Demystifying RPKI-Invalid Prefixes: Hidden Causes and Security Risks

Weitong Li[1], Tao Wan[2], and Tijay Chung[1]

[1]Virginia Tech, [2]CableLabs

Border Gateway Protocol (BGP)

- Each network resource owner (e.g., VT) announces its IP prefixes to the rest of routers, so that they can learn the path towards VT.

- However, BGP does not have builtin security mechanism

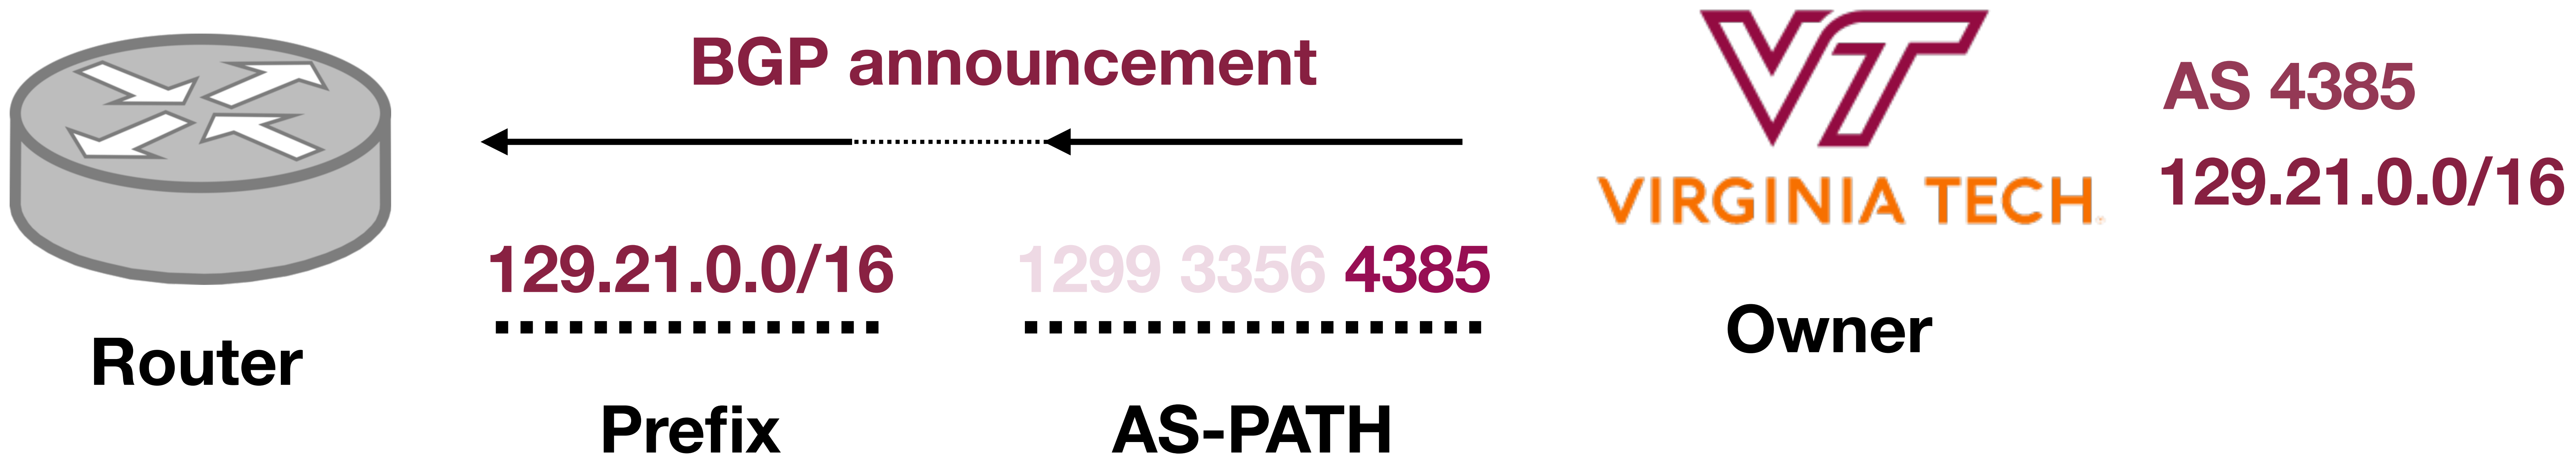

BGP announcement

Router

129.21.0.0/16

Prefix

1299 3356 4385

AS-PATH

VIRGINIA TECH

Owner

AS 4385

129.21.0.0/16

Border Gateway Protocol (BGP)

- An adversary can announce prefixes that not belong to it, thus hijack the resource

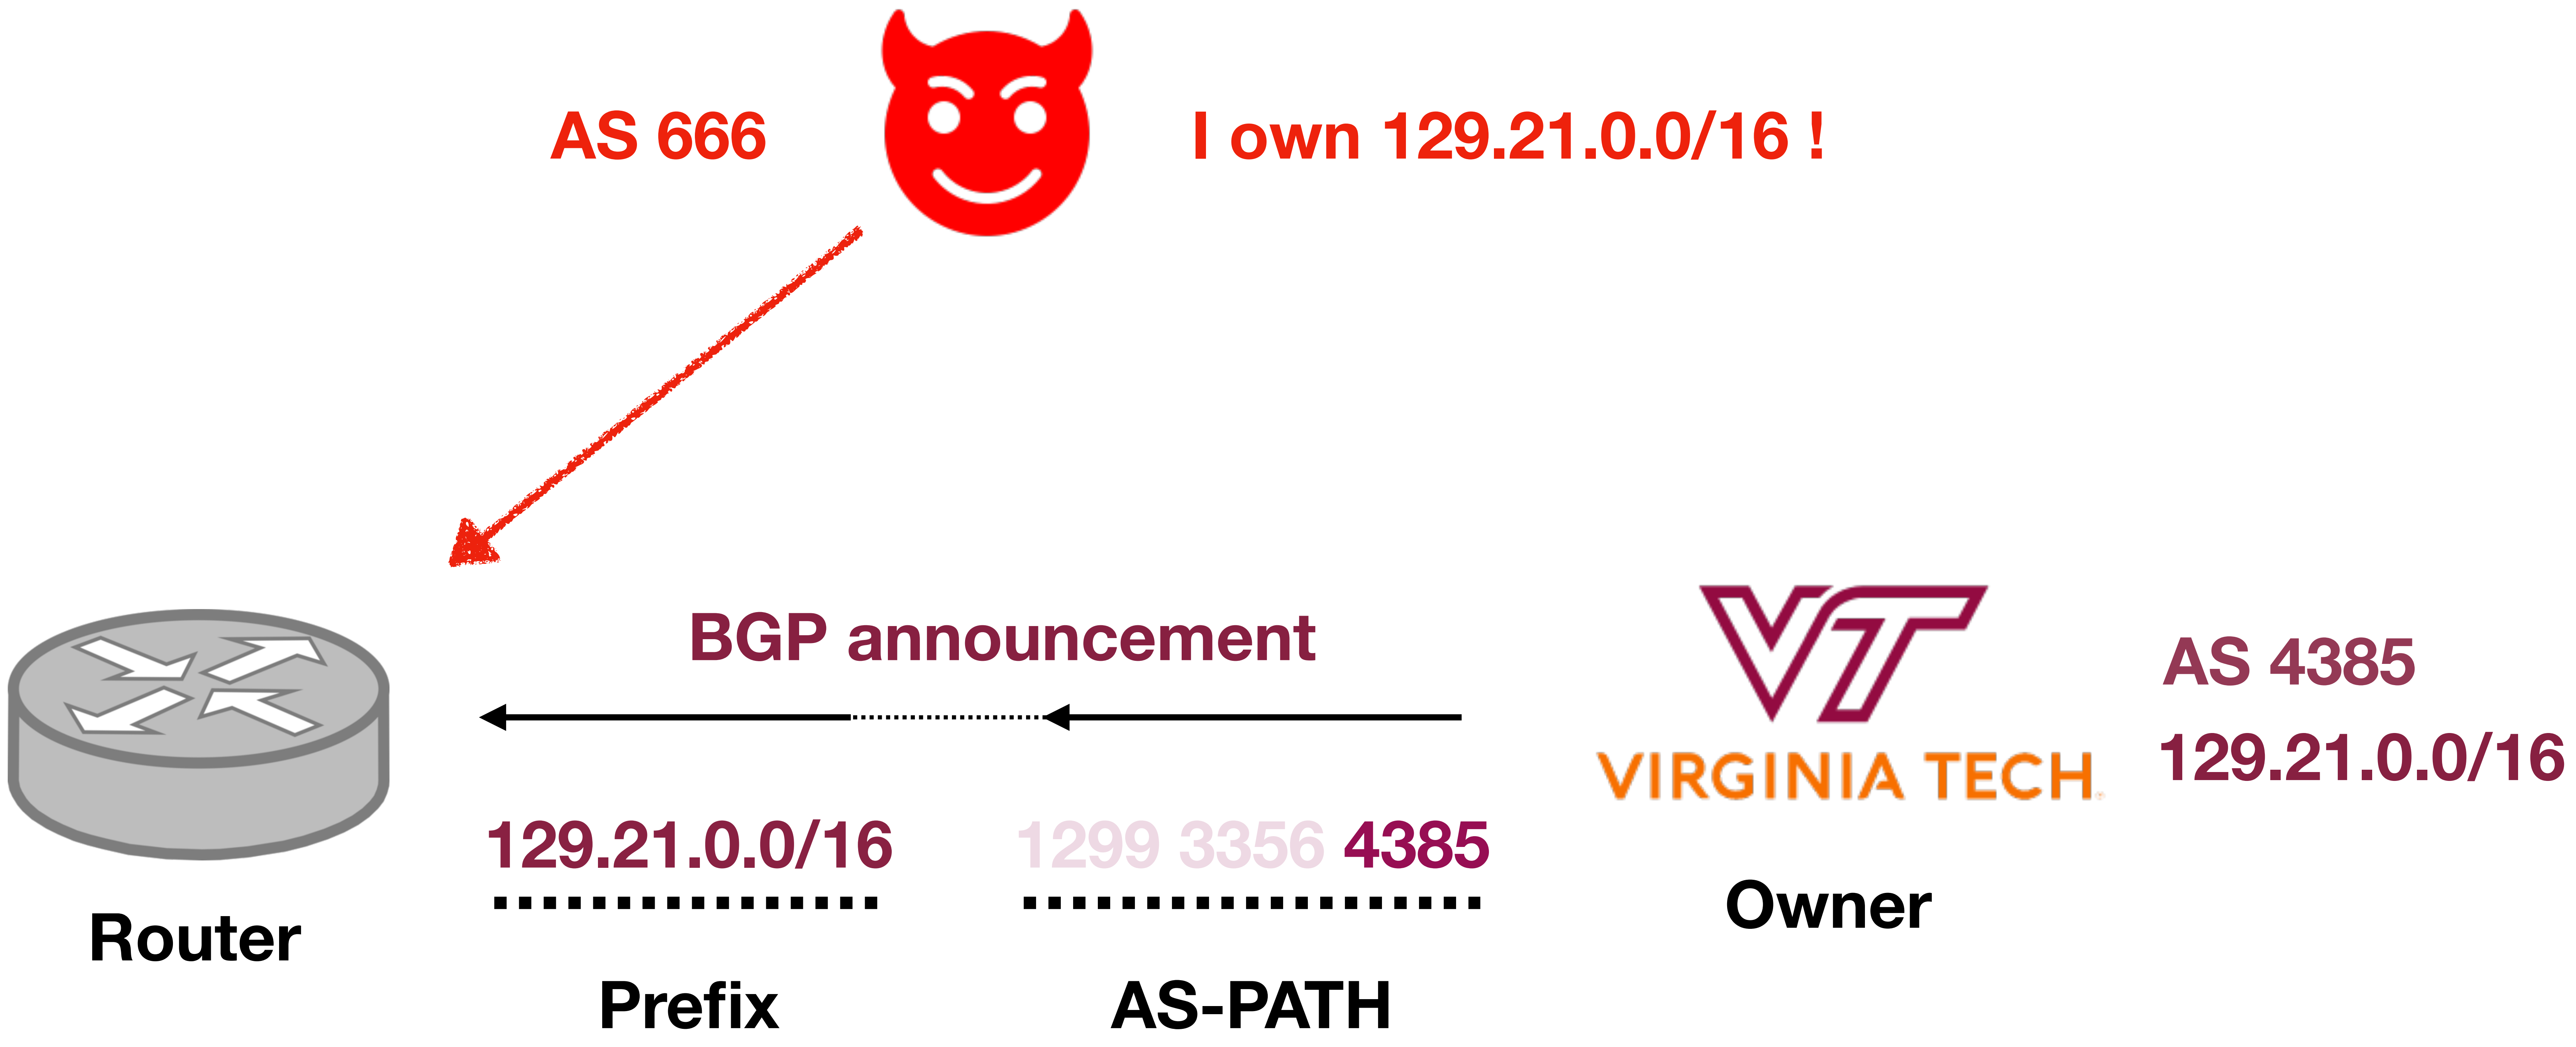

AS 666

I own 129.21.0.0/16 !

BGP announcement

AS 4385

129.21.0.0/16

129.21.0.0/16

1299 3356 4385

Router

Prefix

AS-PATH

Owner

BGP Hijacks

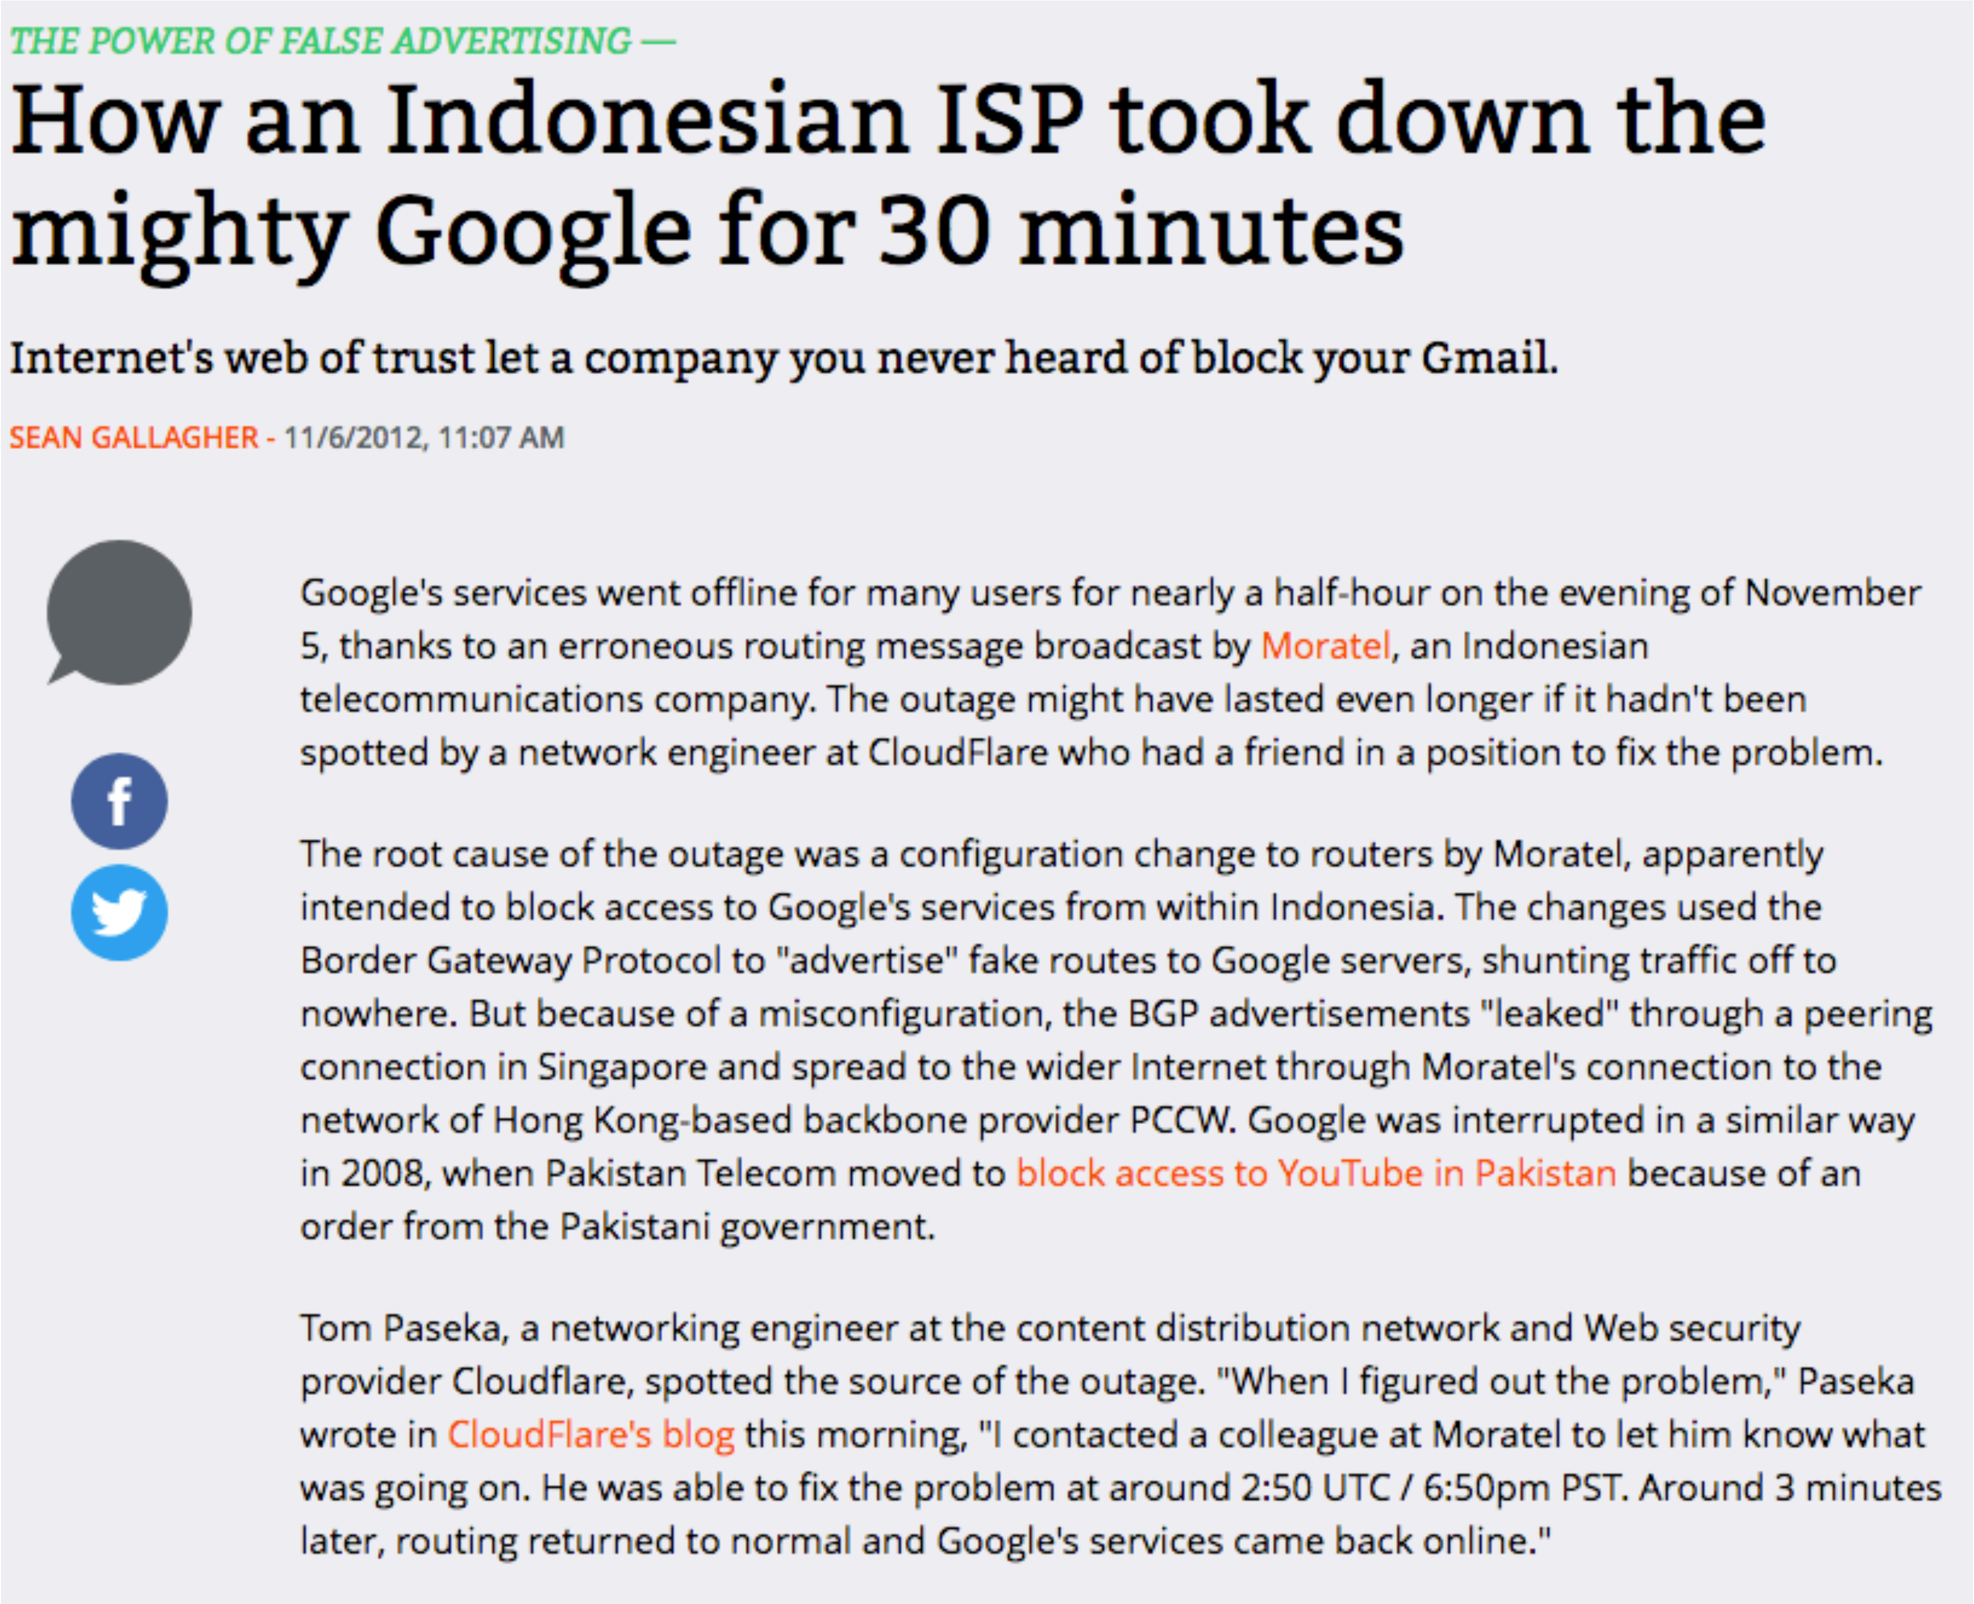

THE POWER OF FALSE ADVERTISING —

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."

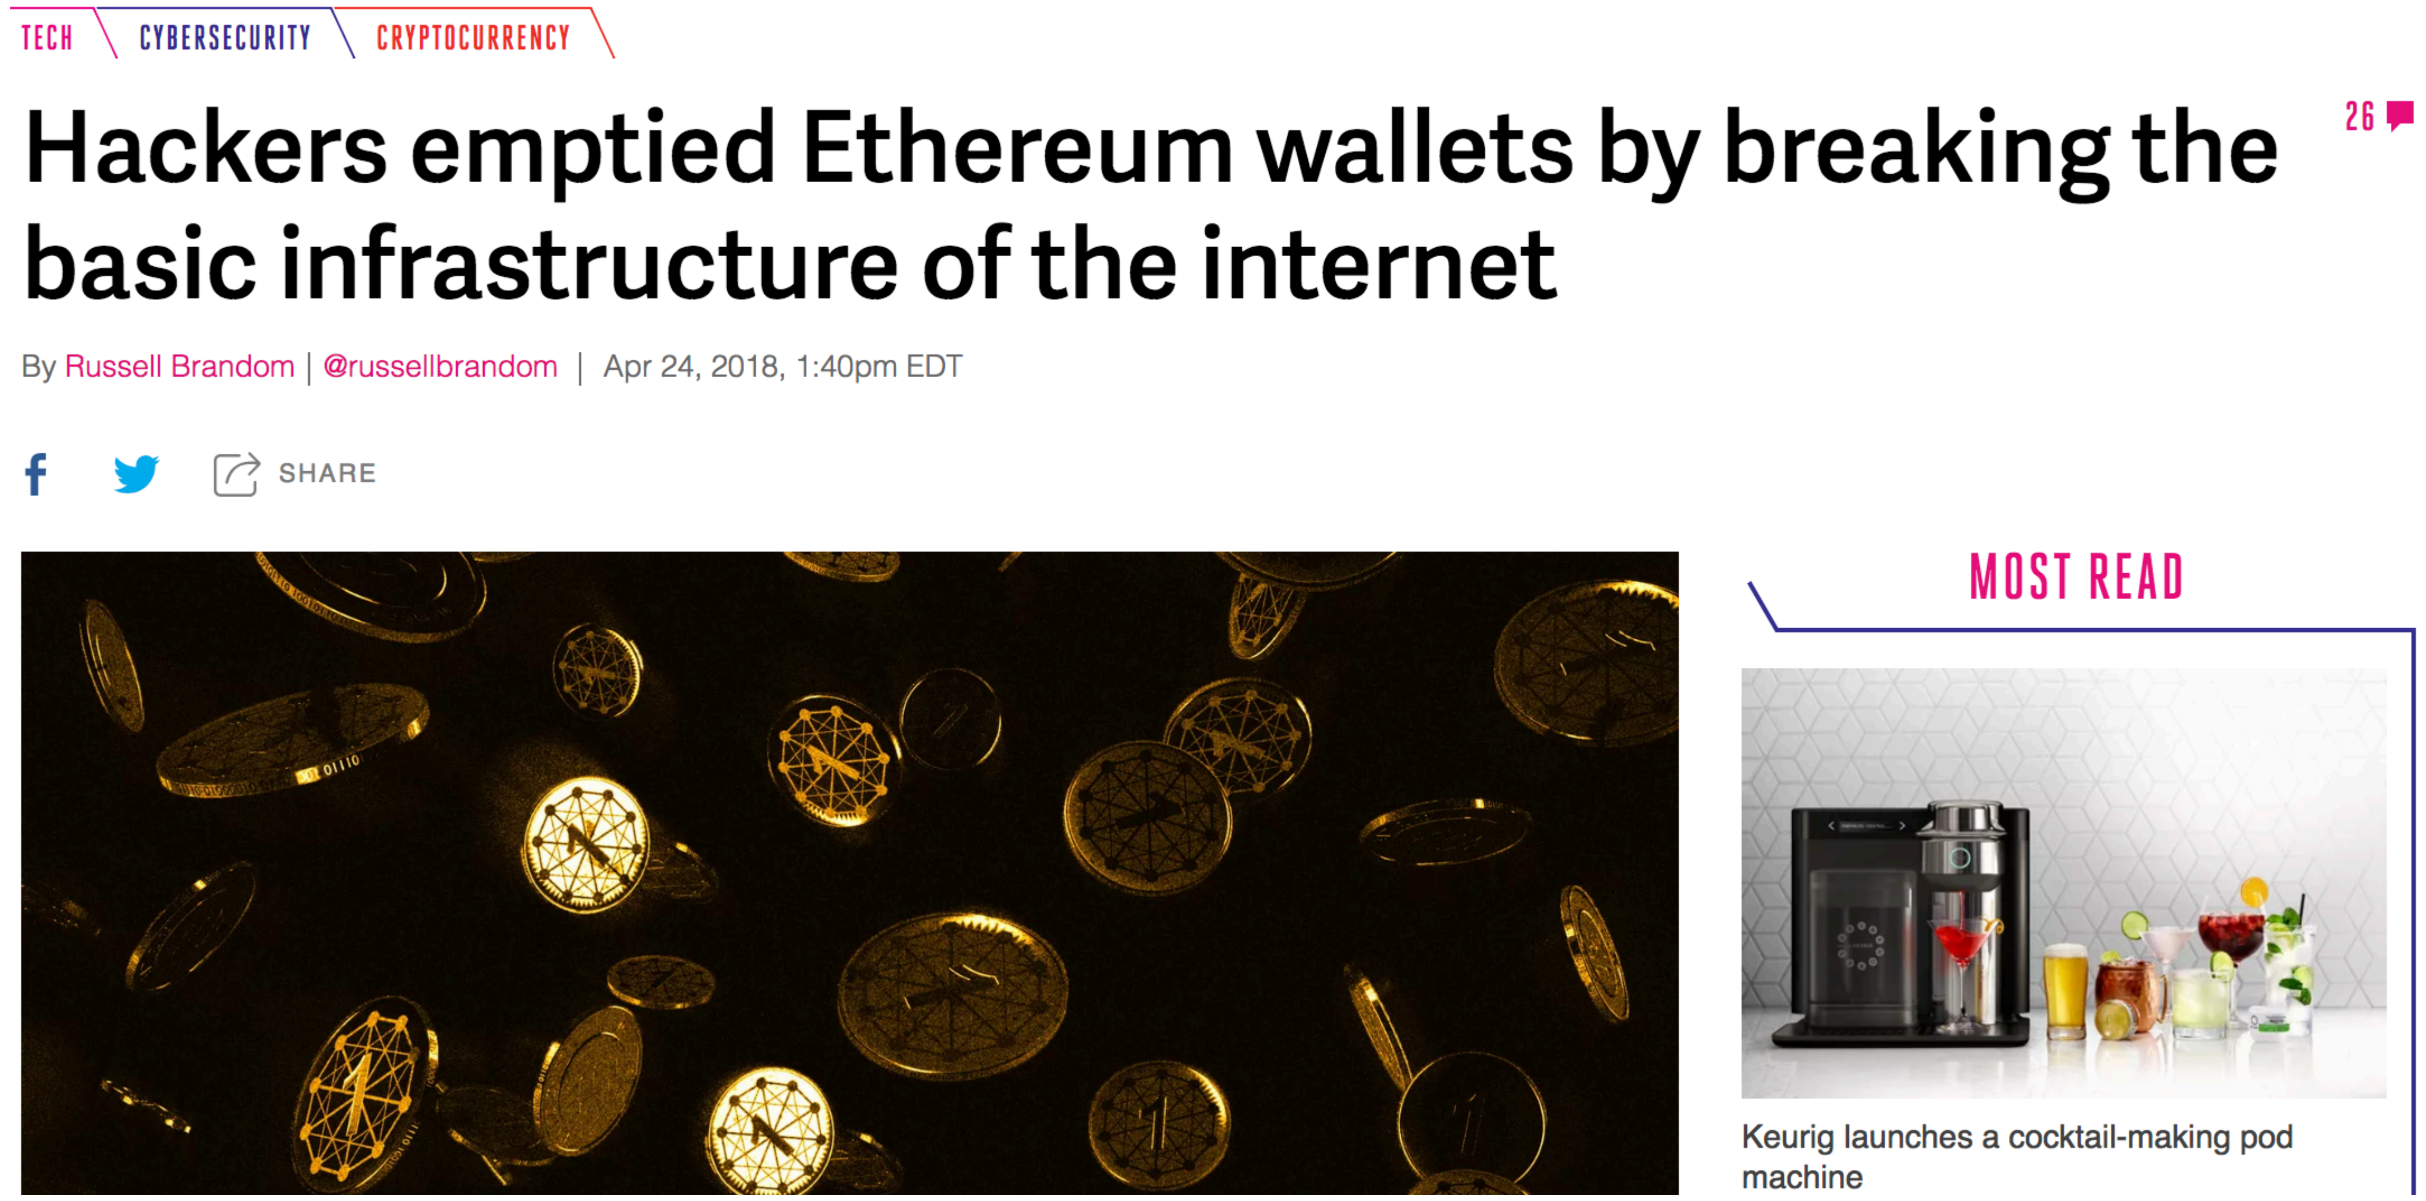

Hackers emptied Ethereum wallets by breaking the basic infrastructure of the internet

By Russell Brandom | @russellbrandom | Apr 24, 2018, 1:40pm EDT

Keurig launches a cocktail-making pod machine

The Solution: RPKI

- Resource Public Key Infrastructure (RPKI) is proposed to secure Internet's routing and prevent hijacks

- The deployment of RPKI starts from 2008 and received more

The Solution: RPKI

- Resource Public Key Infrastructure (RPKI) is proposed to secure Internet's routing and prevent hijacks
- The deployment of RPKI starts from 2008
- RPKI contains two parts:
  - Route Origin Authorization (ROA)
  - Route Origin Validation (ROV)

RPKI Structure: ROA

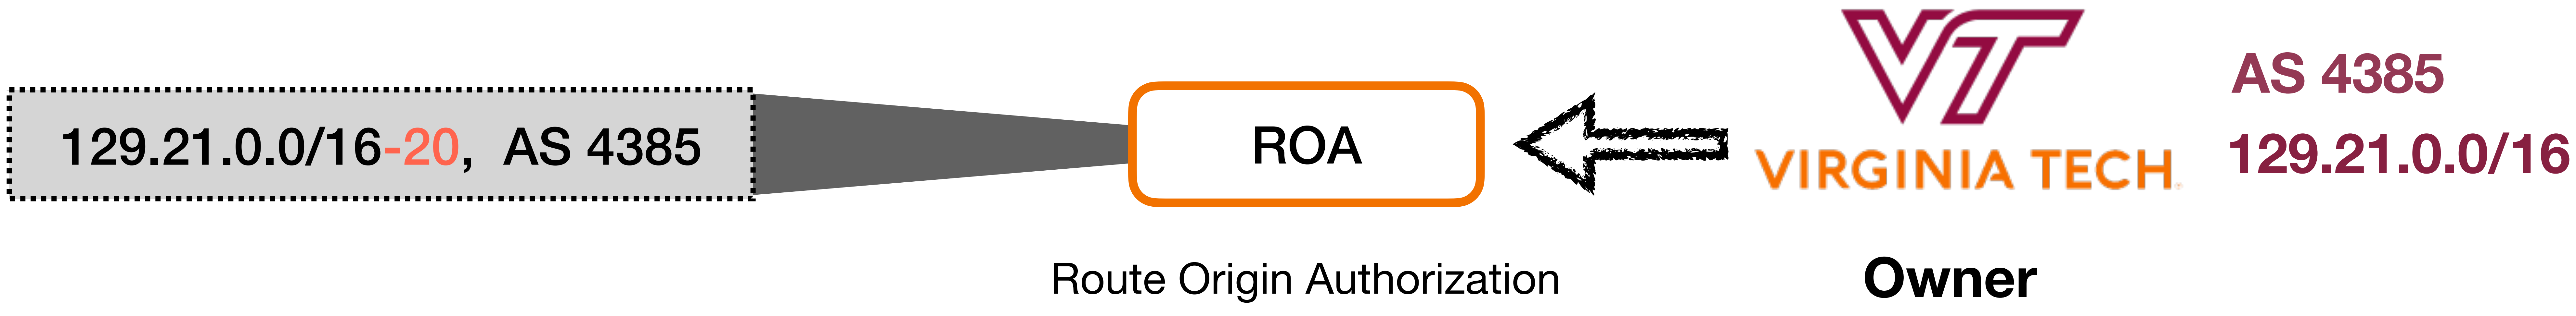

129.21.0.0/16-20, AS 4385

ROA

Route Origin Authorization

Owner

AS 4385

129.21.0.0/16

RPKI Structure: ROA

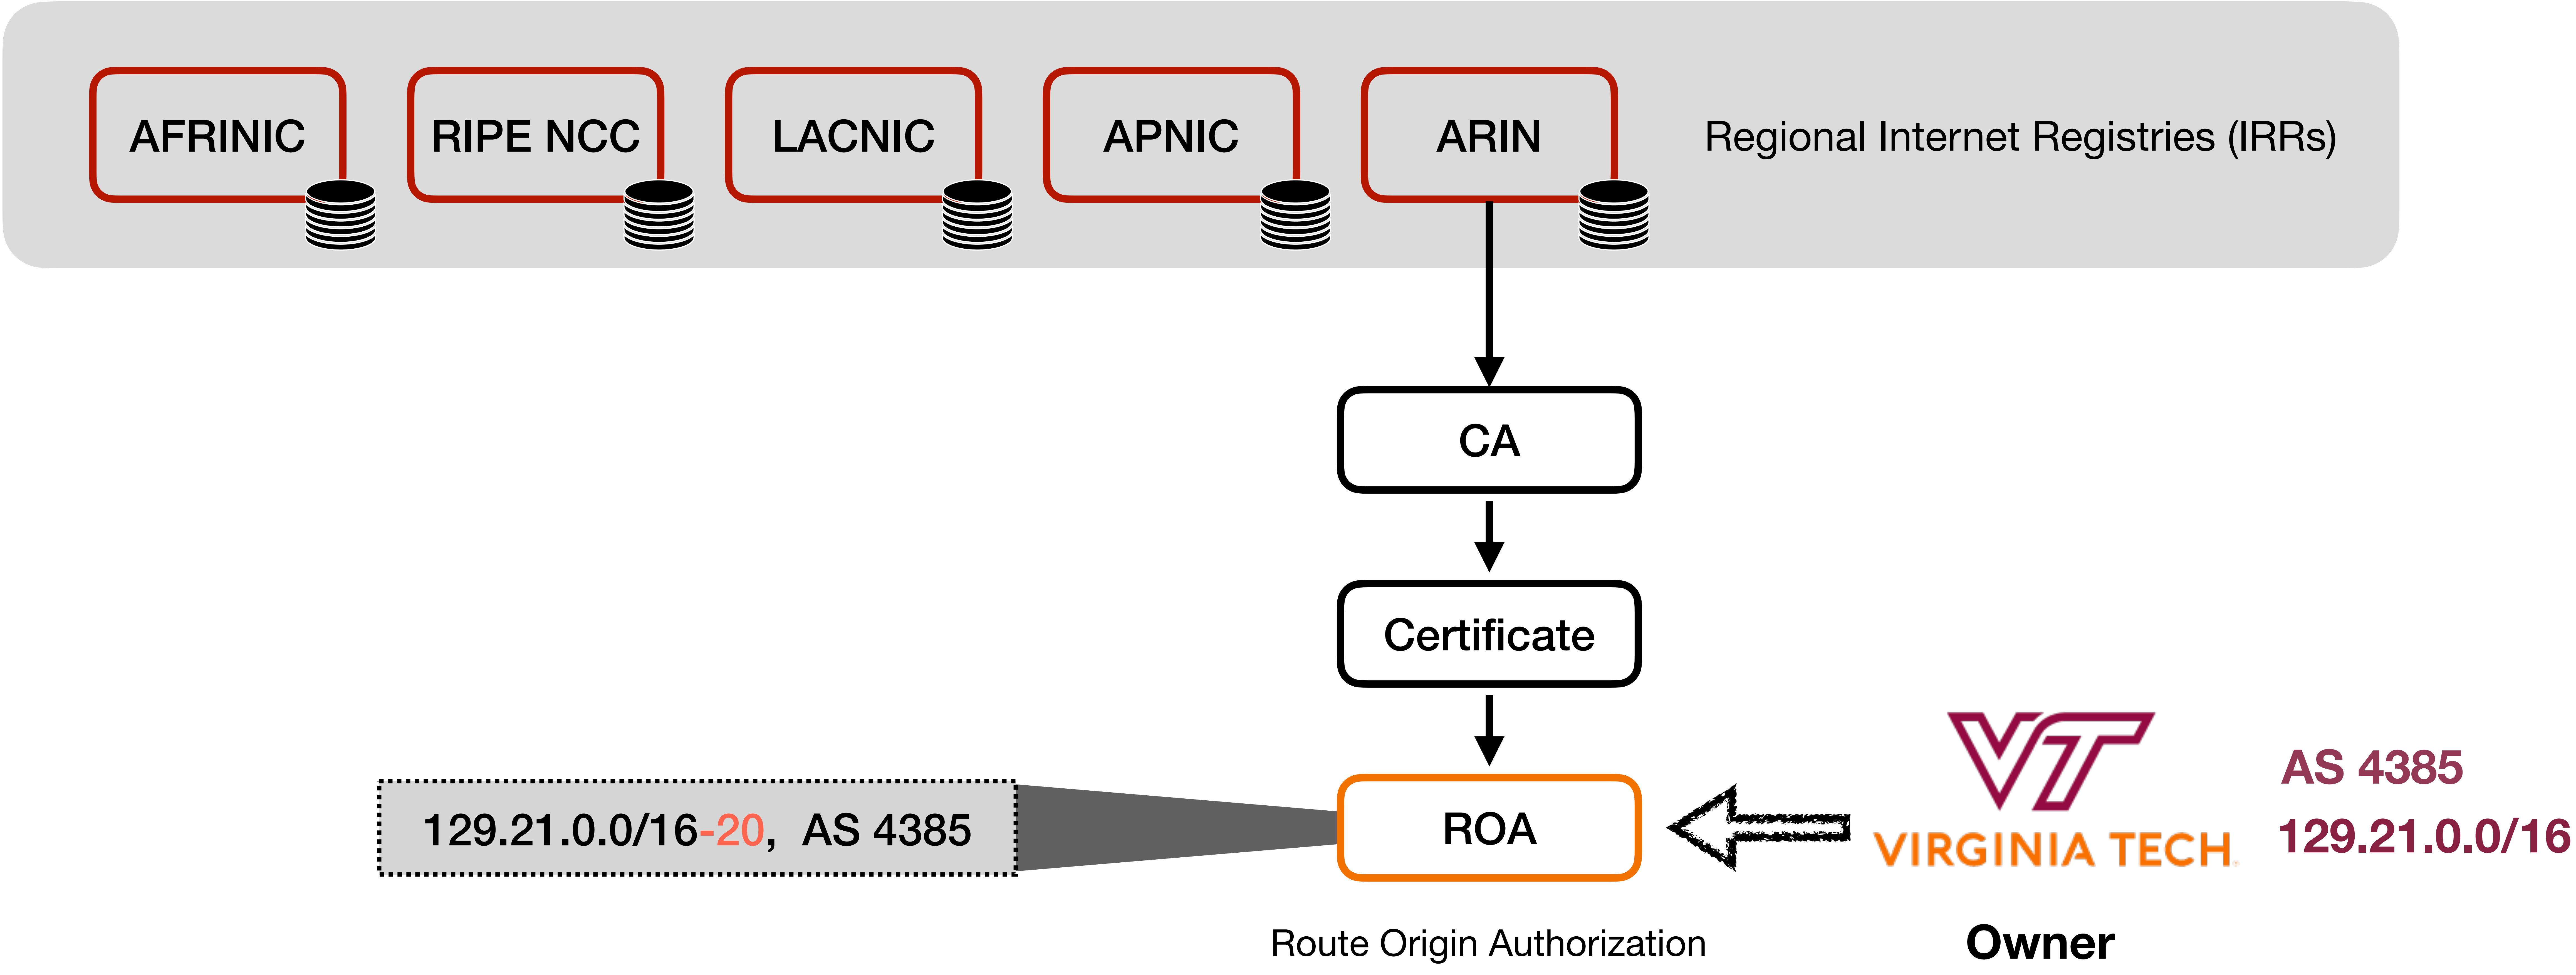
AFRINIC
RIPE NCC
LACNIC
APNIC
ARIN
Regional Internet Registries (IRRs)
CA
Certificate
129.21.0.0/16-20, AS 4385
ROA
Route Origin Authorization
VIRGINIA TECH
AS 4385
129.21.0.0/16
Owner

