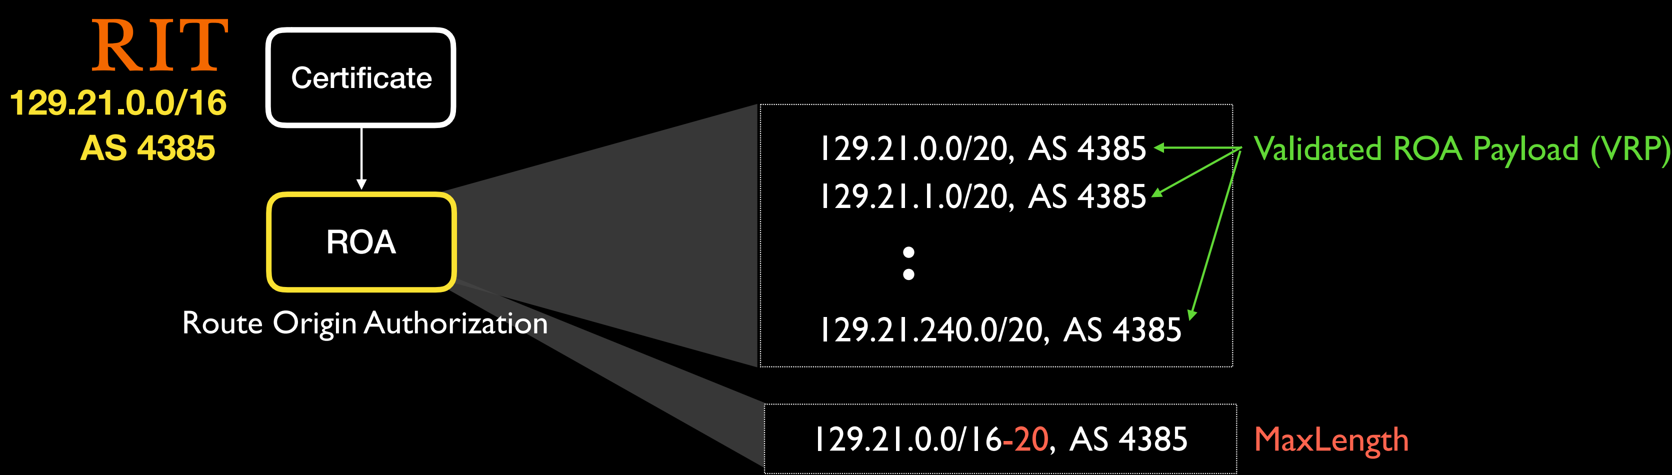

RPKI Structure

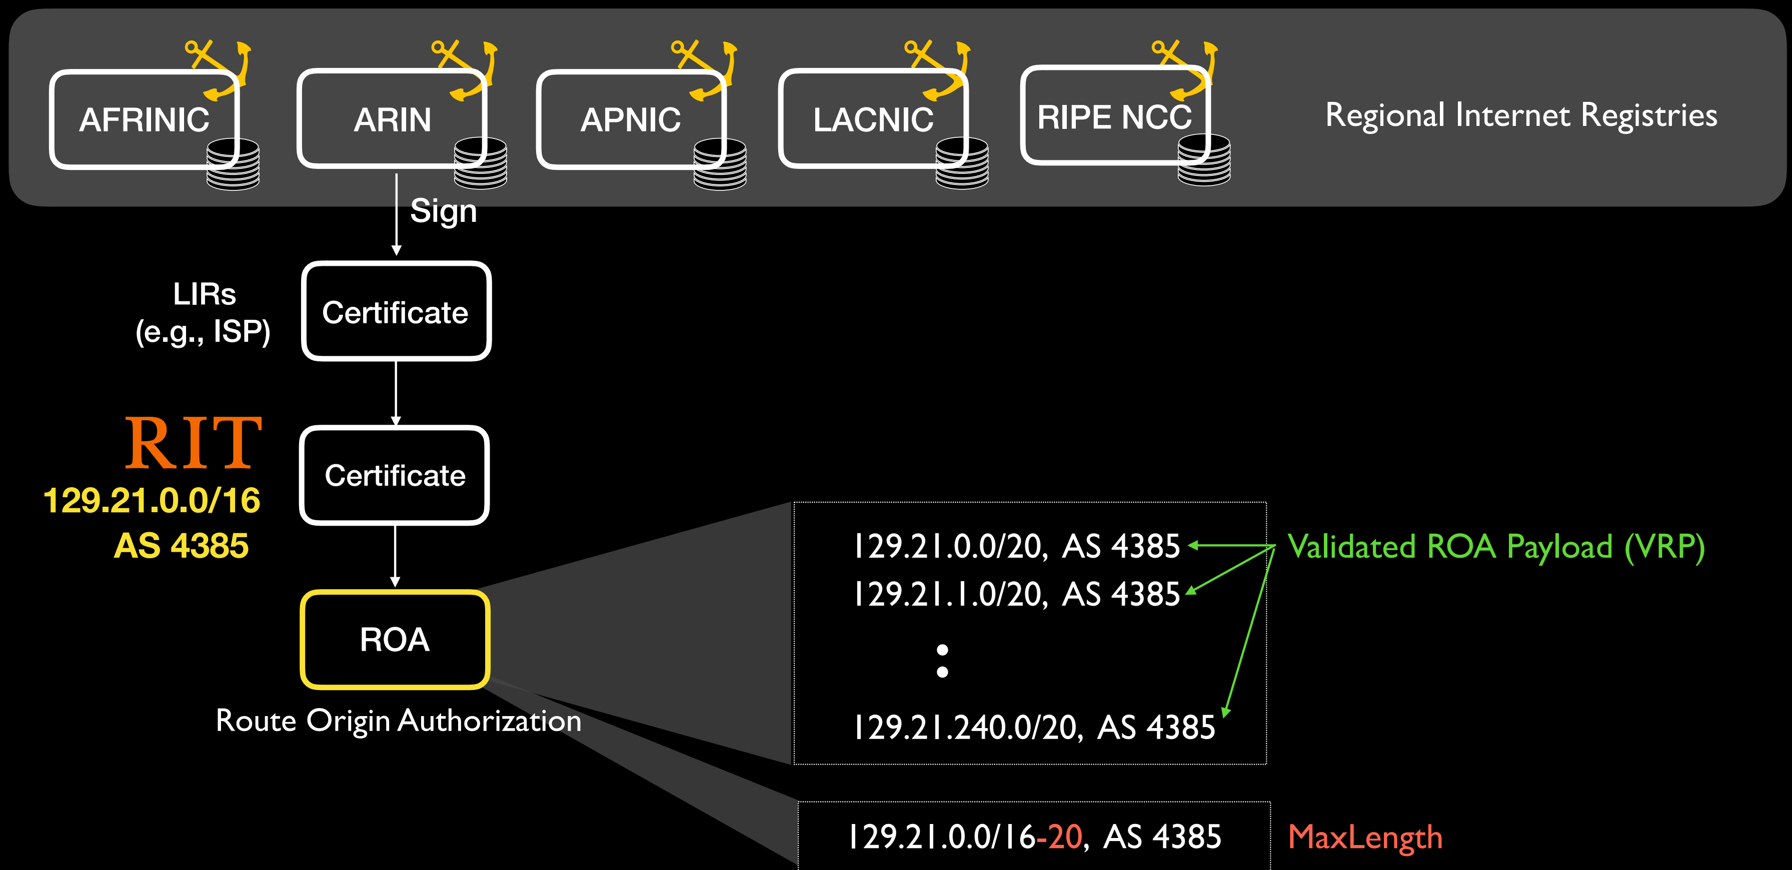

RPKI: How it works?

What does an resource owner needs to do to protect their IP prefixes?

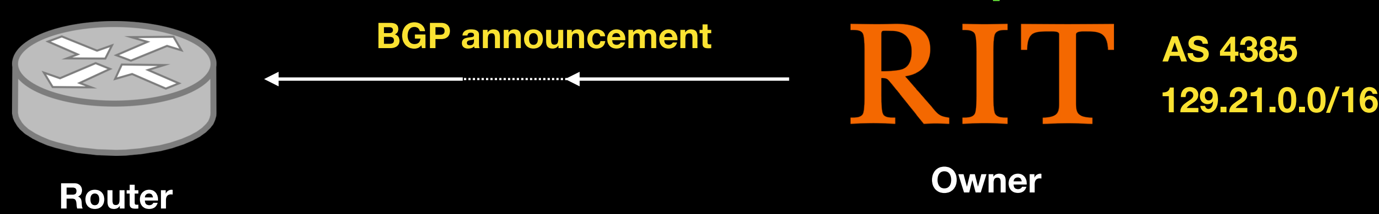

How can a router verify BGP announcements using RPKI?

RPKI: How it works?
Validation process: Valid

Prefix-to-AS Mapping Database

1.1.0.0/16 AS 111
2.0.0.0/8-16 AS 222
3.3.0.0/16 AS 333
4.4.4.0/24 AS 444

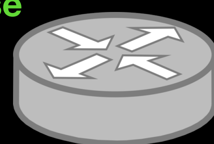

Router

BGP announcement

1.1.0.0/16 AS 111

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

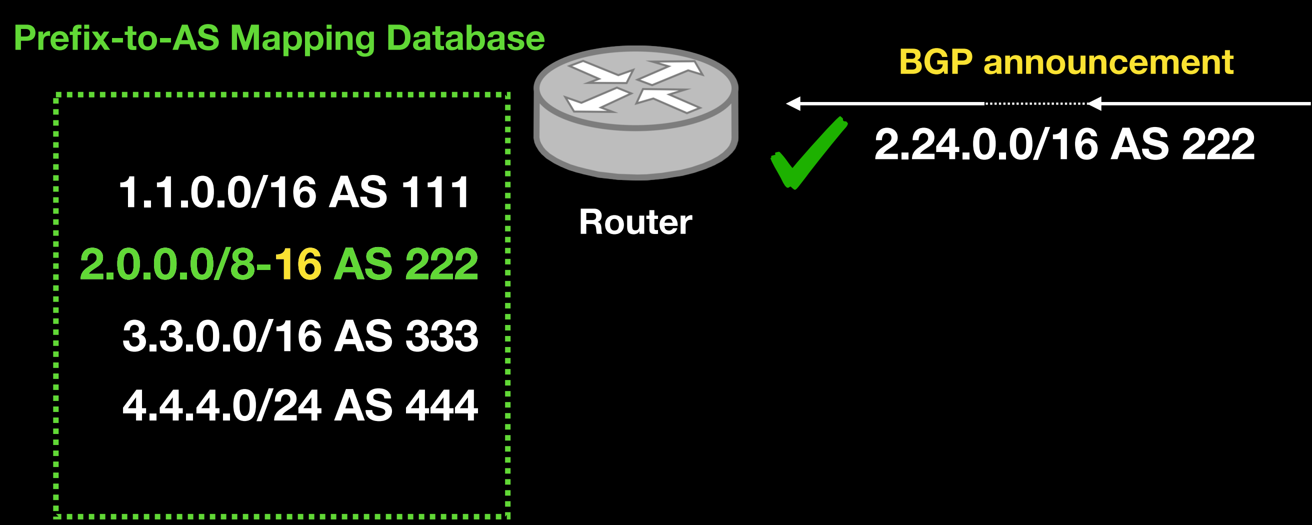

RPKI: How it works?
Validation process: Invalid (too-specific)

Prefix-to-AS Mapping Database

1.1.0.0/16 AS 111
2.0.0.0/8-16 AS 222
3.3.0.0/16 AS 333
4.4.4.0/24 AS 444

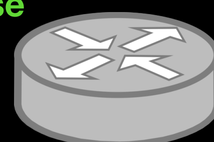

Router

BGP announcement

3.3.3.0/24 AS 333

✗ Covered, but the announcement is too specific

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

Prefix-to-AS Mapping Database

1.1.0.0/16 AS 111
2.0.0.0/8-16 AS 222
3.3.0.0/16 AS 333
4.4.4.0/24 AS 444

Router

BGP announcement

4.4.4.0/24 AS 555

✘ IP prefix is matched, but the ASN is different.