RPKI Structure: Route Origin Validation

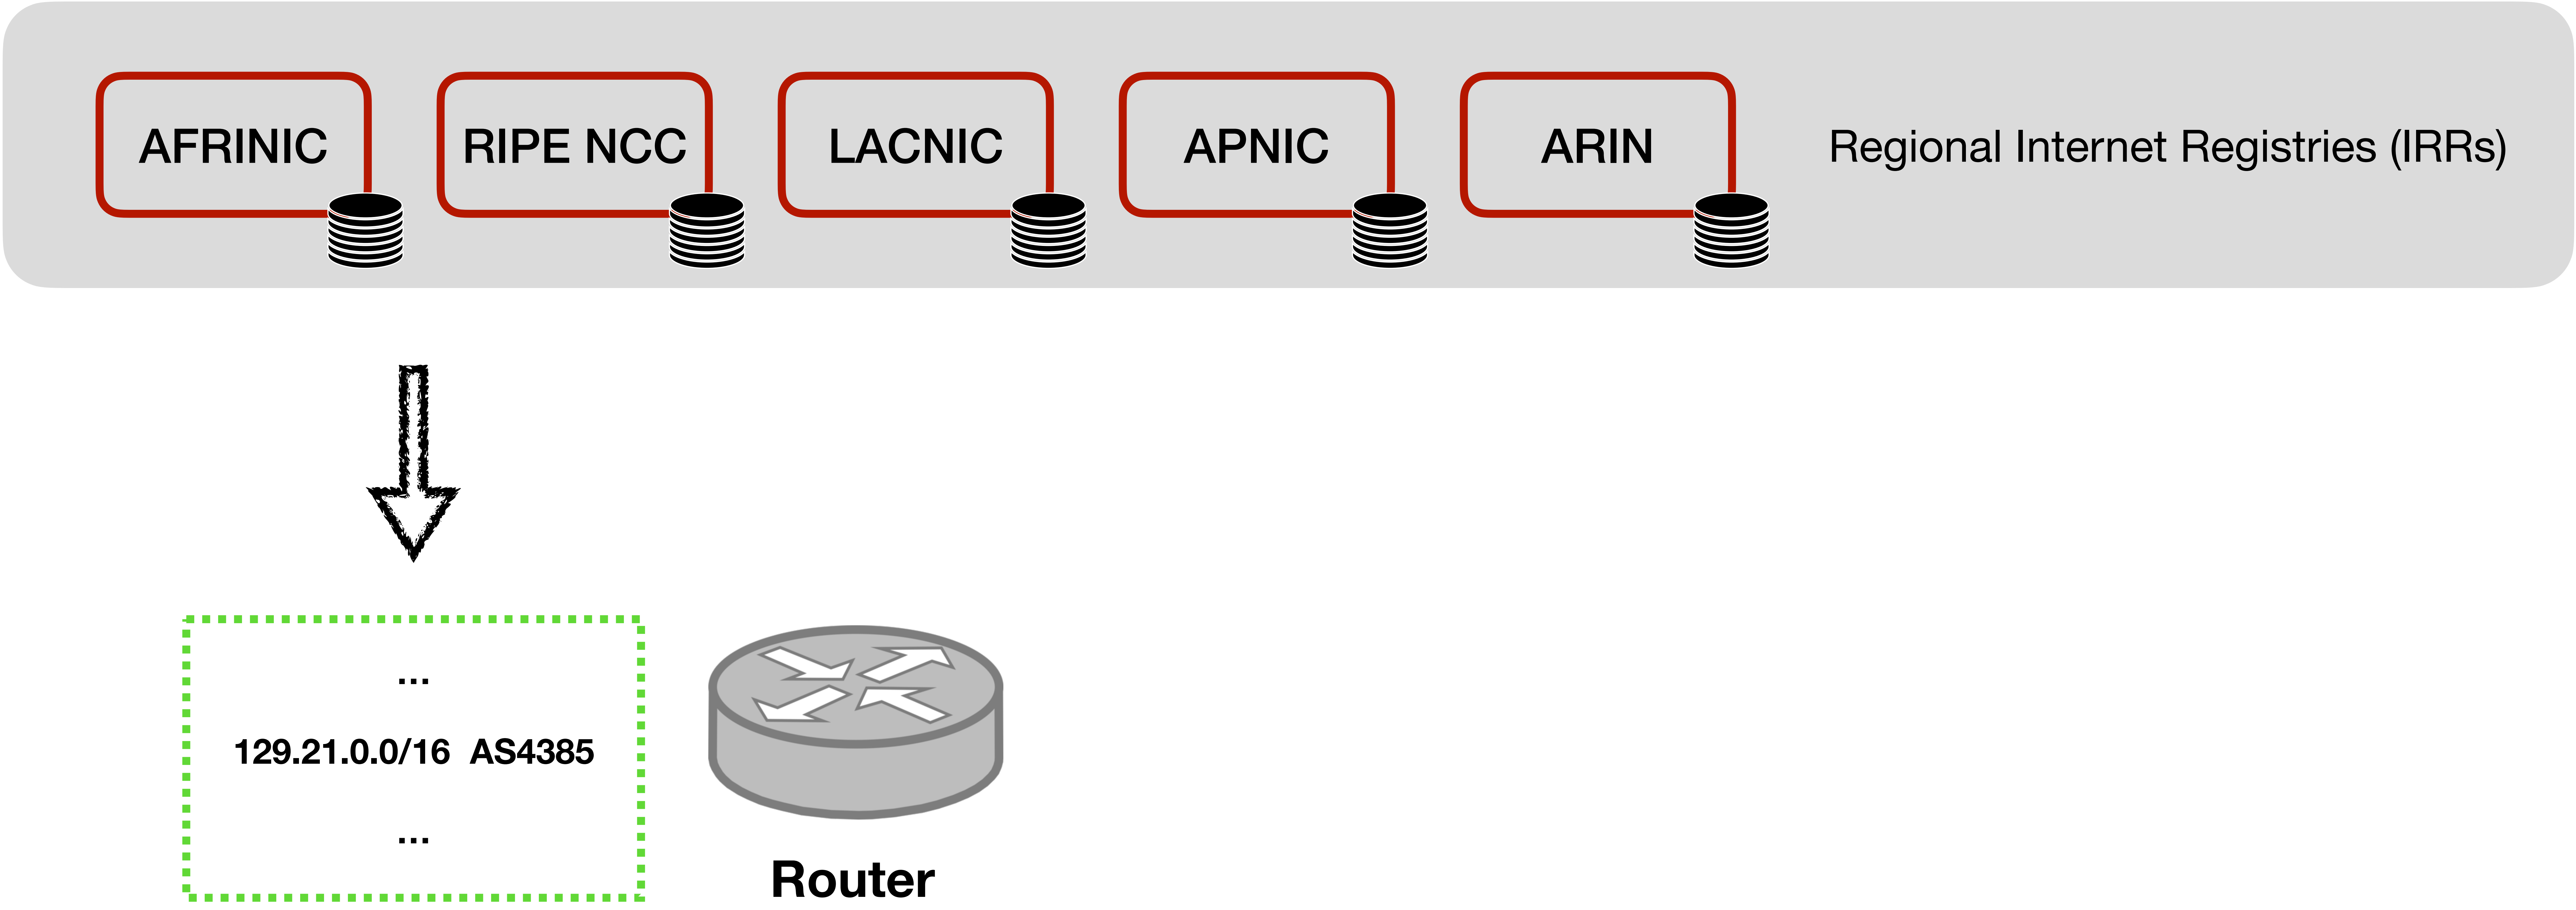

AFRINIC

RIPE NCC

LACNIC

APNIC

ARIN

Regional Internet Registries (IRRs)

...

129.21.0.0/16 AS4385

...

Router

(Cryptographically verifiable)
Prefix-to-AS Mapping Database

RPKI Structure: Route Origin Validation

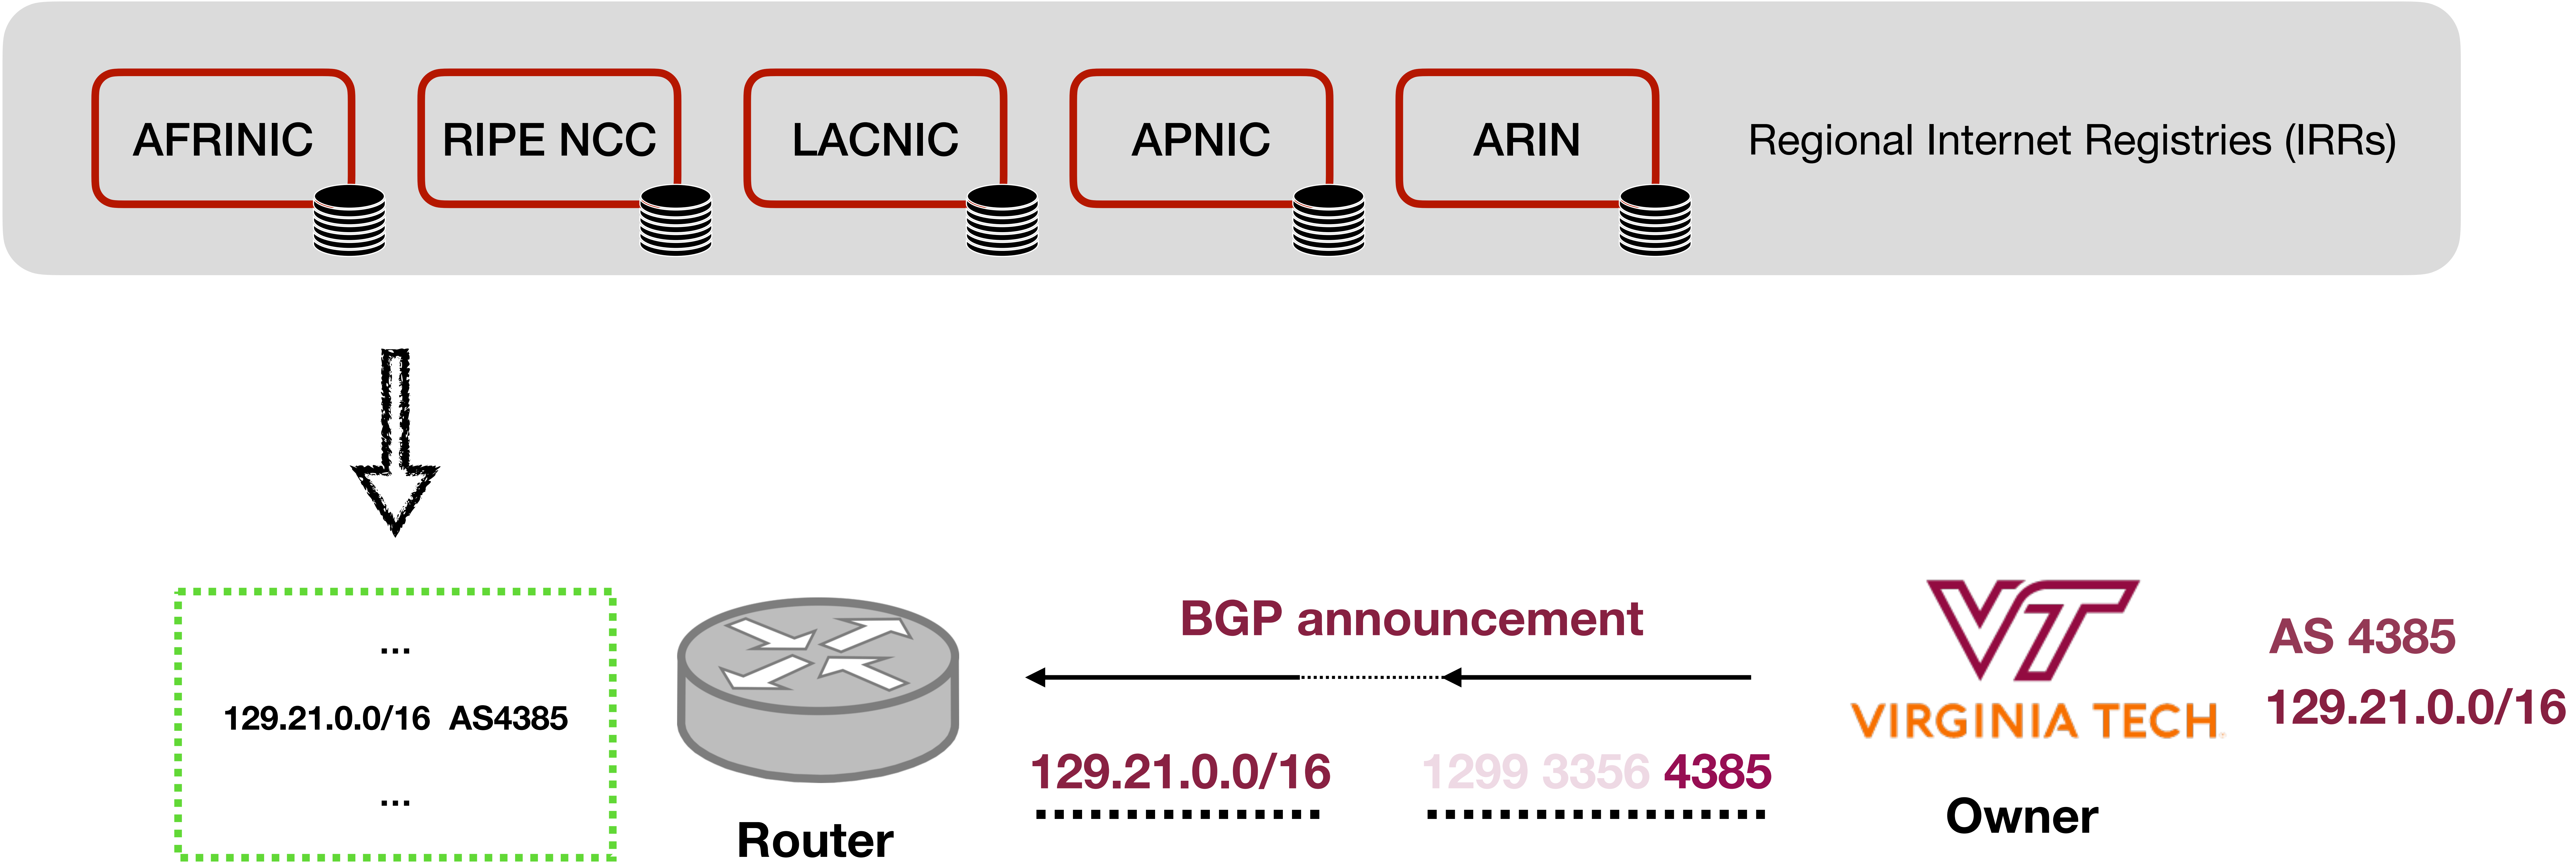

AFRINIC

RIPE NCC

LACNIC

APNIC

ARIN

Regional Internet Registries (IRRs)

...

129.21.0.0/16 AS4385

...

Router

BGP announcement

129.21.0.0/16

Prefix

1299 3356 4385

AS-PATH

AS 4385

129.21.0.0/16

Owner

(Cryptographically verifiable)
Prefix-to-AS Mapping Database

RPKI Structure: ROV

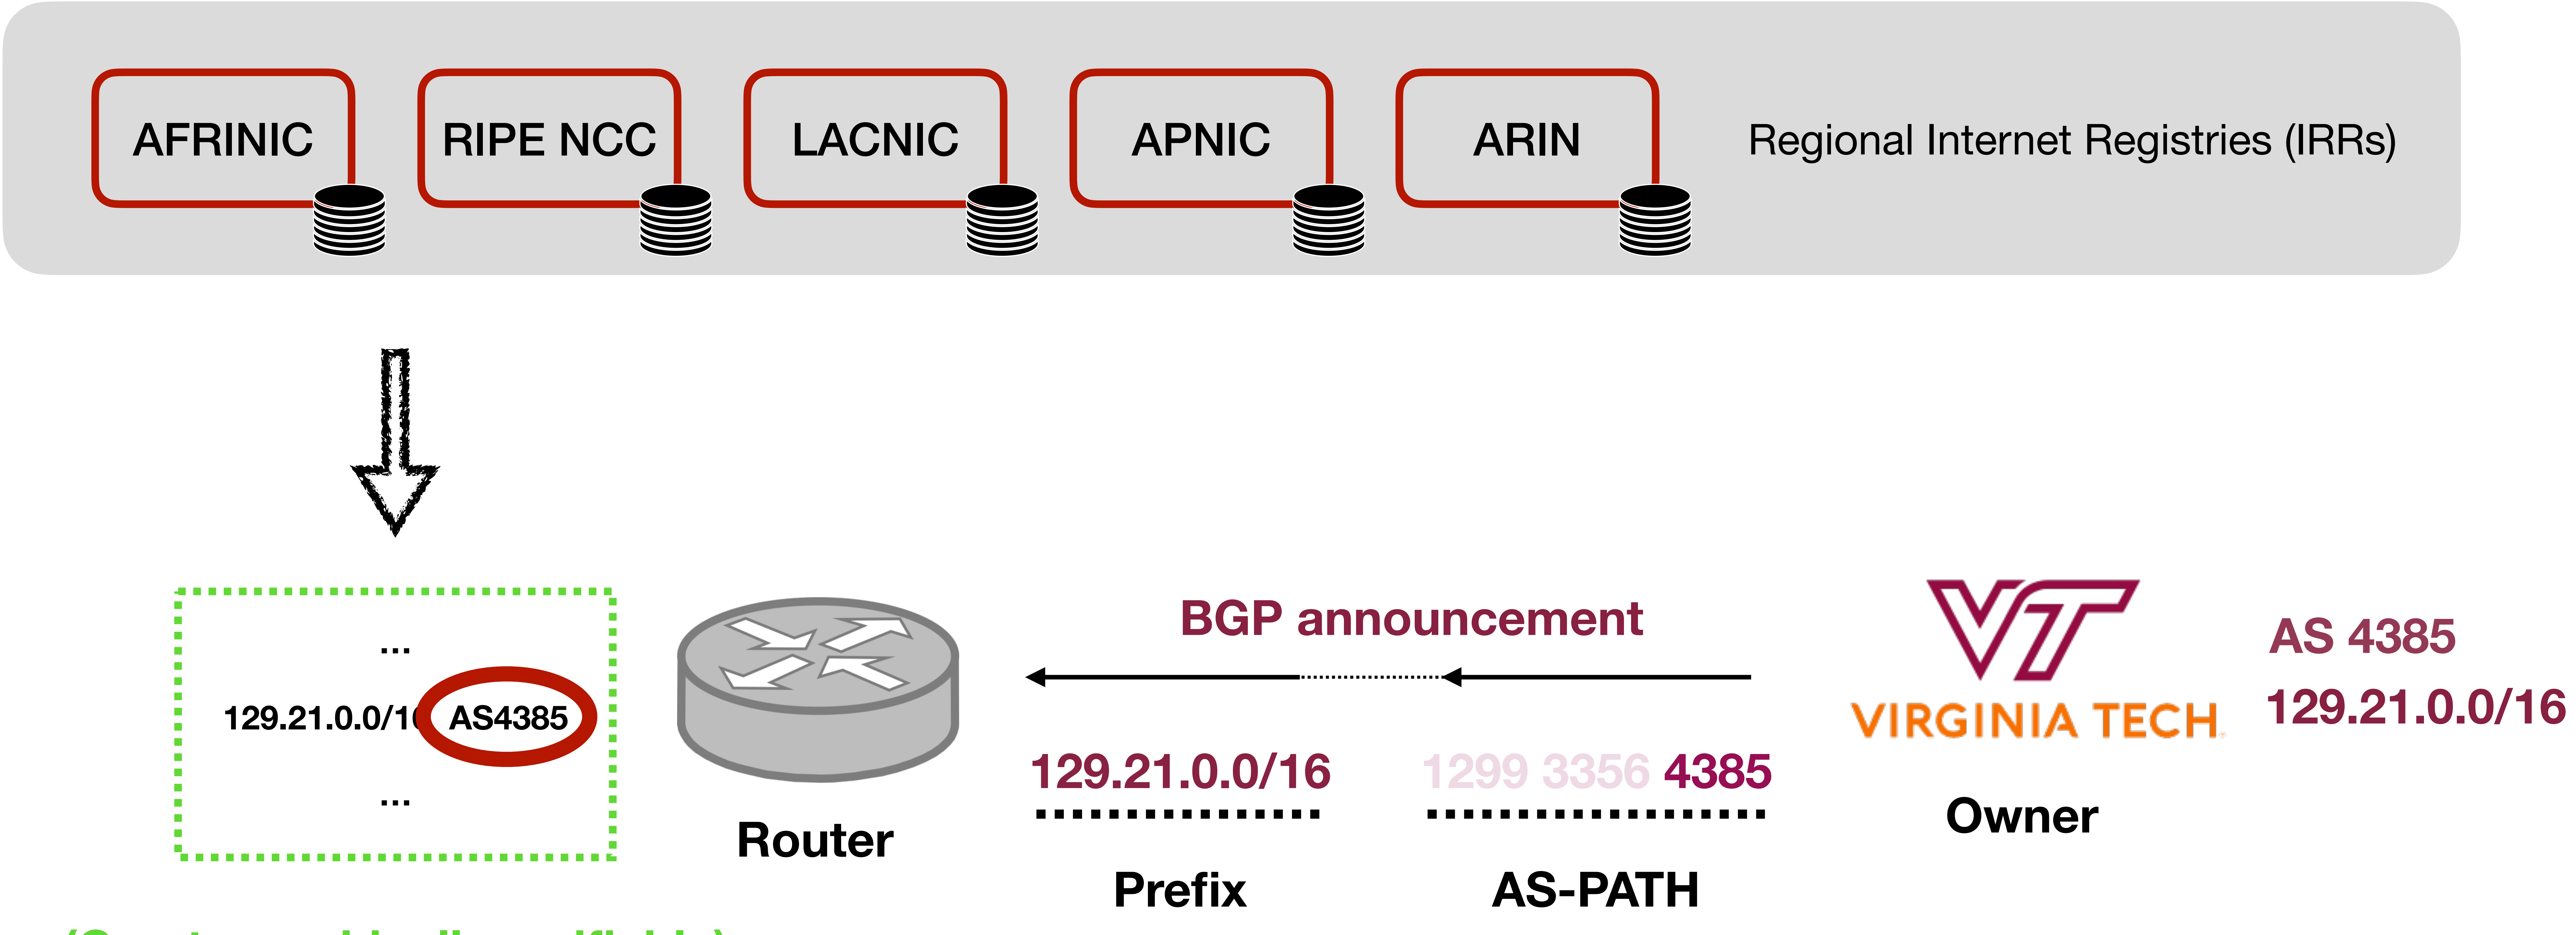

AFRINIC

RIPE NCC

LACNIC

APNIC

ARIN

Regional Internet Registries (IRRs)

...

129.21.0.0/1(AS4385

...

Router

BGP announcement

129.21.0.0/16

1299 3356 4385

Prefix

AS-PATH

Owner

AS 4385

129.21.0.0/16

(Cryptographically verifiable)
Prefix-to-AS Mapping Database

RPKI Structure: ROV

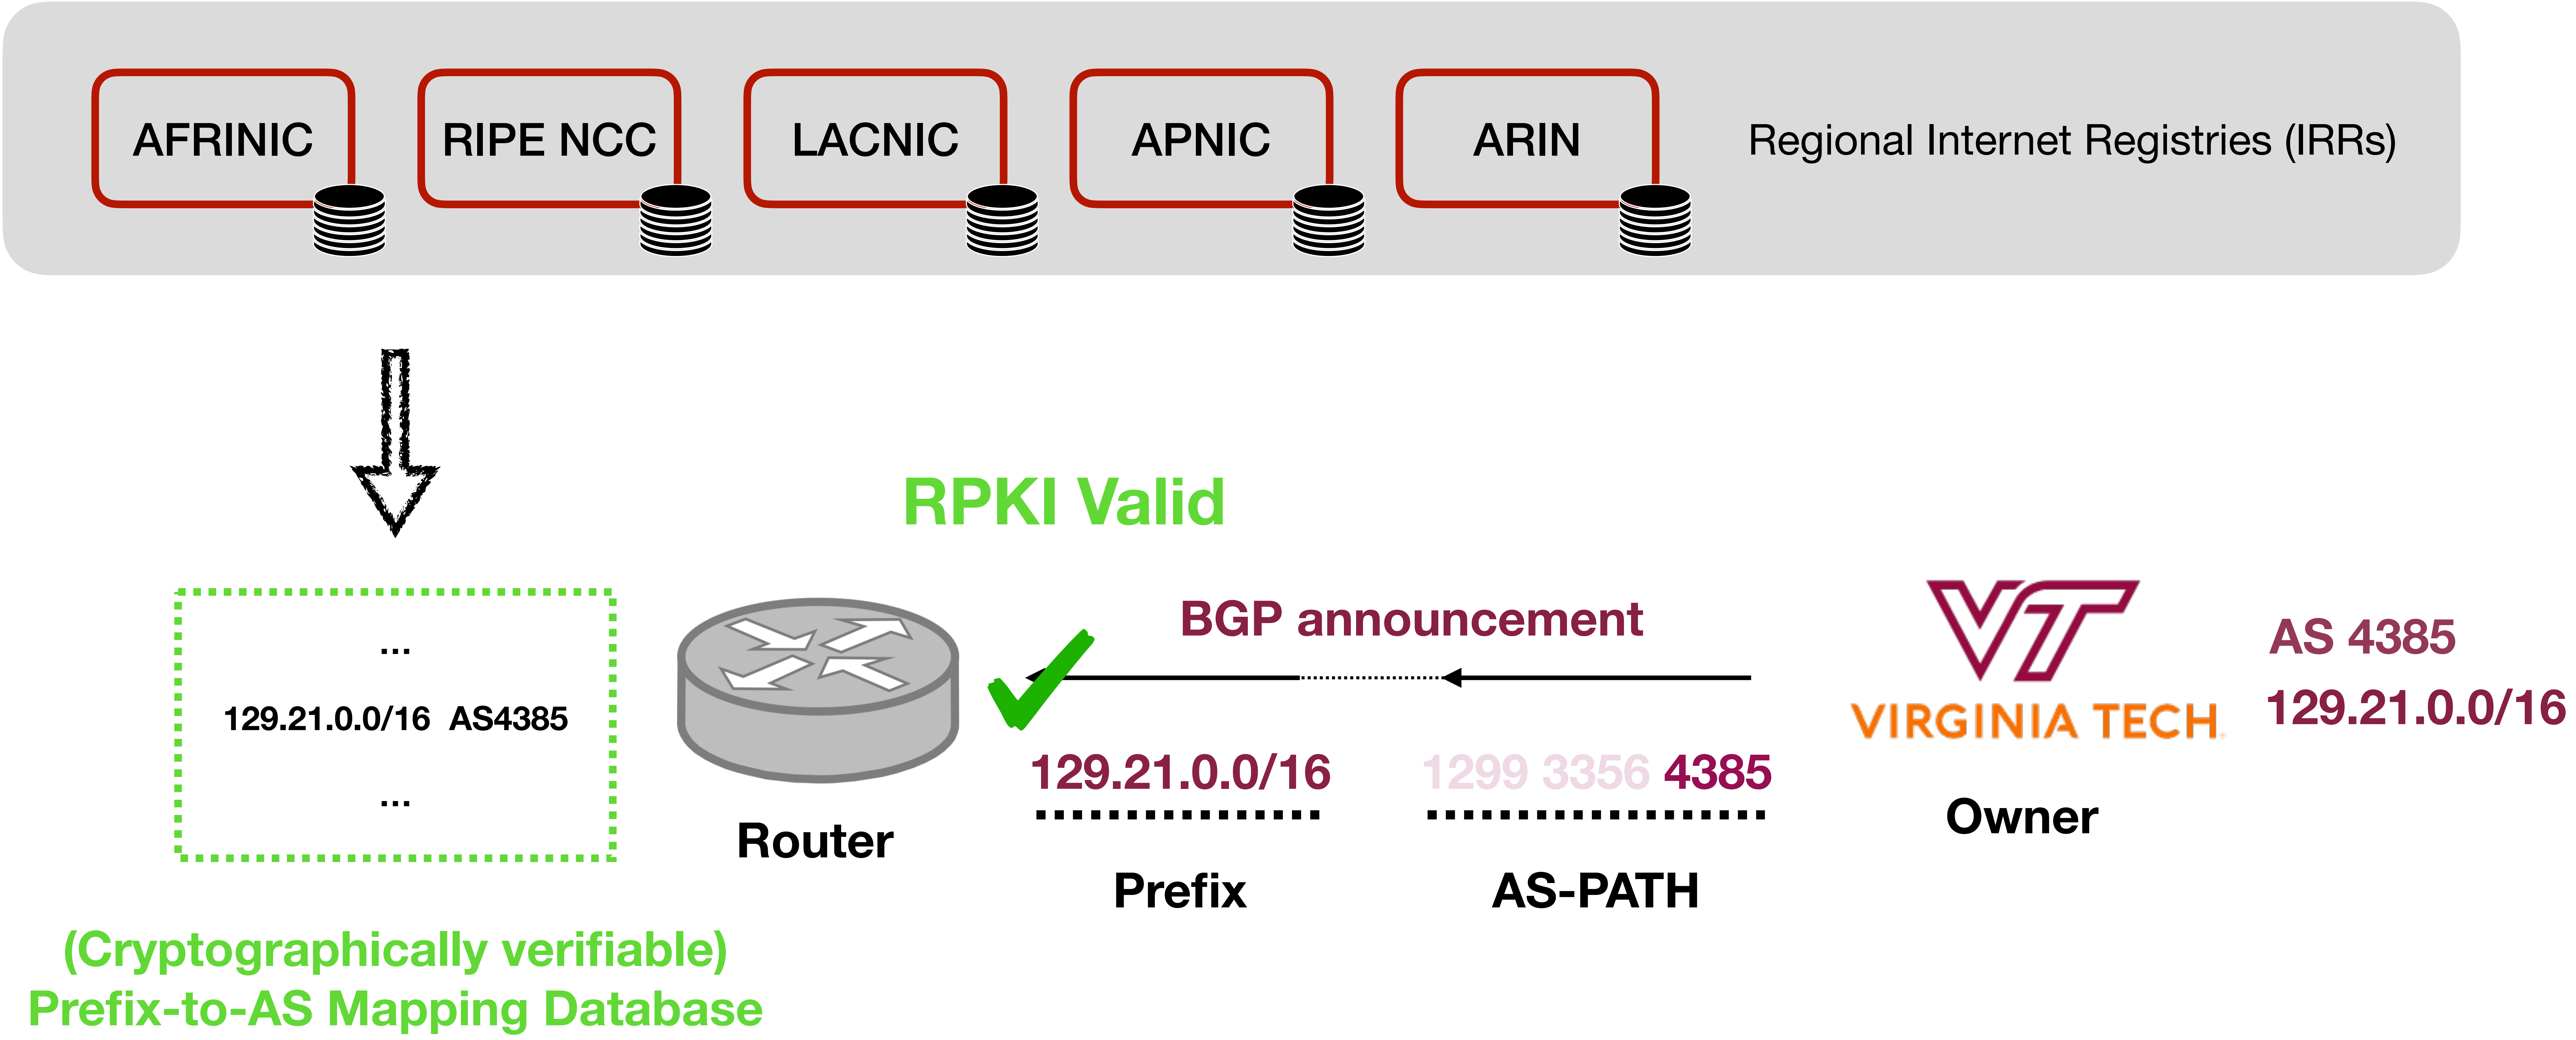

AFRINIC

RIPE NCC

LACNIC

APNIC

ARIN

Regional Internet Registries (IRRs)

RPKI Valid

...

129.21.0.0/16 AS4385

...

(Cryptographically verifiable)
Prefix-to-AS Mapping Database

Router

BGP announcement

129.21.0.0/16

Prefix

1299 3356 4385

AS-PATH

AS 4385
129.21.0.0/16

Owner

RPKI Structure: ROV

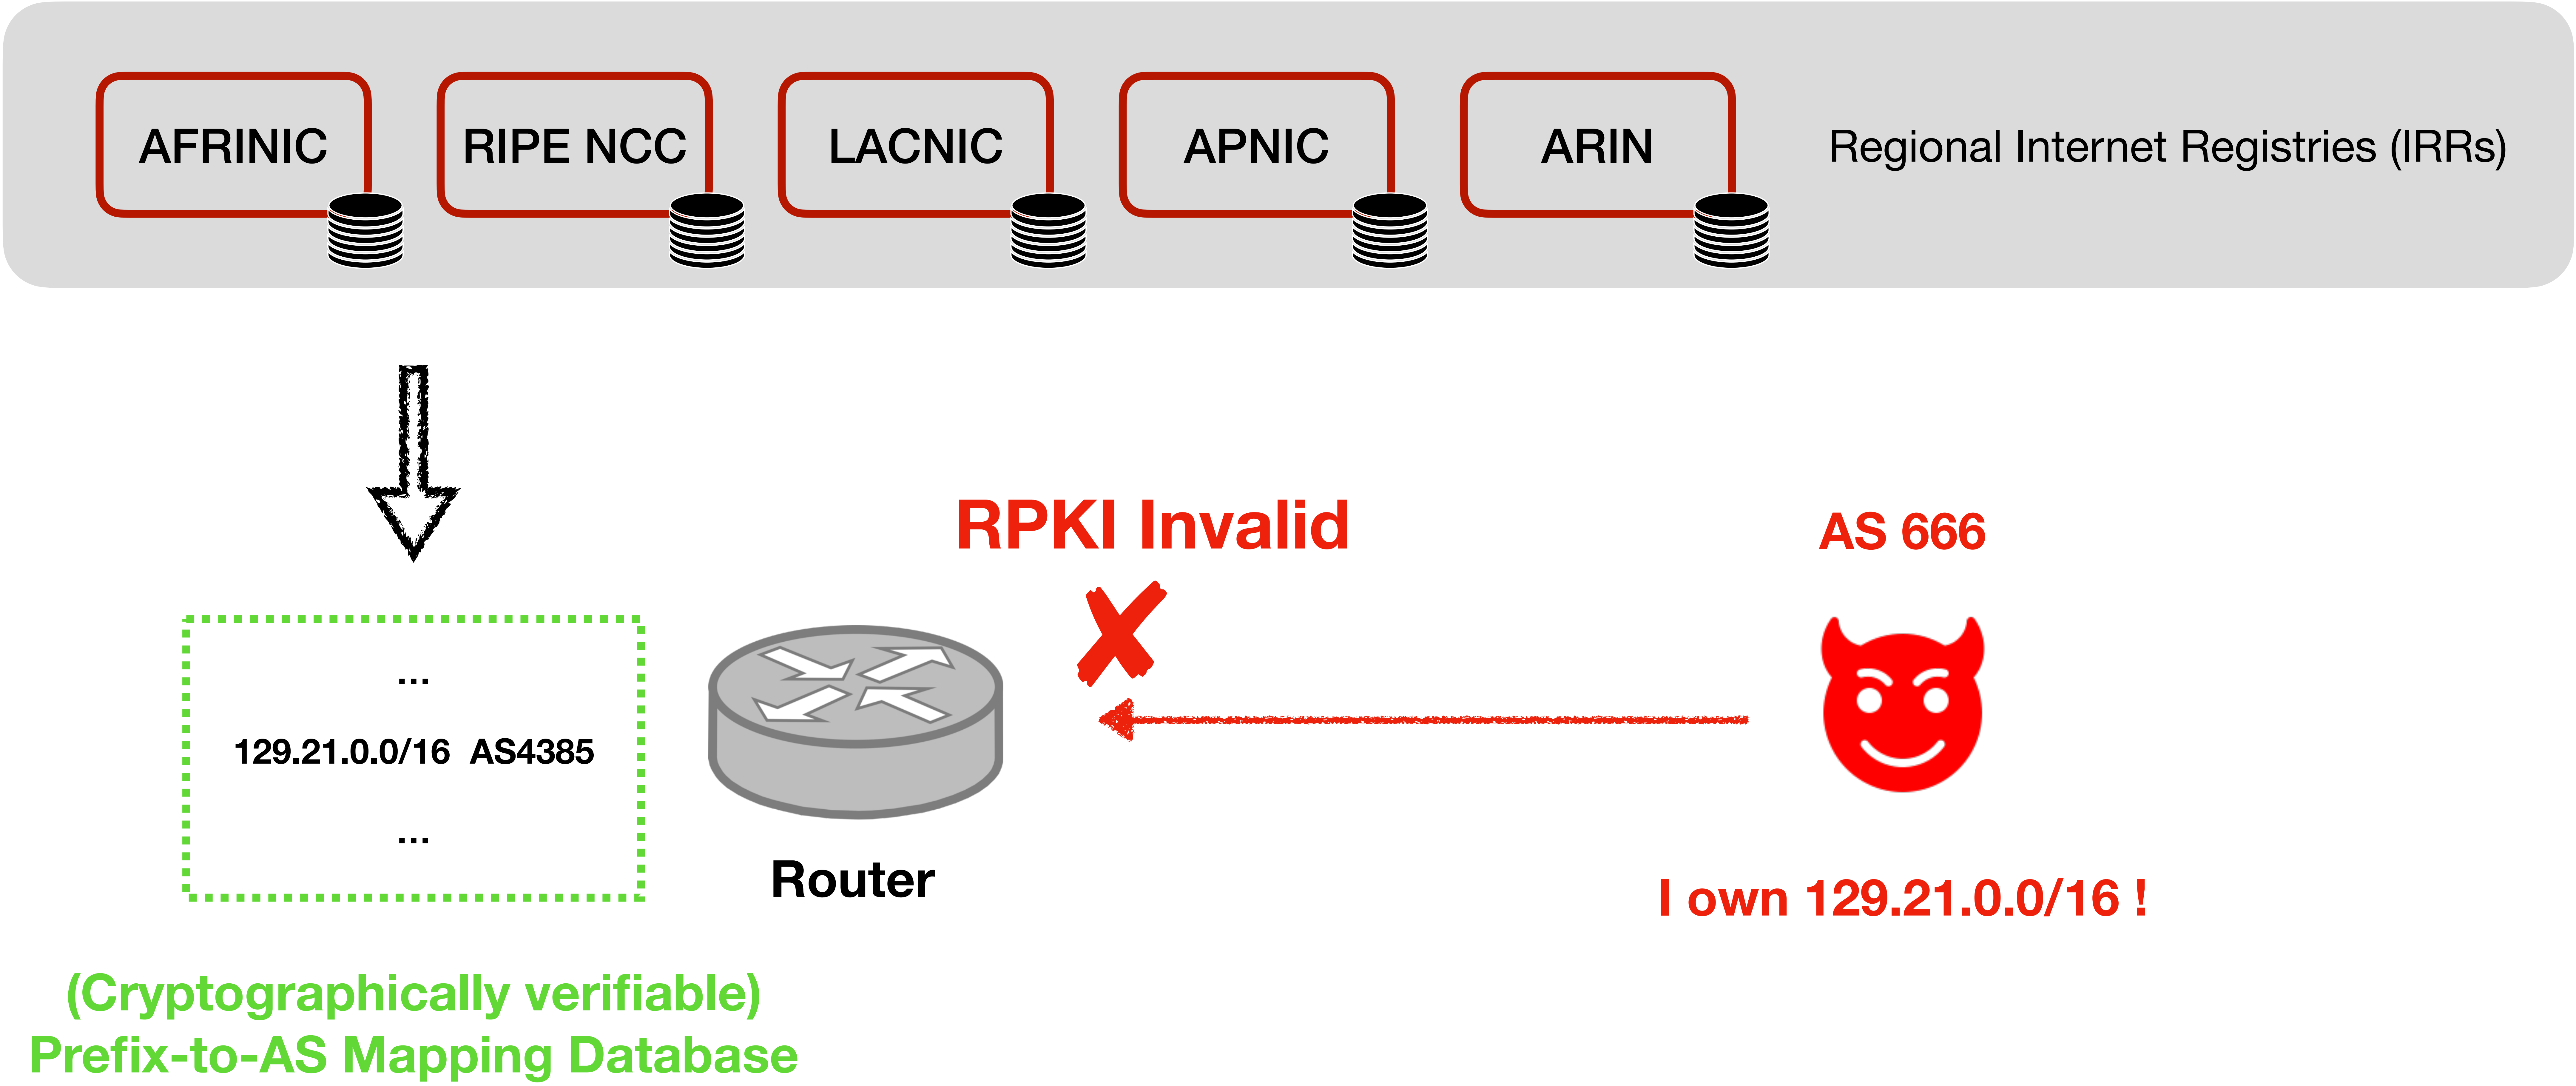

AFRINIC

RIPE NCC

LACNIC

APNIC

ARIN

Regional Internet Registries (IRRs)

...

129.21.0.0/16 AS4385

...

Router

RPKI Invalid

AS 666

I own 129.21.0.0/16 !

(Cryptographically verifiable)
Prefix-to-AS Mapping Database

The Deployment of ROA

- We analyze BGP routing tables from all collectors of RouteViews and RIPE RIS during a 18-month period

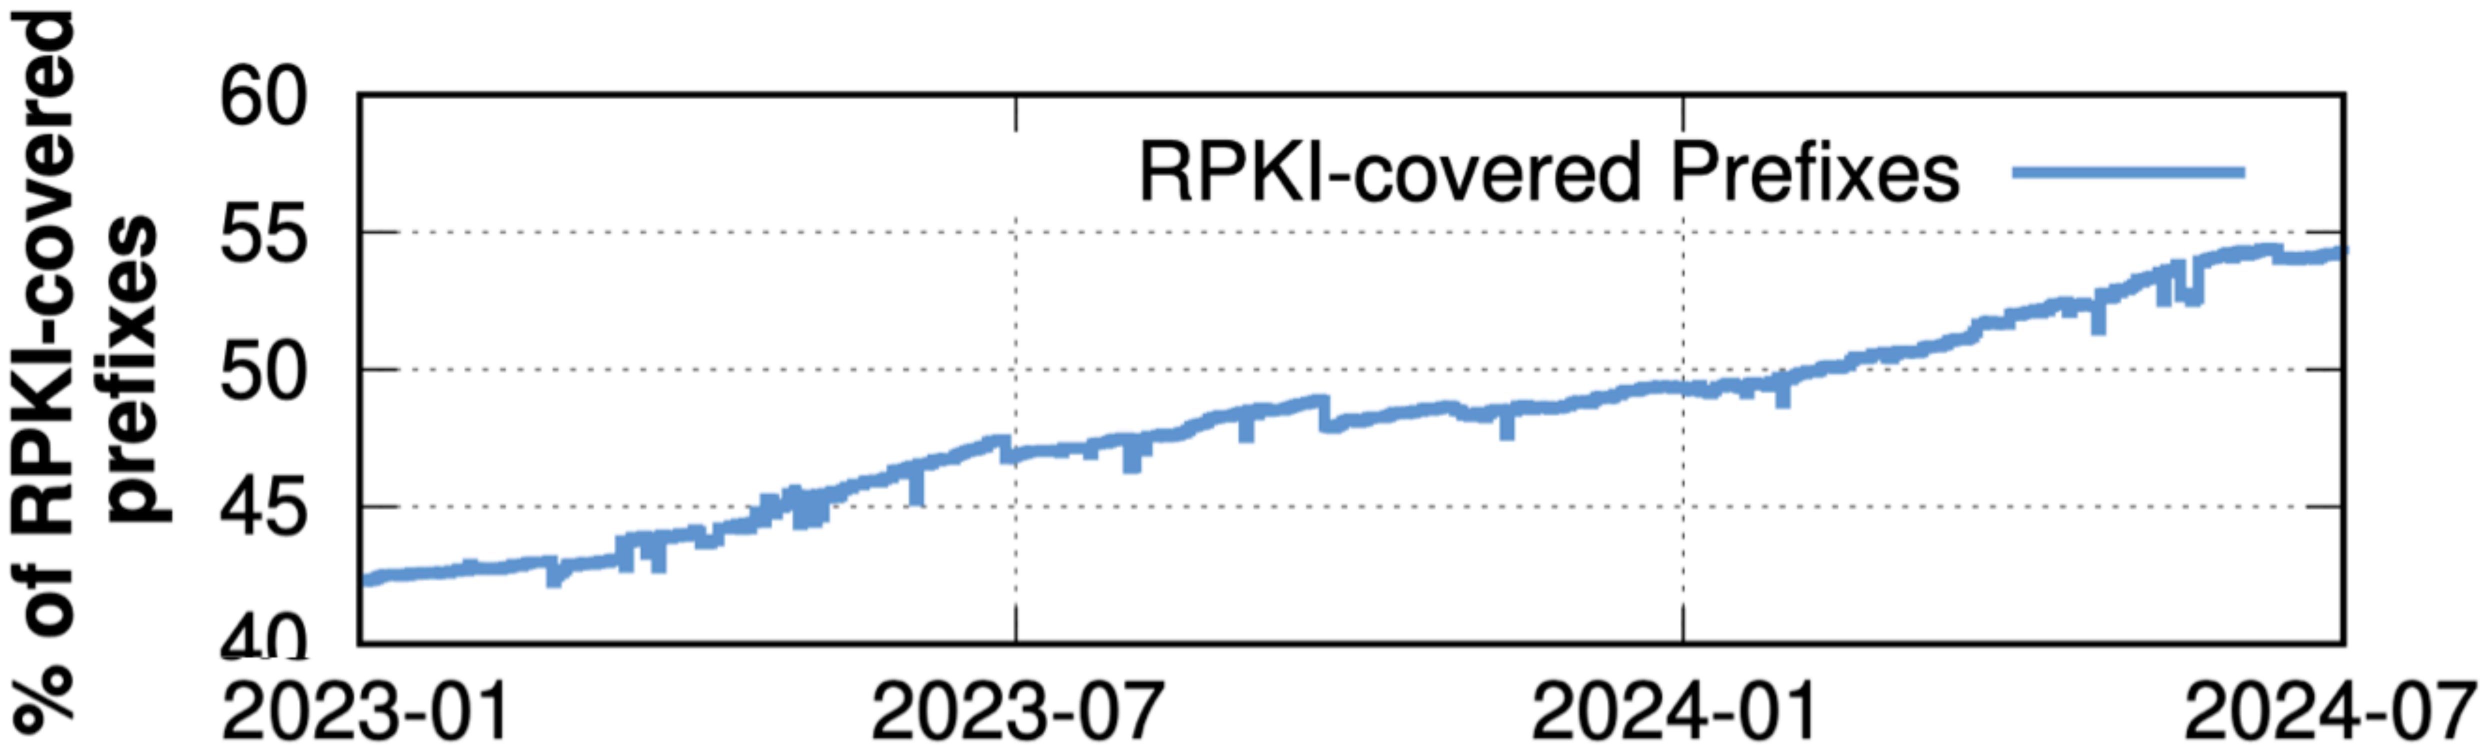

The Deployment of ROA

- We analyze BGP routing tables from all collectors of RouteViews and RIPE RIS during a 18-month period

- The coverage of ROAs are growing, 54% of the IPv4 spaces are covered by ROAs. More networks are deploying ROA, but ….

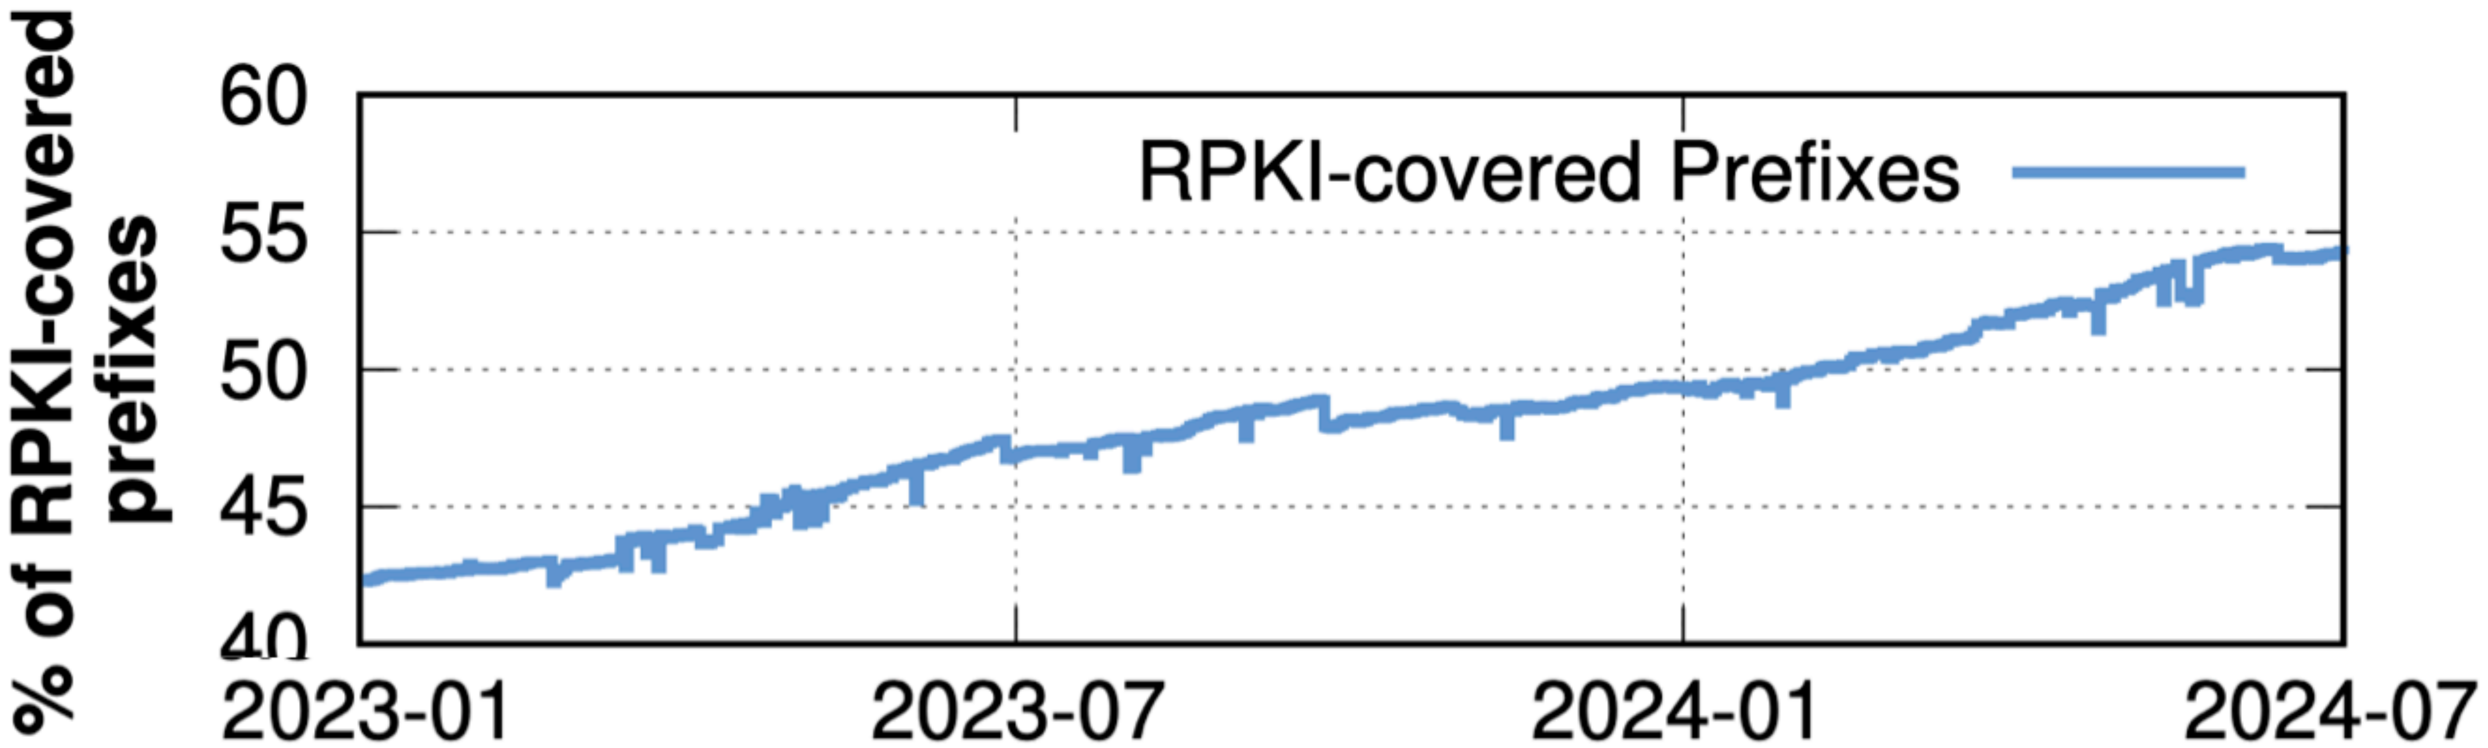

RPKI-Invalid routes never go away

- The coverage of ROAs are growing, 54% of the IPv4 spaces are covered by ROAs. More networks are deploying ROA, but we can still see 7000 RPKI-invalid routes everyday!

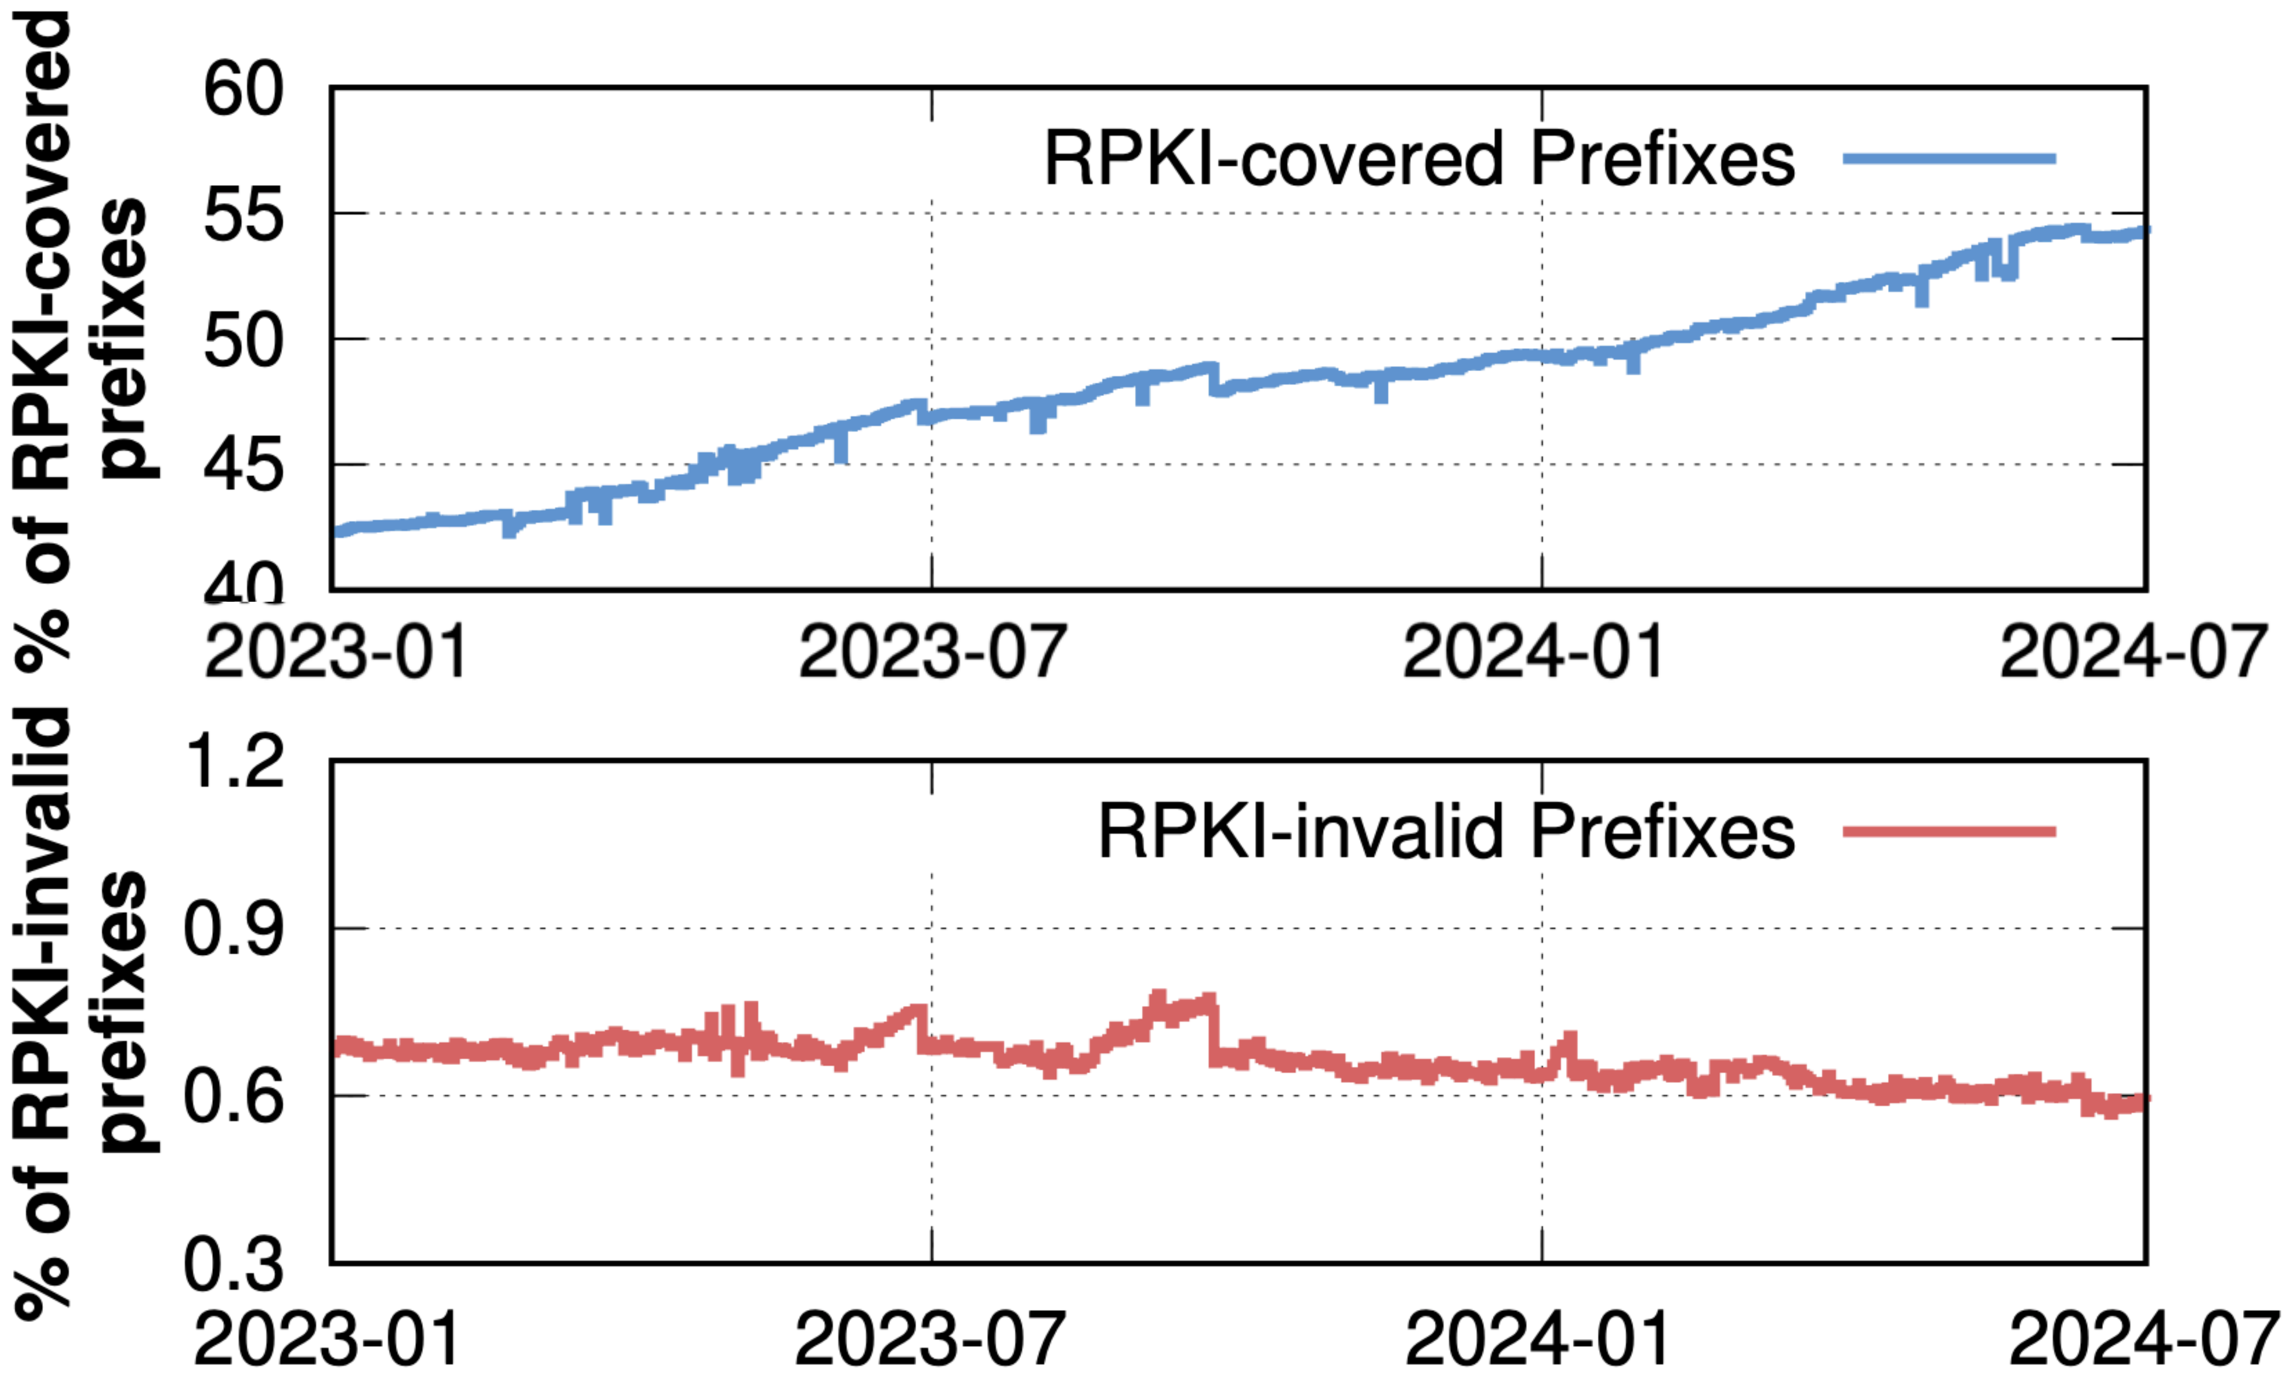

RPKI-Invalid routes never go away

- The coverage of ROAs are growing, 54% of the IPv4 spaces are covered by ROAs. More networks are deploying ROA, but we can still see 7000 RPKI-invalid routes everyday!

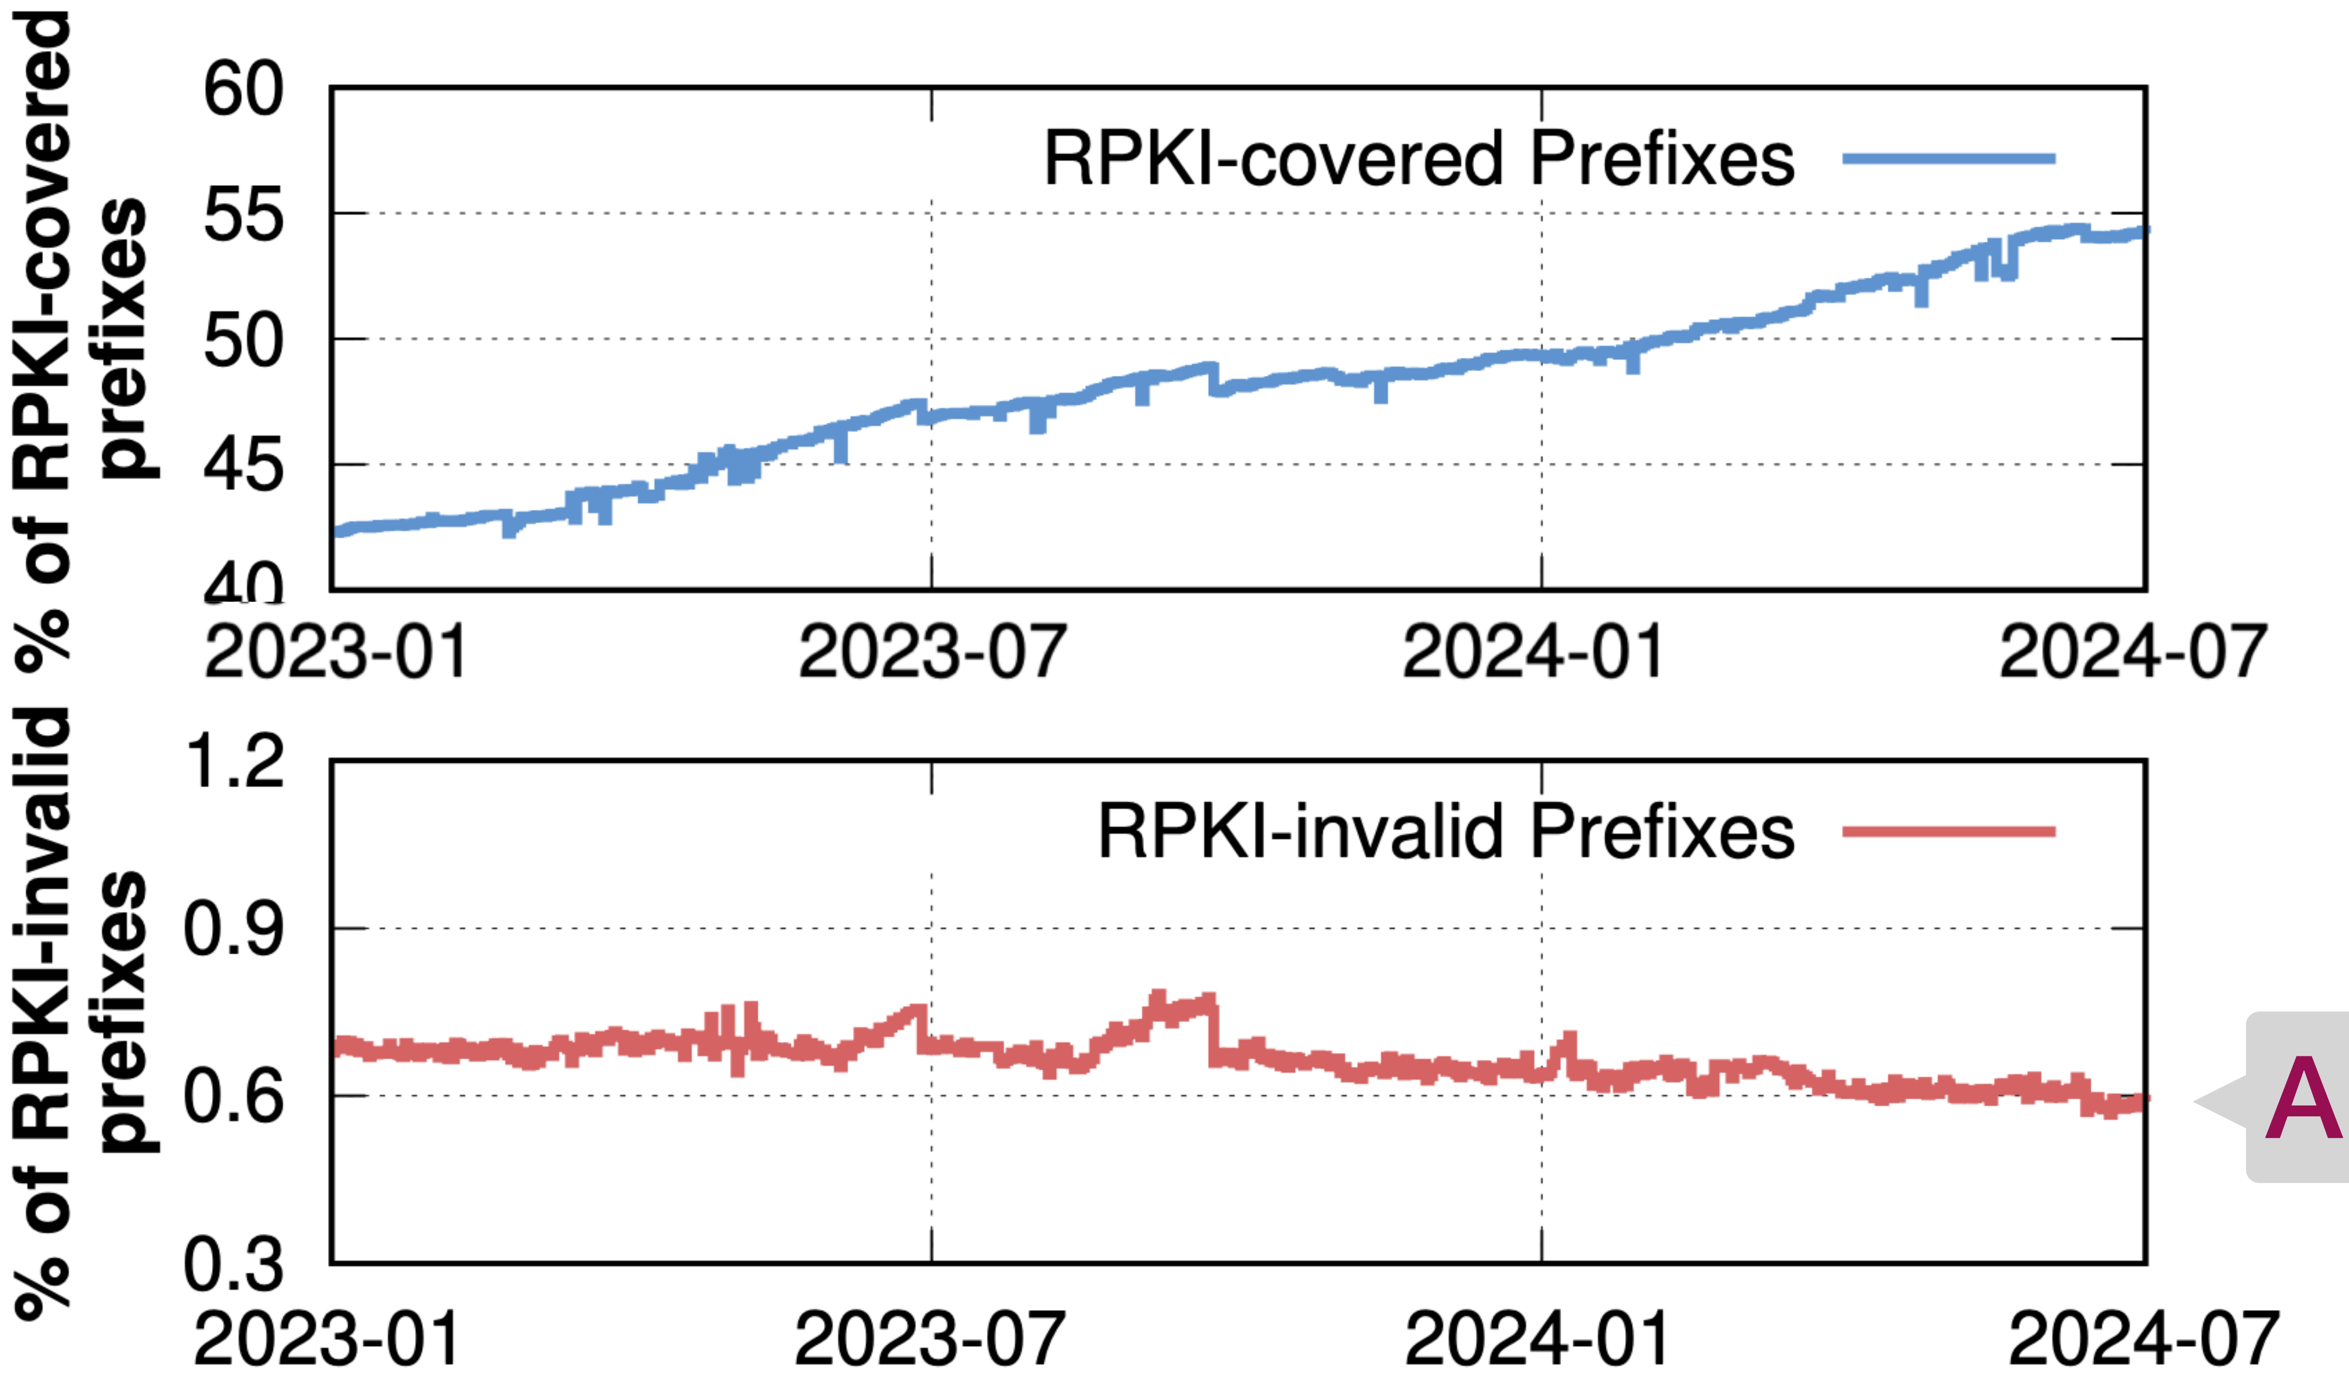

Are they all hijacks?

Question 1:
Why RPKI-Invalid Happens?