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

Prefix-to-AS Mapping Database

1.1.0.0/16 AS 111

2.0.0.0/8-16 AS 222

3.3.0.0/16 AS 333

4.4.4.0/24 AS 555

Router

BGP announcement

5.5.0.0/16 AS 555

? Uncovered, thus unknown

What RPKI can vs. can't do

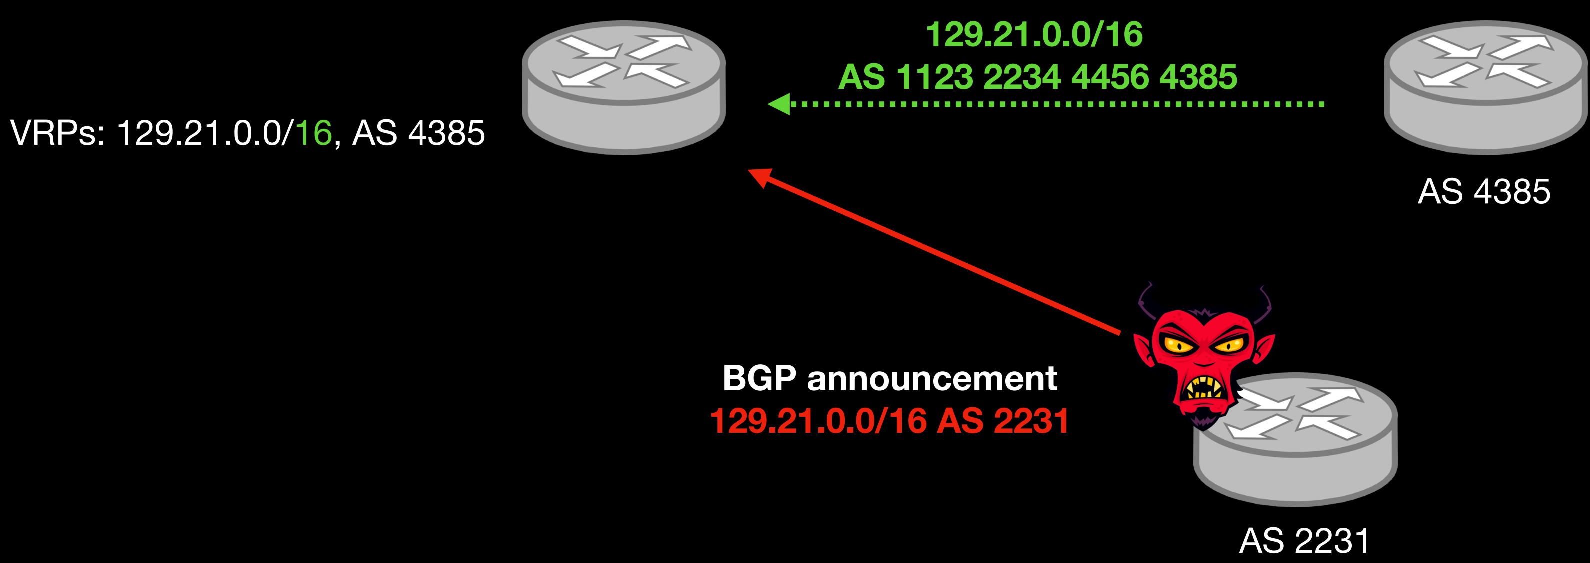

RPKI attests that the origin AS number is authorized to announce the prefix(es)