Known Misconfigurations

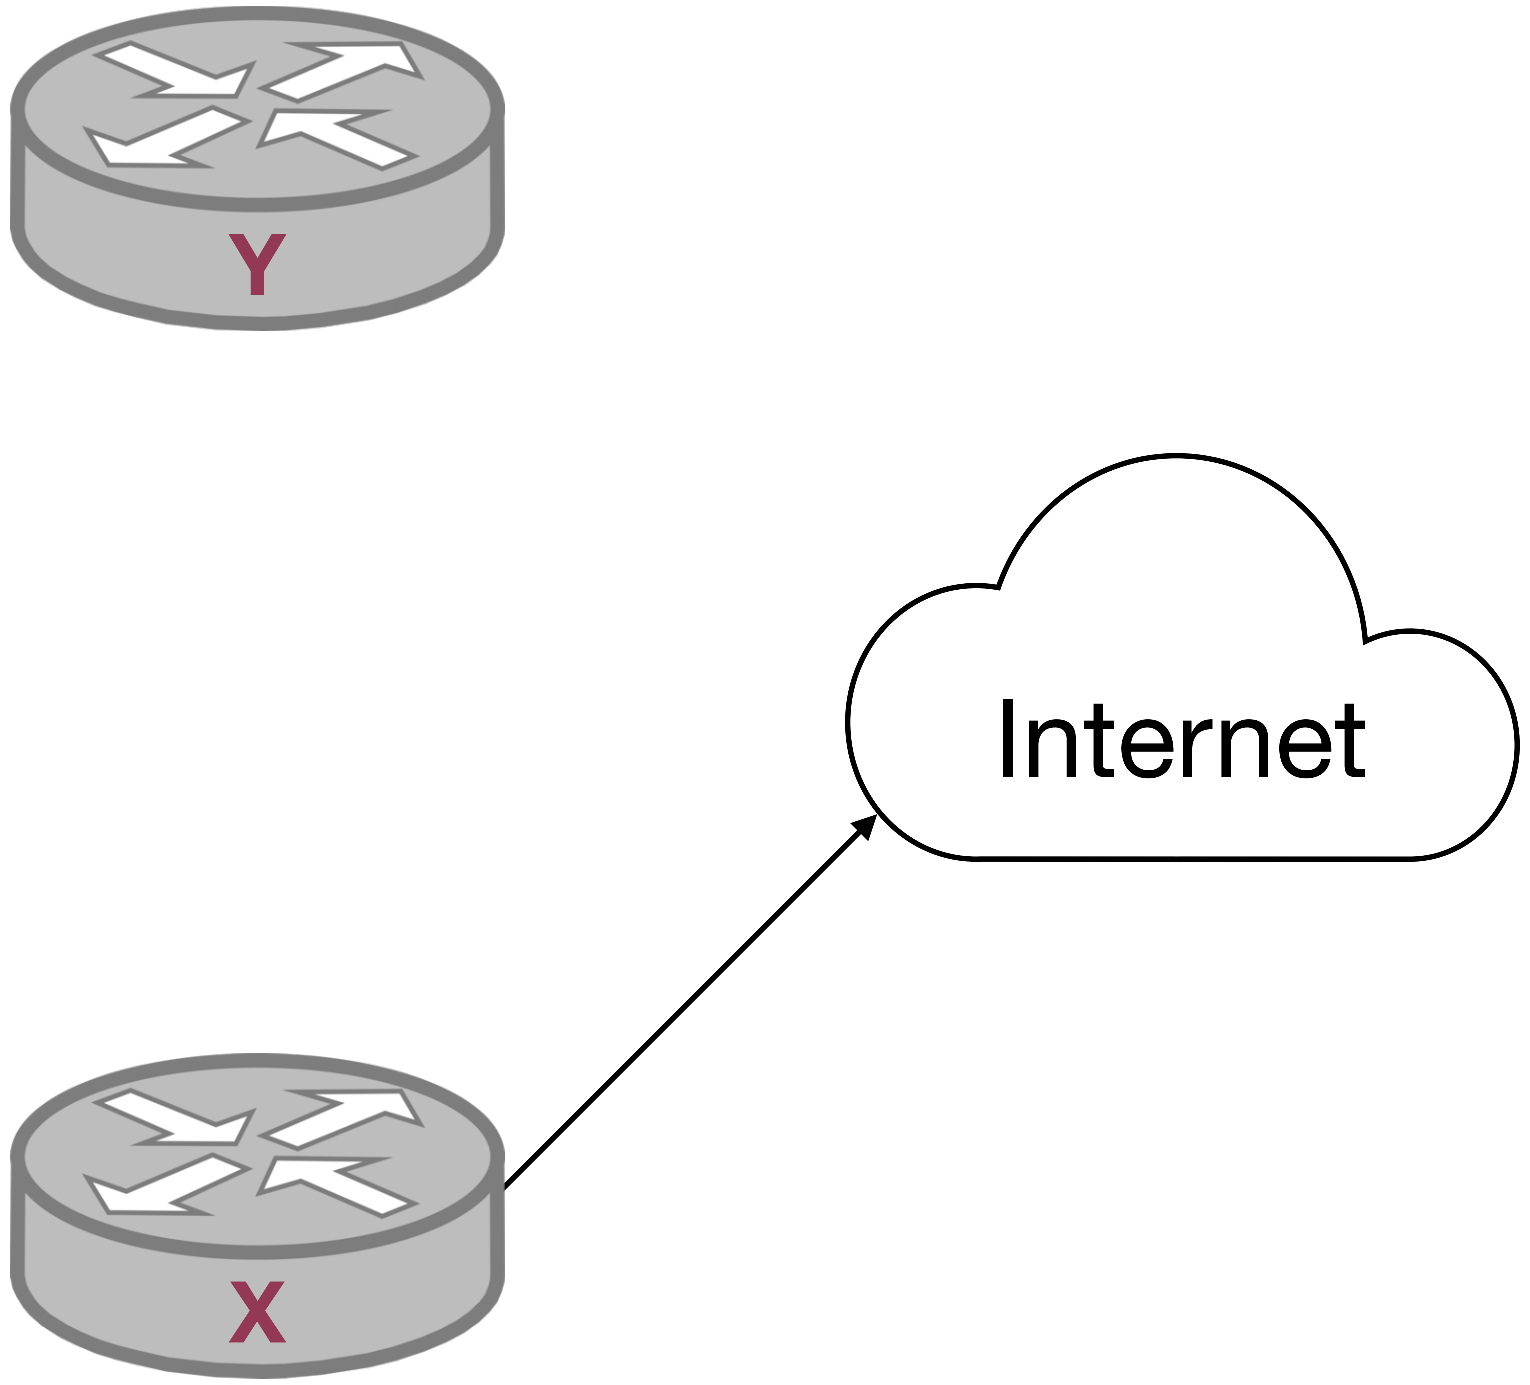

BGP Announcement:

129.21.0.0/20, AS X

ROA:

129.21.0.0/16-20, AS Y

Known Misconfigurations

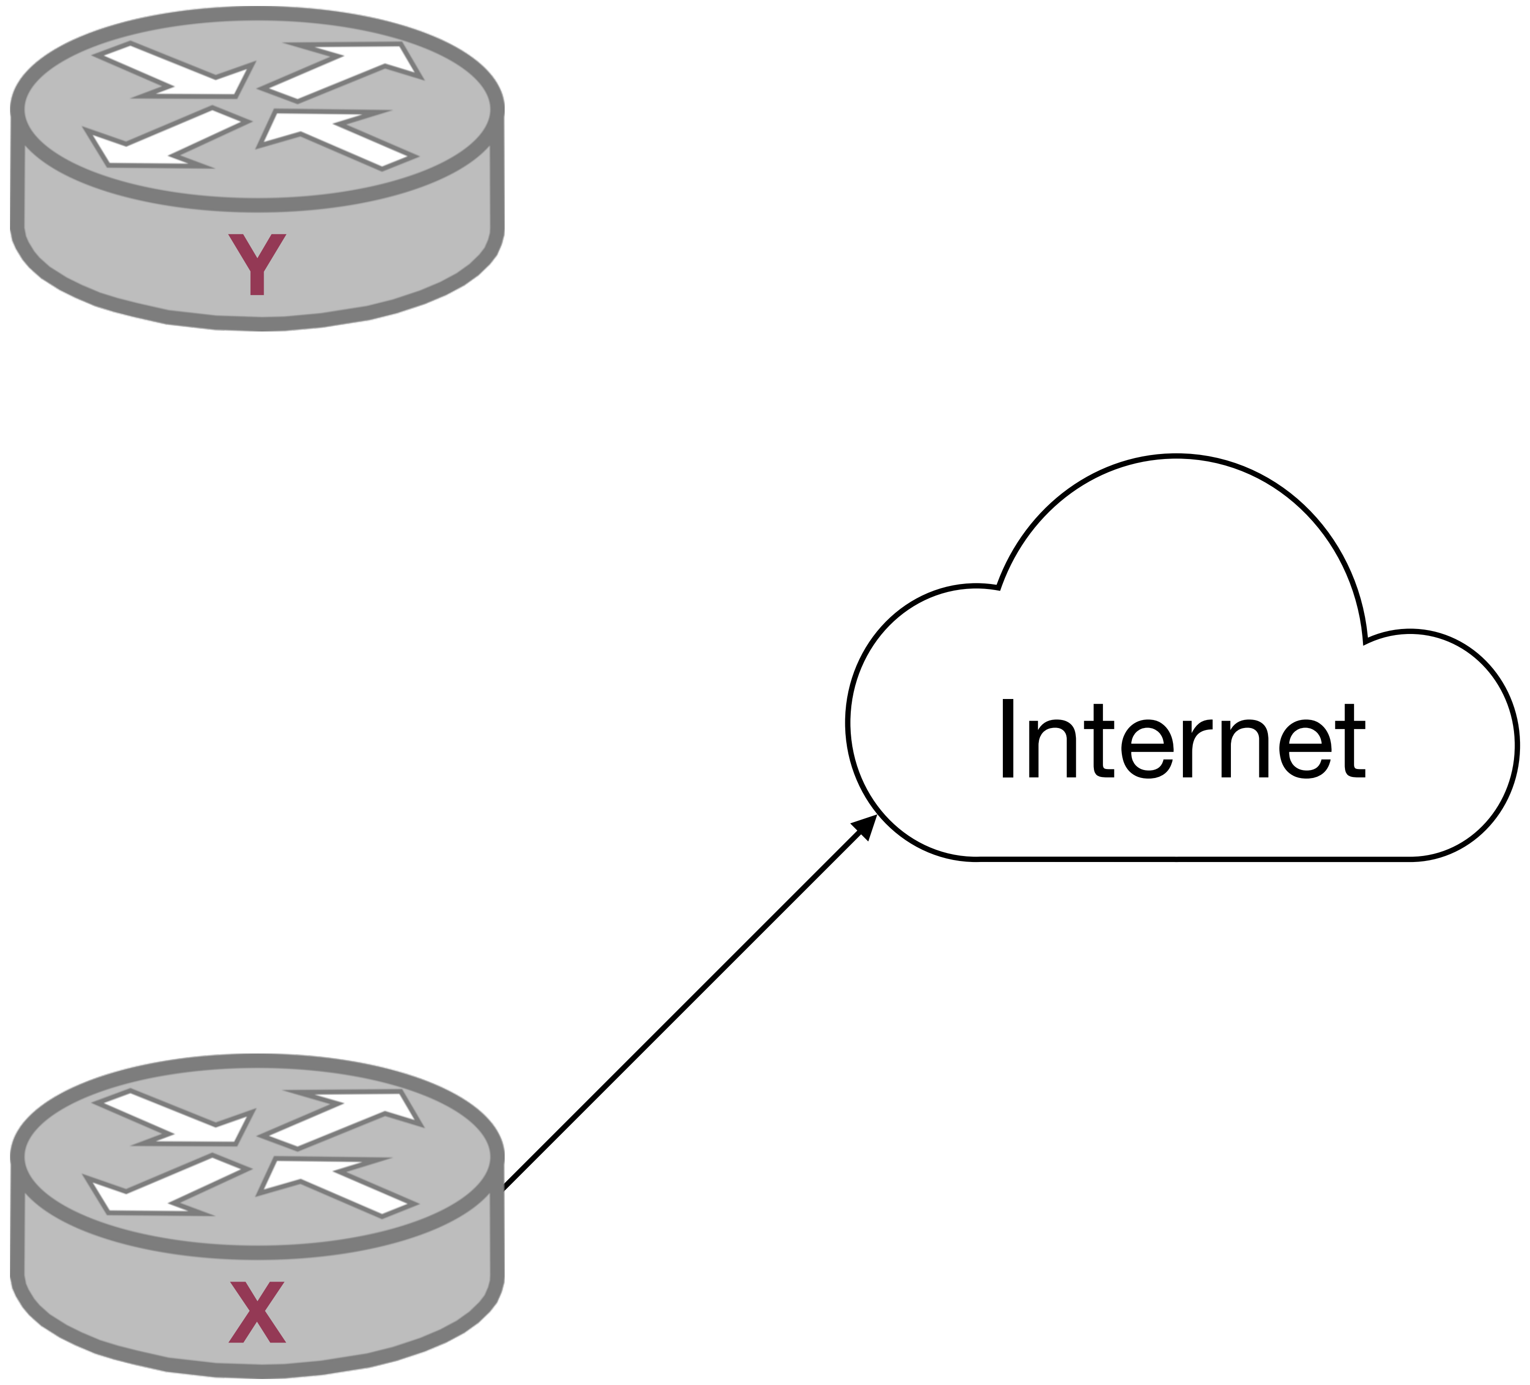

BGP Announcement:

129.21.0.0/20, AS X ✗ **RPKI Invalid**

ROA:

129.21.0.0/16-20, AS Y

Known Misconfigurations

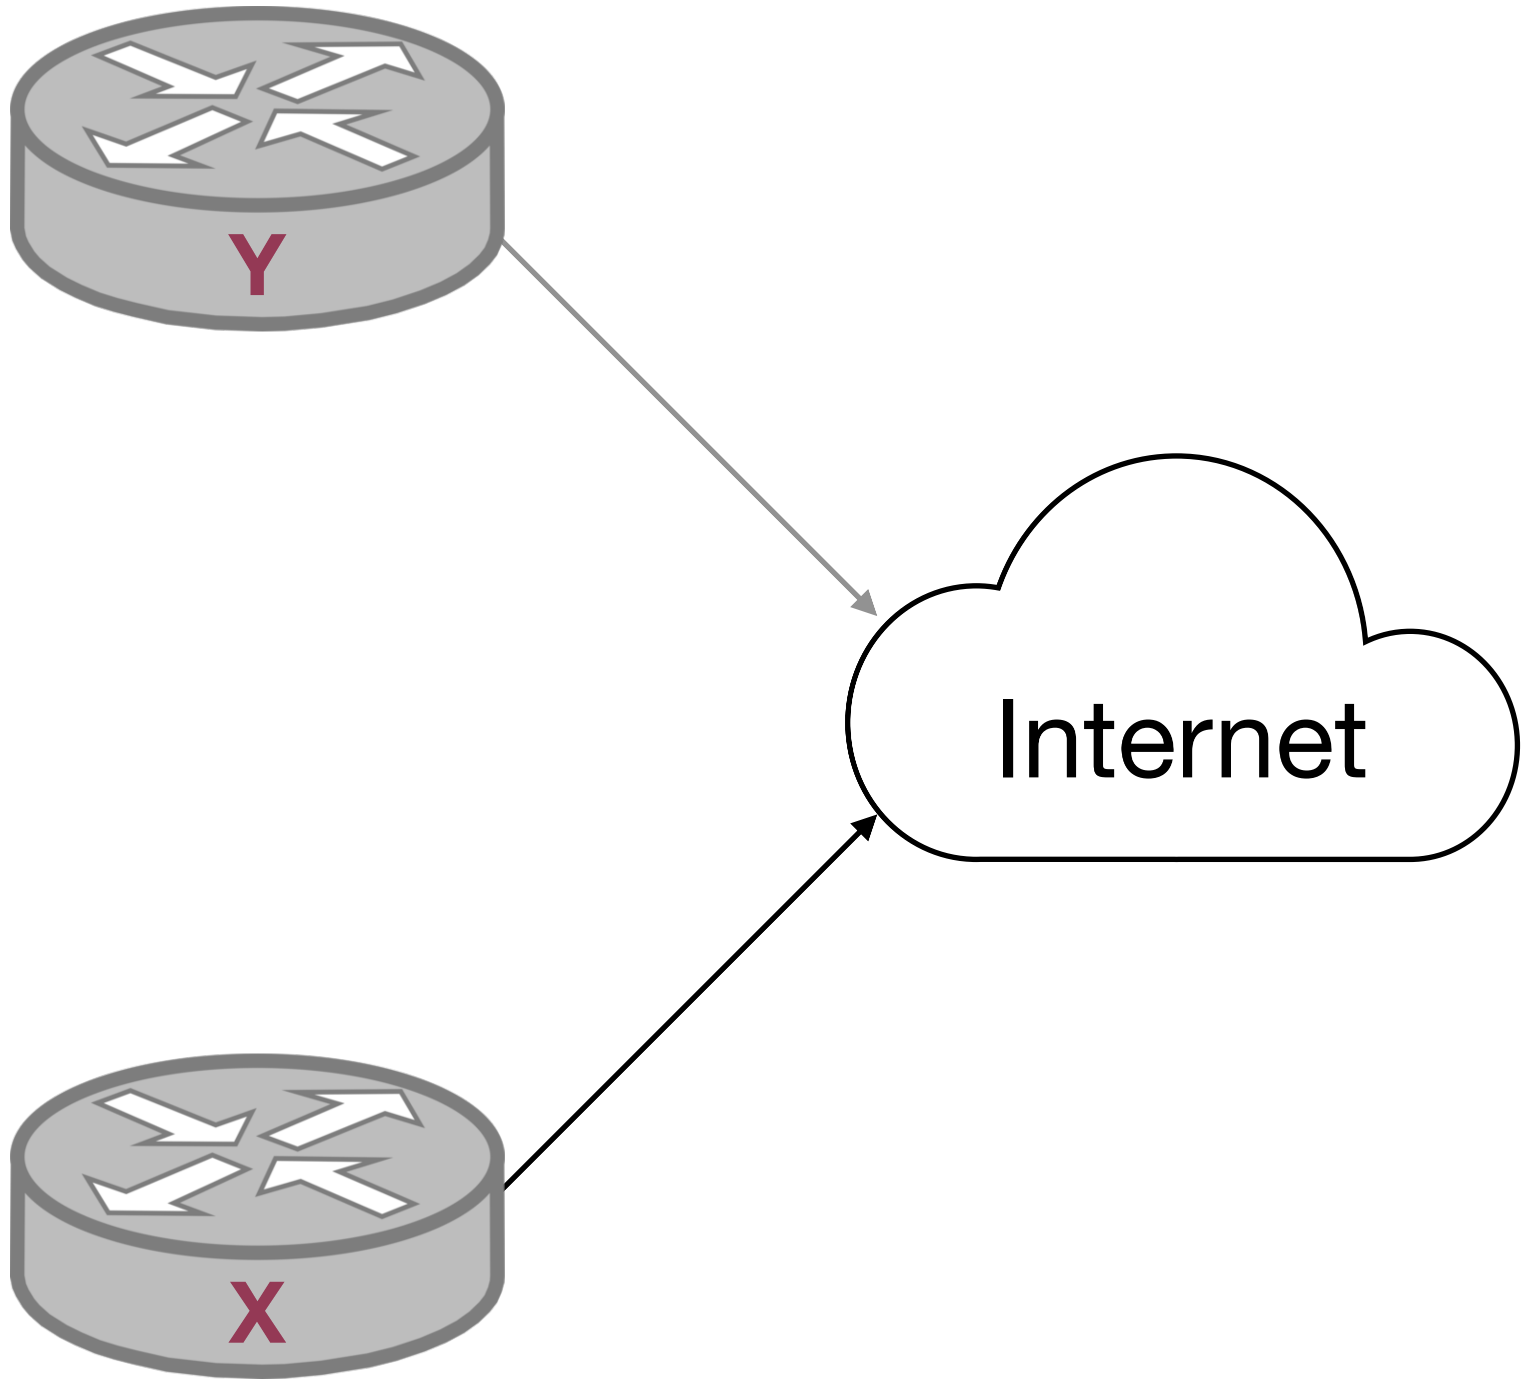

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

Known Misconfigurations in Same Org

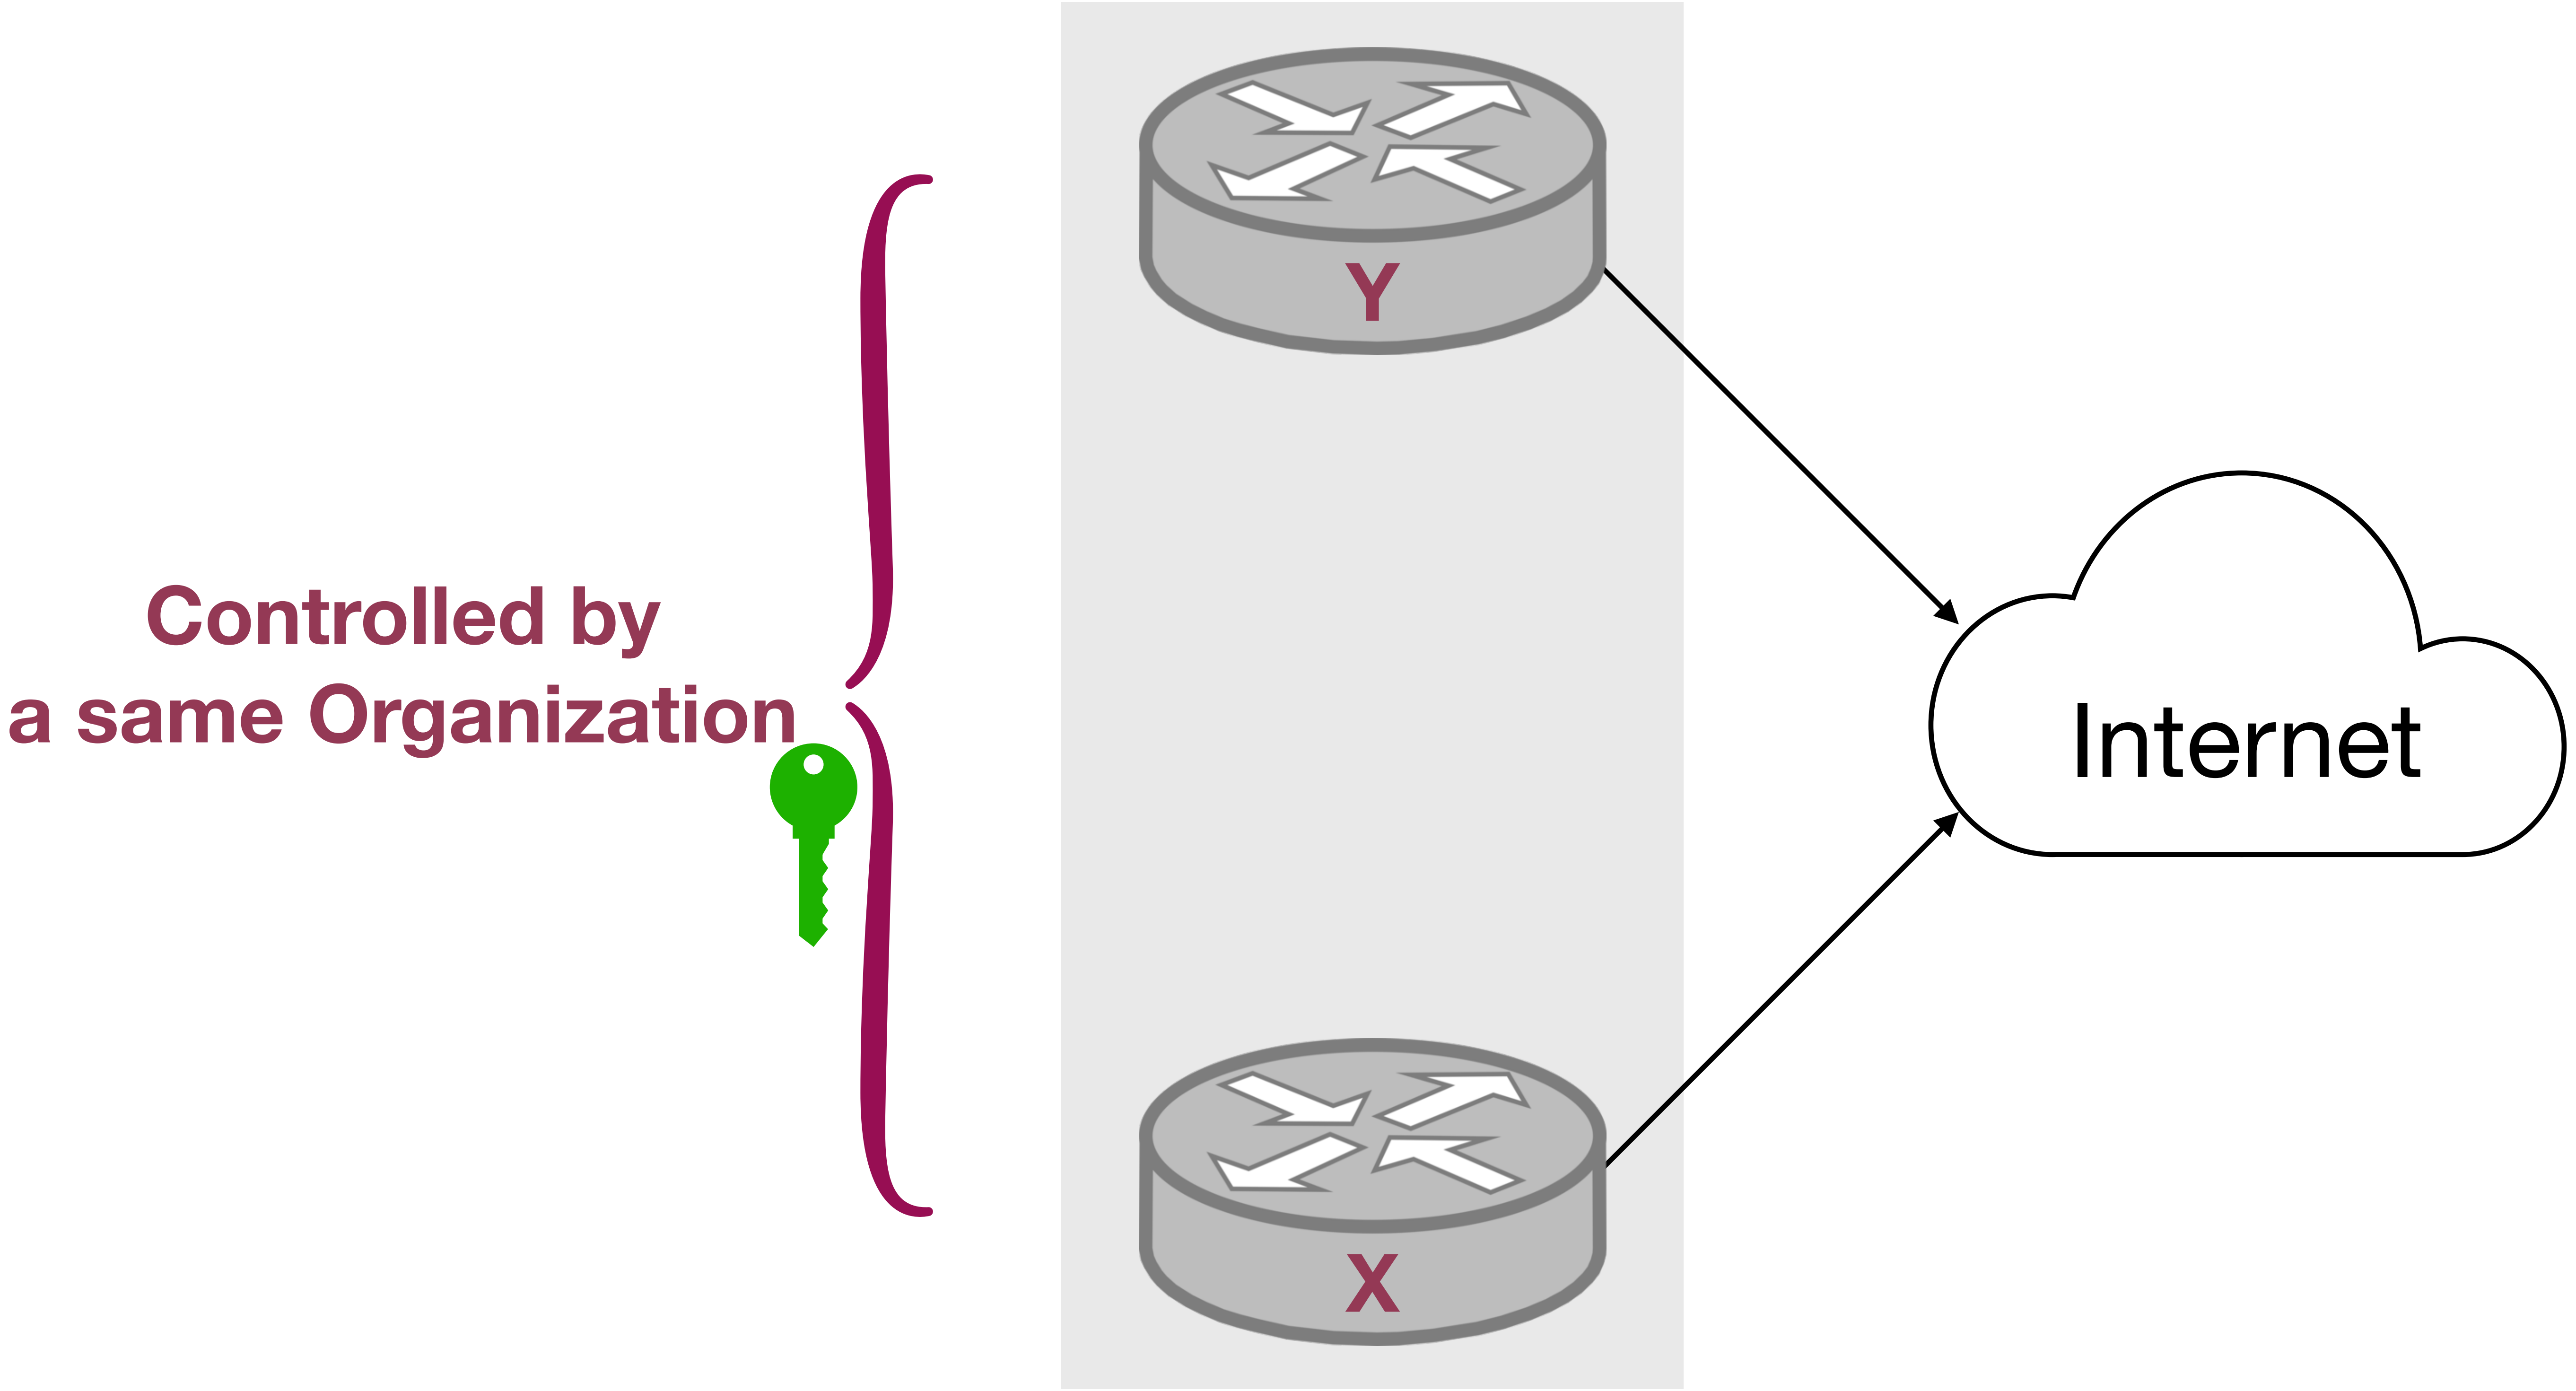

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 **: Access to ROA**

Known Misconfigurations in IP Transit

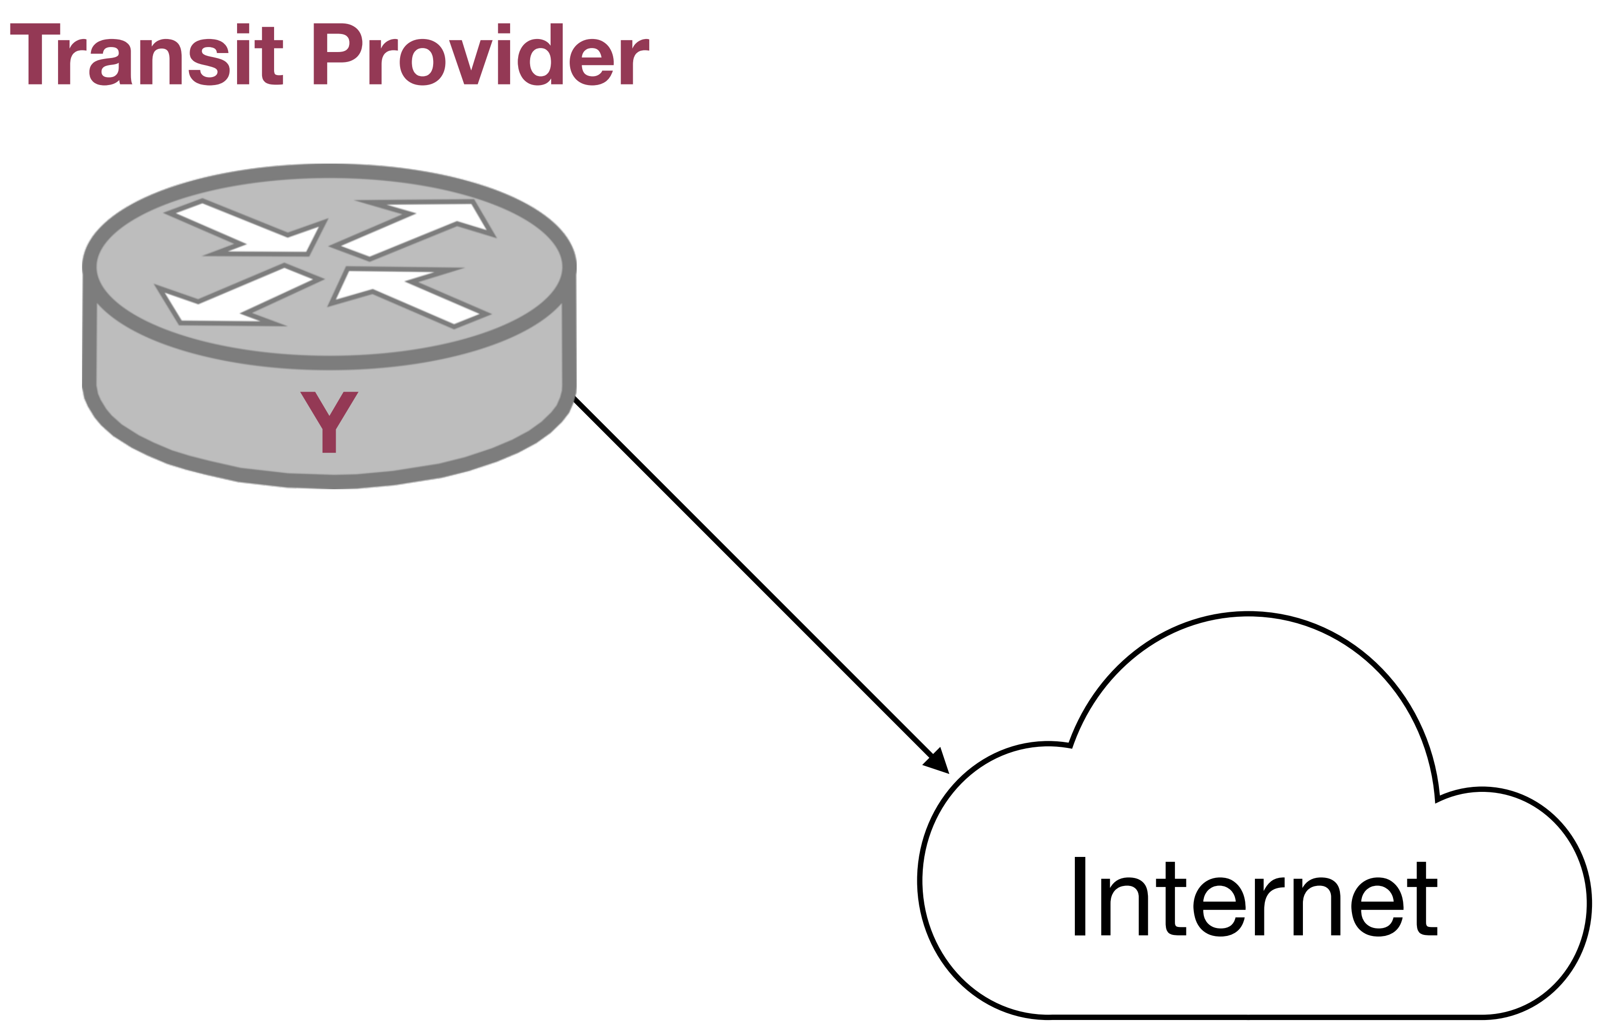

BGP Announcement:

129.21.0.0/16, AS Y

ROA:

129.21.0.0/16-20, AS Y

Known Misconfigurations in IP Transit

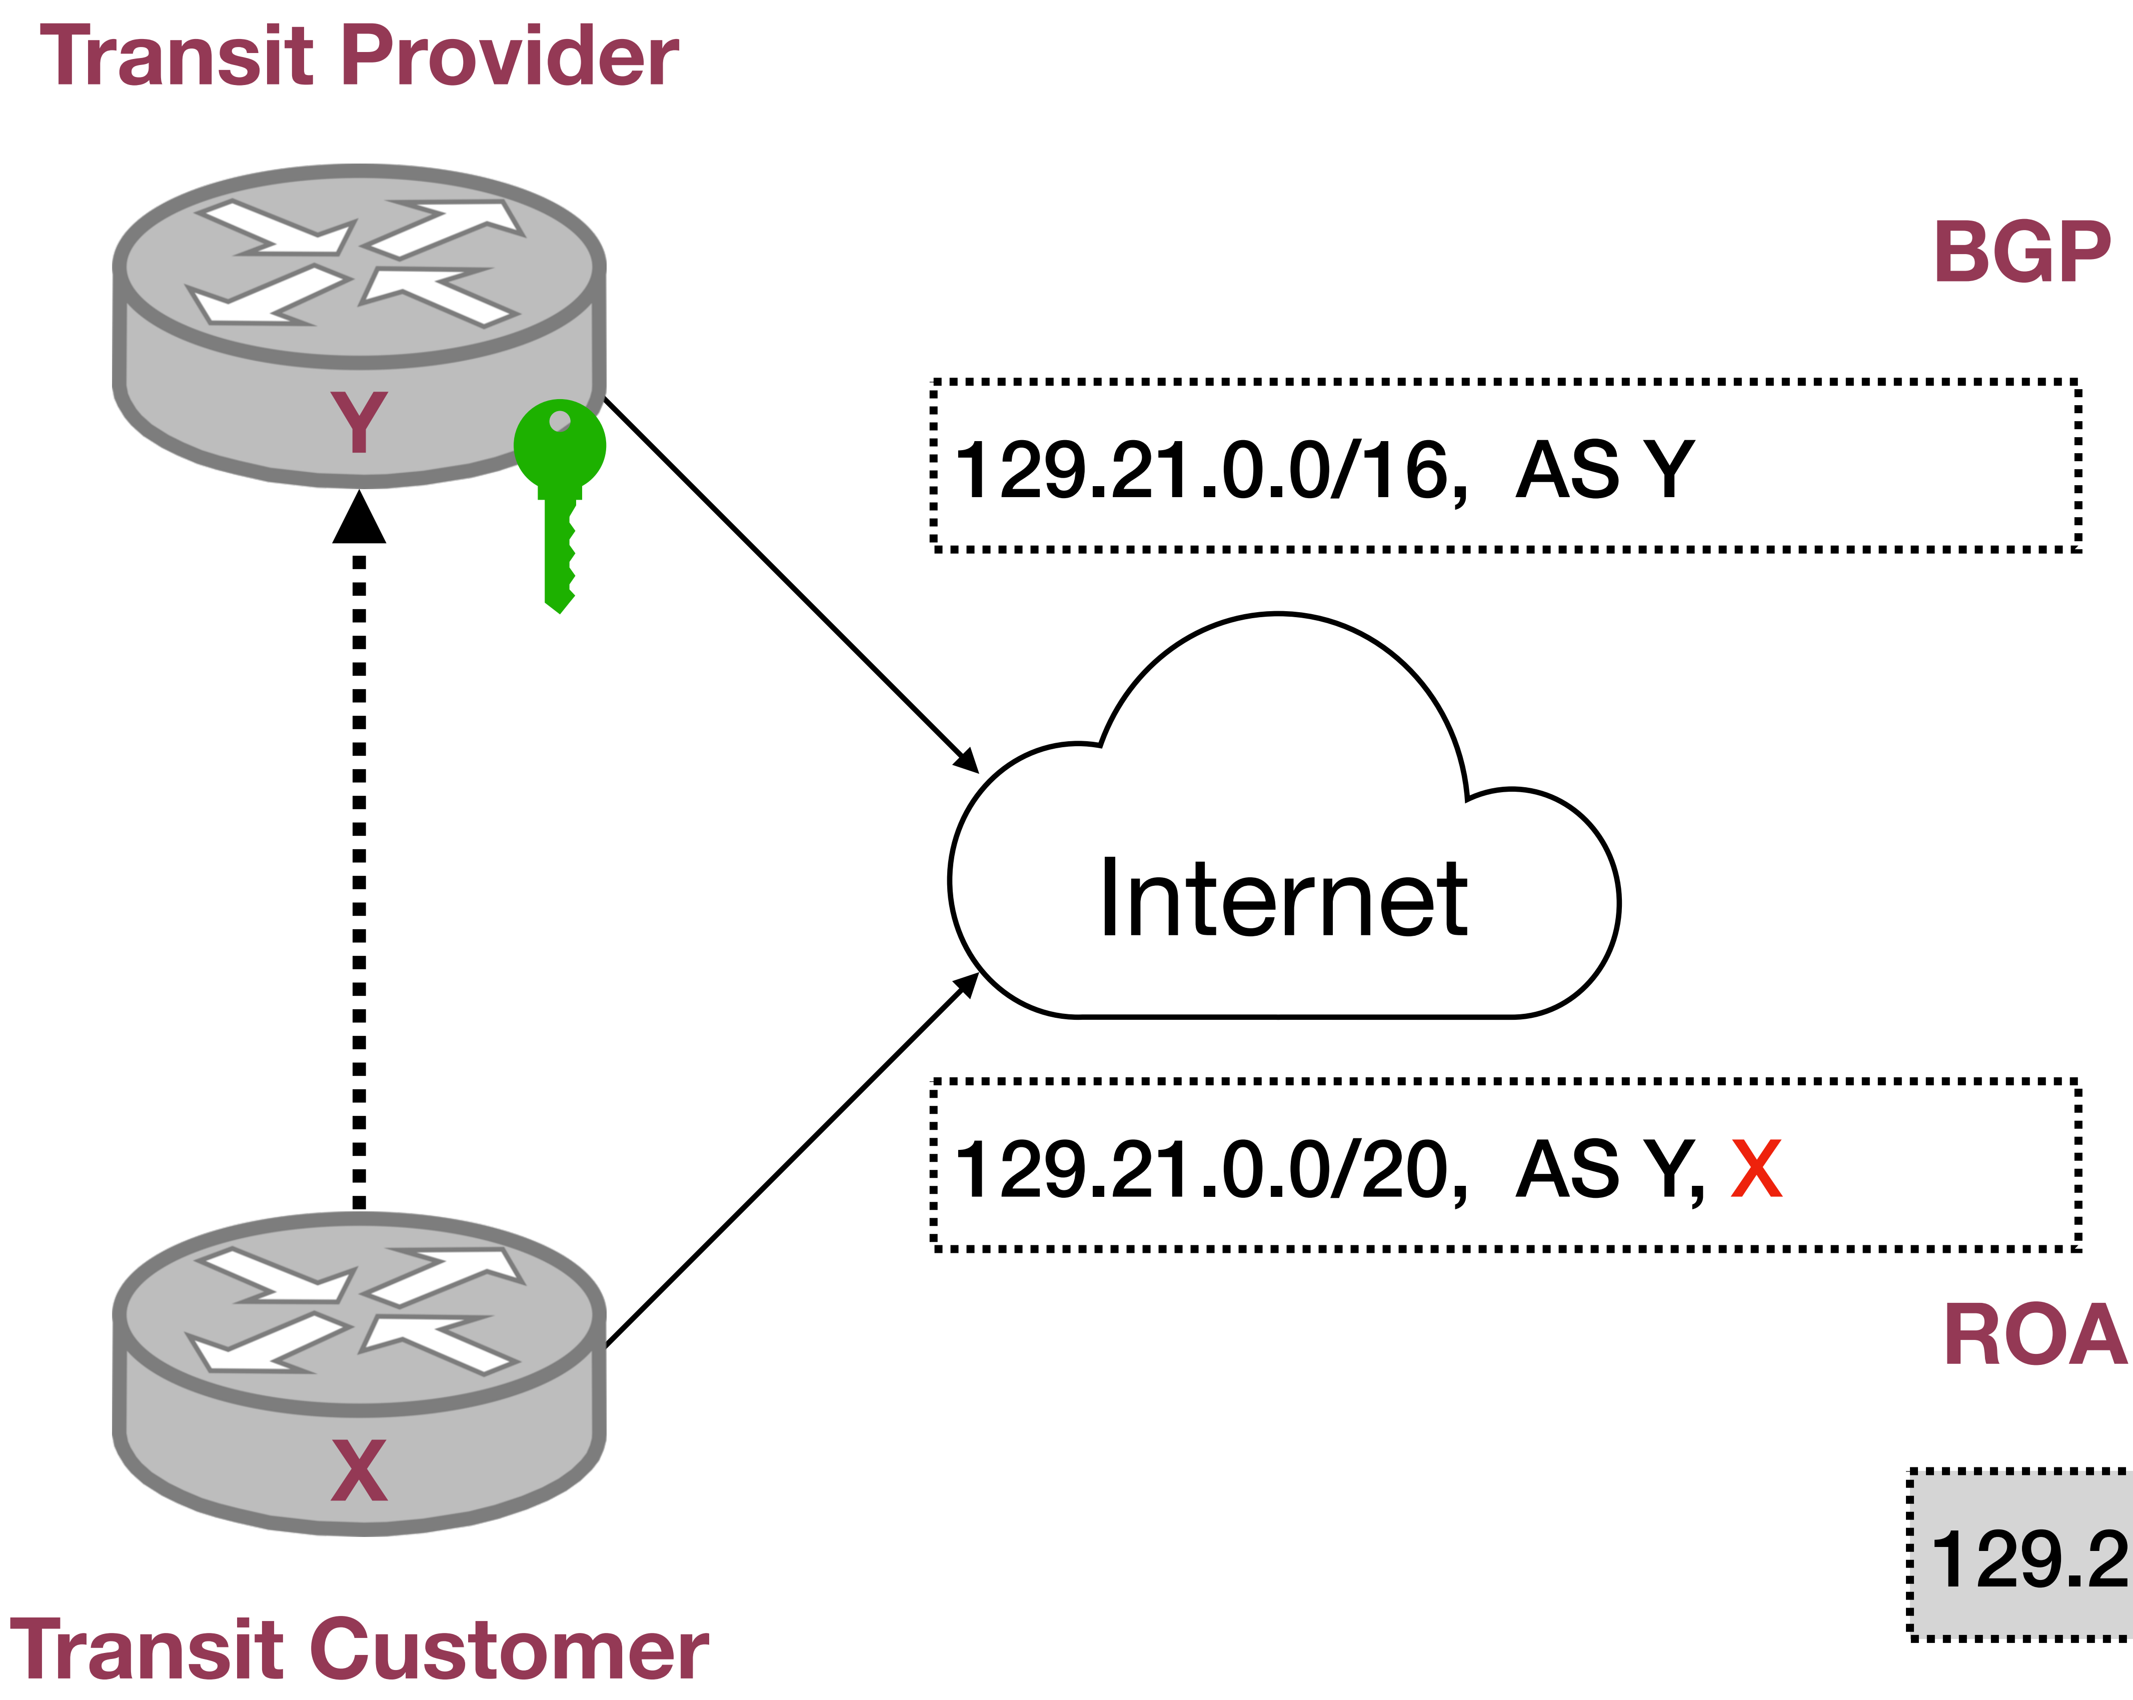

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS Y, X

ROA:

129.21.0.0/16-20, AS Y

🔑 : Access to ROA

Known Misconfigurations in IP Transit

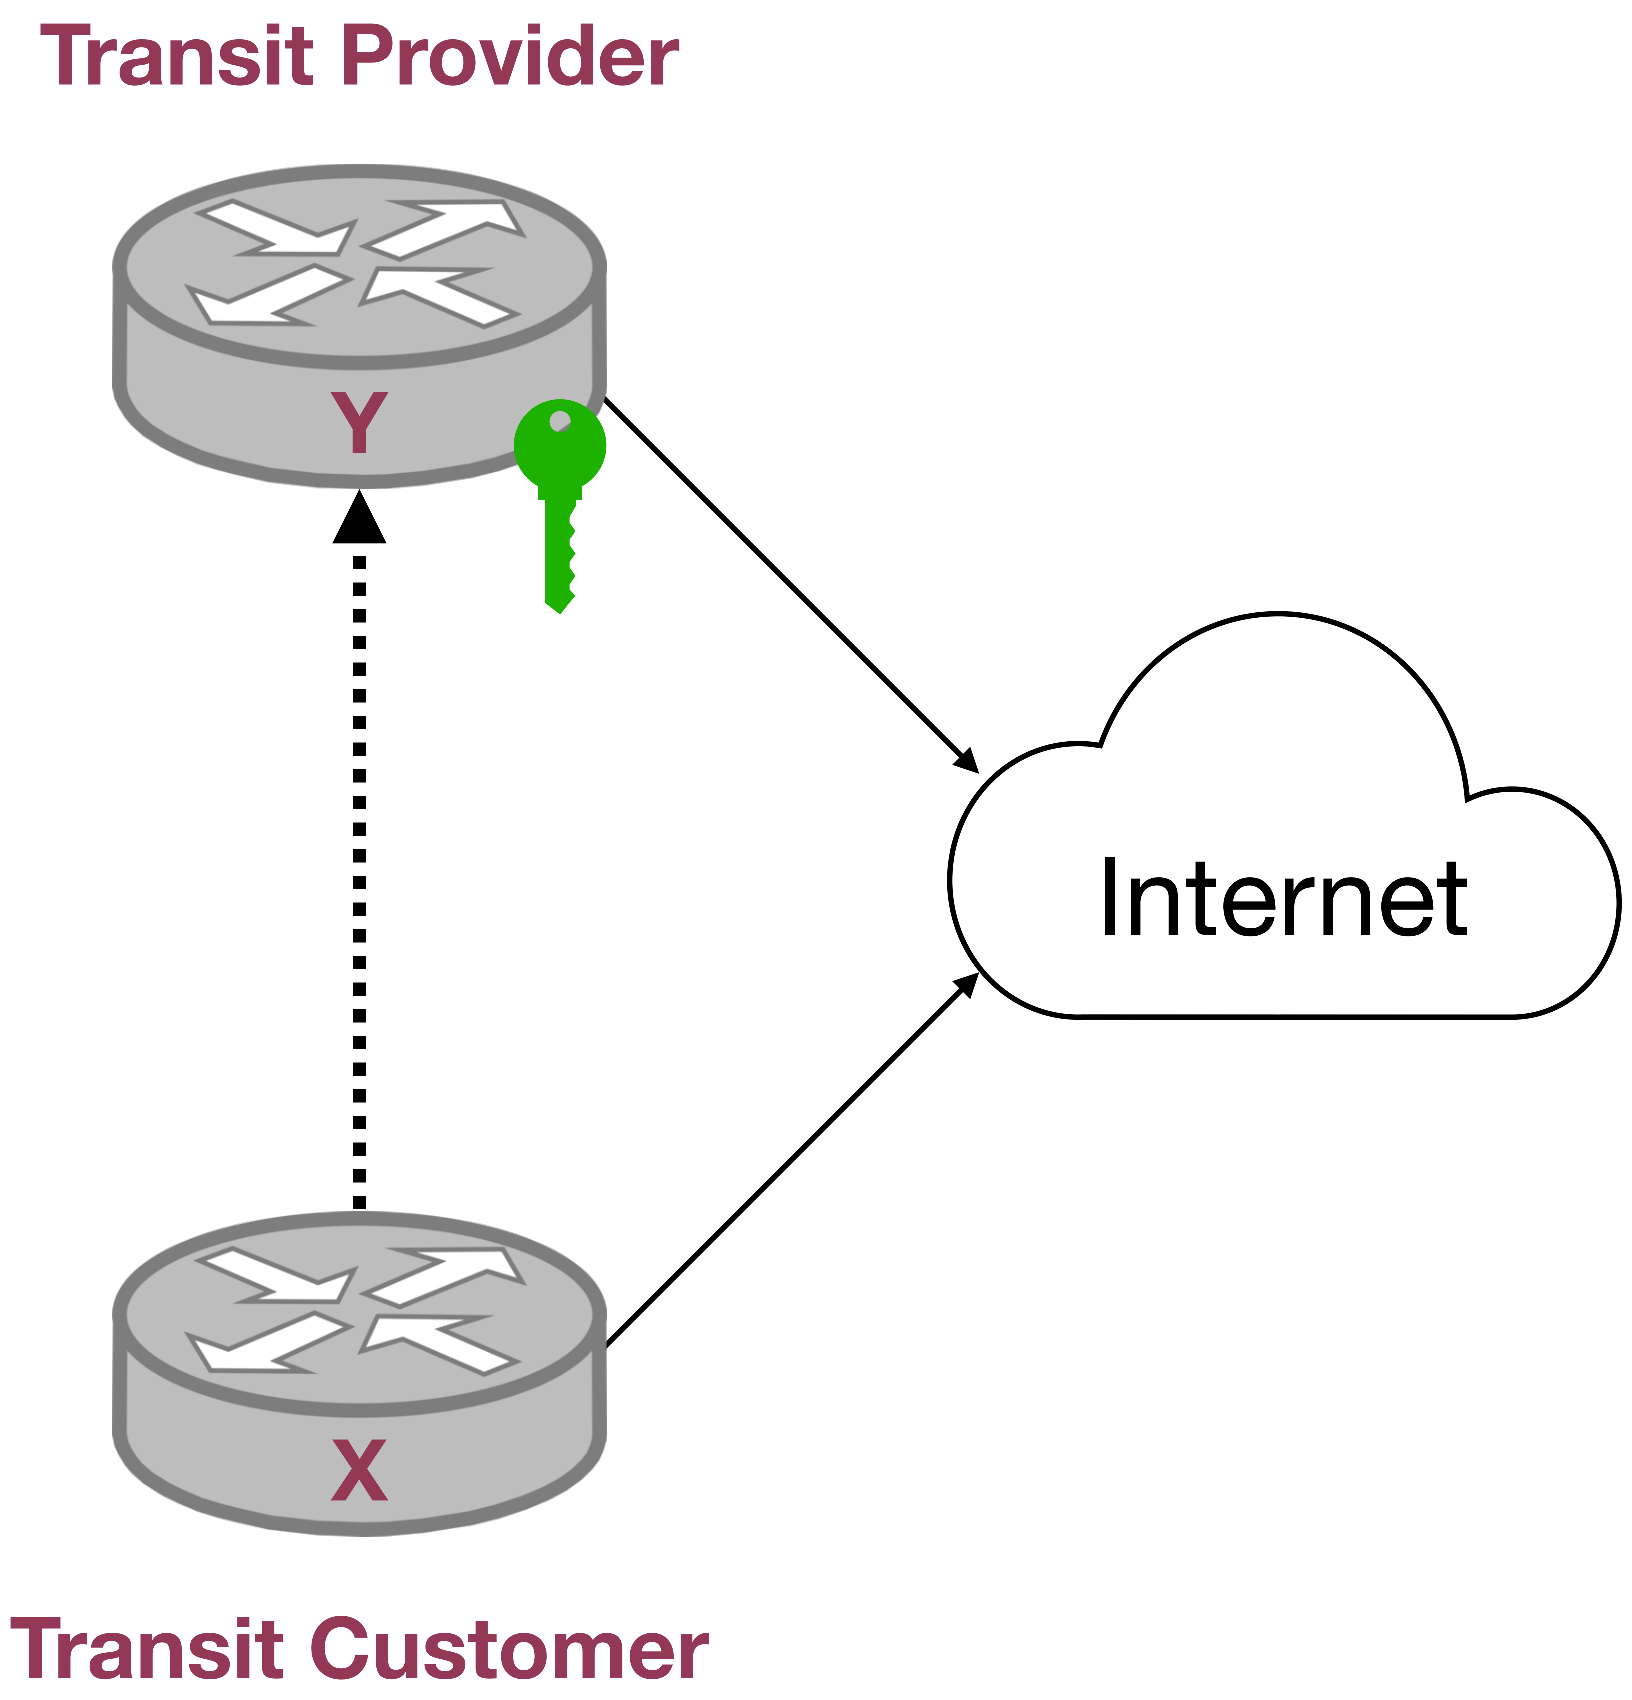

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS Y, X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 **: Access to ROA**

Known Misconfigurations

How to identify:

1. **Transit**: checking BGP path from BGP datasets from Routeviews and RIS
2. **Same org**: using AS to Org mapping database CAIDA AS2Org (We also have one AS2Org mapping dataset using LLM called ASINT!)

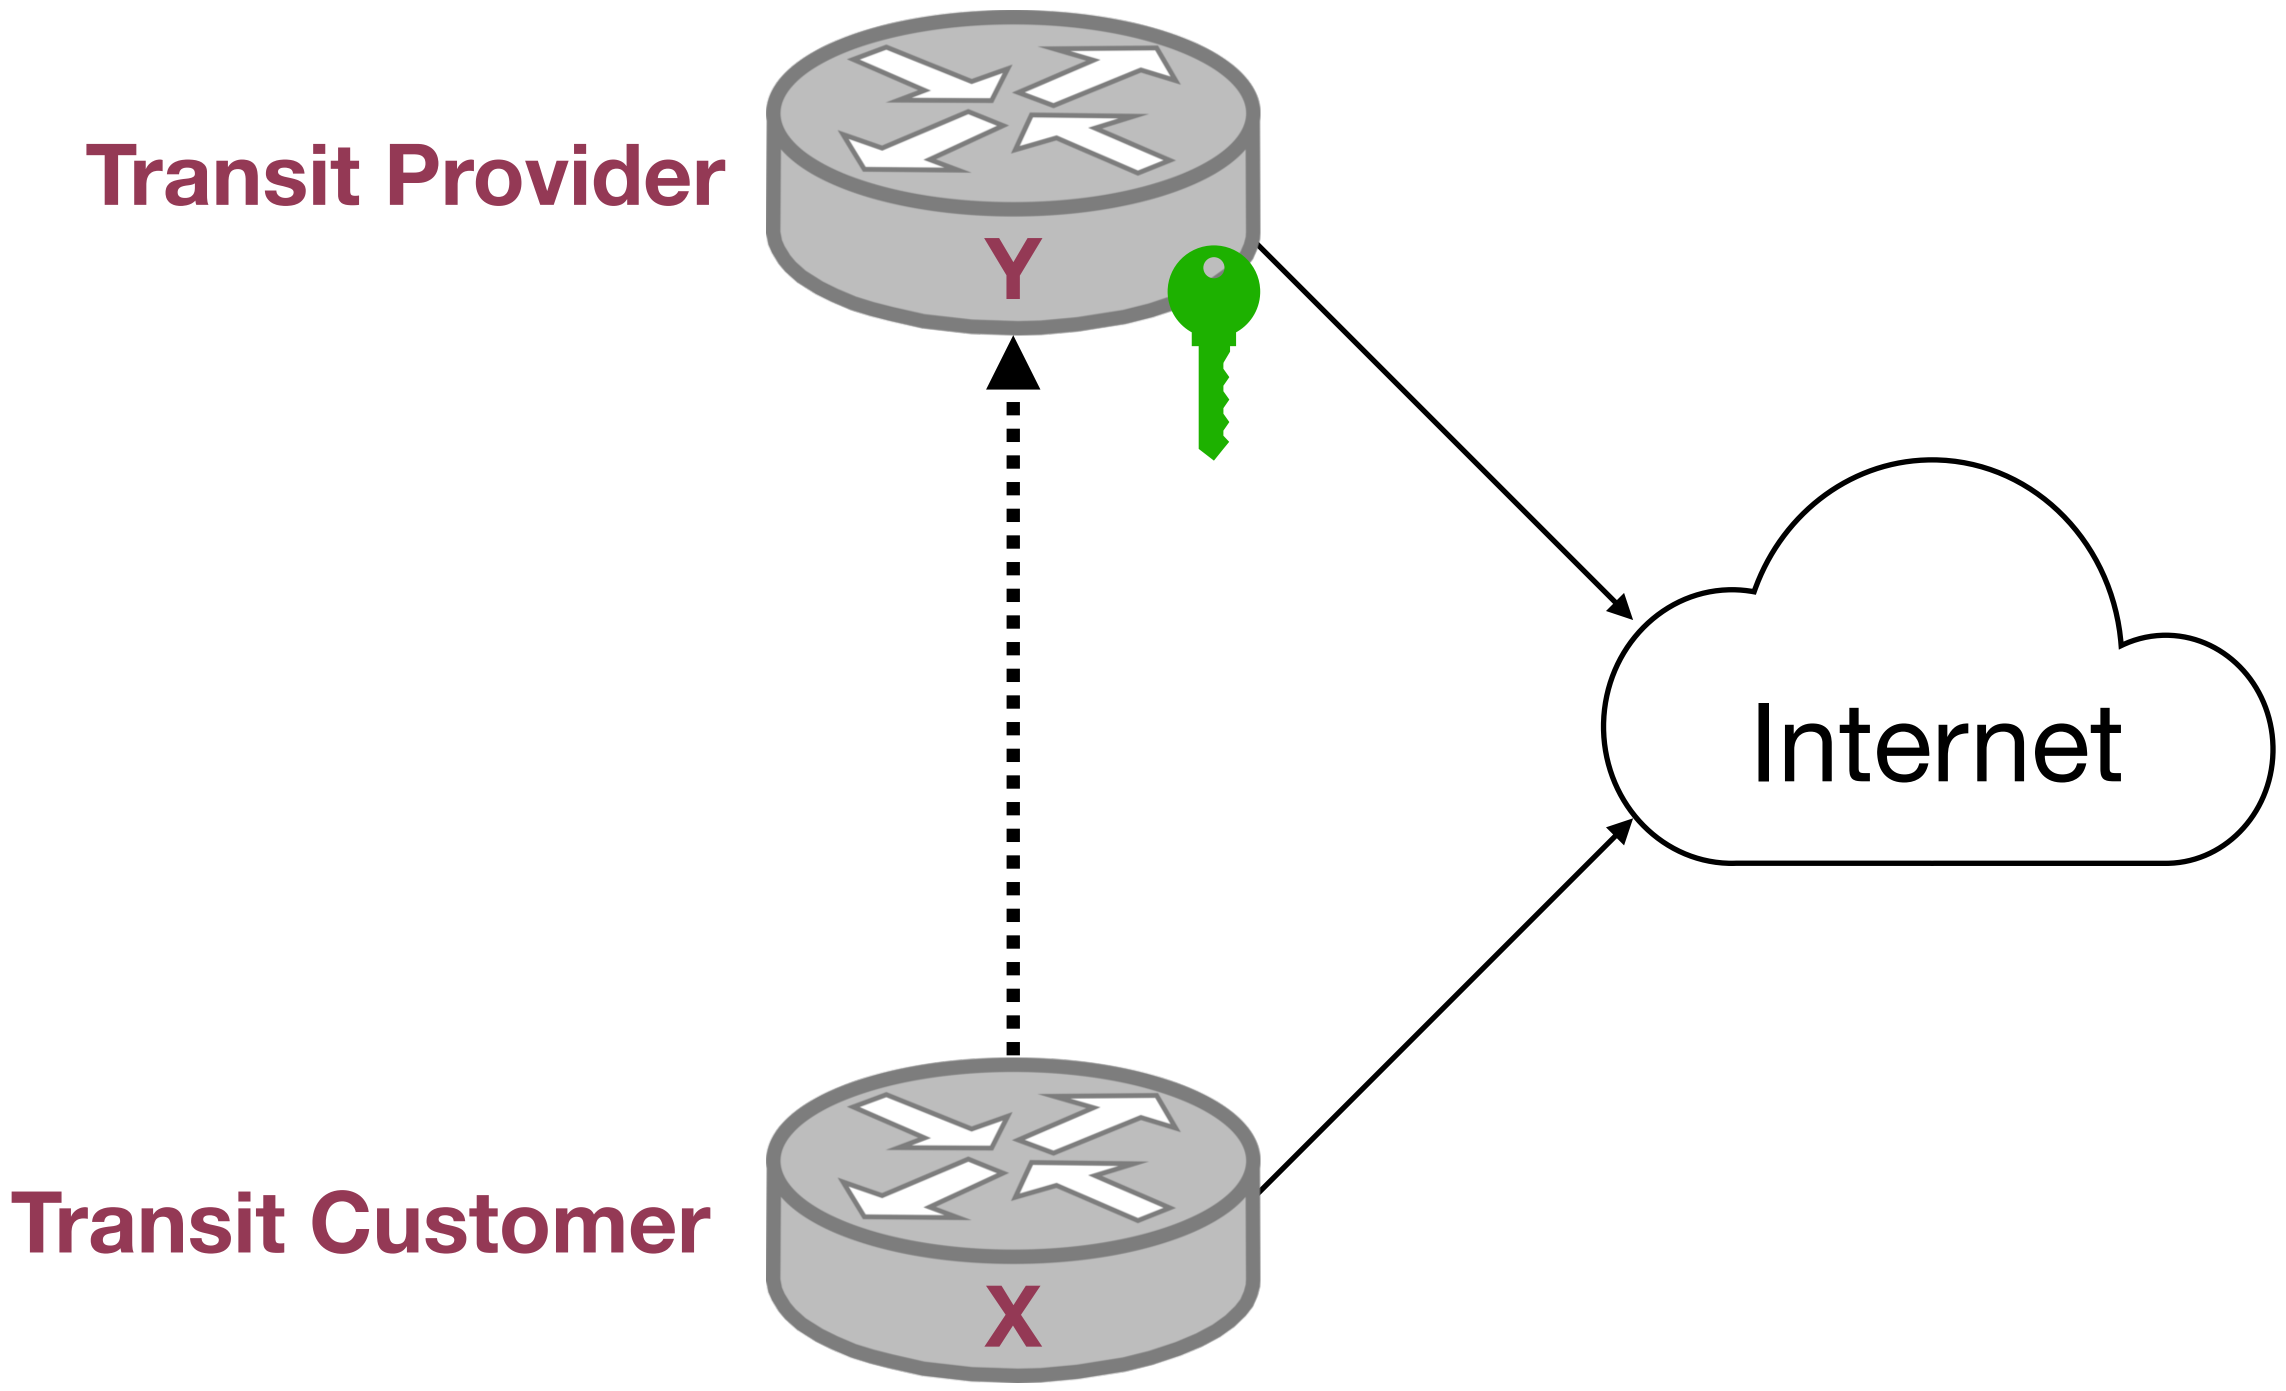

Known Misconfigurations

How to identify:

1. **Transit**: checking BGP path from BGP datasets from Routeviews and RIS

2. **Same org**: using AS to Org mapping database CAIDA AS2Org (We also have one AS2Org mapping dataset using LLM called ASINT!)

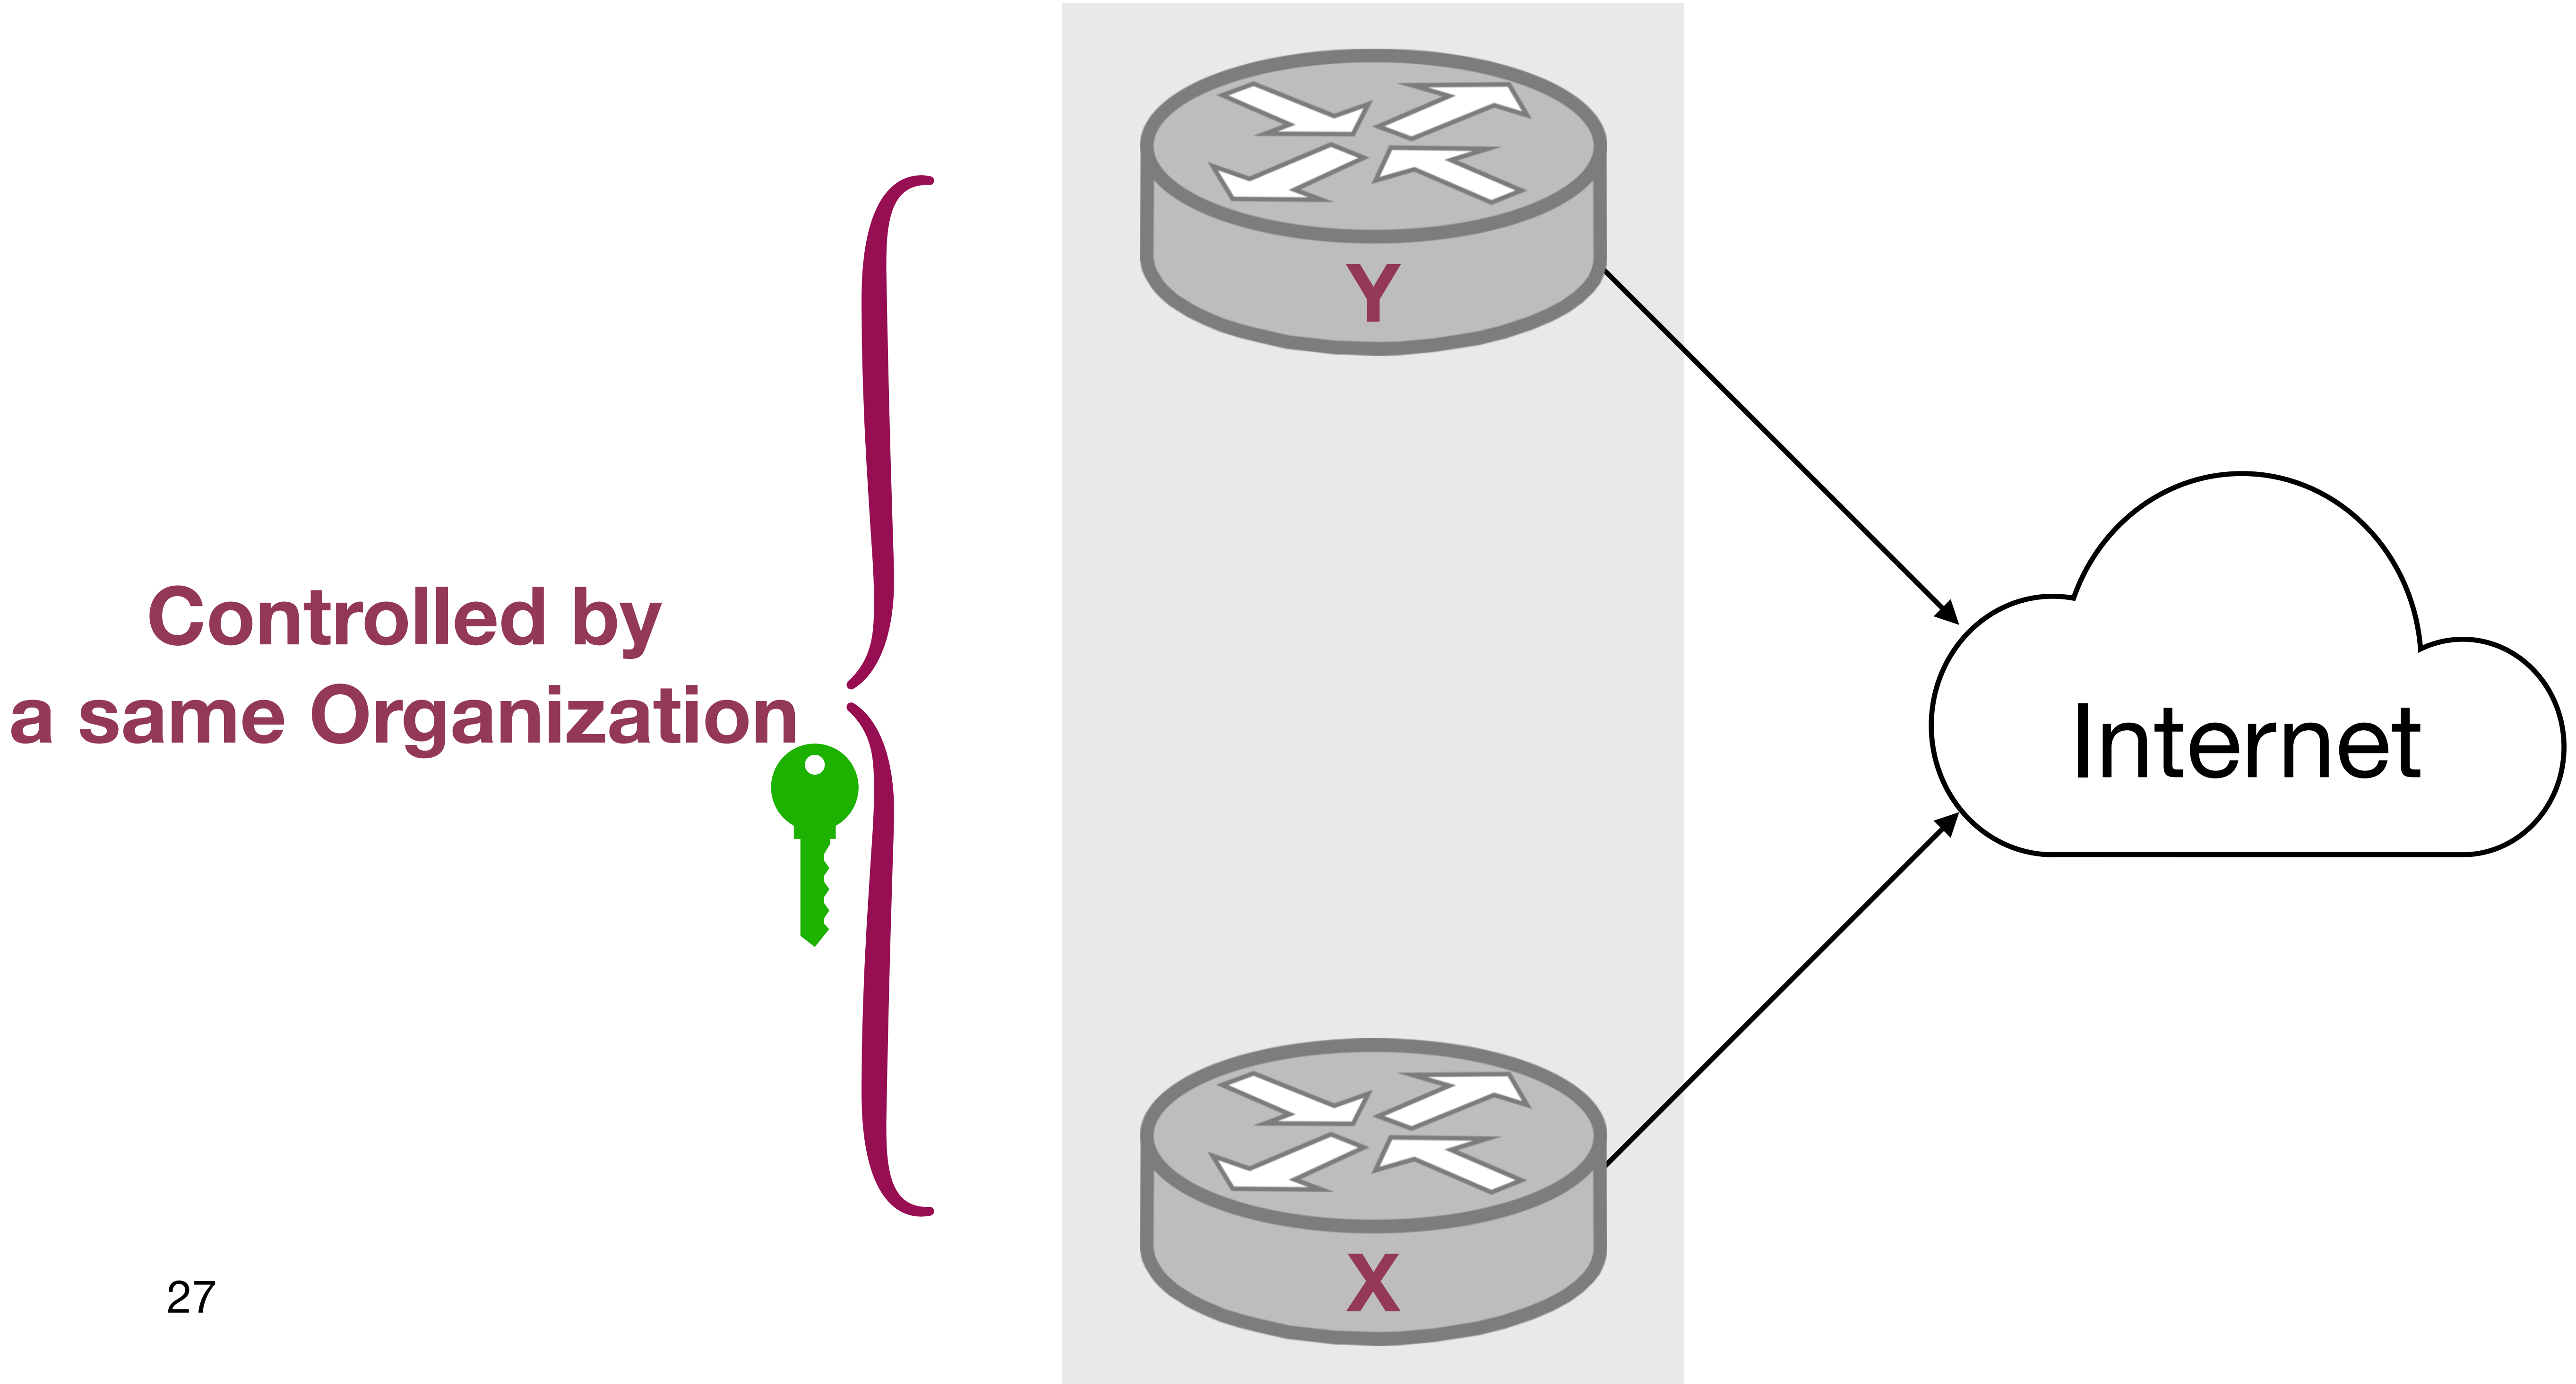

Known Misconfigurations

How to identify [1]:

1. **Transit**: checking BGP path from BGP datasets from RouteViews and RIS
2. **Same org**: using AS to Org mapping database CAIDA AS2Org (We also have one AS2Org mapping dataset using LLM called ASINT!)

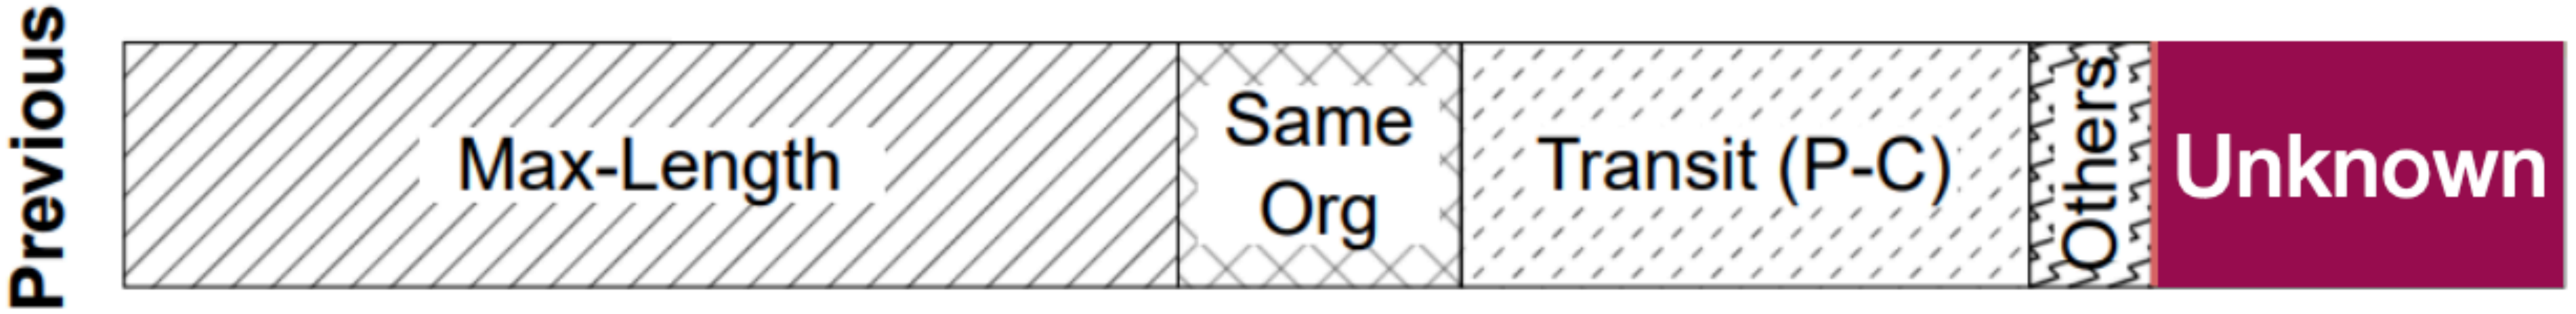

Are they all hijacks?

[1] IMC'19 RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins

Known Misconfigurations

How to identify [1]:

1. **Transit**: checking BGP path from BGP datasets from RouteViews and RIS
2. **Same org**: using AS to Org mapping database CAIDA AS2Org (We also have one AS2Org mapping dataset using LLM called ASINT!)

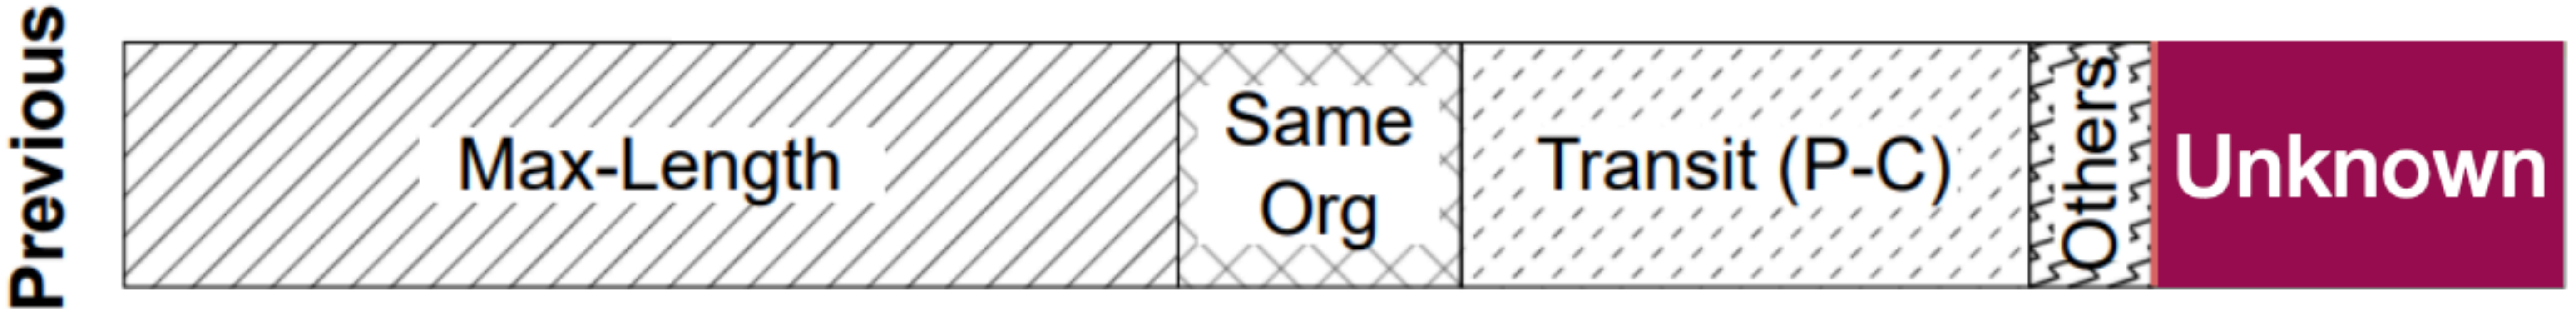

Are they all hijacks?