What RPKI can vs. can't do

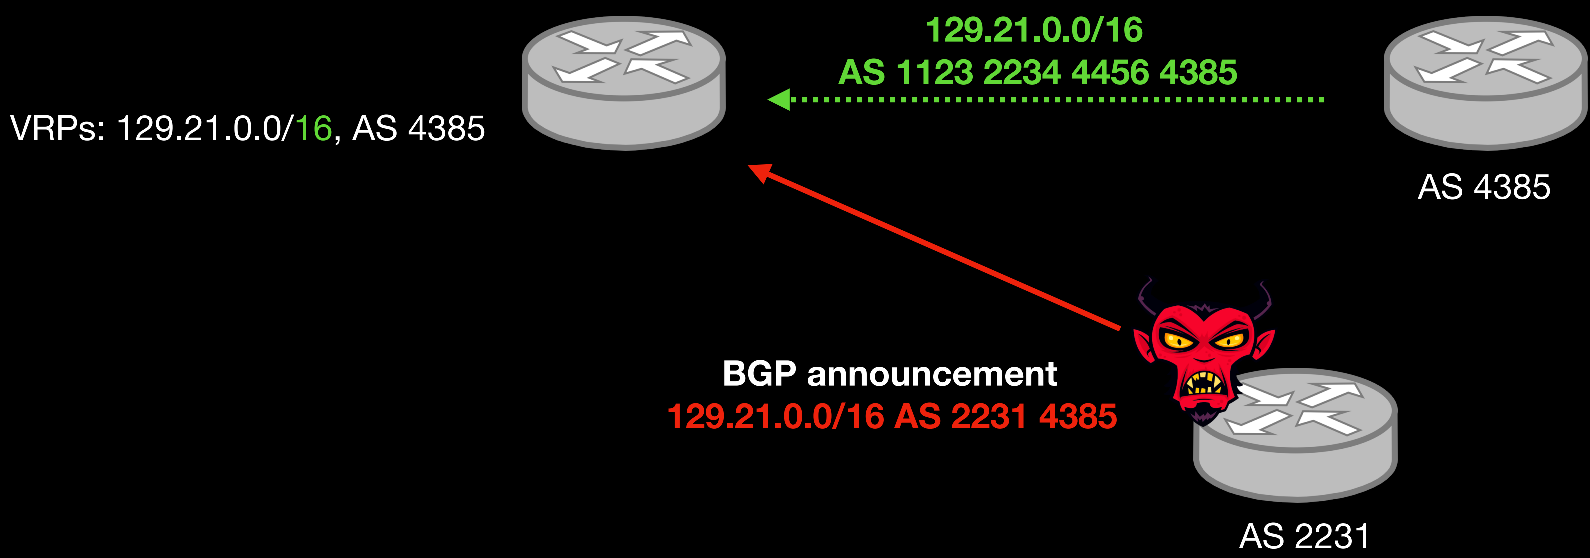

RPKI does not protect from path-shortening attacks

What RPKI can vs. can't do

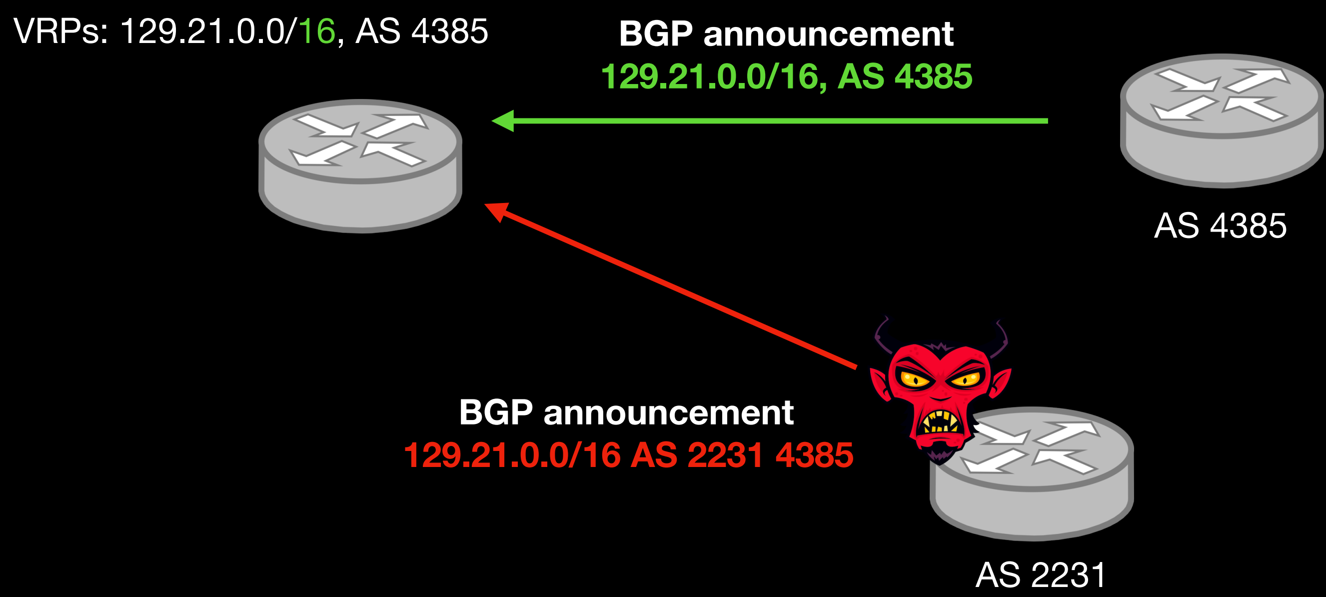

RPKI does not provide "Path" validation

What RPKI can vs. can't do

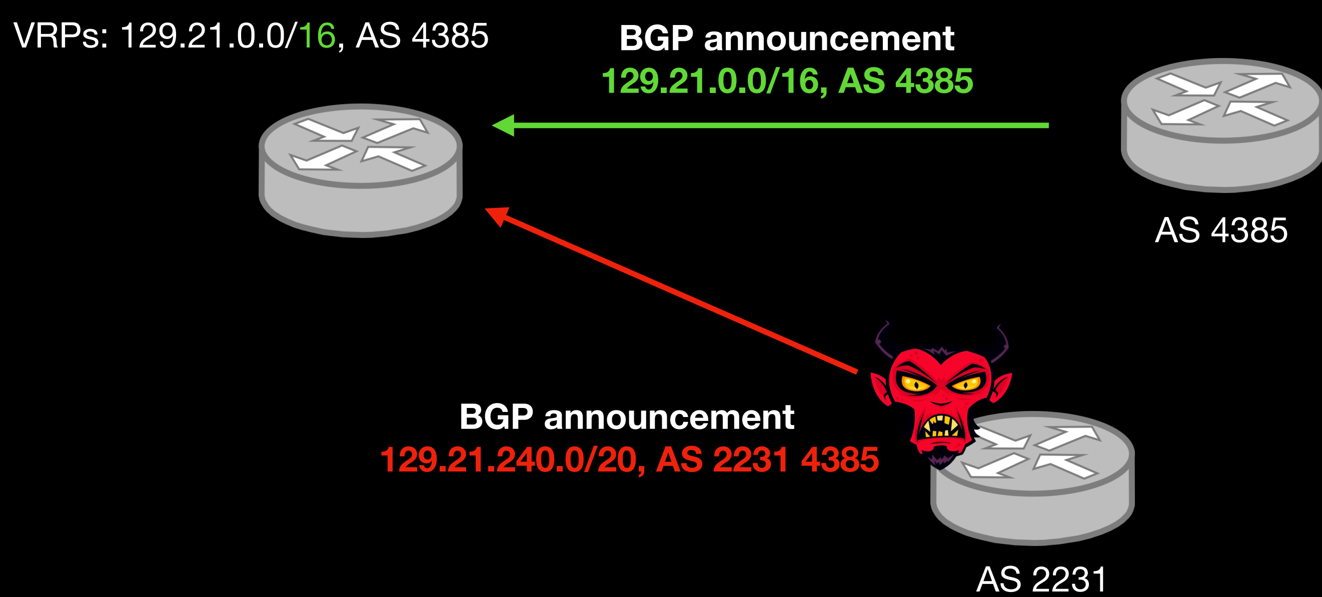

RPKI can protect from sub-prefix hijacking

RPKI: How it works?
Validation Process

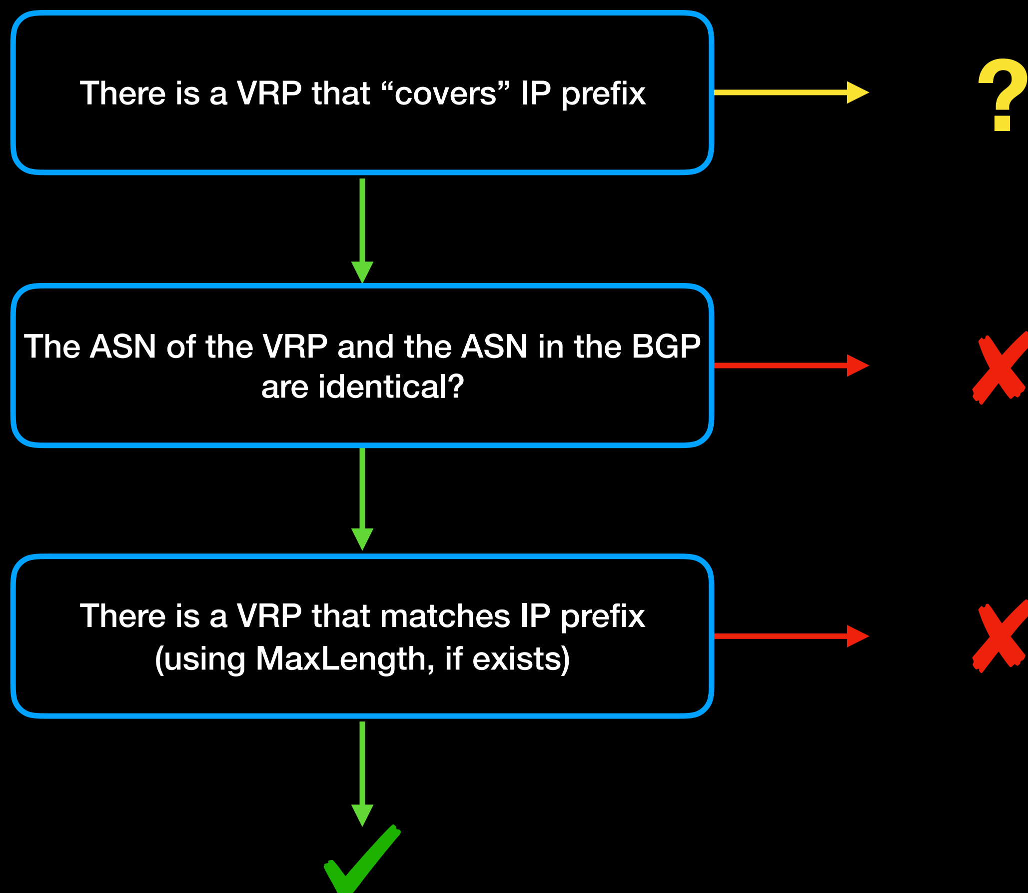

Why do we study RPKI?

It is relatively new

It works differently

It is easy to deploy

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

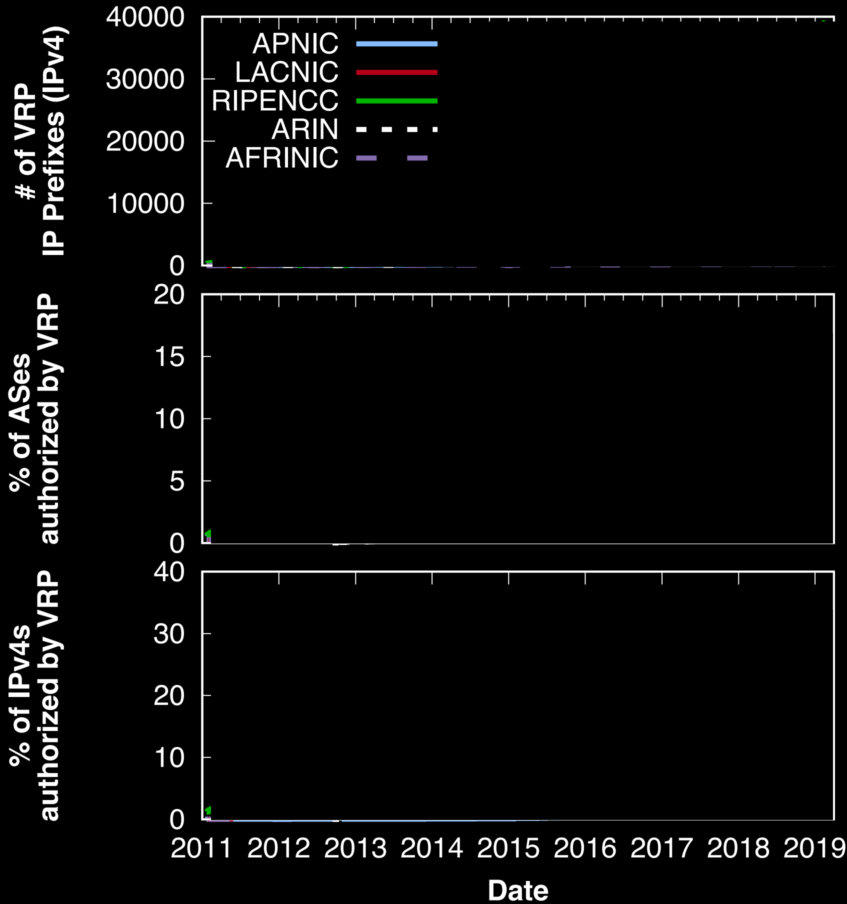

A general increasing trend in adoption of RPKI!

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

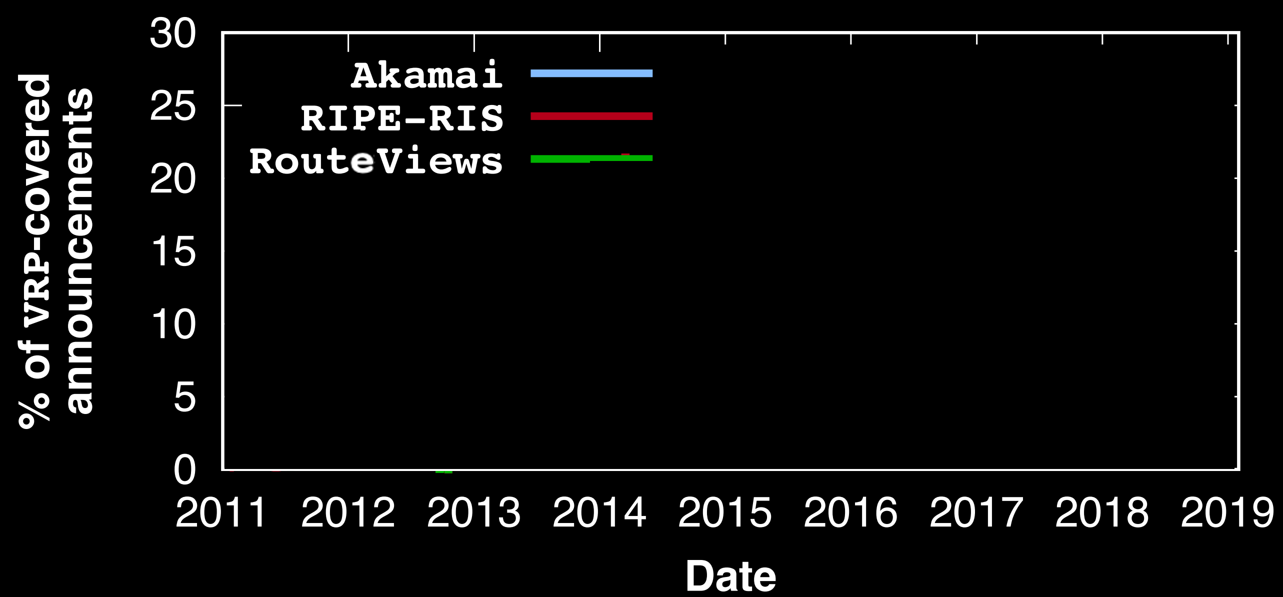

Deployment

RPKI-enabled BGP announcements are consistently increasing

RPKI validation over BGP announcements

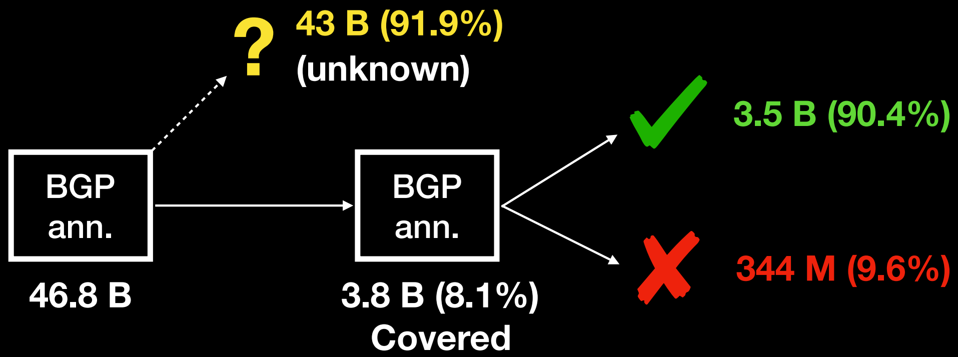

RPKI validation over BGP announcements

During 2011, 48.92% covered announcements were invalid; 27.47% of invalid were due to announced IP prefixes being covered, but not matched with VRPs

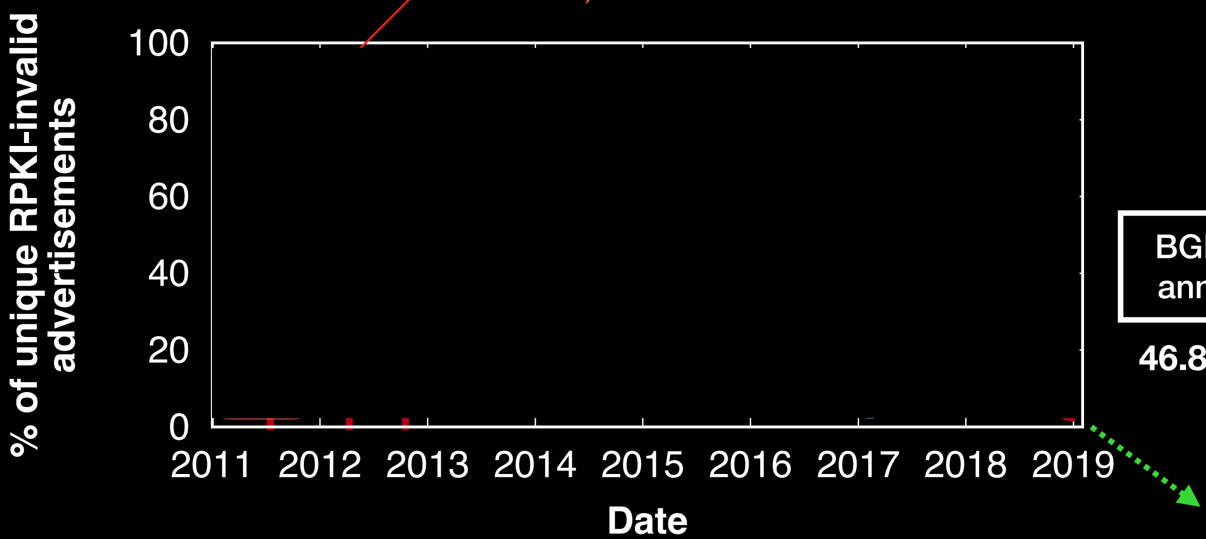

Only 2~4%

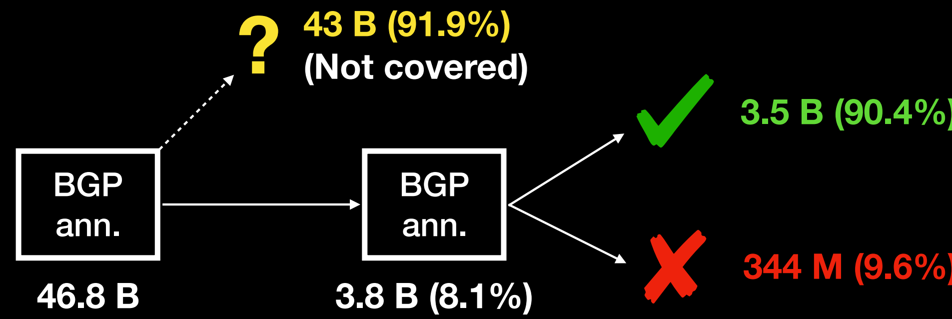

27

Then, why are they invalid?

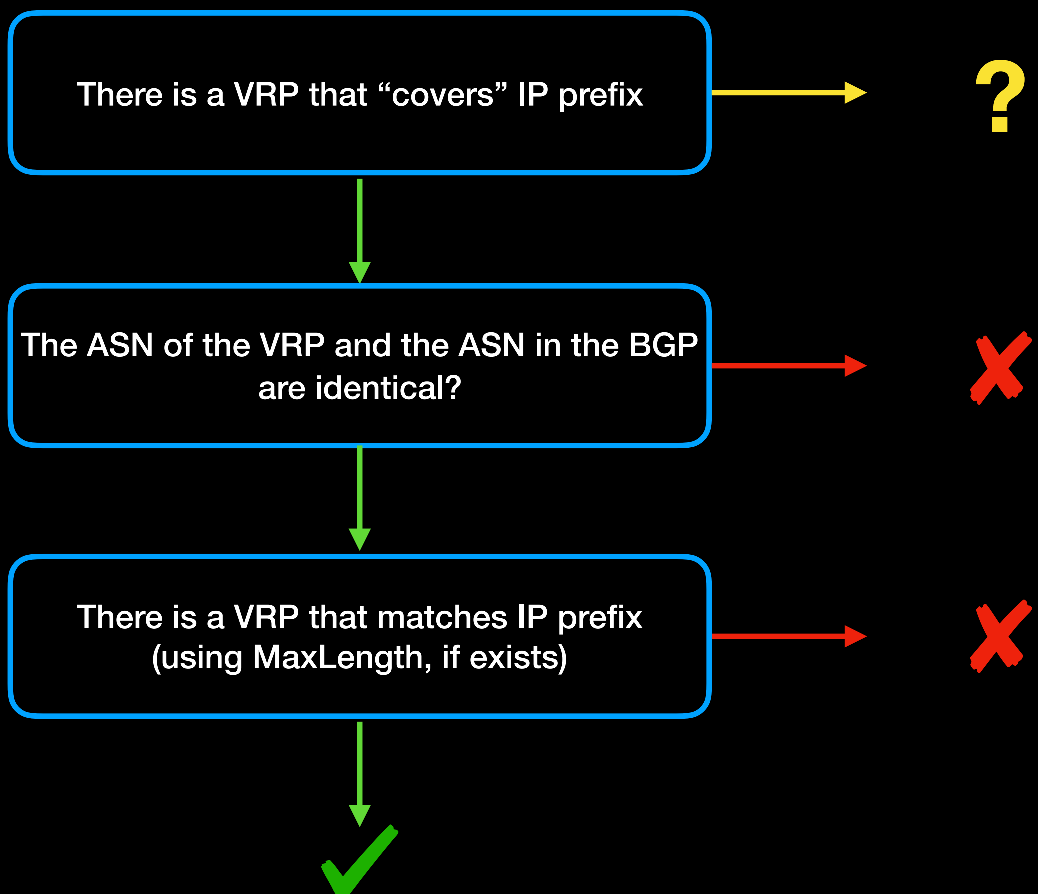

Then, why are they invalid?

There is a VRP that “covers” IP prefix

?

The ASN of the VRP and the ASN in the BGP are identical?

Wrong ASN

Potential Reasons:

- Malicious hijacking attacks?

There is a VRP that matches IP prefix (using MaxLength, if exists)