[1] IMC’19 RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins

Hidden Relationships: Transit + Org

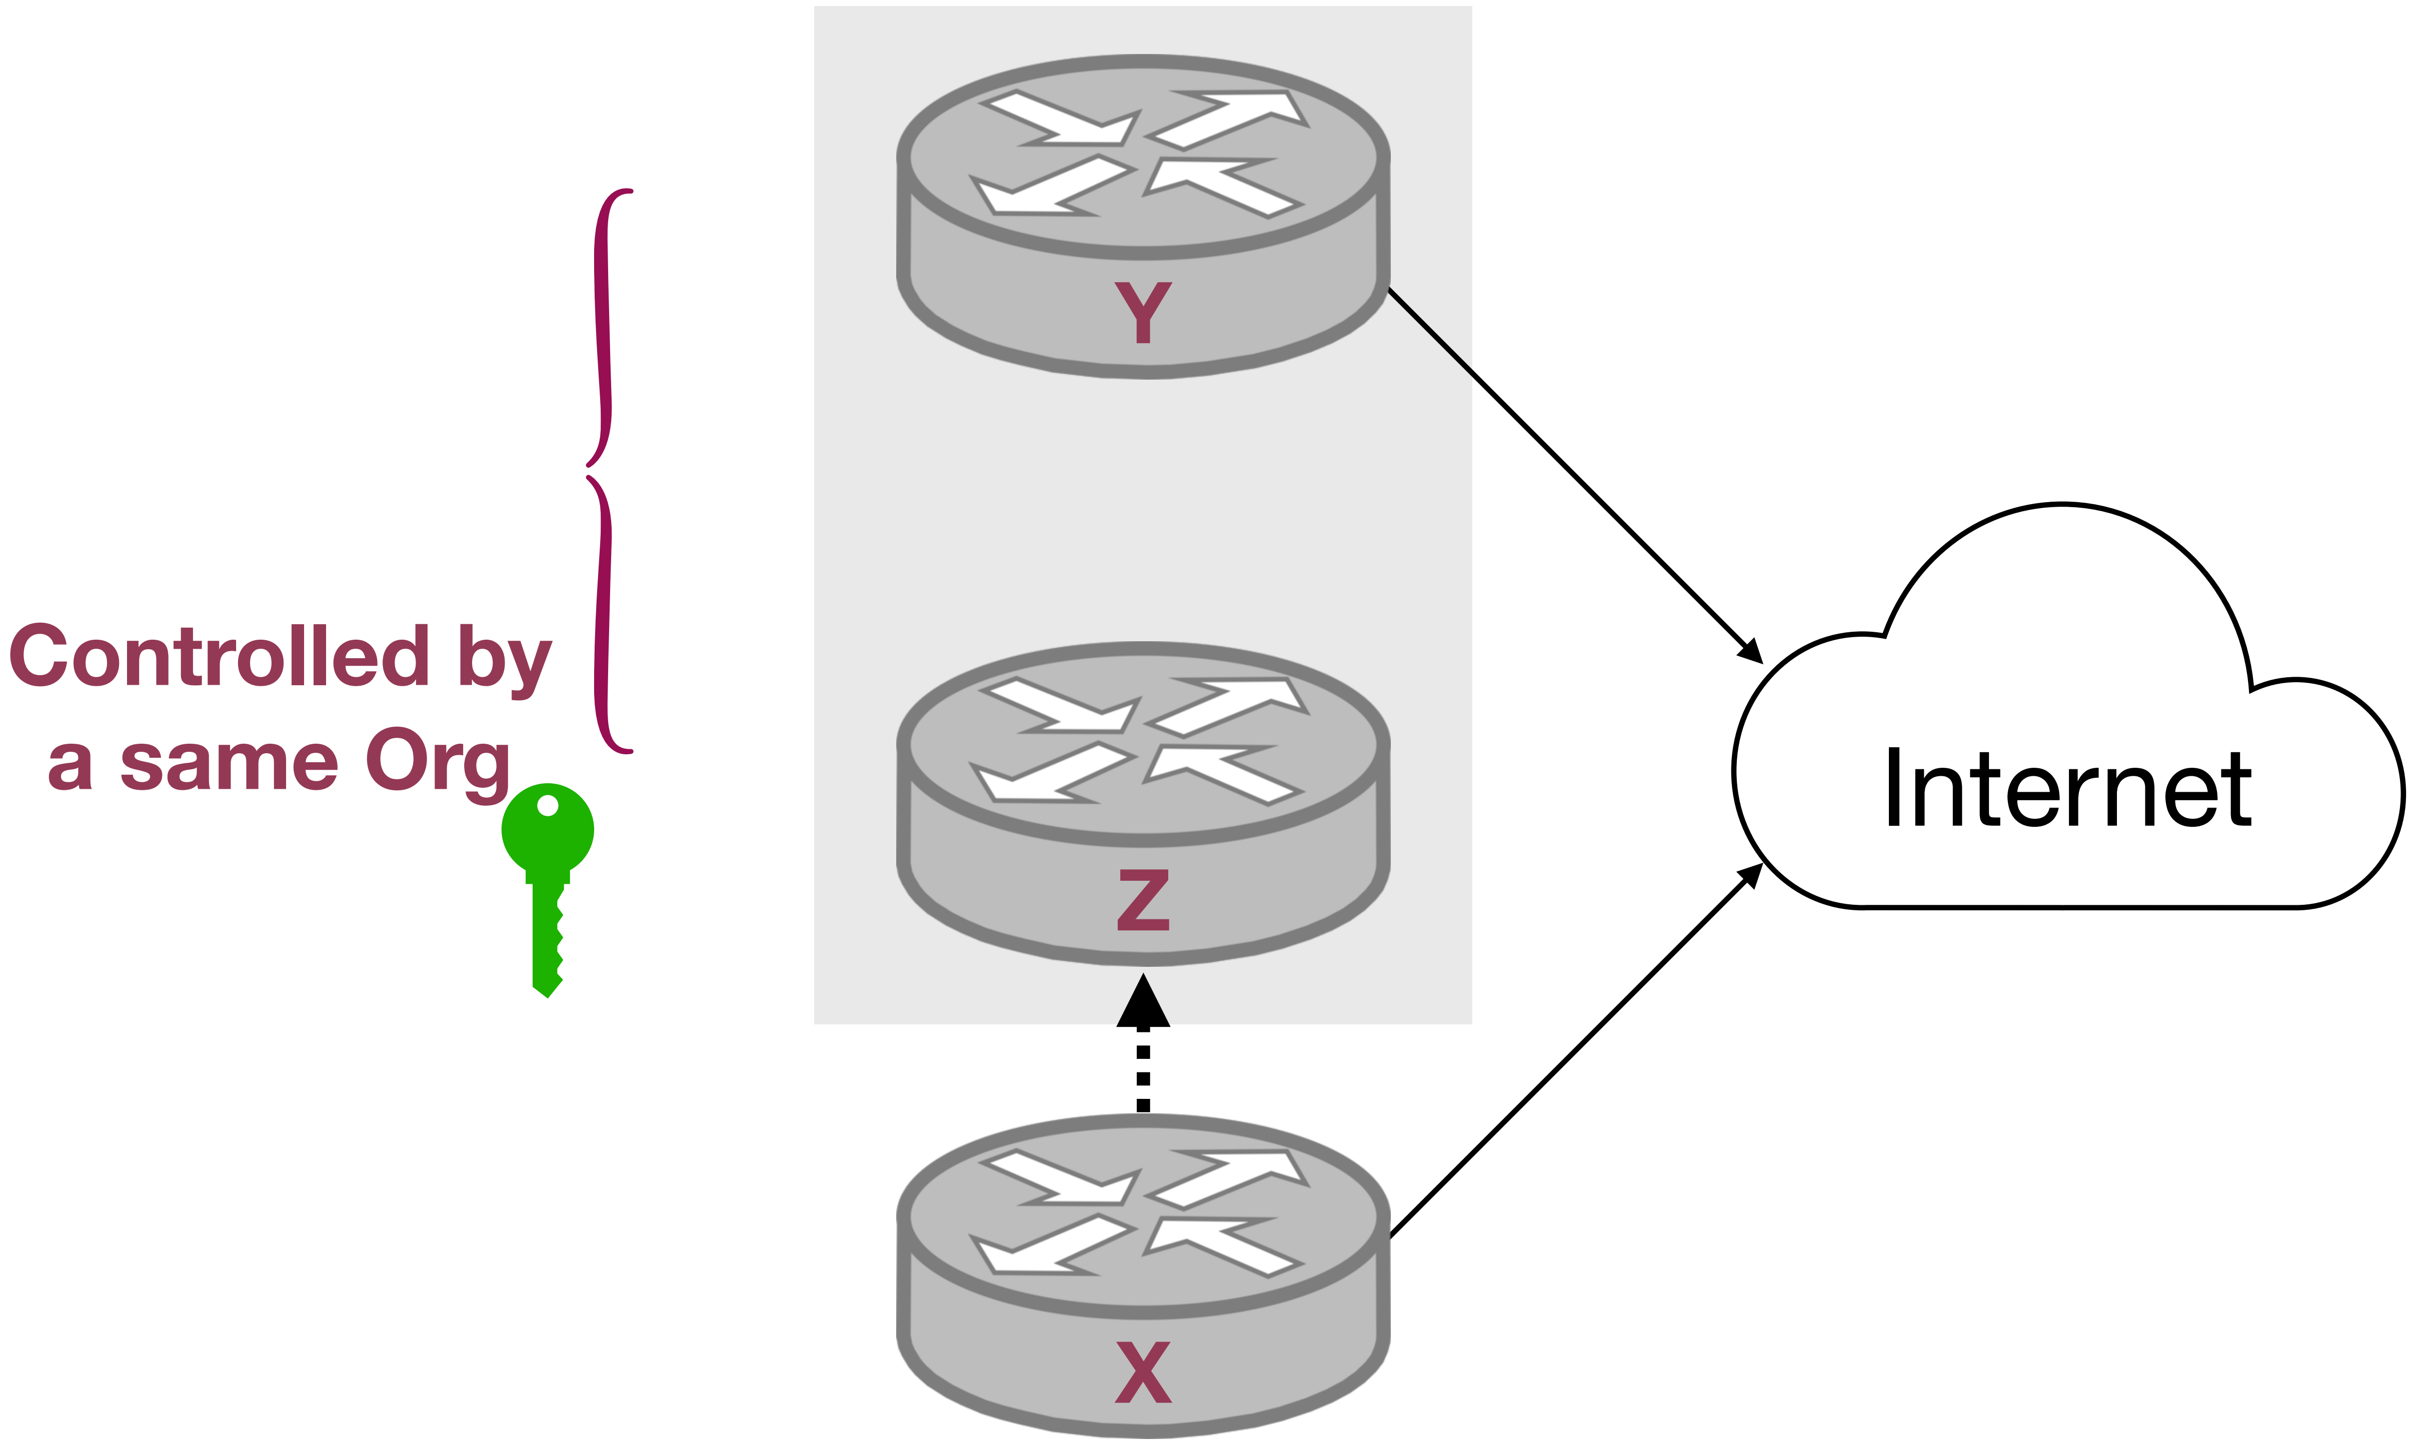

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS Z X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 **: Access to ROA**

Hidden Relationships

- What if we can’t find any relationship between X and Y in BGP or AS2ORG mapping?

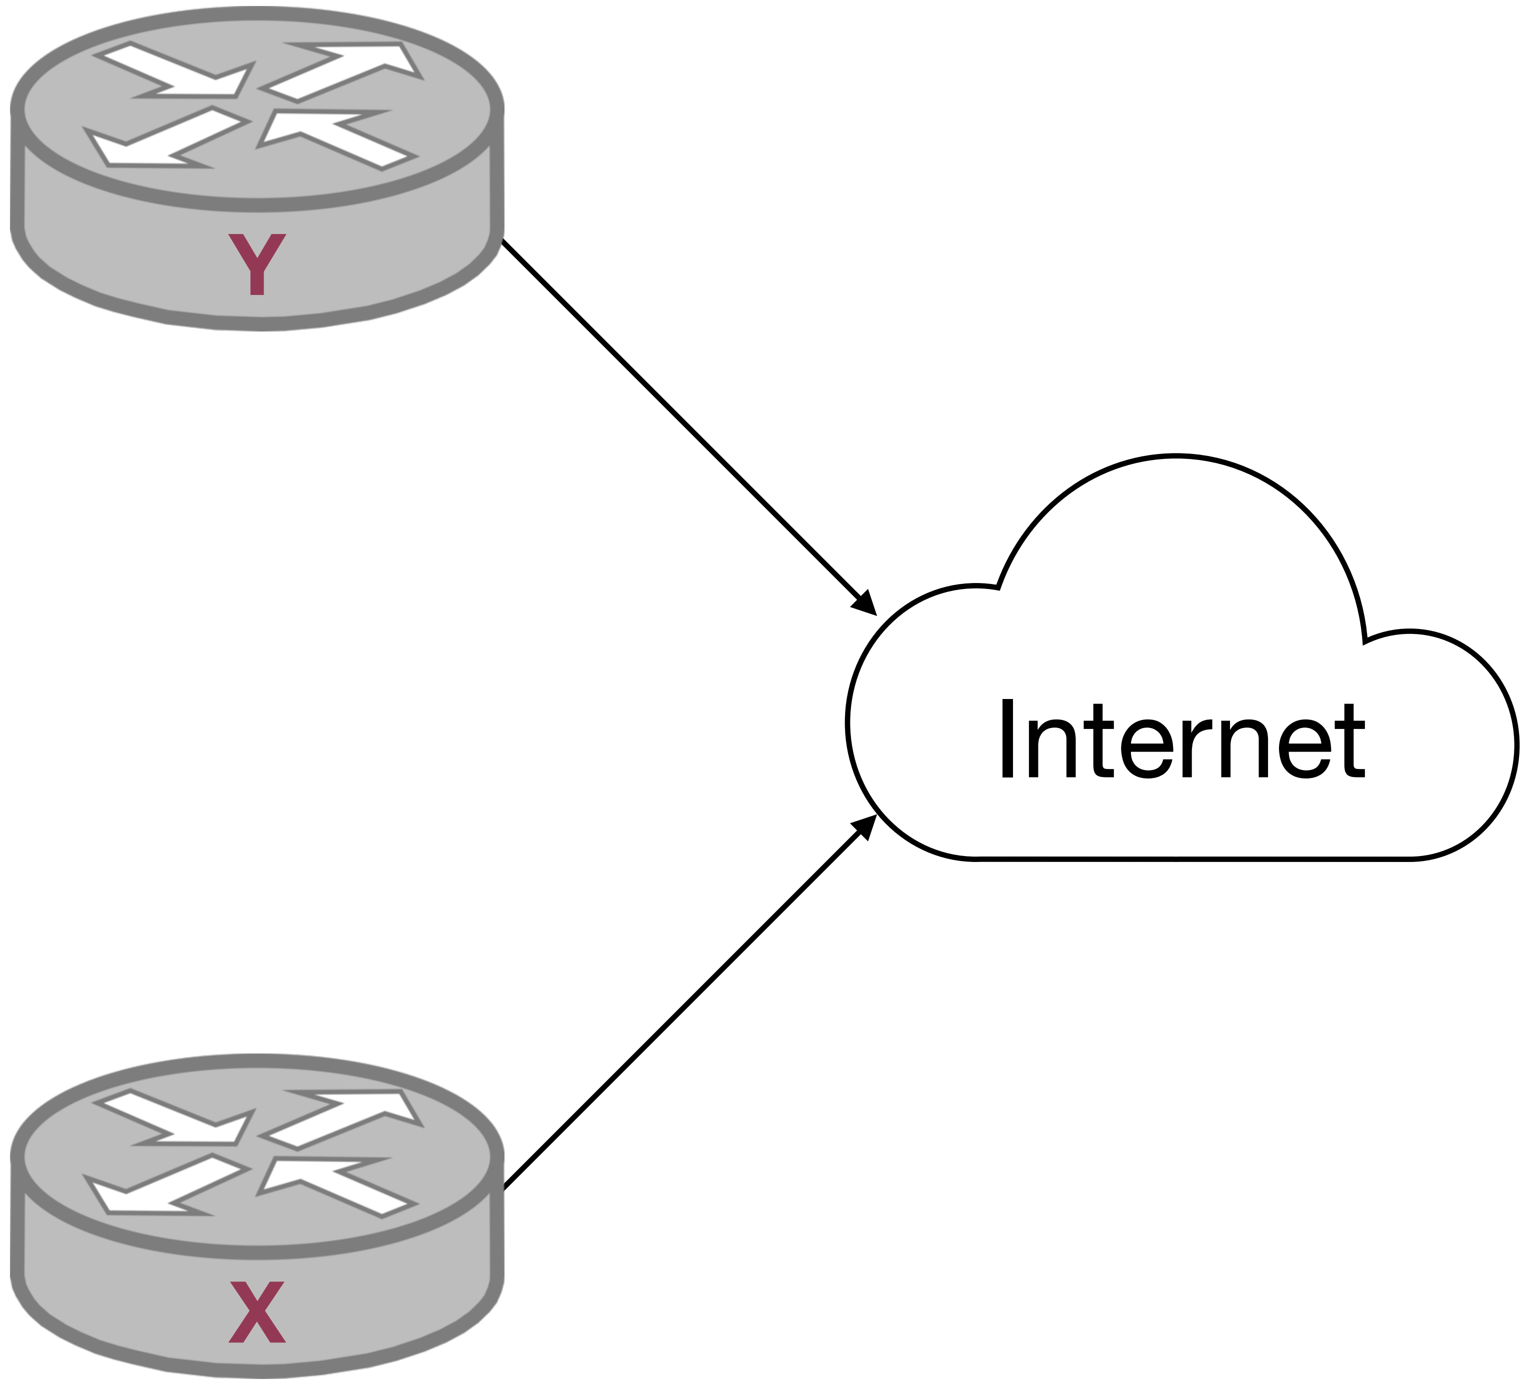

Hidden Relationships: Transit + Tunnel

Transit services may not show up in the BGP path!

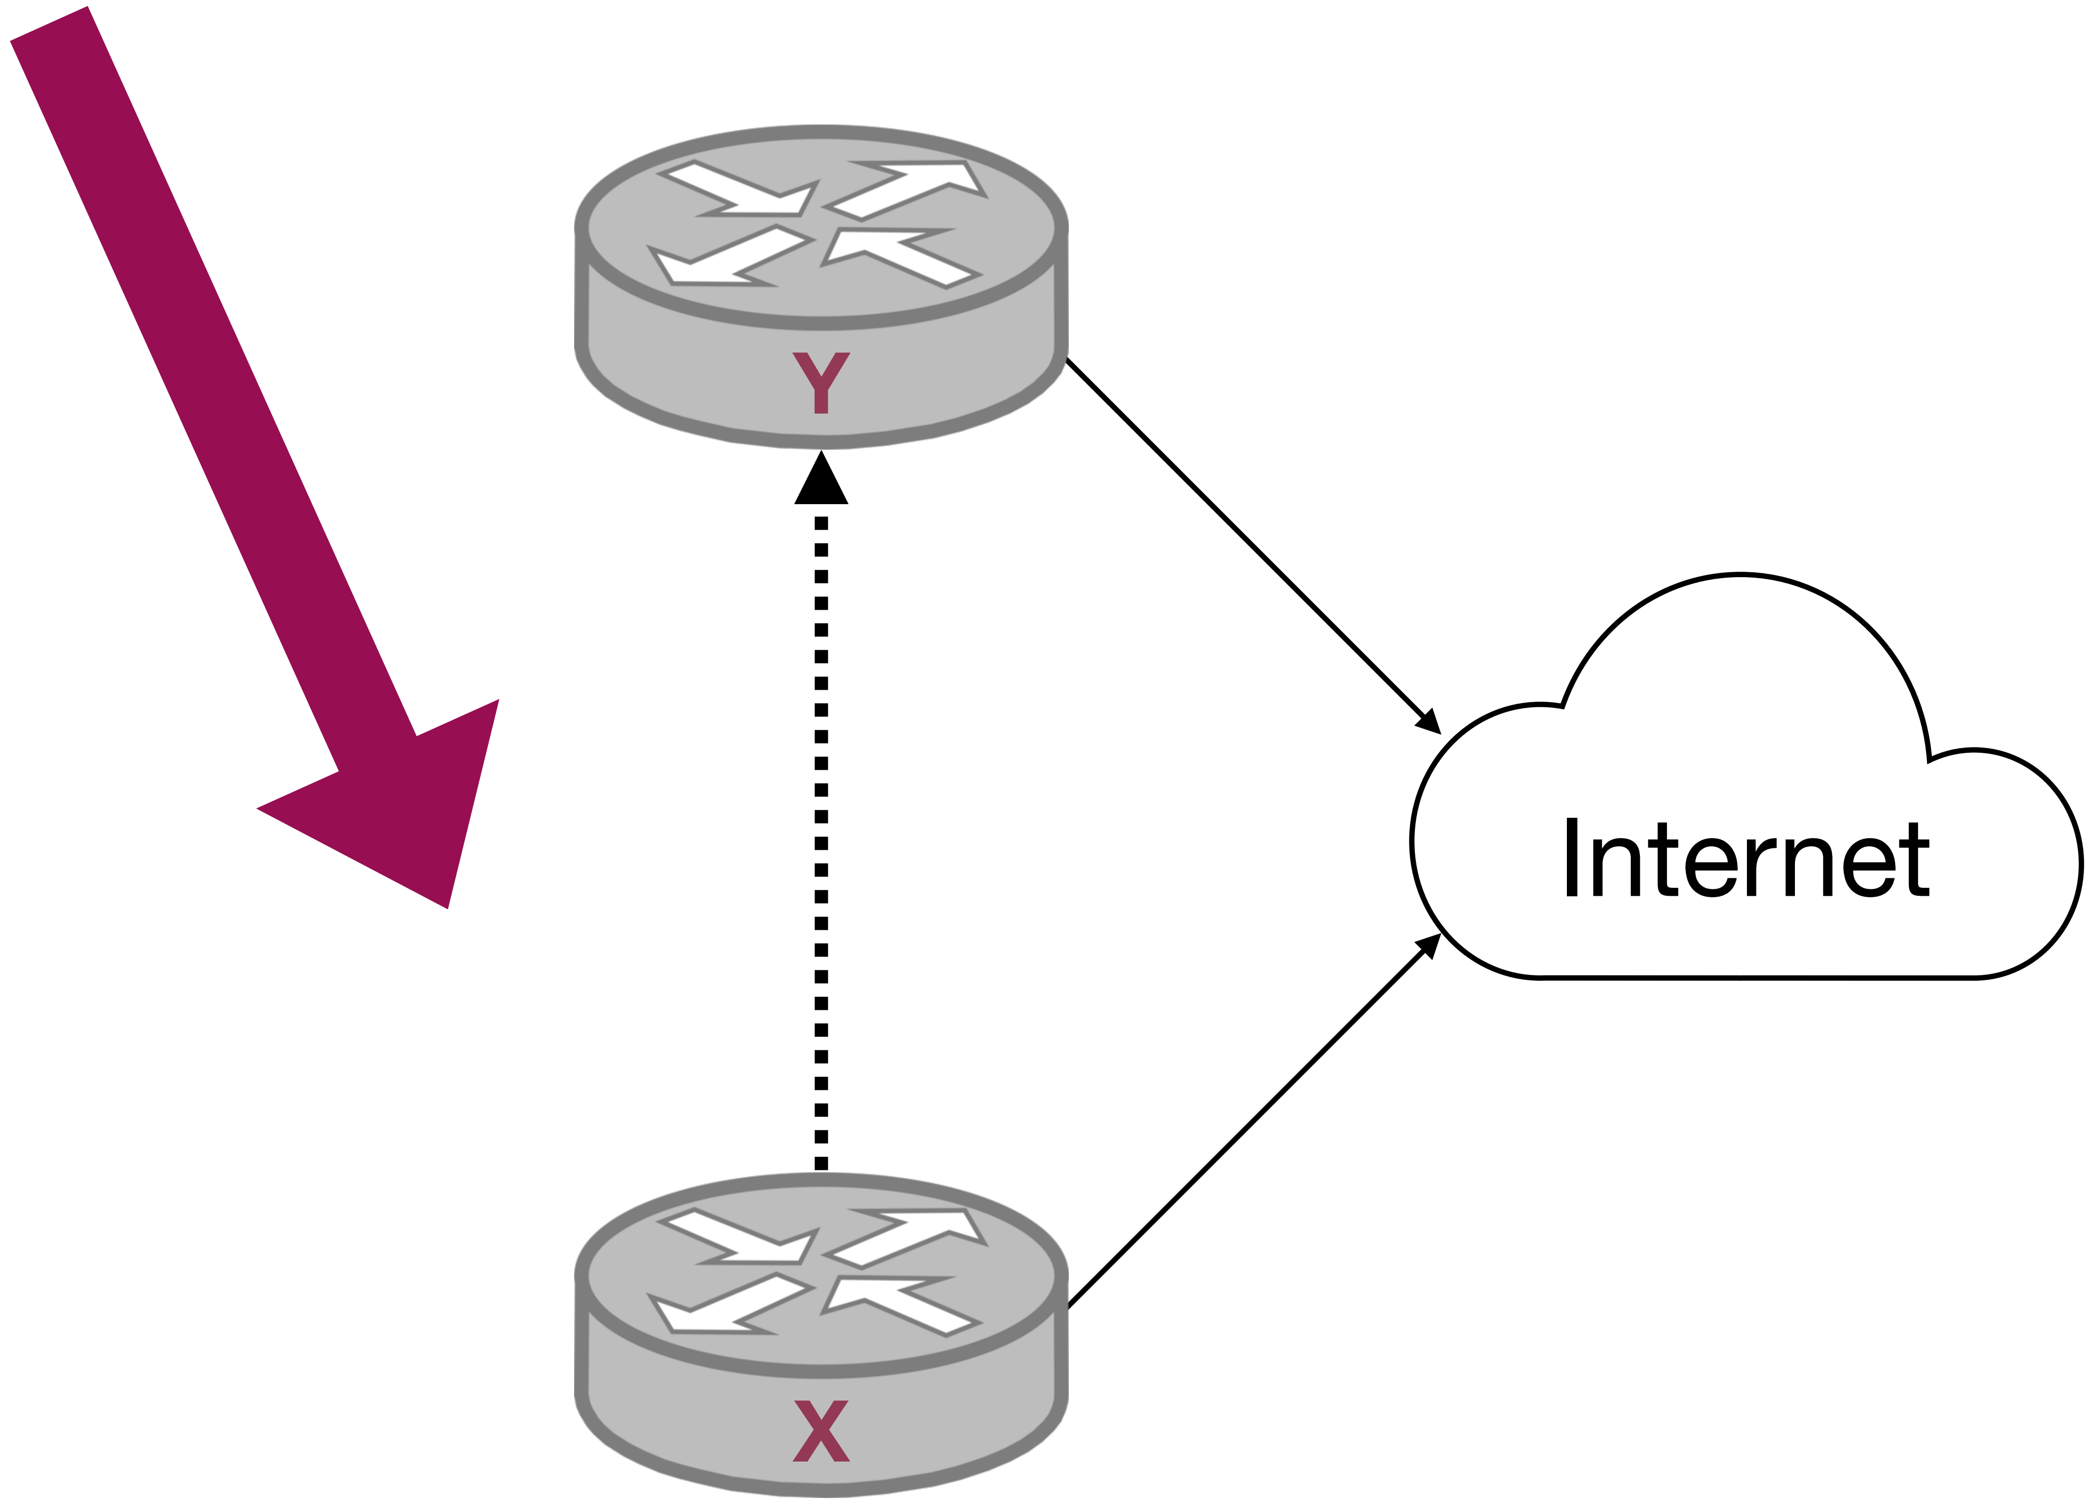

Hidden Relationships: Transit + Tunnel

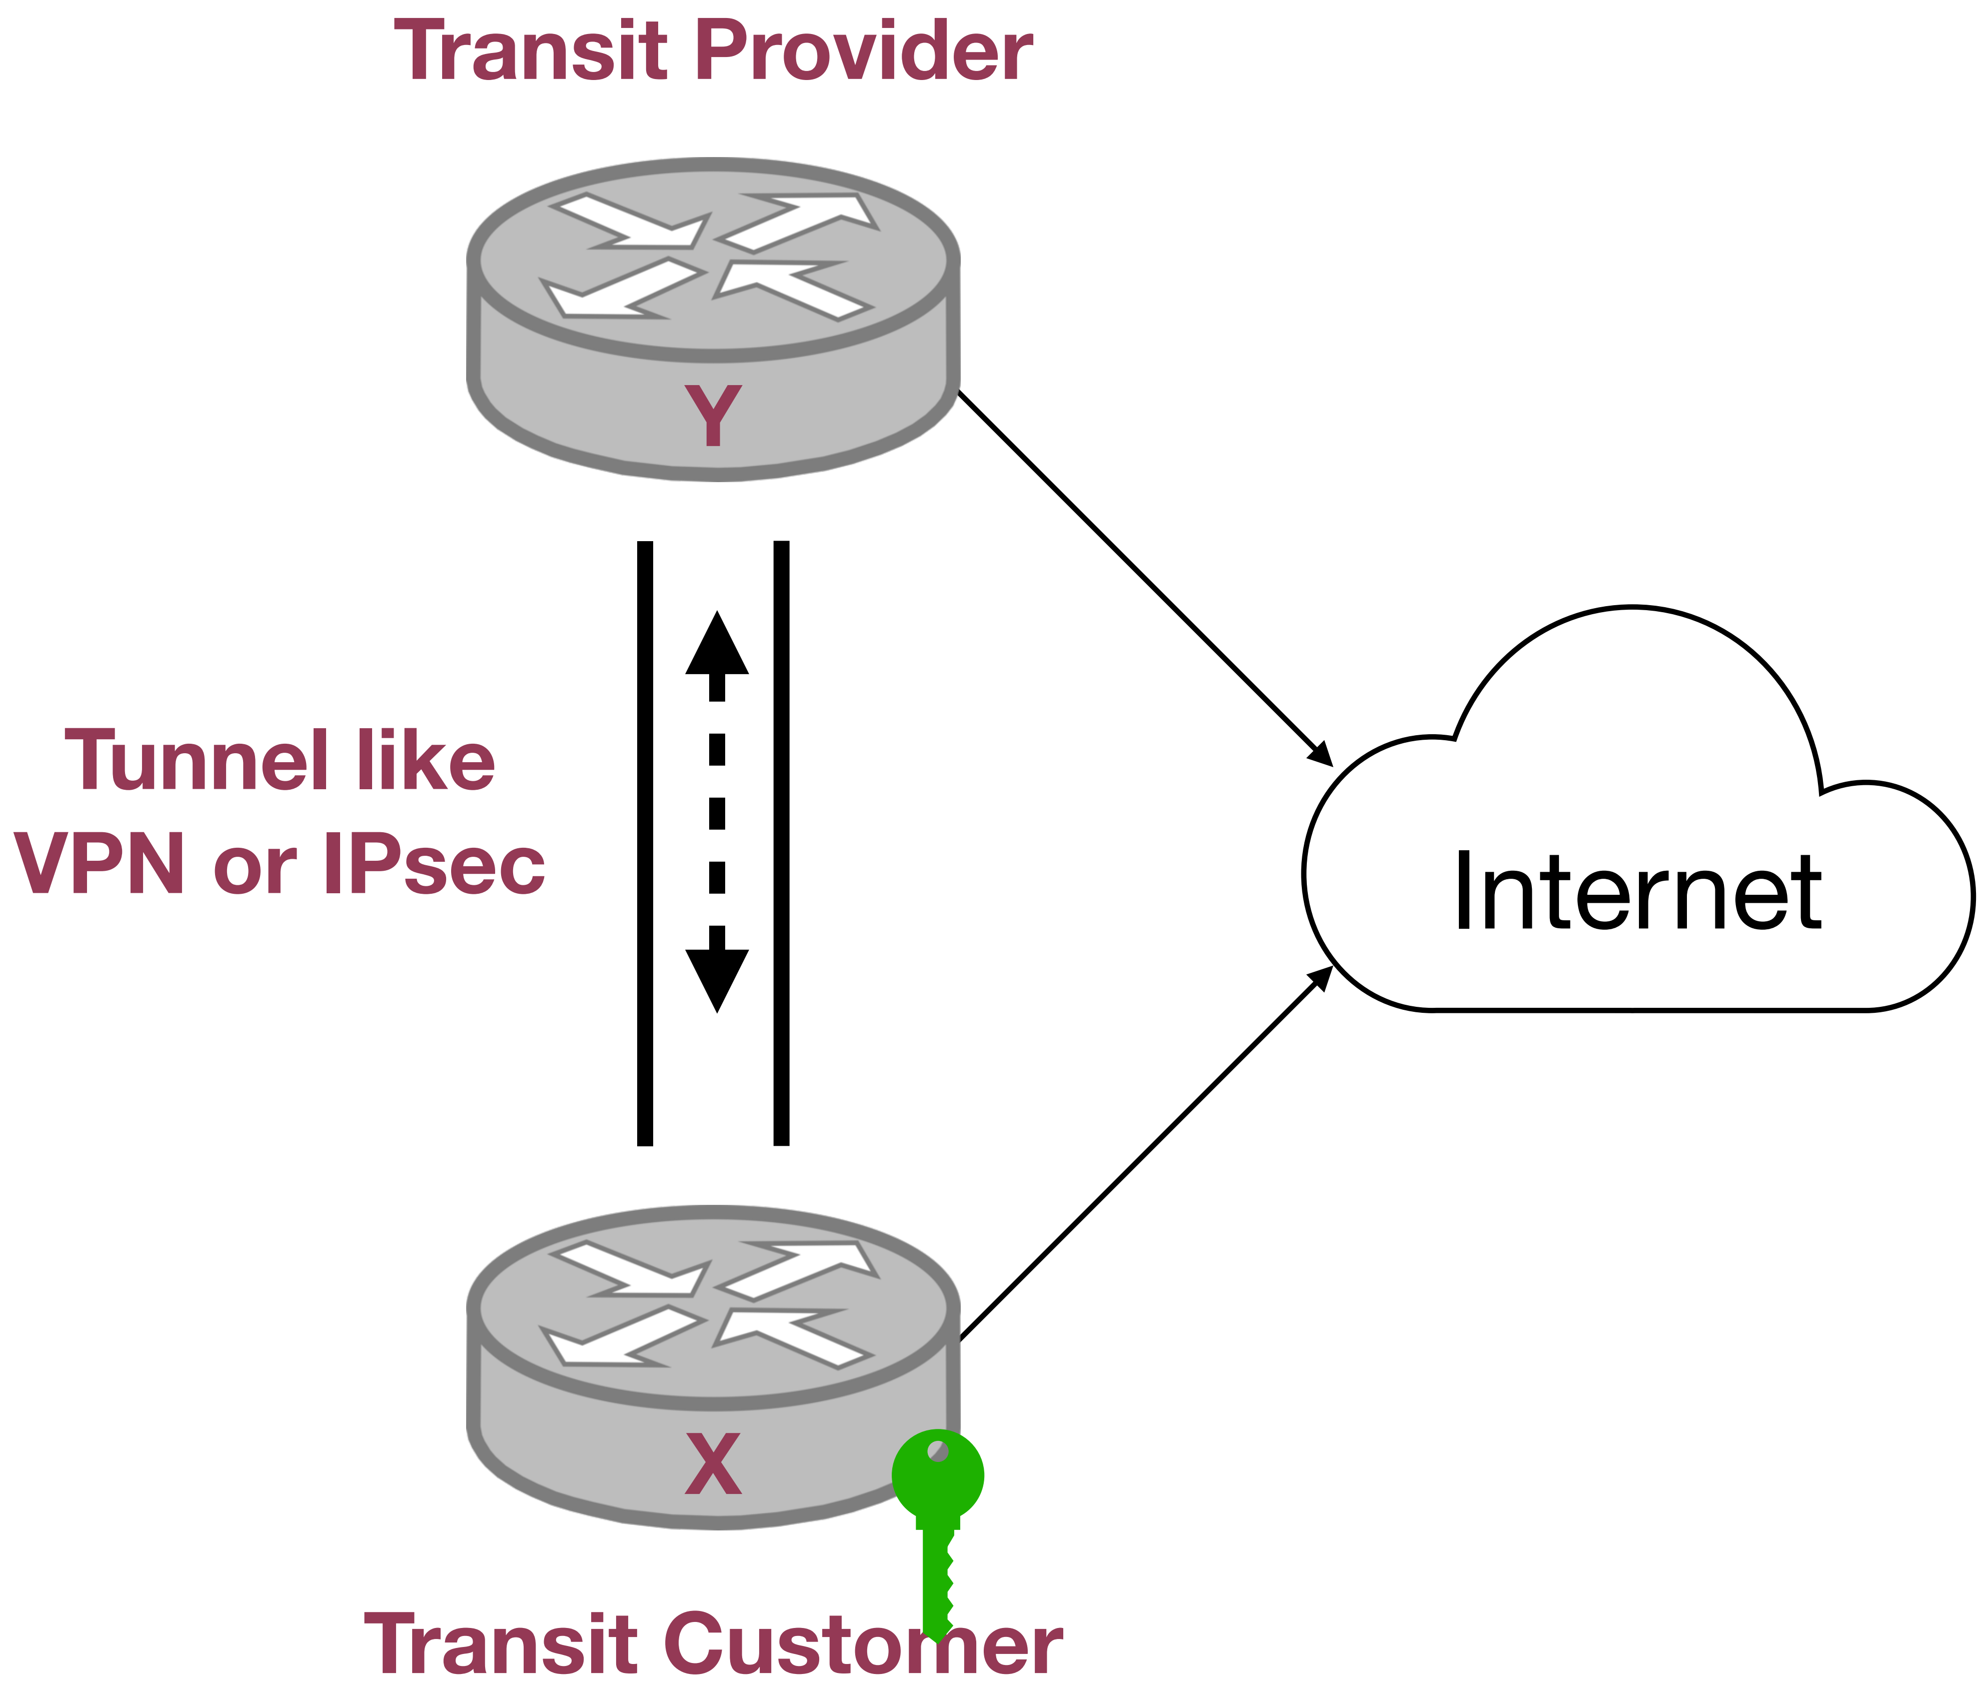

BGP Announcement:

129.21.0.0/16, AS X

129.21.0.0/20, AS Y ✗

ROA:

129.21.0.0/16-20, AS X

🔑 : Access to ROA

Hidden Relationships: Transit + Tunnel

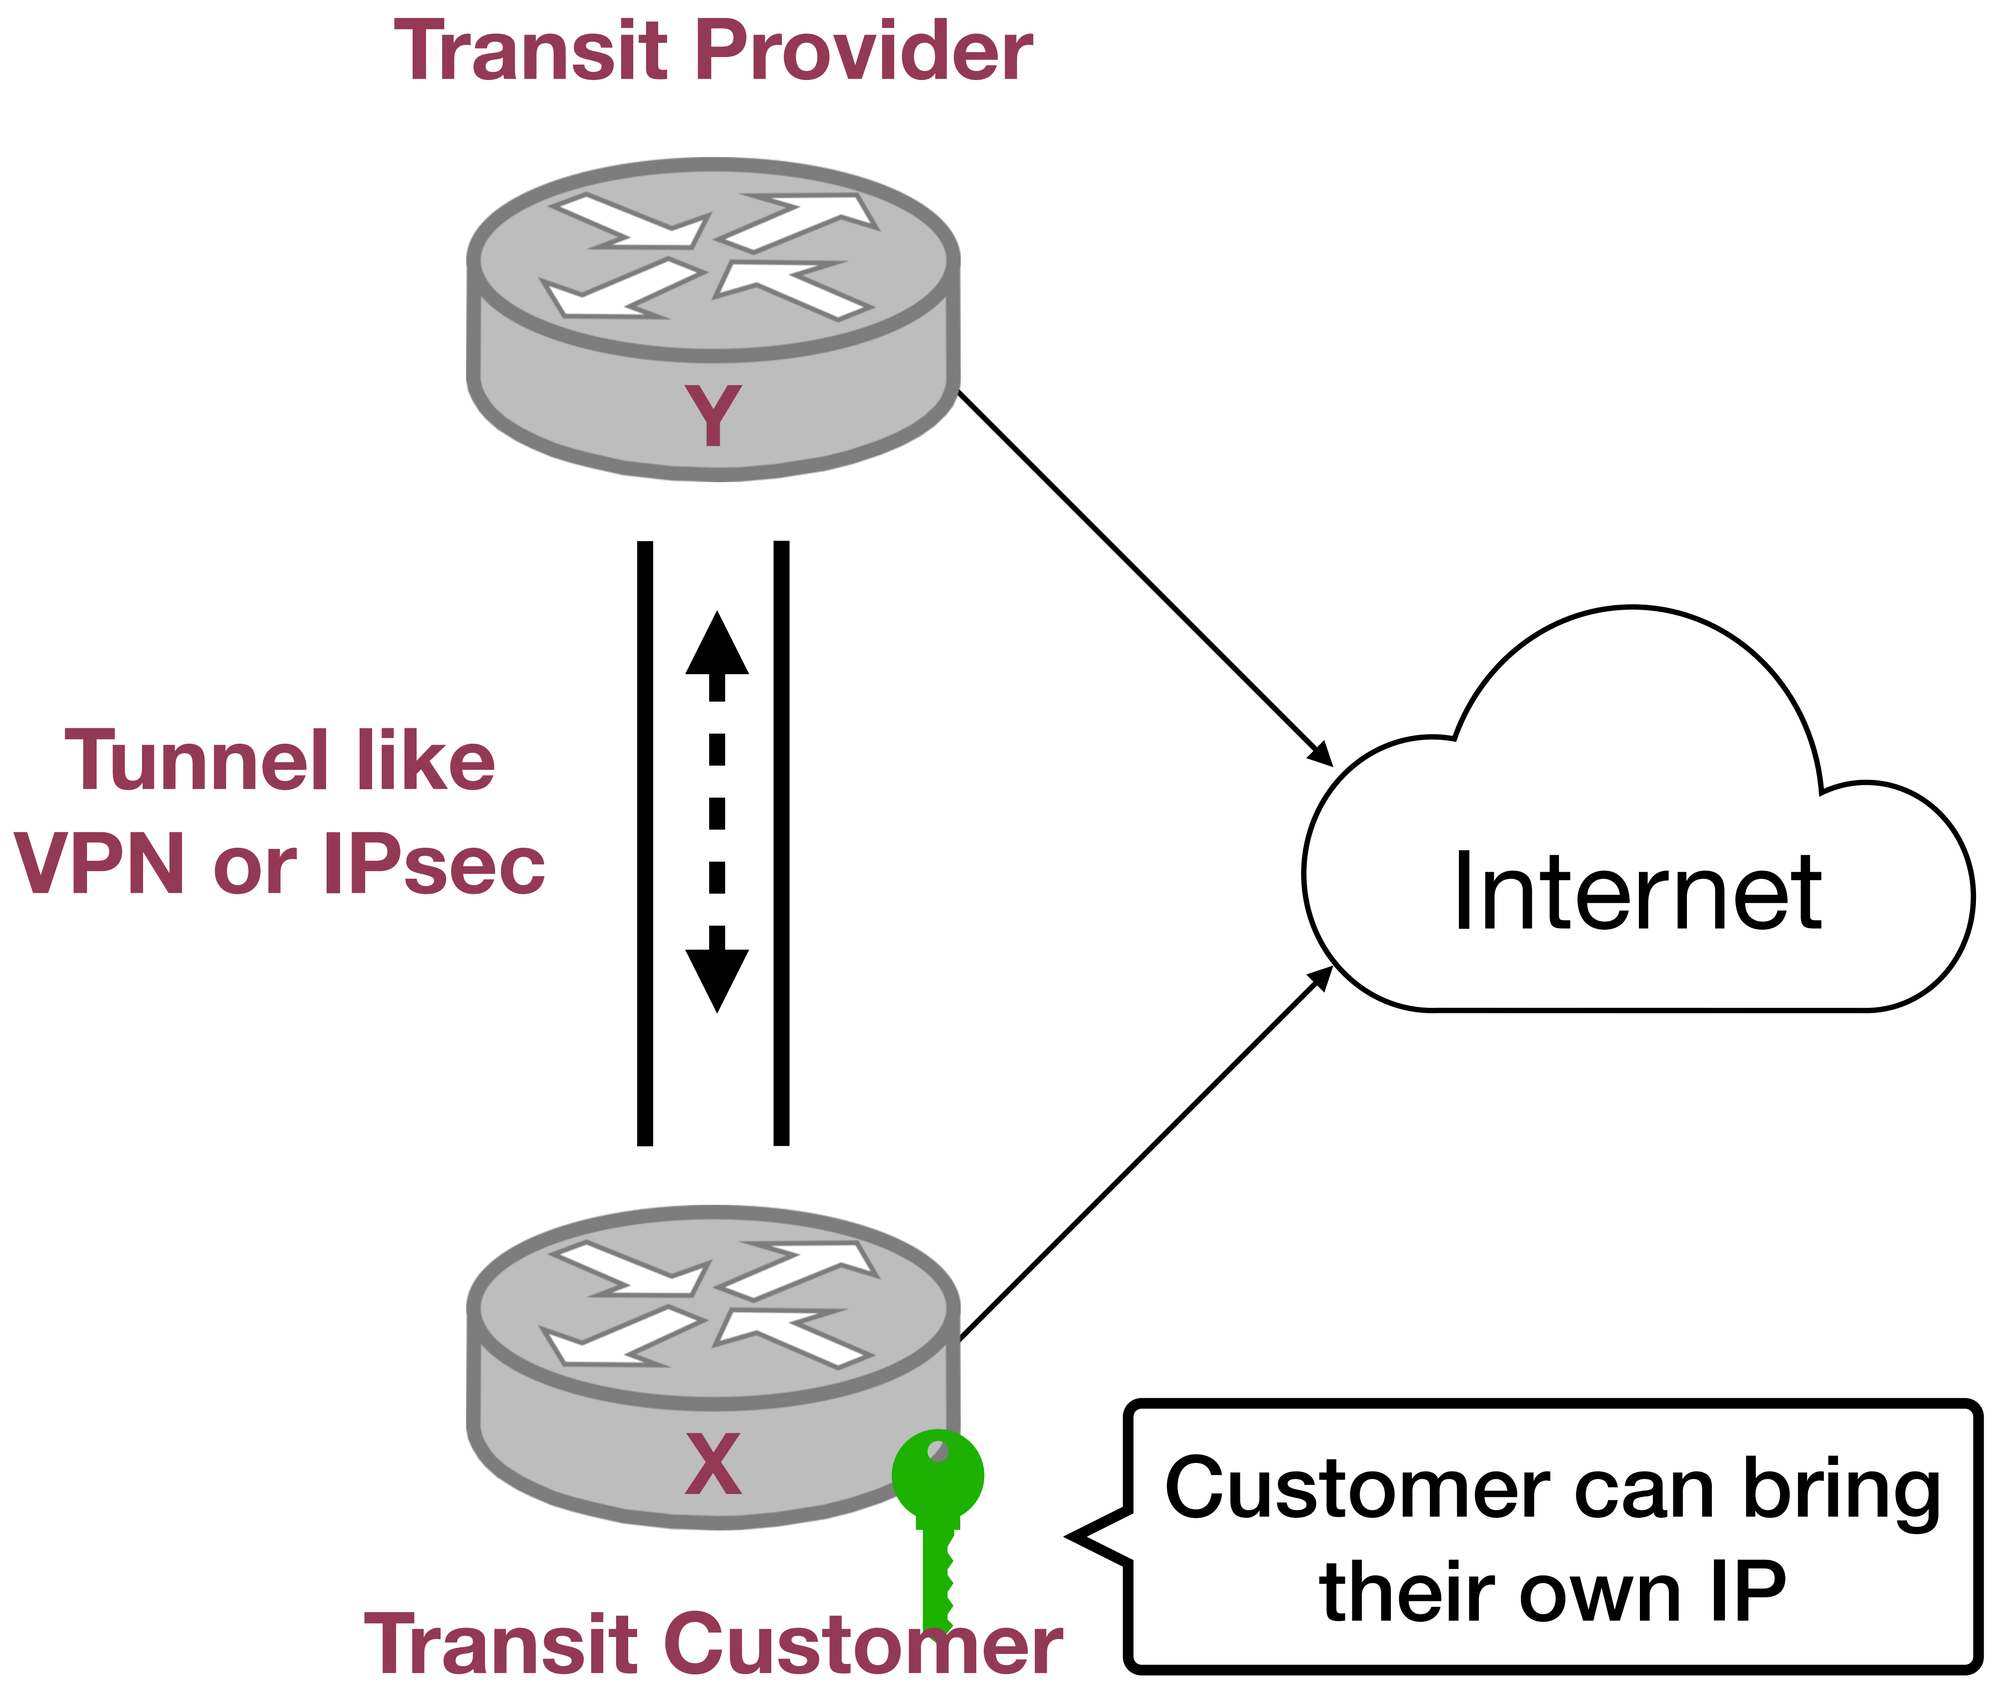

BGP Announcement:

129.21.0.0/16, AS X

129.21.0.0/20, AS Y ✗

ROA:

129.21.0.0/16-20, AS X

🔑 **: Access to ROA**

Hidden Relationships

- Transit service is not the only reason!

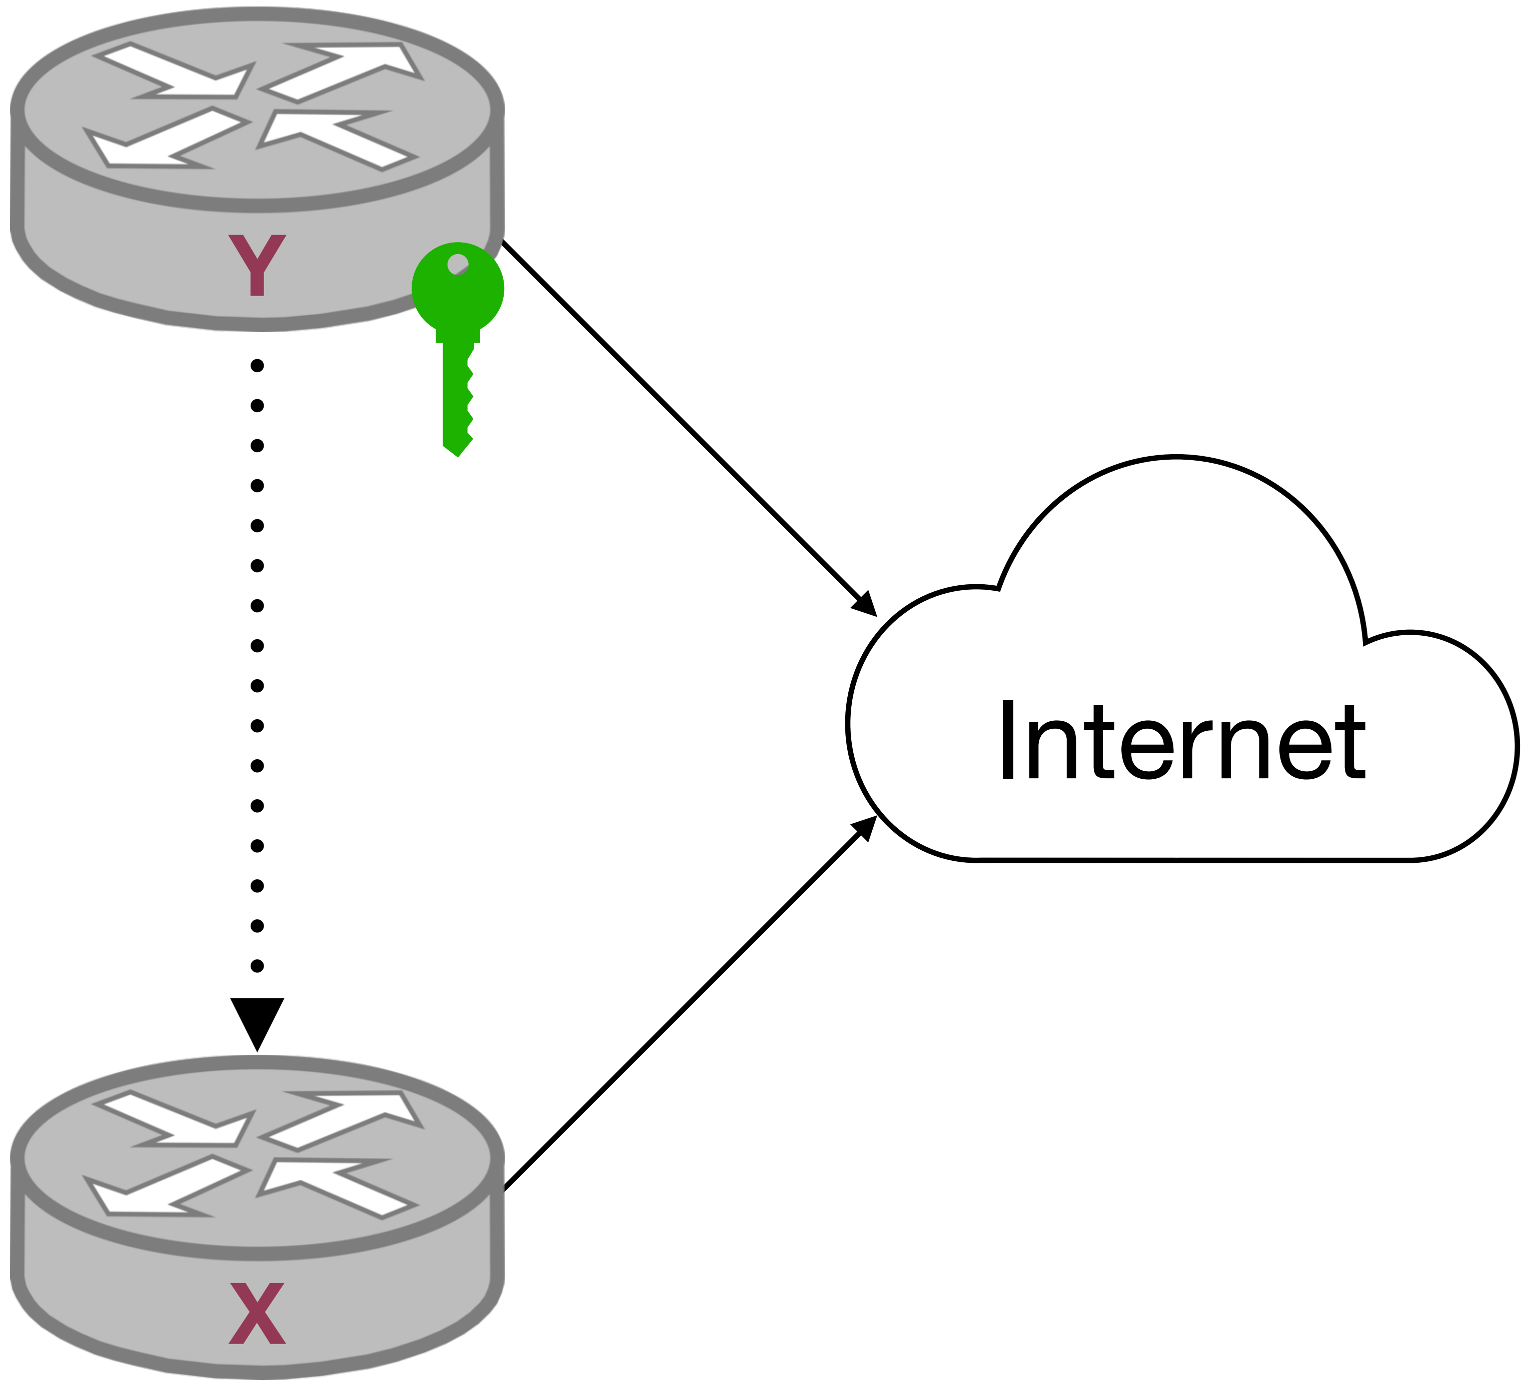

Hidden Relationships: Leasing

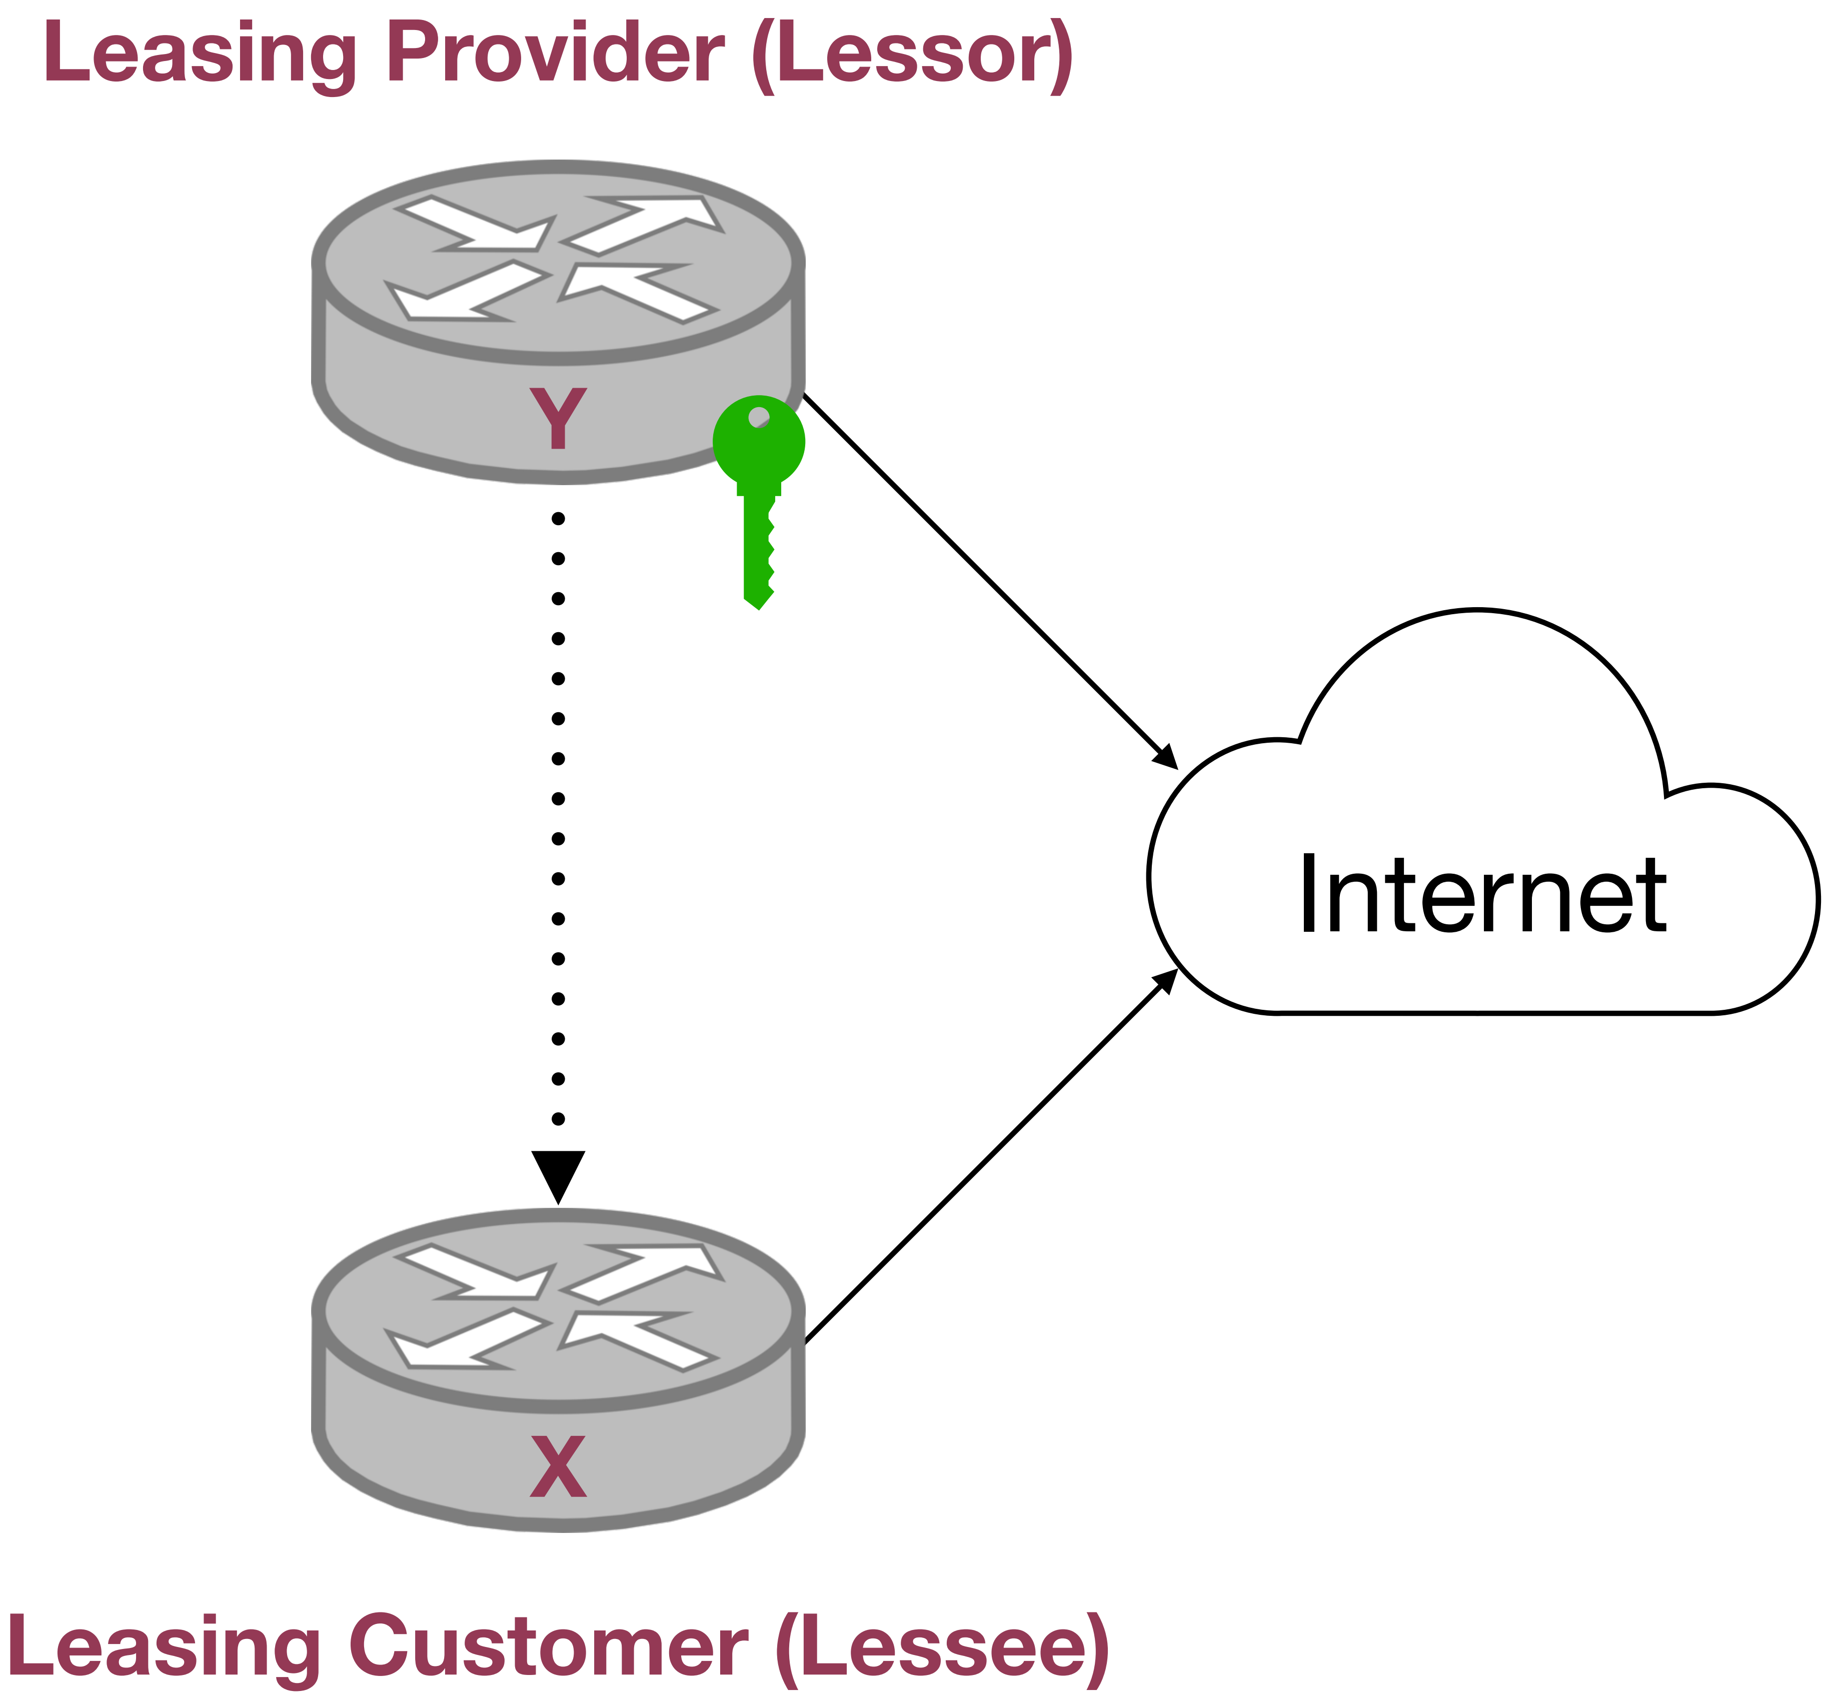

Hidden Relationships: Leasing (direct)

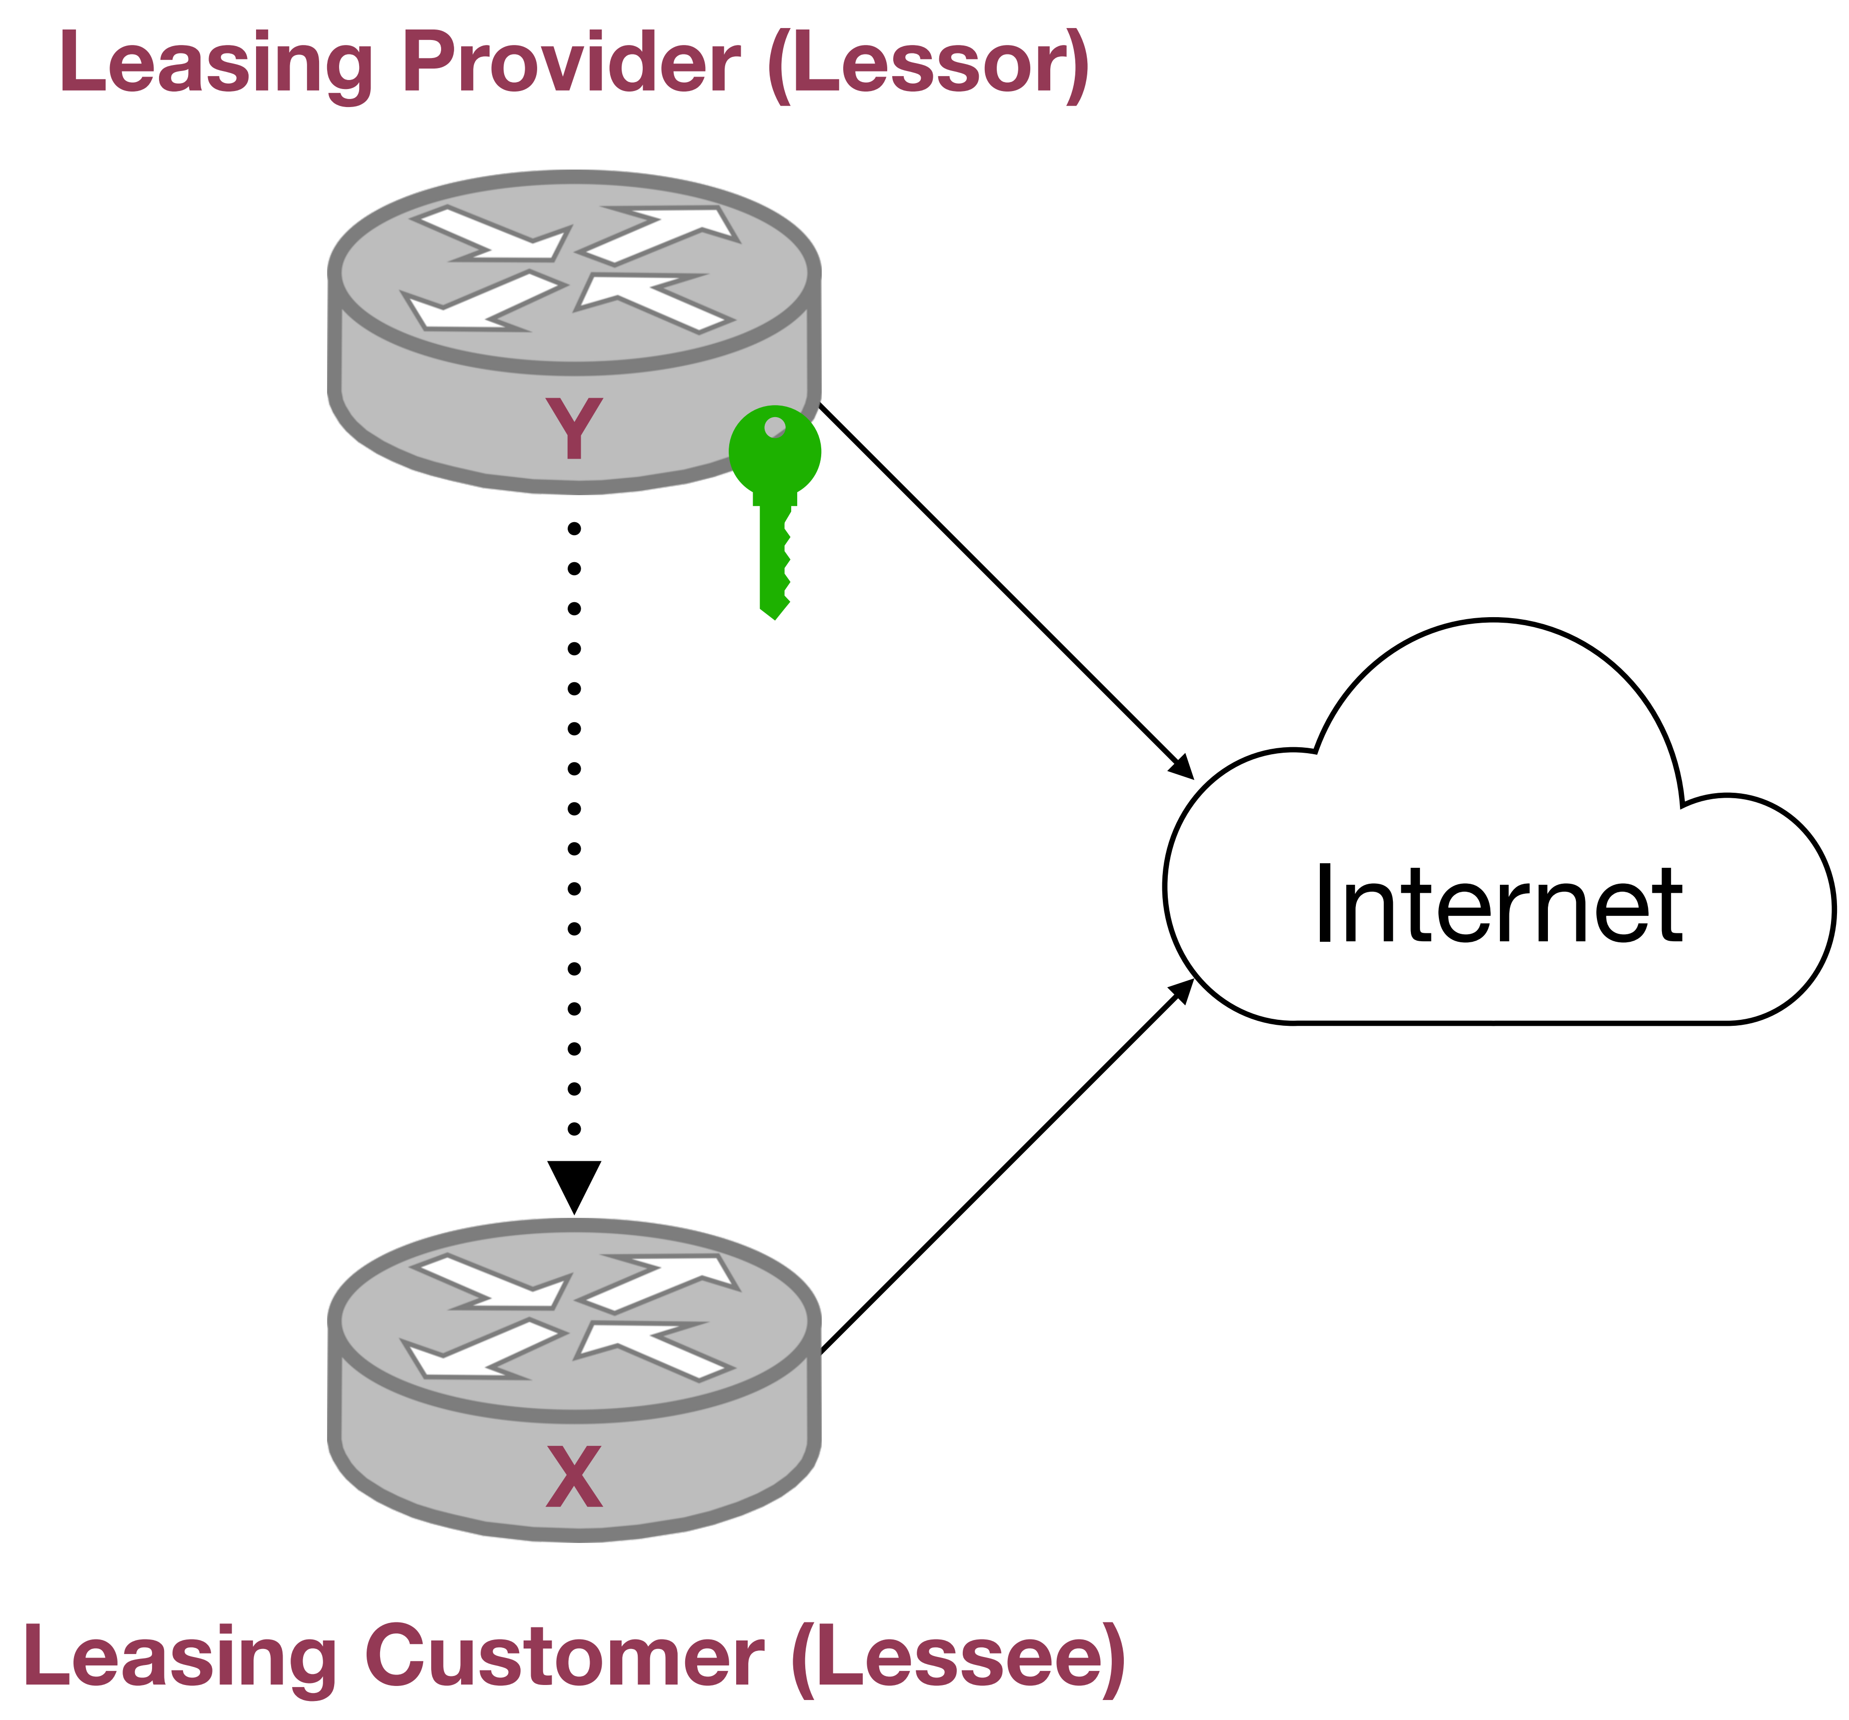

Leasing Provider (Lessor)

Y

Internet

X

Leasing Customer (Lessee)

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 **: Access to ROA**

Hidden Relationships: Leasing (w/ Broker)