Too-specific

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

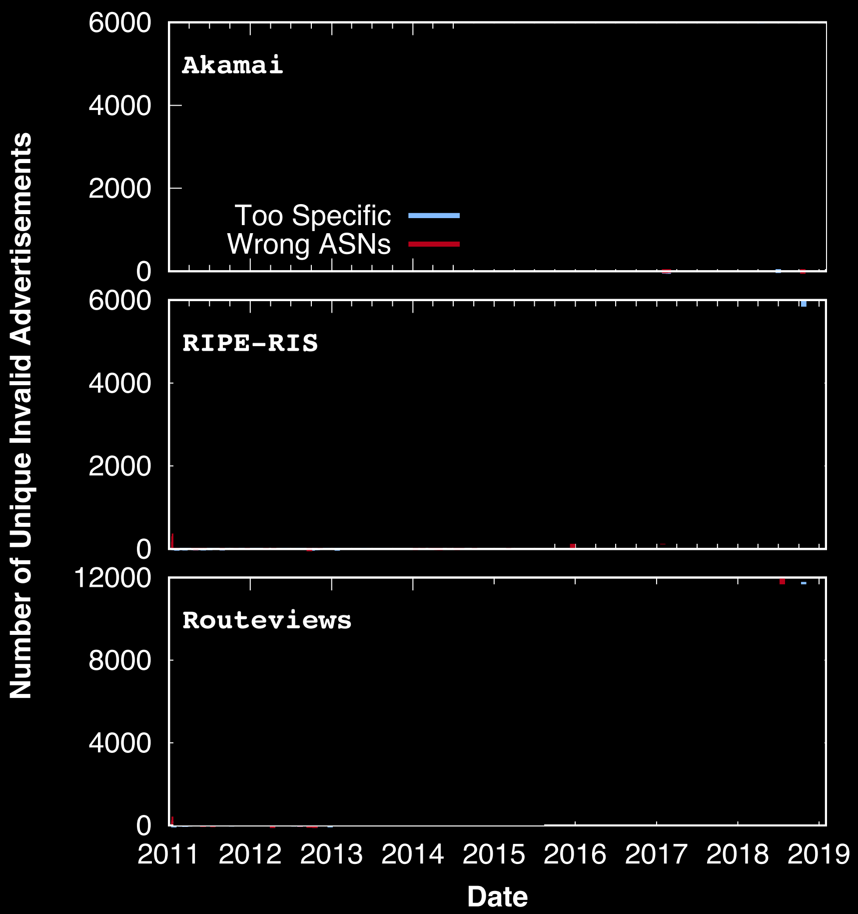

Too specific vs. Wrong ASNs

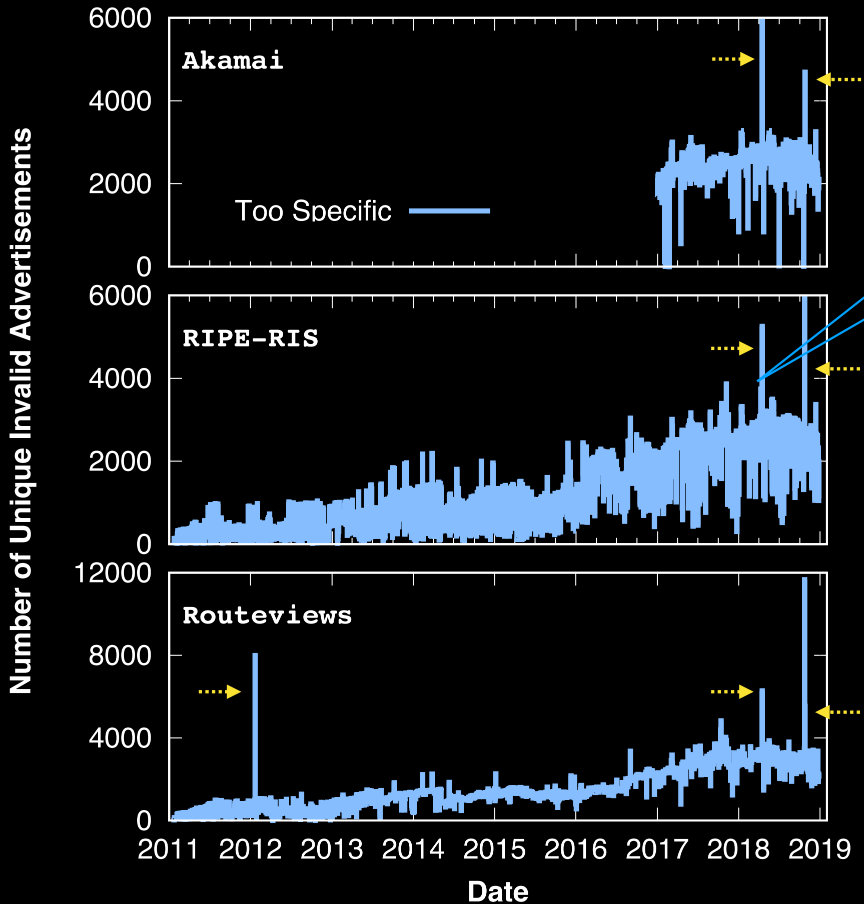

AS 5089 (Virgin Media Limited)

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

October 23, 2018

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

Too-specific and MaxLength attribute

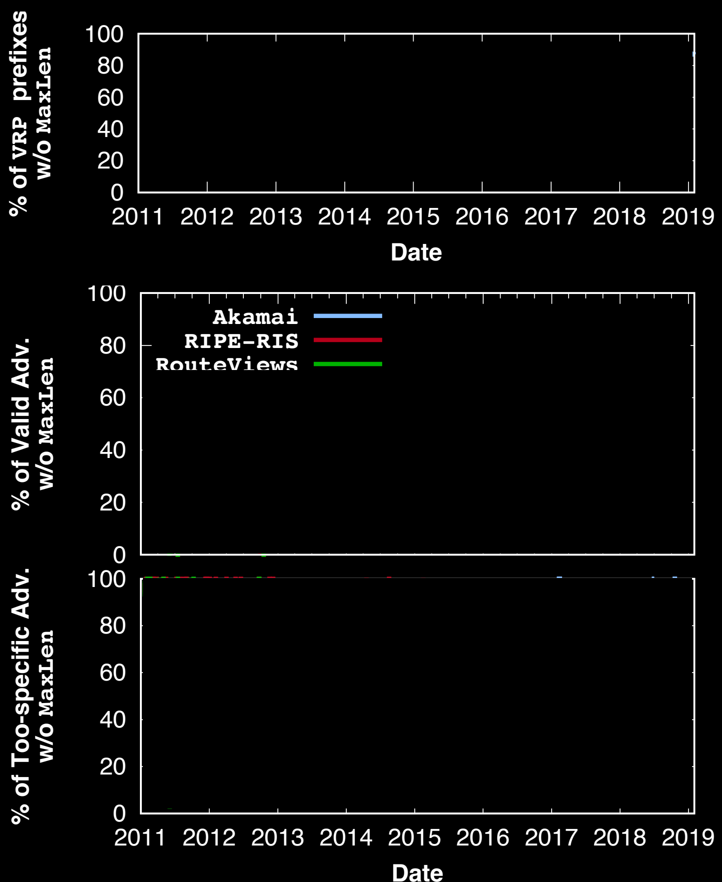

The use of MaxLength has been decreasing

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

Wrong ASN

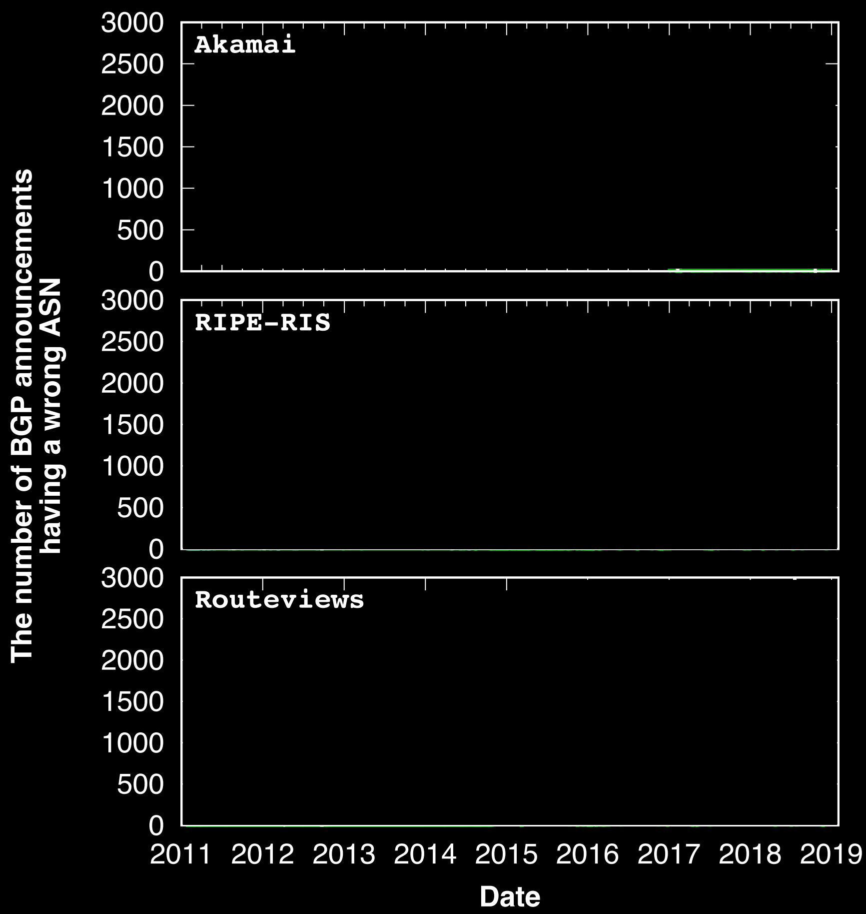

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource "scrubbing" of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don't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

Wrong ASN: Same ISP

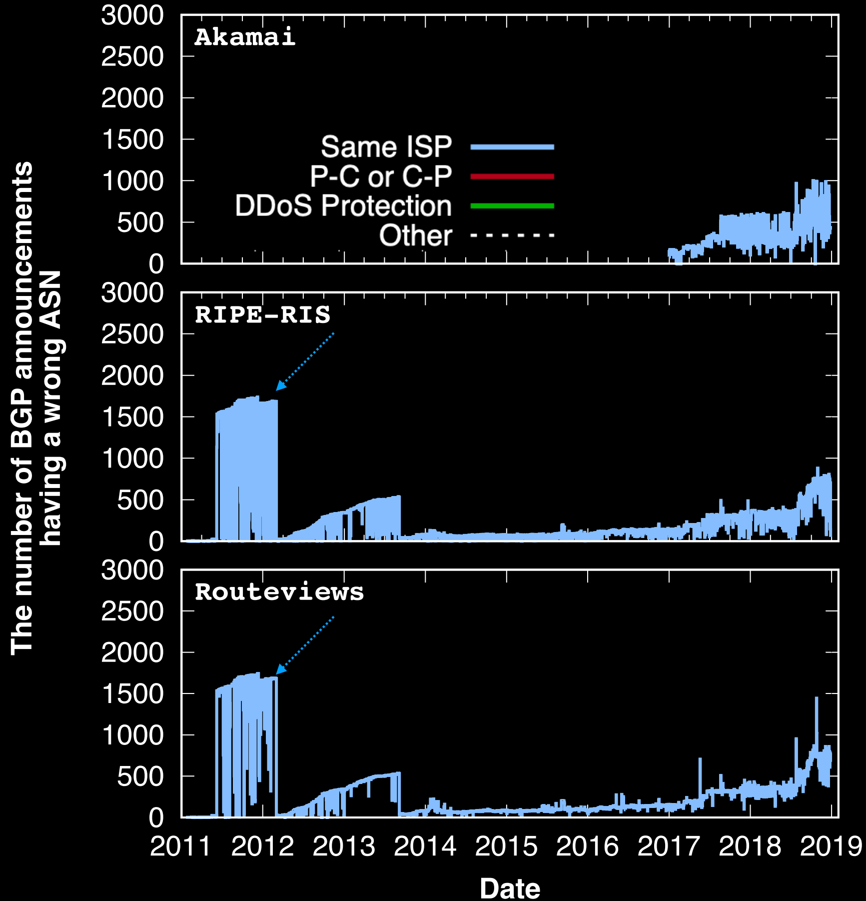

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource "scrubbing" of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don't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

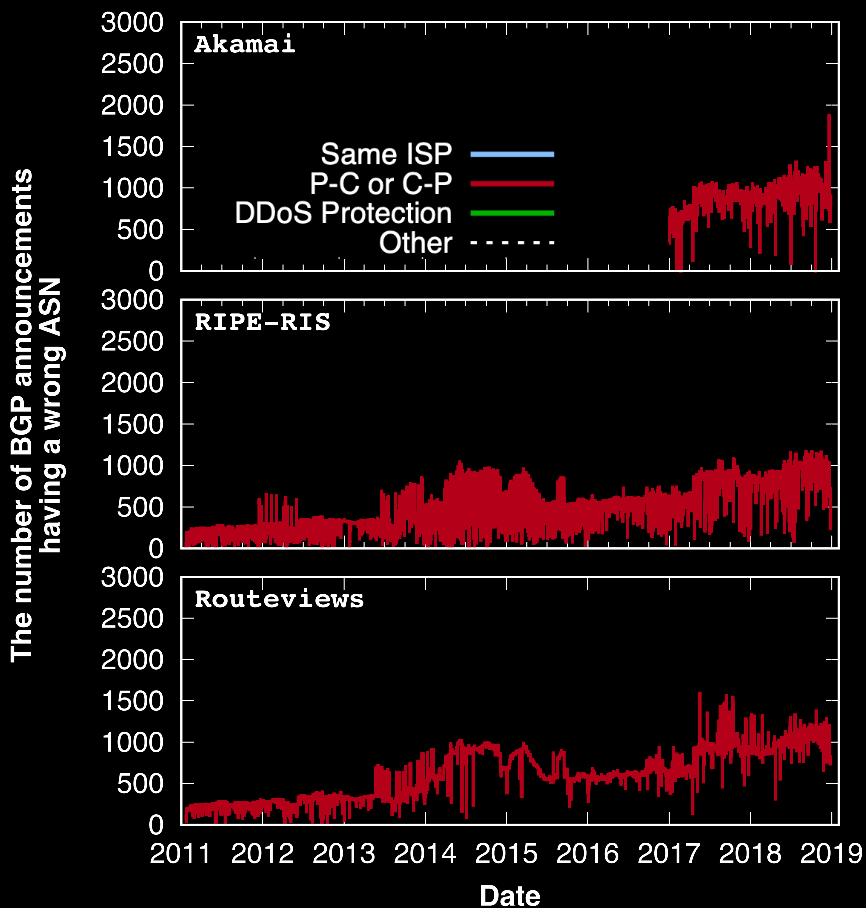

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource "scrubbing" of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don't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

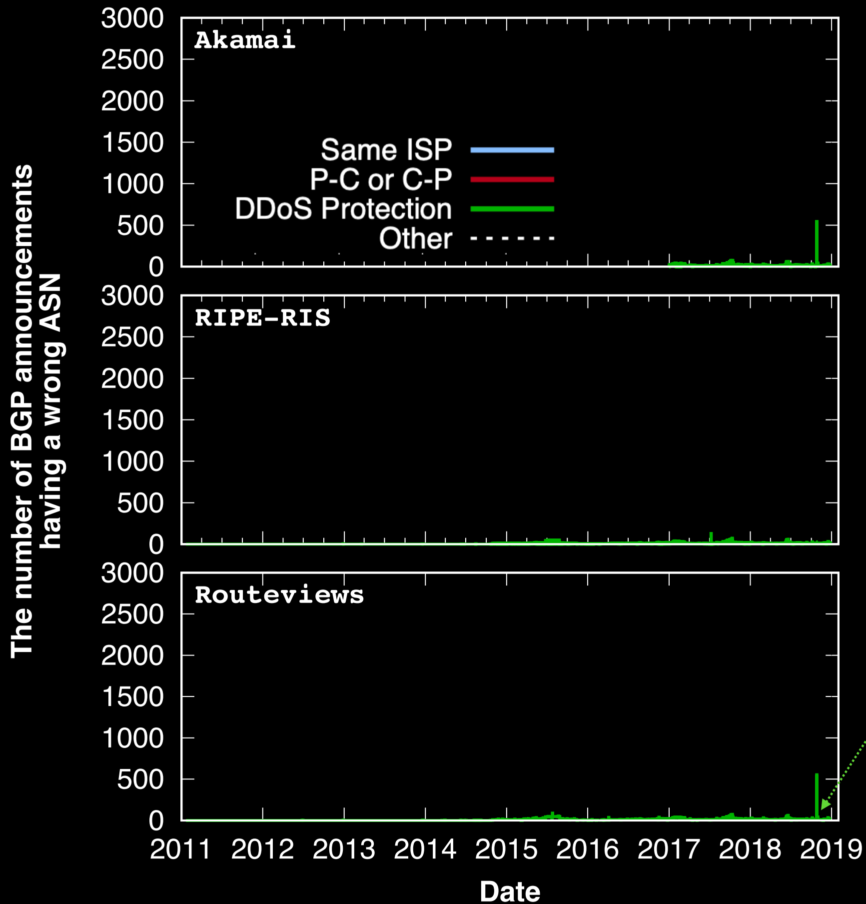

Same ISP	Two different ASNs are managed by the same operator
Provider—Customer Relationship	An AS can sub-allocate part of its IP prefixes to its customer
DDoS Protection	Origin ASes may outsource “scrubbing” of their traffic by using traffic diversion to a DDoS protection service (DPS)
Other	We don’t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

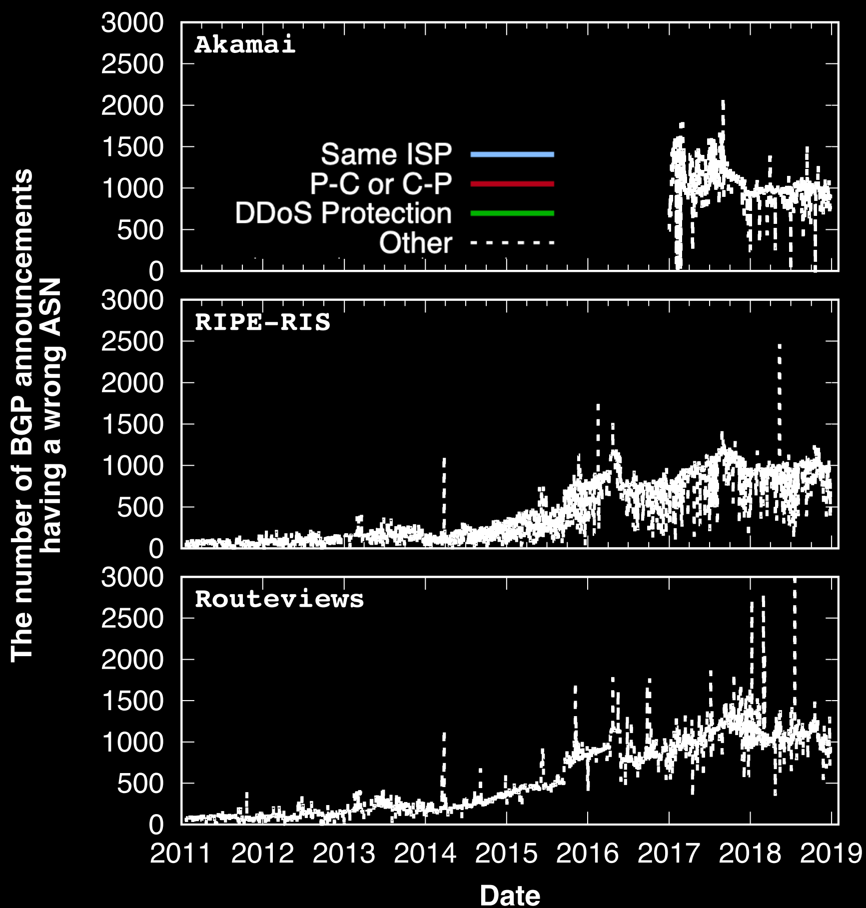

(1) AS 37468 (Angola Cables) announced more than 2,500 IP prefixes owned by 82 ASes on May 11, 2018 and 15,000 IP prefixes owned by 1,554 ASes on July 19, 2018

(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?

Case-study: BGPStream

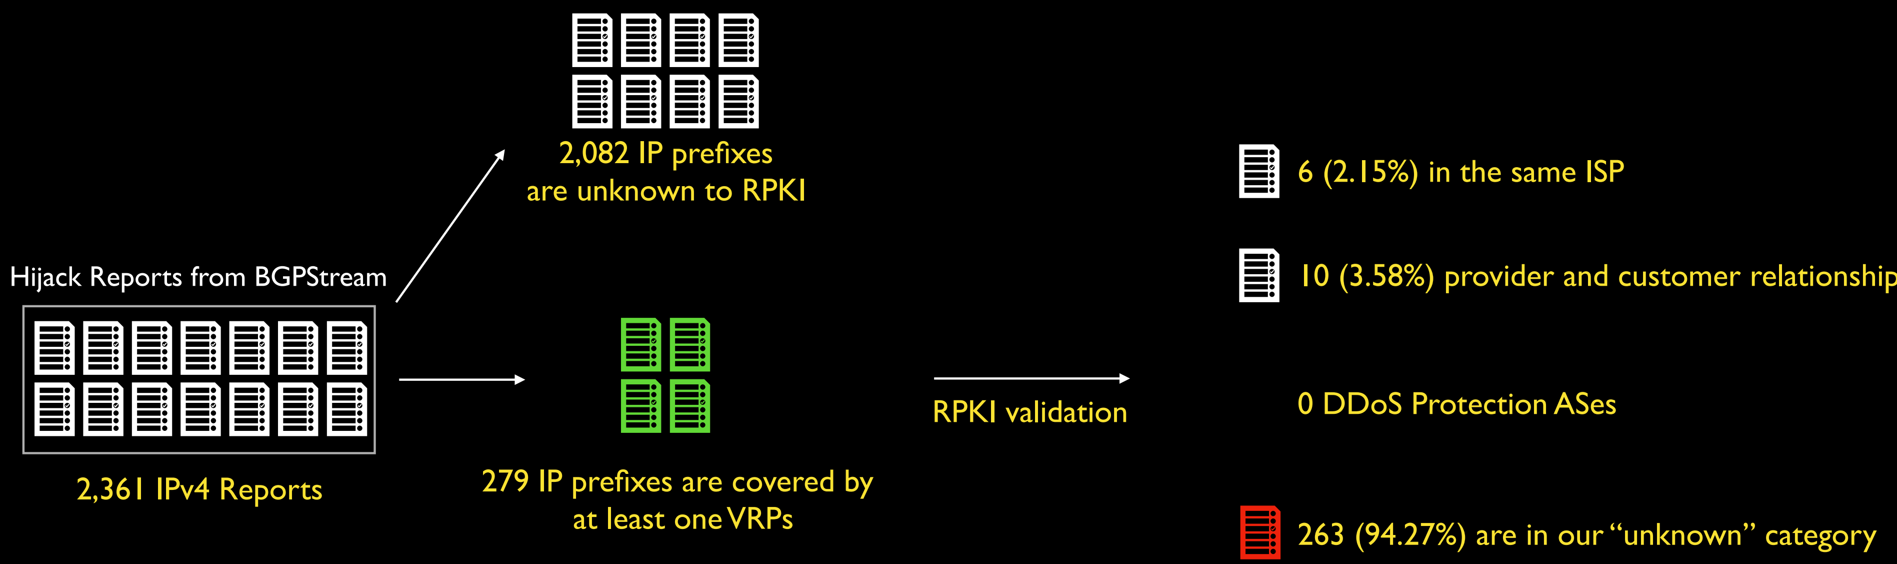

Conclusion and Discussion

- 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
- Open Question: how can we identify hijacking attempt with high confidence?

Deployment: VRPs

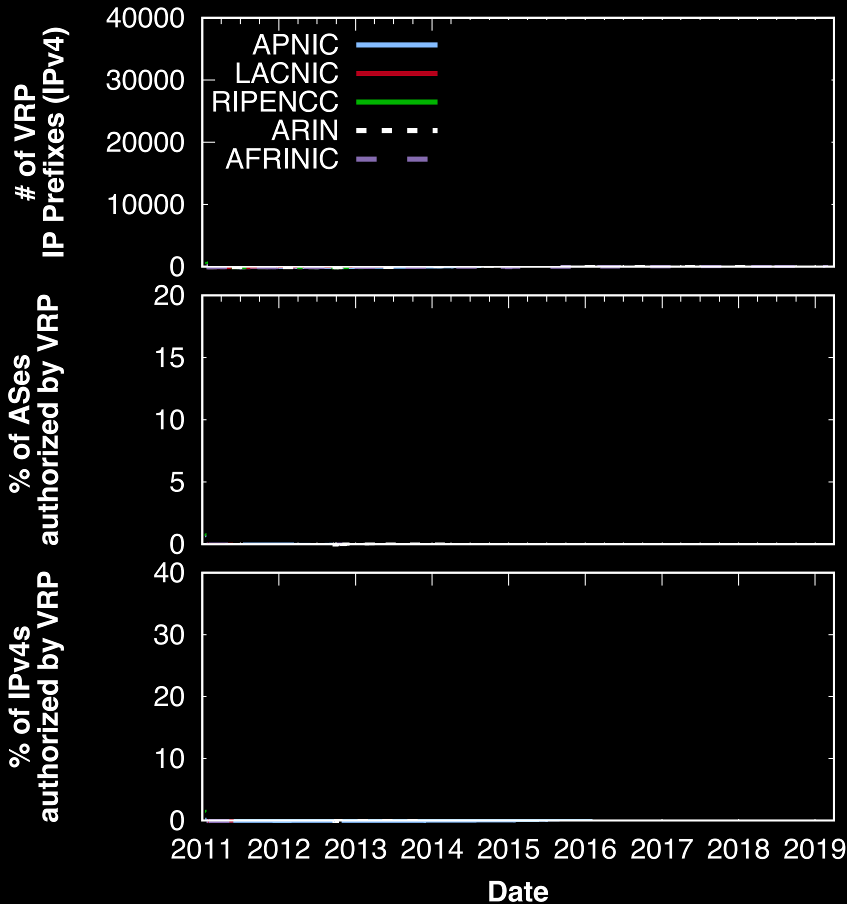

A general increasing trend in adoption of RPKI

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

ROAs with MaxLength attributes were disabled and those VRPs were separately introduced without MaxLength (June 6th), but rolled back on June 19th, 2017

RPKI validation over BGP announcements

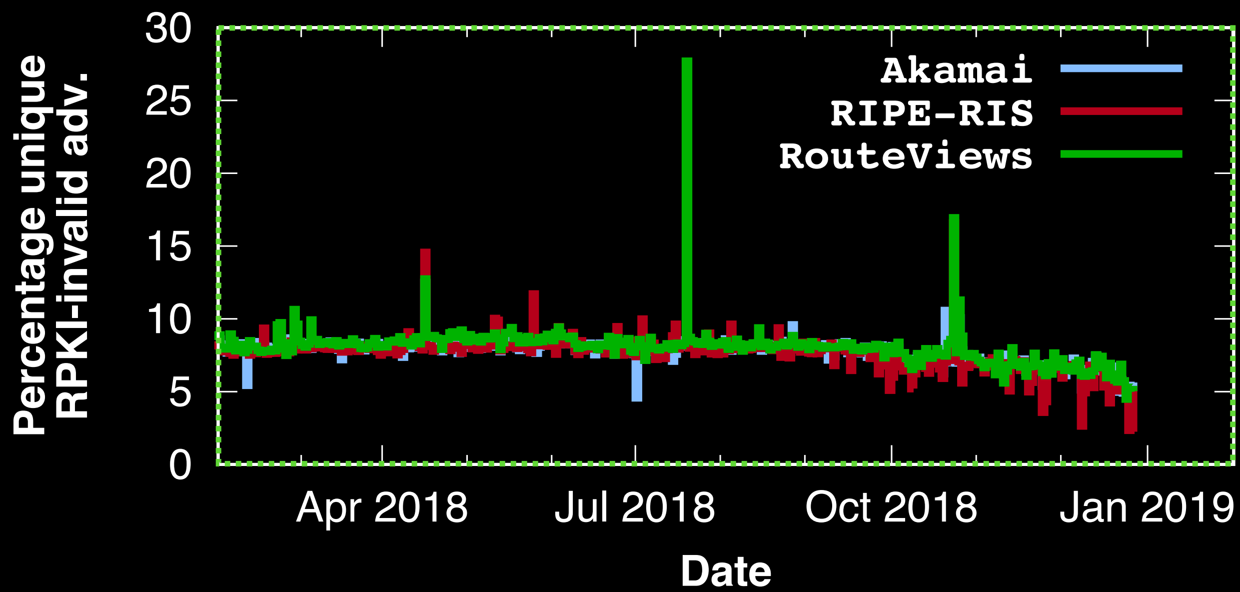

Quality of BGP announcements

Overall percentage of invalid prefixes has been decreasing rapidly

Why Covering is not valid?

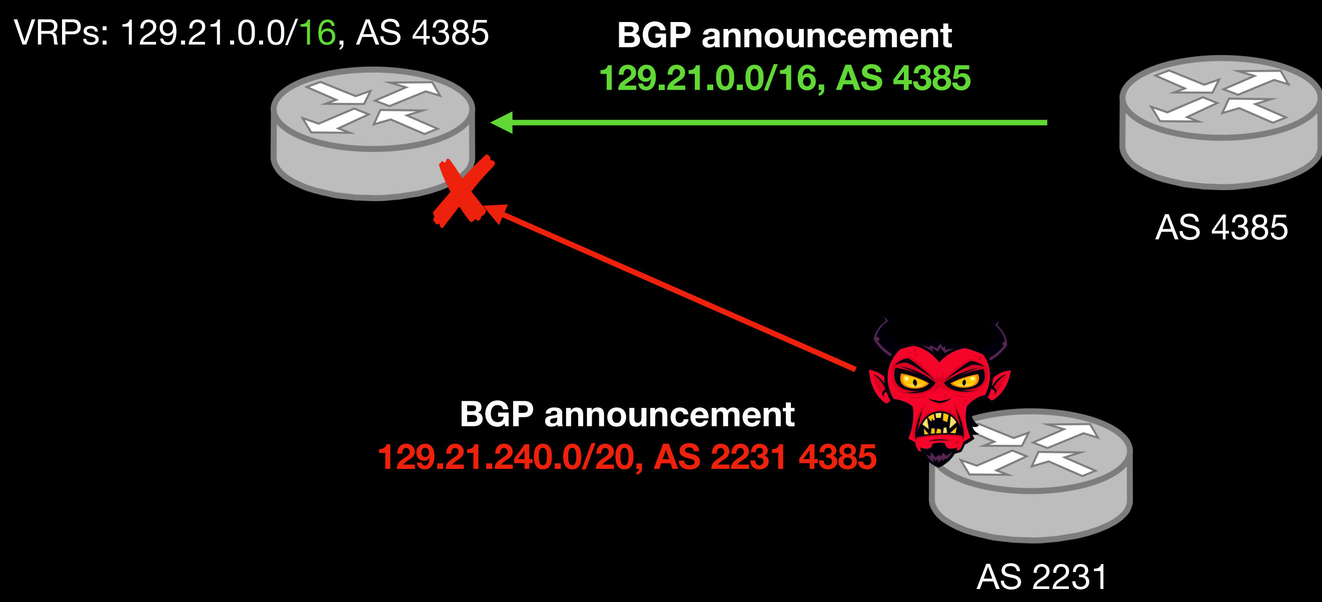

Traffic from the “other” category

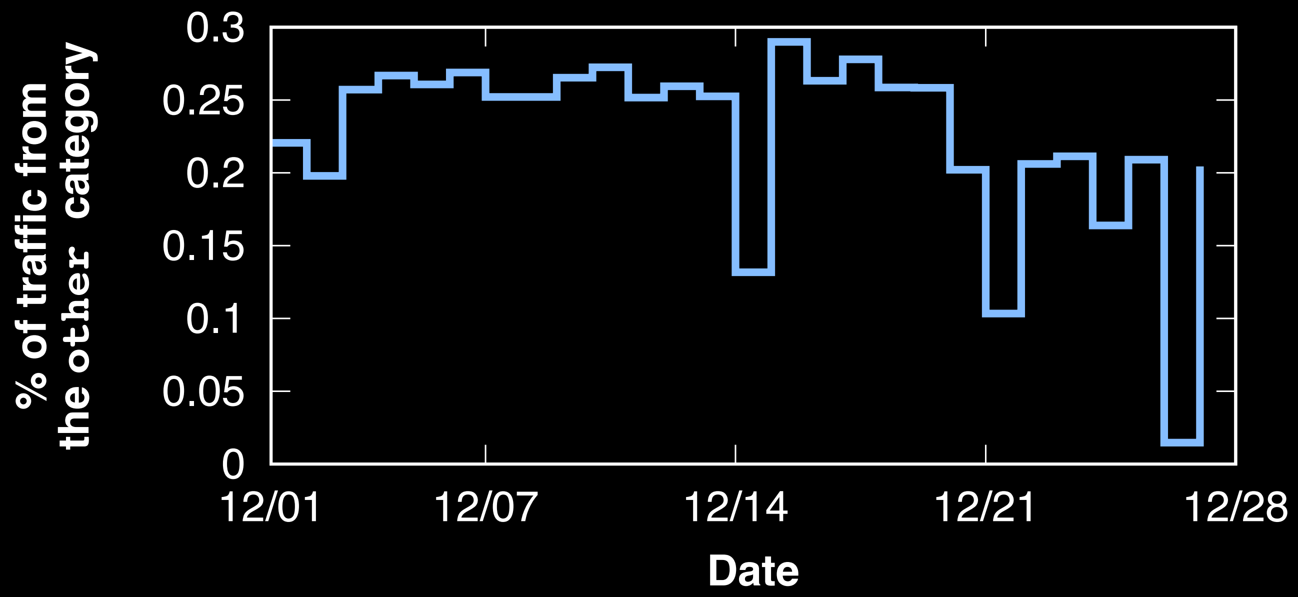

Amount of Traffic

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

How a Router Uses RPKI

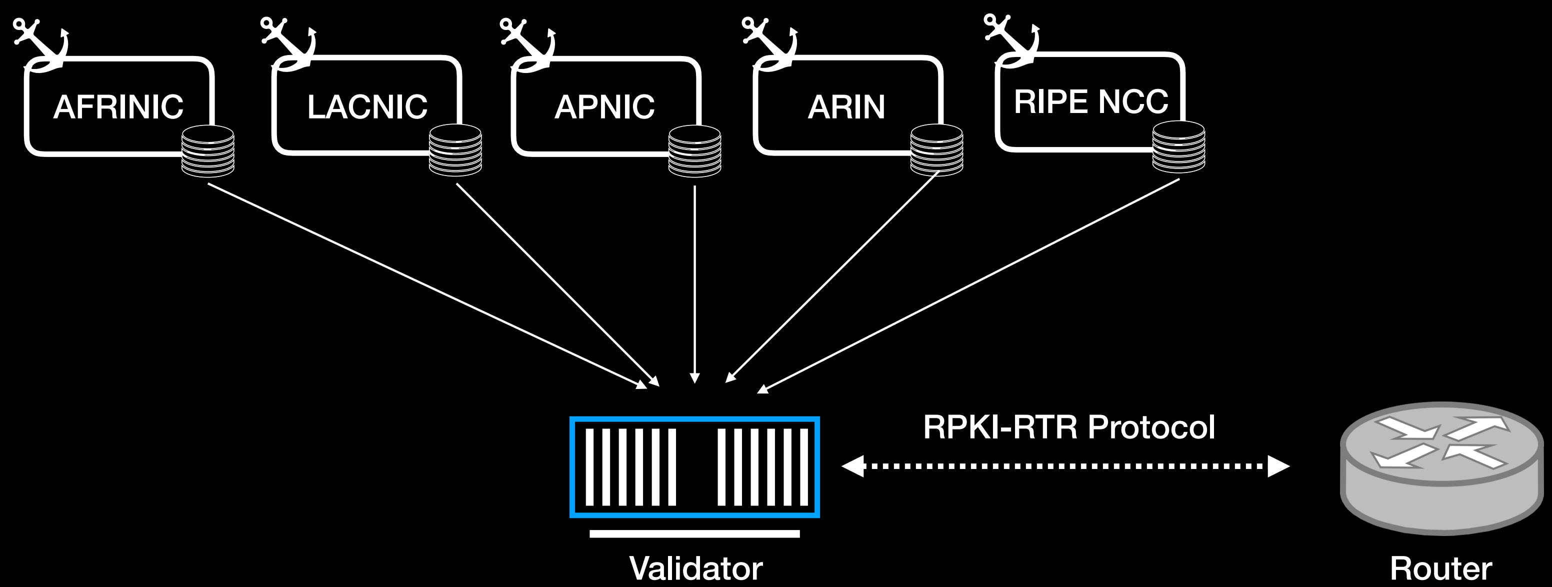

Routinator (NLNetLabs)
OctoRPKI (Cloudflare)
RPKI Validator (RIPE NCC)
...

ROV (Route Origin Validation)
A route prefix is "covered"

An IP prefix is **covered**	The IP prefix address and VRP IP prefix address are identical for all bits specified by the VRP IP prefix length

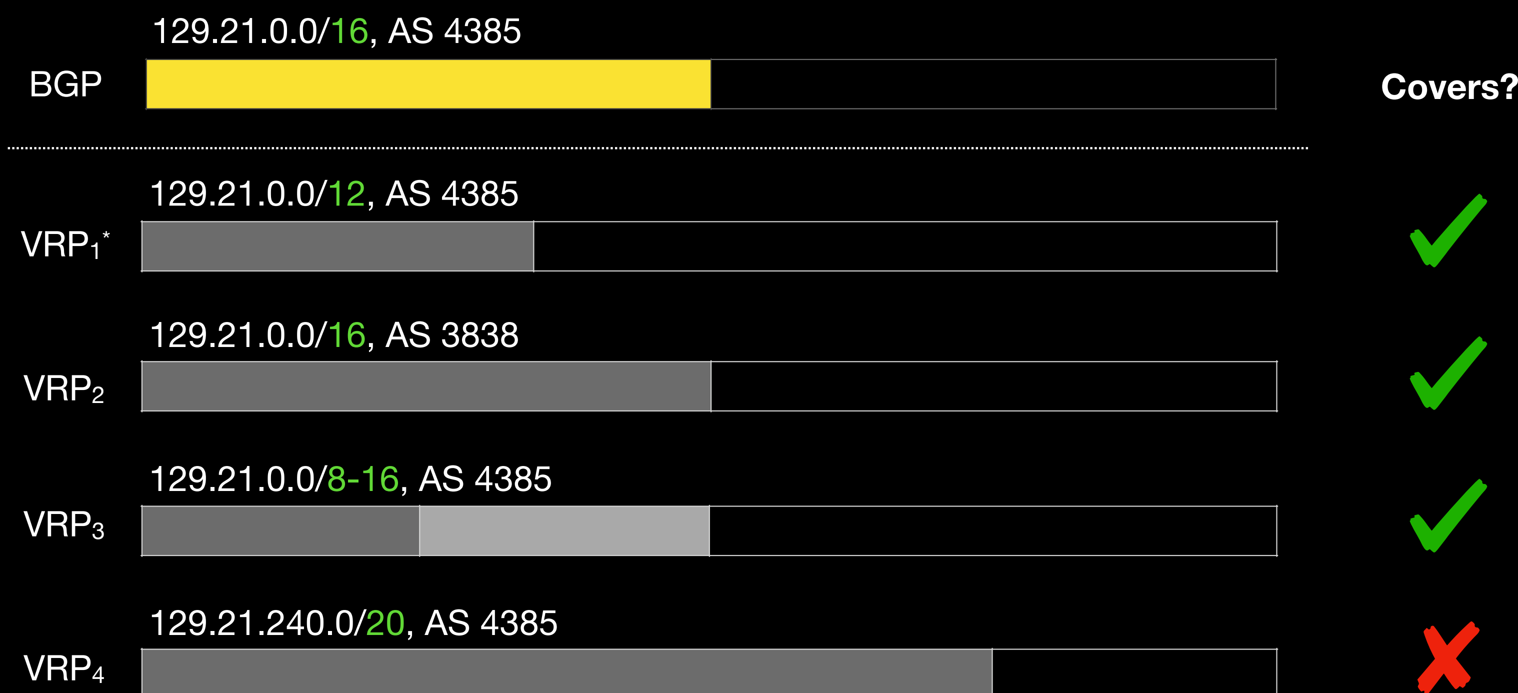

*Validated ROA Payloads

ROV (Route Origin Validation)
A route prefix is “matched”

An IP prefix is **matched**	1. VRP IP prefix covers the announced IP prefix 2. VRP’s ASN == Announced ASN 3. Announced IP prefix length <= VRP’s IP prefix length (including MaxLength)

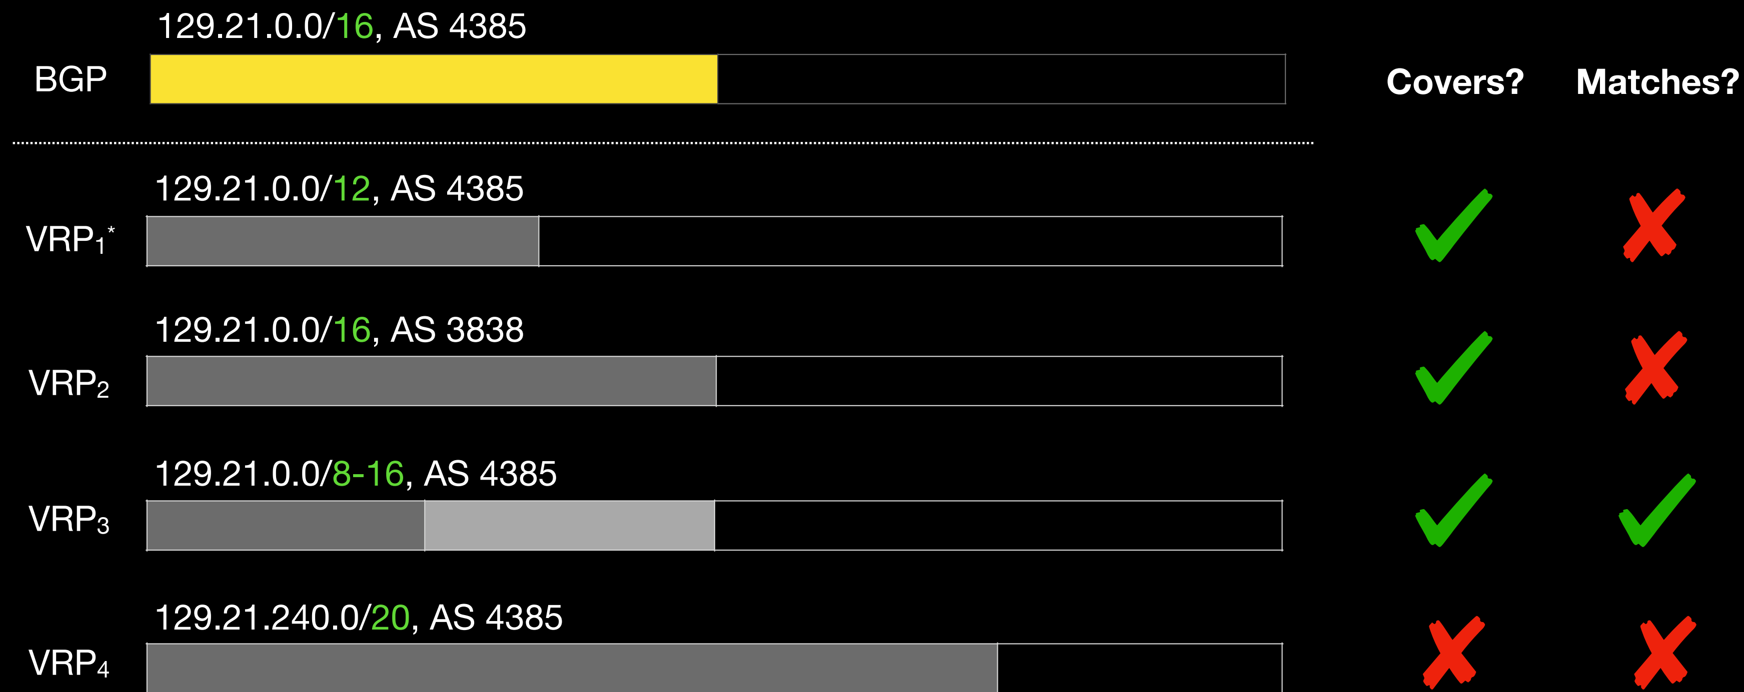

*Validated ROA Payloads

ROV (Route Origin Validation) Validation

?	Unknown	No VRP Covers the Route Prefix
✓	Valid	At least one VRP Matches the Route Prefix.
✗	Invalid	At least one VRP Covers the Route Prefix, but no VRP

BGP: 129.21.0.0/16, AS 4385

	Prefix	Covers?	Matches?	Status
VRP$_1$	129.21.0.0/12, AS 4385	✓	✗	✗
VRP$_2$	129.21.0.0/16, AS 3838	✓	✗	✗
VRP$_3$	129.21.0.0/8-16, AS 4385	✓	✓	✓
VRP$_3$	129.21.240.0/20, AS 4385	✗	✗	?

% of VRP-covered announcements: IPv4 vs. IPv6