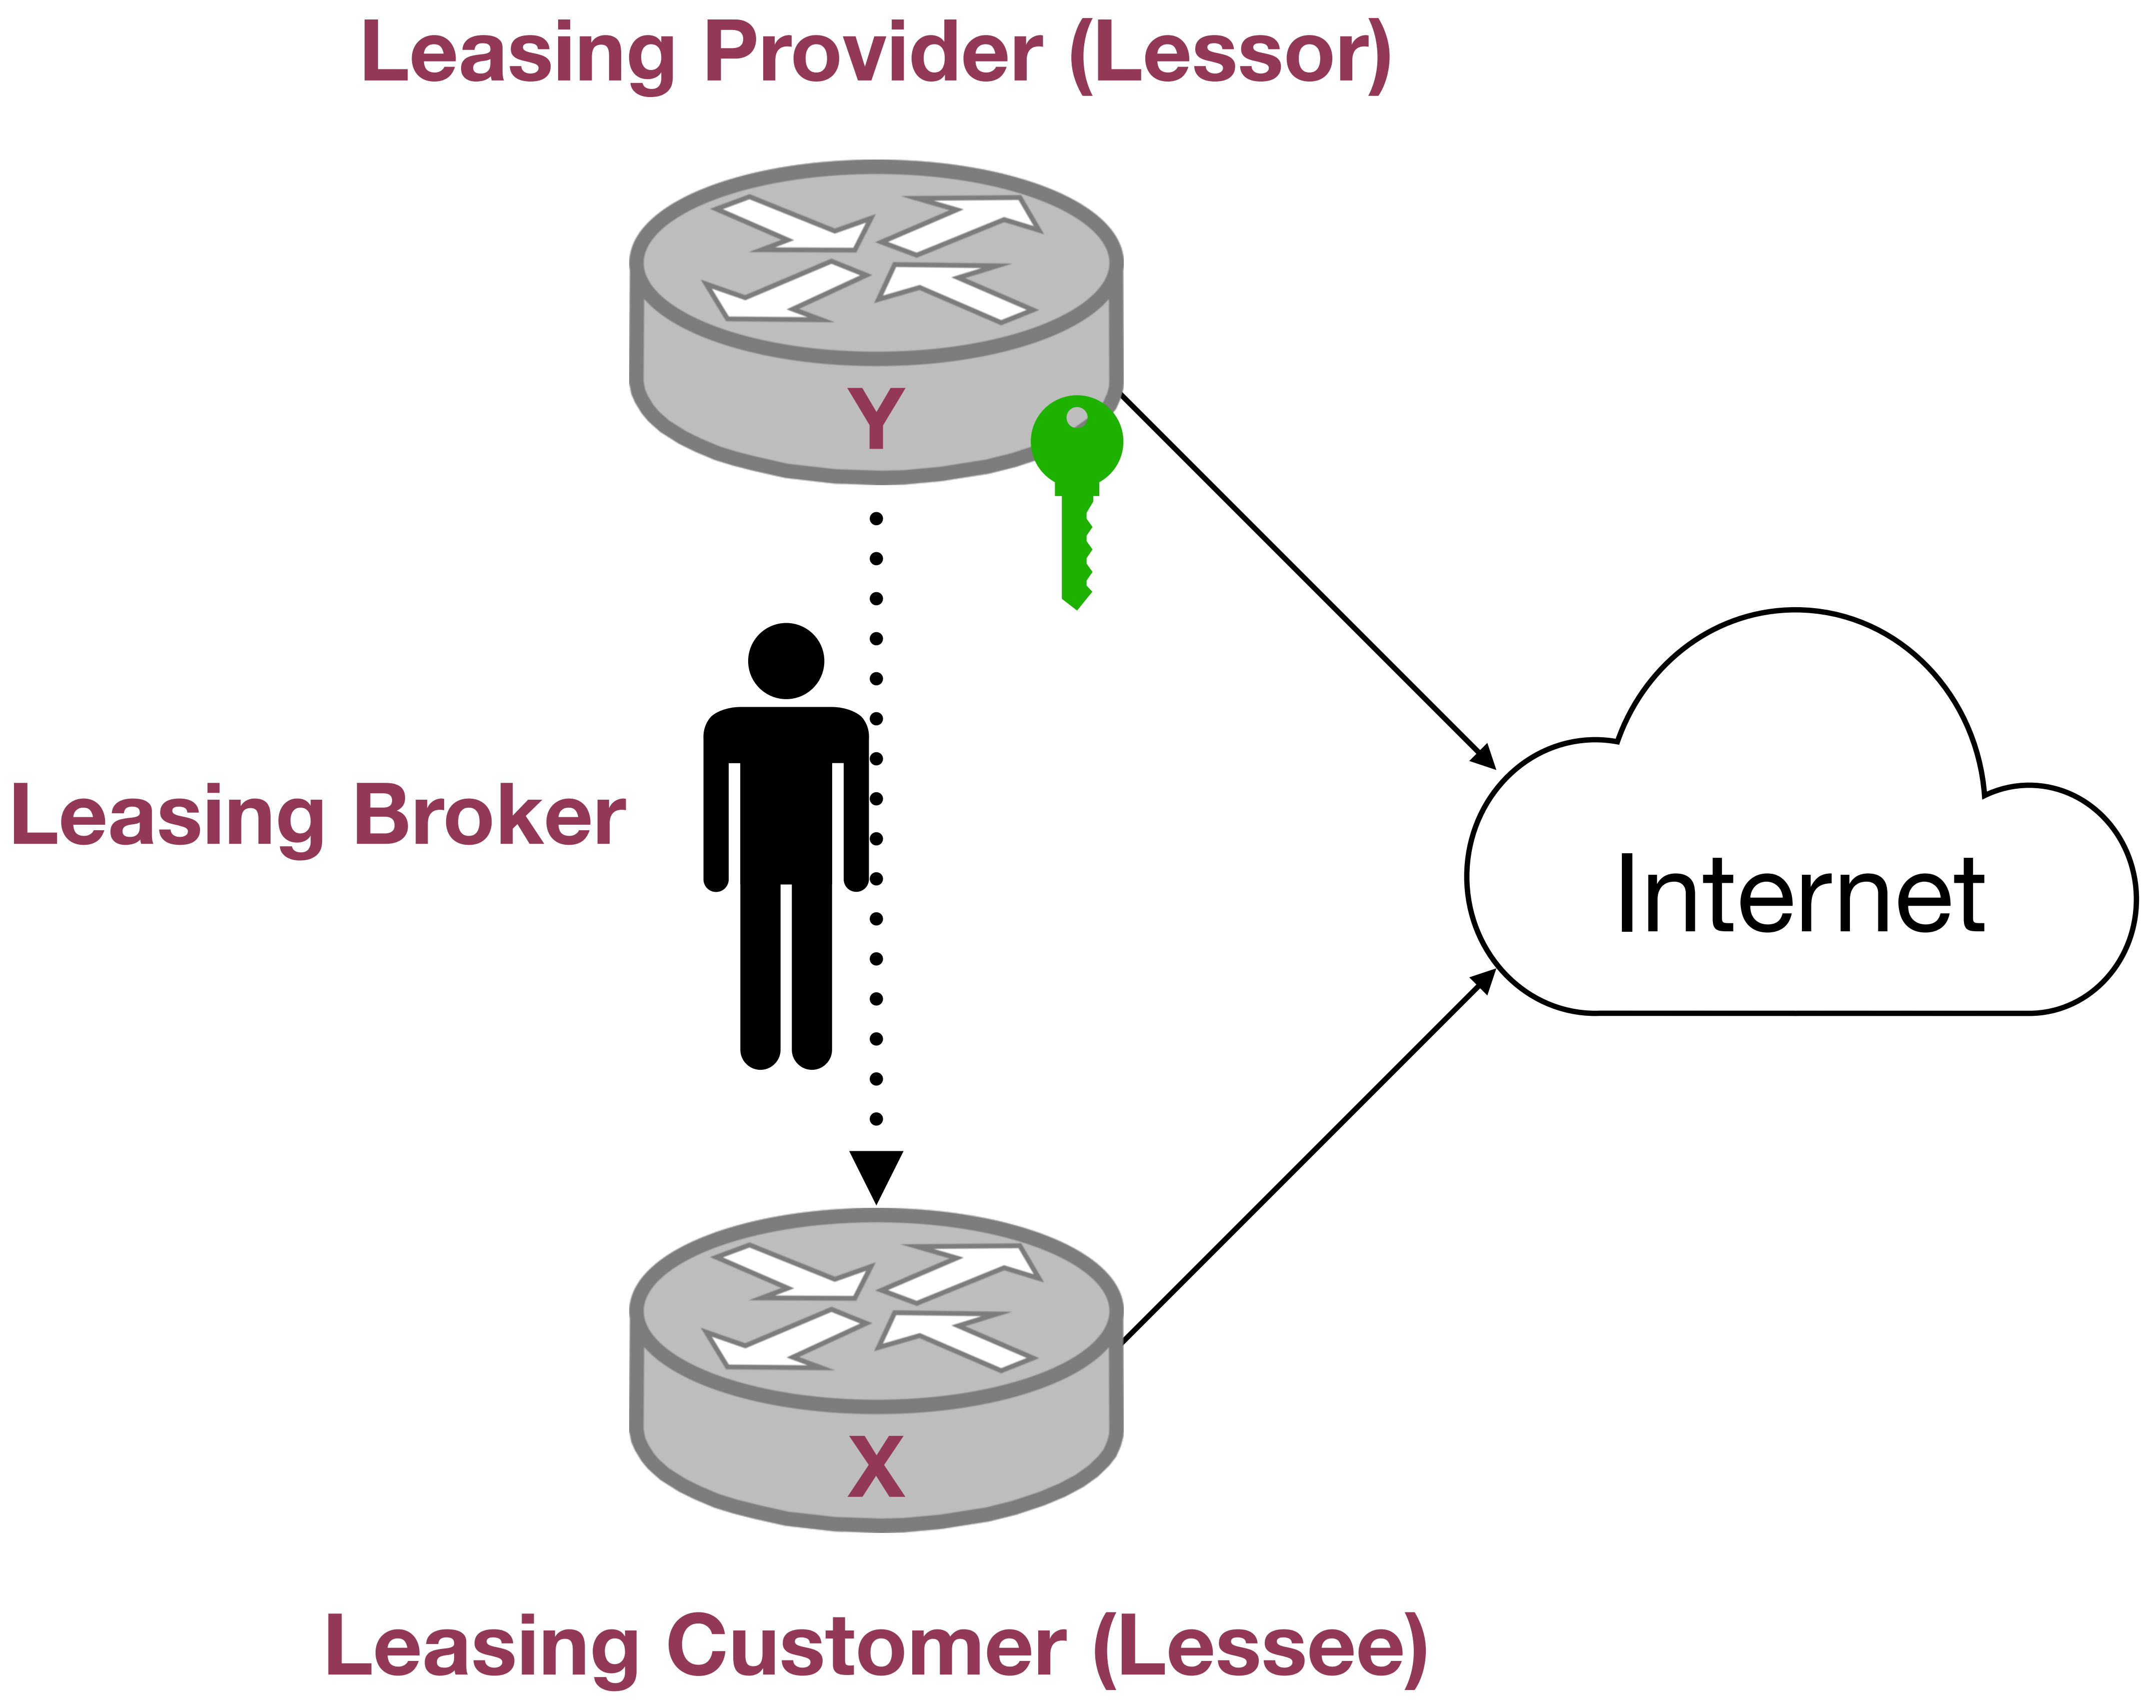

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 **: Access to ROA**

Classification of RPKI-Invalid

Easy ones:

1. **Transit**: checking BGP path from BGP datasets from RouteViews and RIS

2. **Same org**: using AS to Org mapping database CAIDA AS2Org (We also have one AS2Org mapping dataset using LLM called ASINT!)

3. **Transit + Org**: combine BGP and AS2Org

Classification of RPKI-Invalid

Easy ones:

1. **Transit**: checking BGP path from BGP datasets from Routeviews and RIS
2. **Same org**: using AS to Org mapping database CAIDA AS2Org (We also have one AS2Org mapping dataset using LLM called ASINT!)
3. **Transit + Org**: combine BGP and AS2Org
4. **Broker leasing**: Broker usually will leave registration information in WHOIS

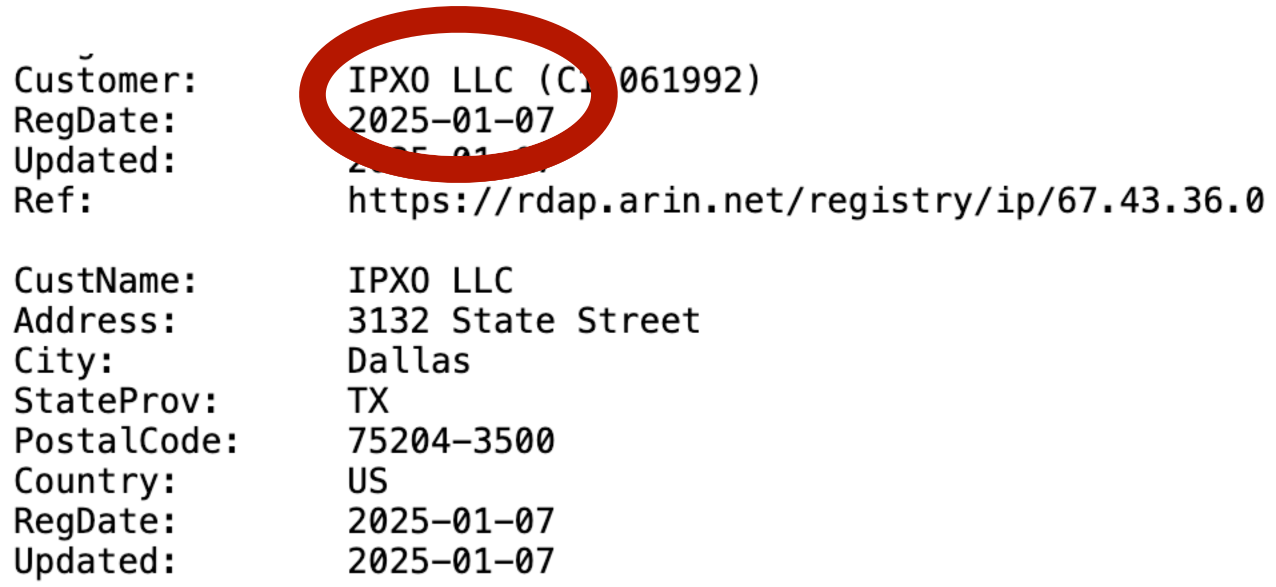

```
Customer:       IPXO LLC (C10061992)
RegDate:        2025-01-07
Updated:        2025-01-07
Ref:            https://rdap.arin.net/registry/ip/67.43.36.0

CustName:       IPXO LLC
Address:        3132 State Street
City:           Dallas
StateProv:      TX
PostalCode:     75204-3500
Country:        US
RegDate:        2025-01-07
Updated:        2025-01-07
```

Classification for Hidden Relationships

Transit with tunneling and direct leasing are hard to identify:

- No evidence in BGP since the BGP origin will be the provider
- No evidence in WHOIS since the WHOIS record will remain lessor

Classification for Hidden Relationships

Transit with tunneling and direct leasing are hard to identify:

- No evidence in BGP since the BGP origin will be the provider
- No evidence in WHOIS since the WHOIS record will remain lessor

Thus we need to guess, but how?

Classification for Hidden Relationships

- Instead of trying to find individual prefix involved (which is hard to identify), we try to find the ASes who provide such hidden services (direct leasing and hidden transit).

Classification for Hidden Relationships

- Instead of trying to find individual prefix involved (which is hard to identify), we try to find the ASes who provide such hidden services (direct leasing and hidden transit).
- Our hypothesis are:
    1. These service providers usually will have many customers
    2. Mistakes in ROA shouldn't be a common situation

Classification for Hidden Relationships

- We collect all RPKI Valid prefixes from RouteViews and RIS BGP tables in the past 18 months.
- Identifying transit providers and lessors ASes and counting the involved prefix per AS
- Testing our methods with ground-truth leasing prefixes and BGP transit providers.

Classification Results

- We attribute 96.9% of the RPKI-Invalid prefixes to these misconfigurations

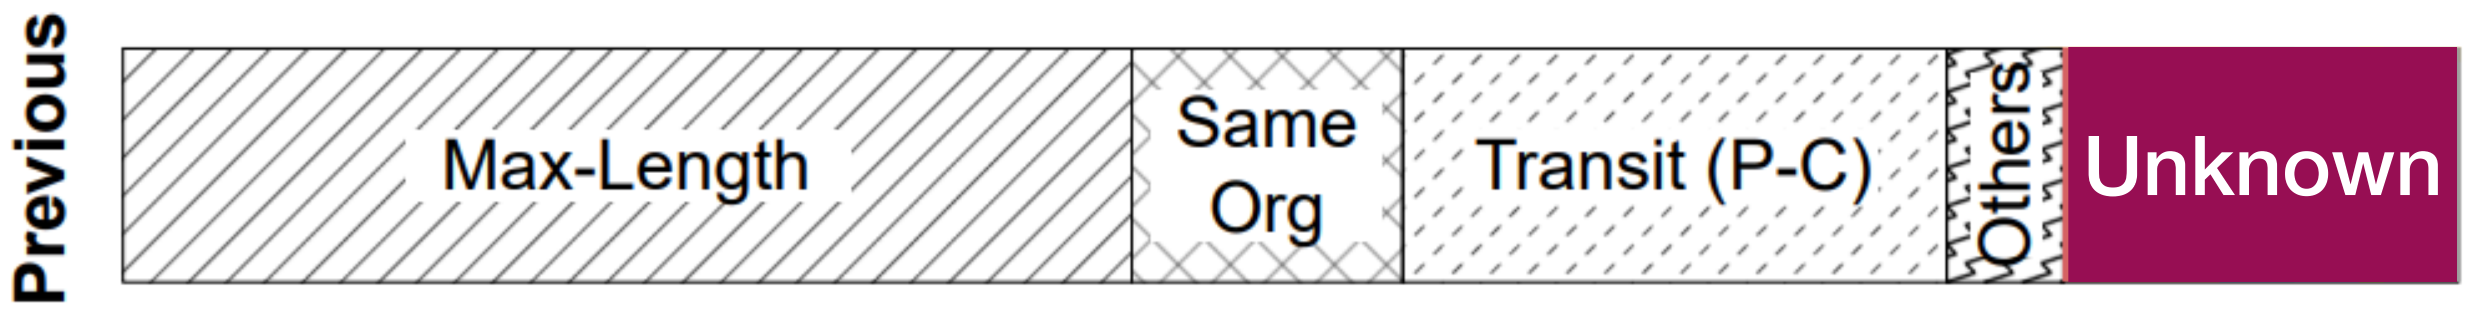

Classification Results

- We attribute 96.9% of the RPKI-Invalid prefixes to these misconfigurations

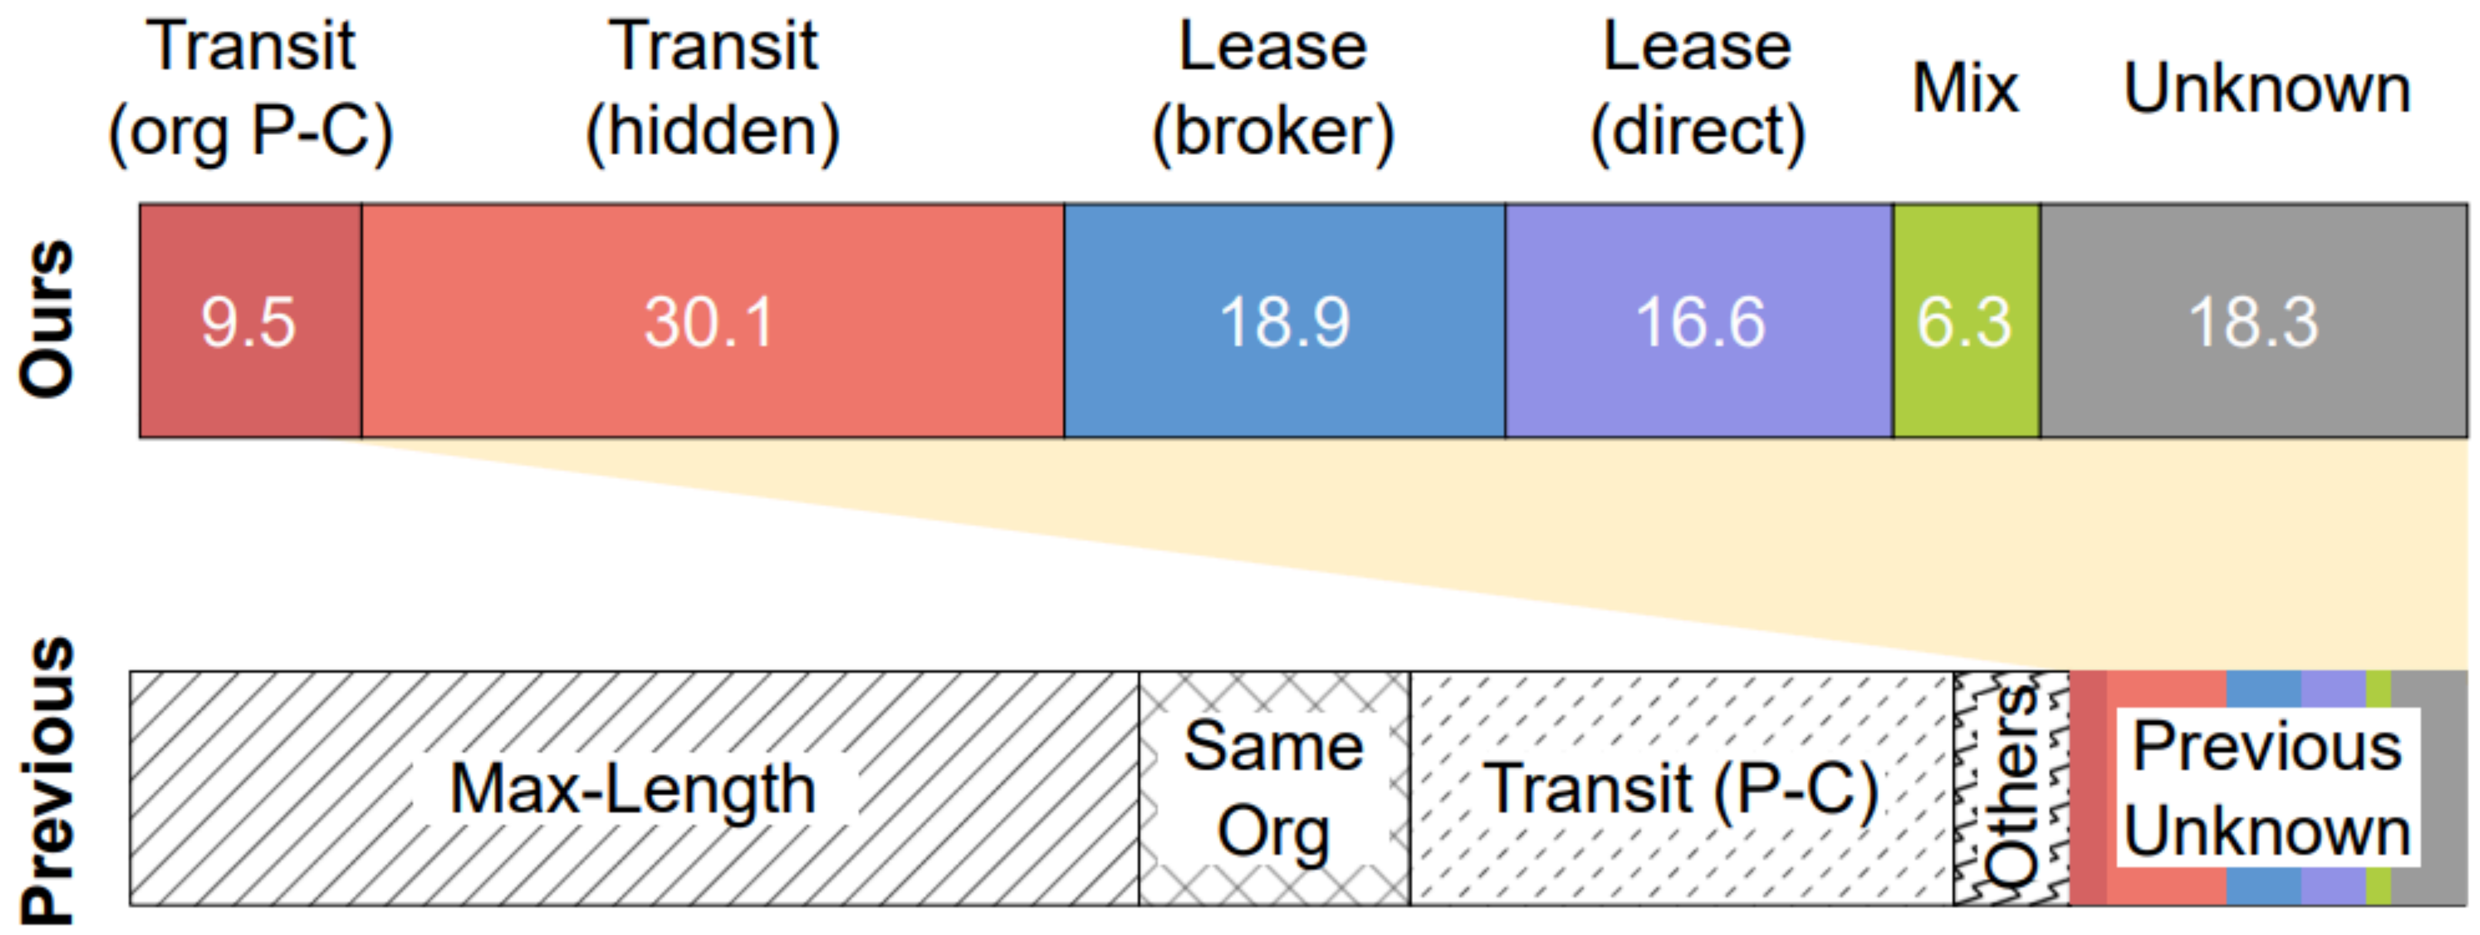

Question 2:
What’s the impact of RPKI-Invalid?

While people keep making mistakes in RPKI,
it’s fine if there’s no impact on daily operation

But we find it’s not.

Impact of RPKI-Invalid: Disconnection

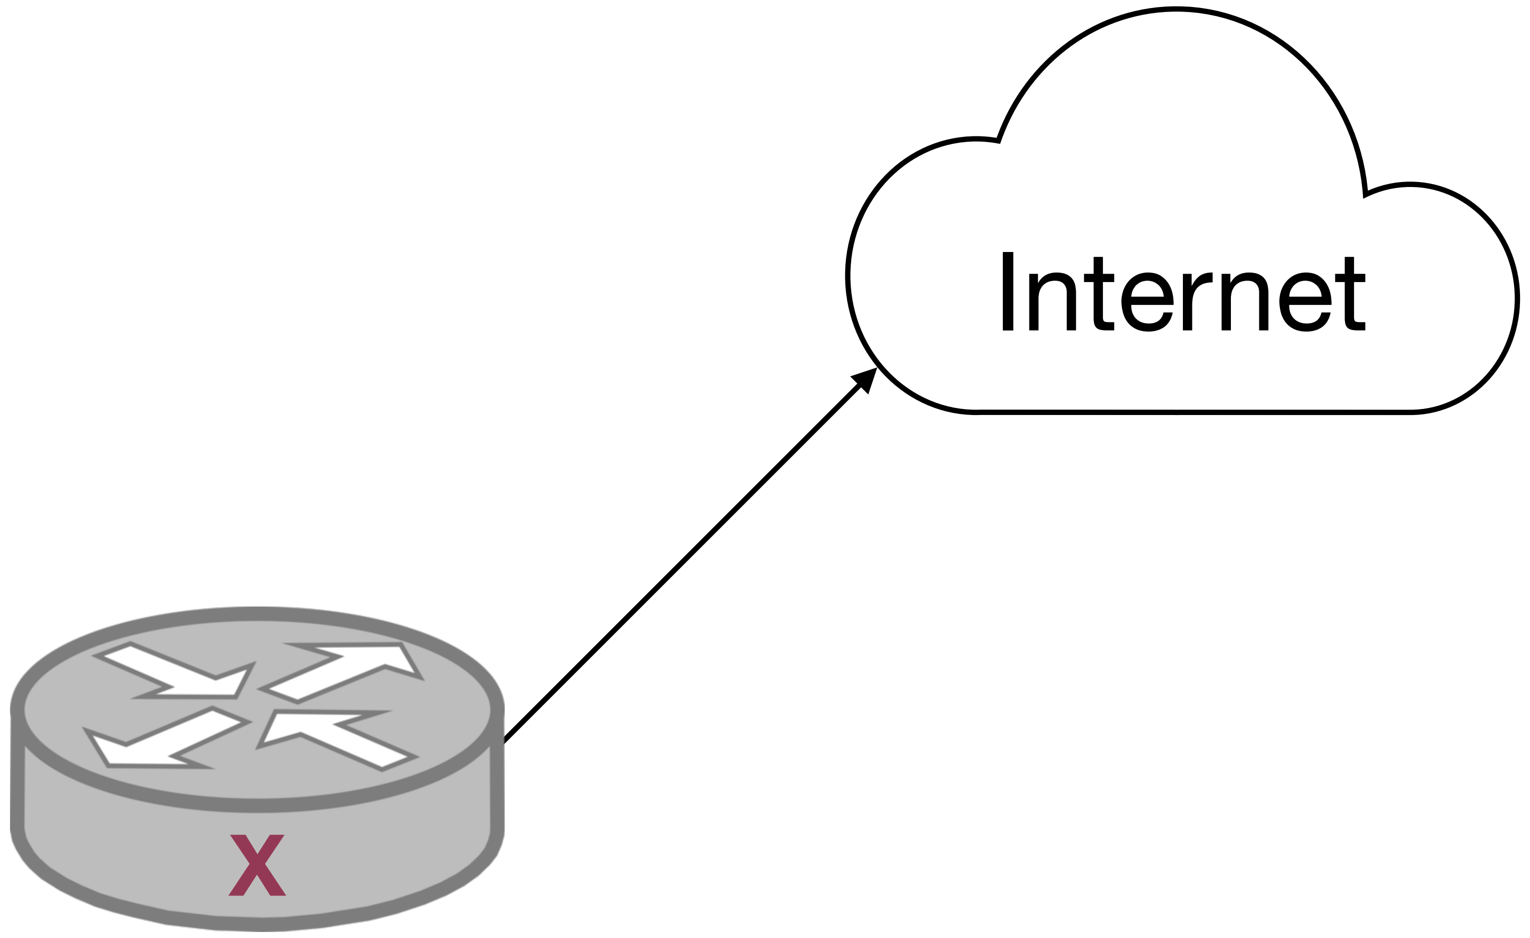

BGP Announcement:

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

Impact of RPKI-Invalid: Disconnection

Disconnected from the Internet

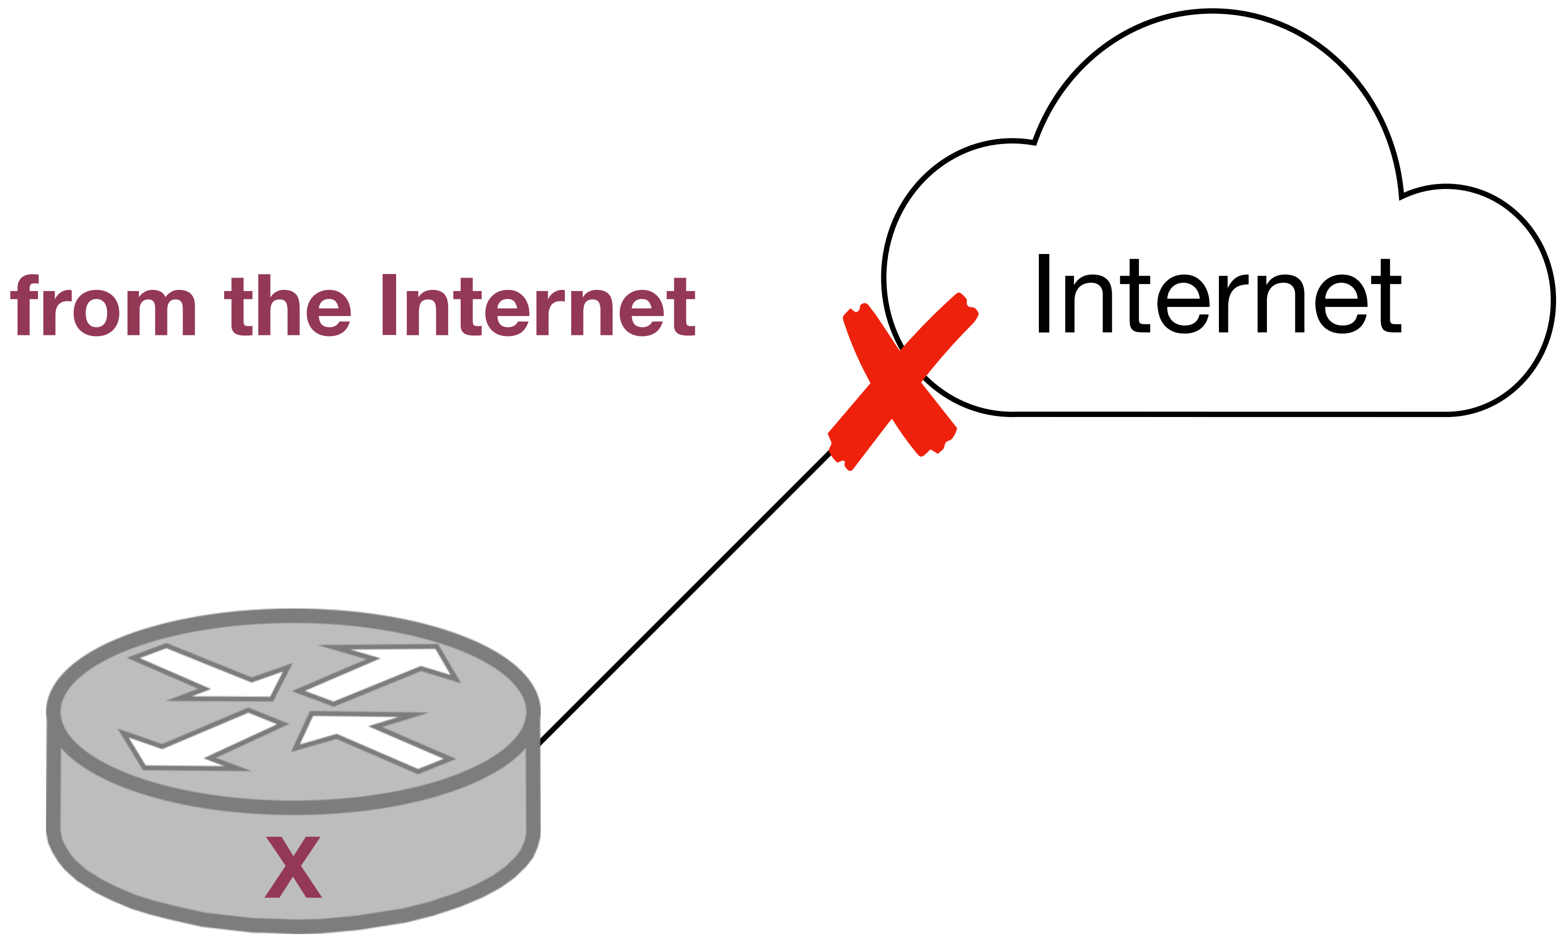

BGP Announcement:

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

Impact of RPKI-Invalid: Disconnection

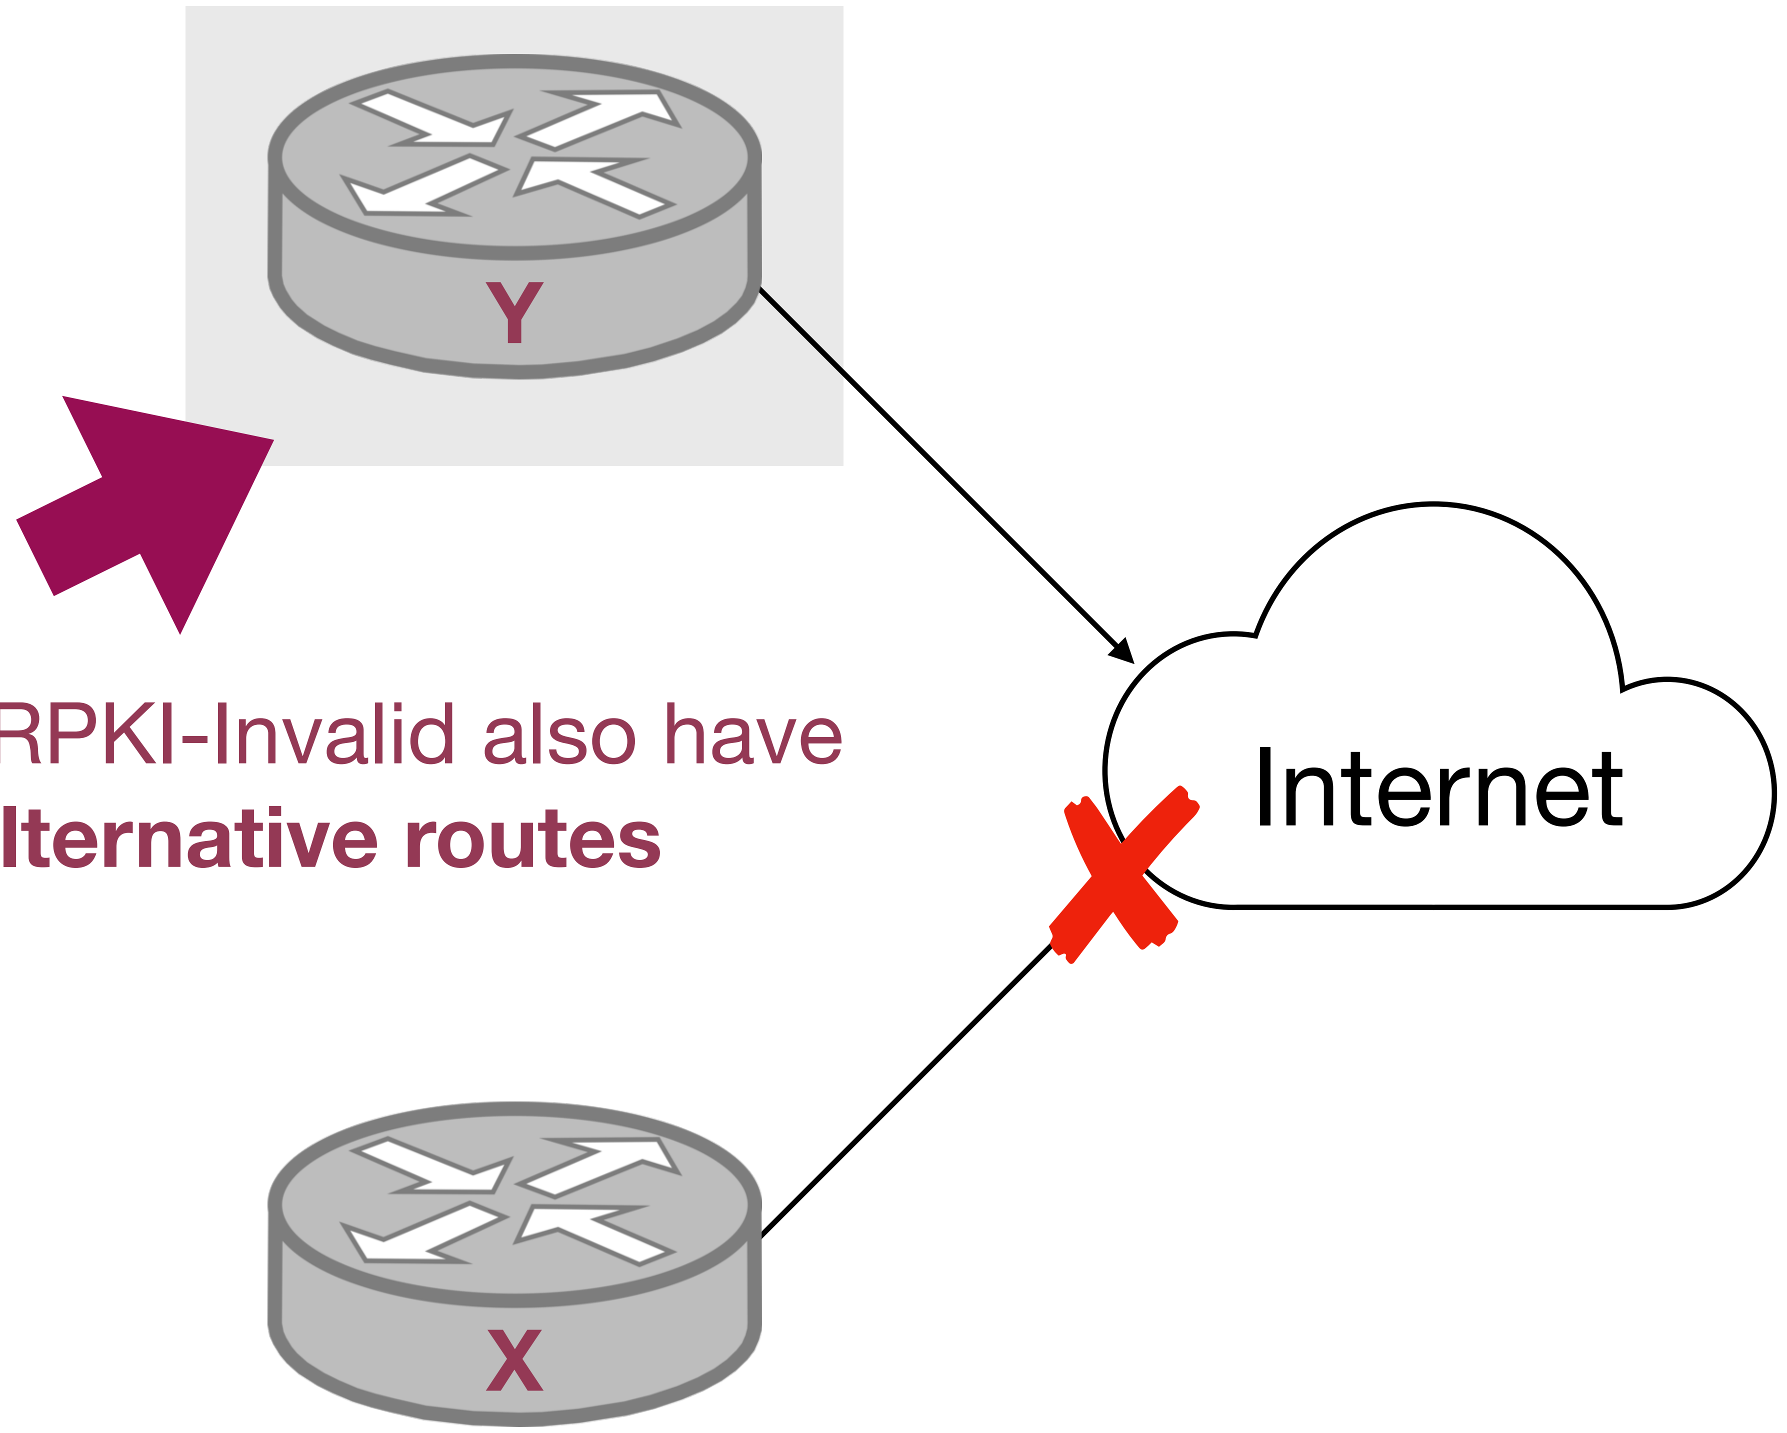

But, more than 95% RPKI-Invalid also have valid/unknown **alternative routes**

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

Impact of RPKI-Invalid: Disconnection

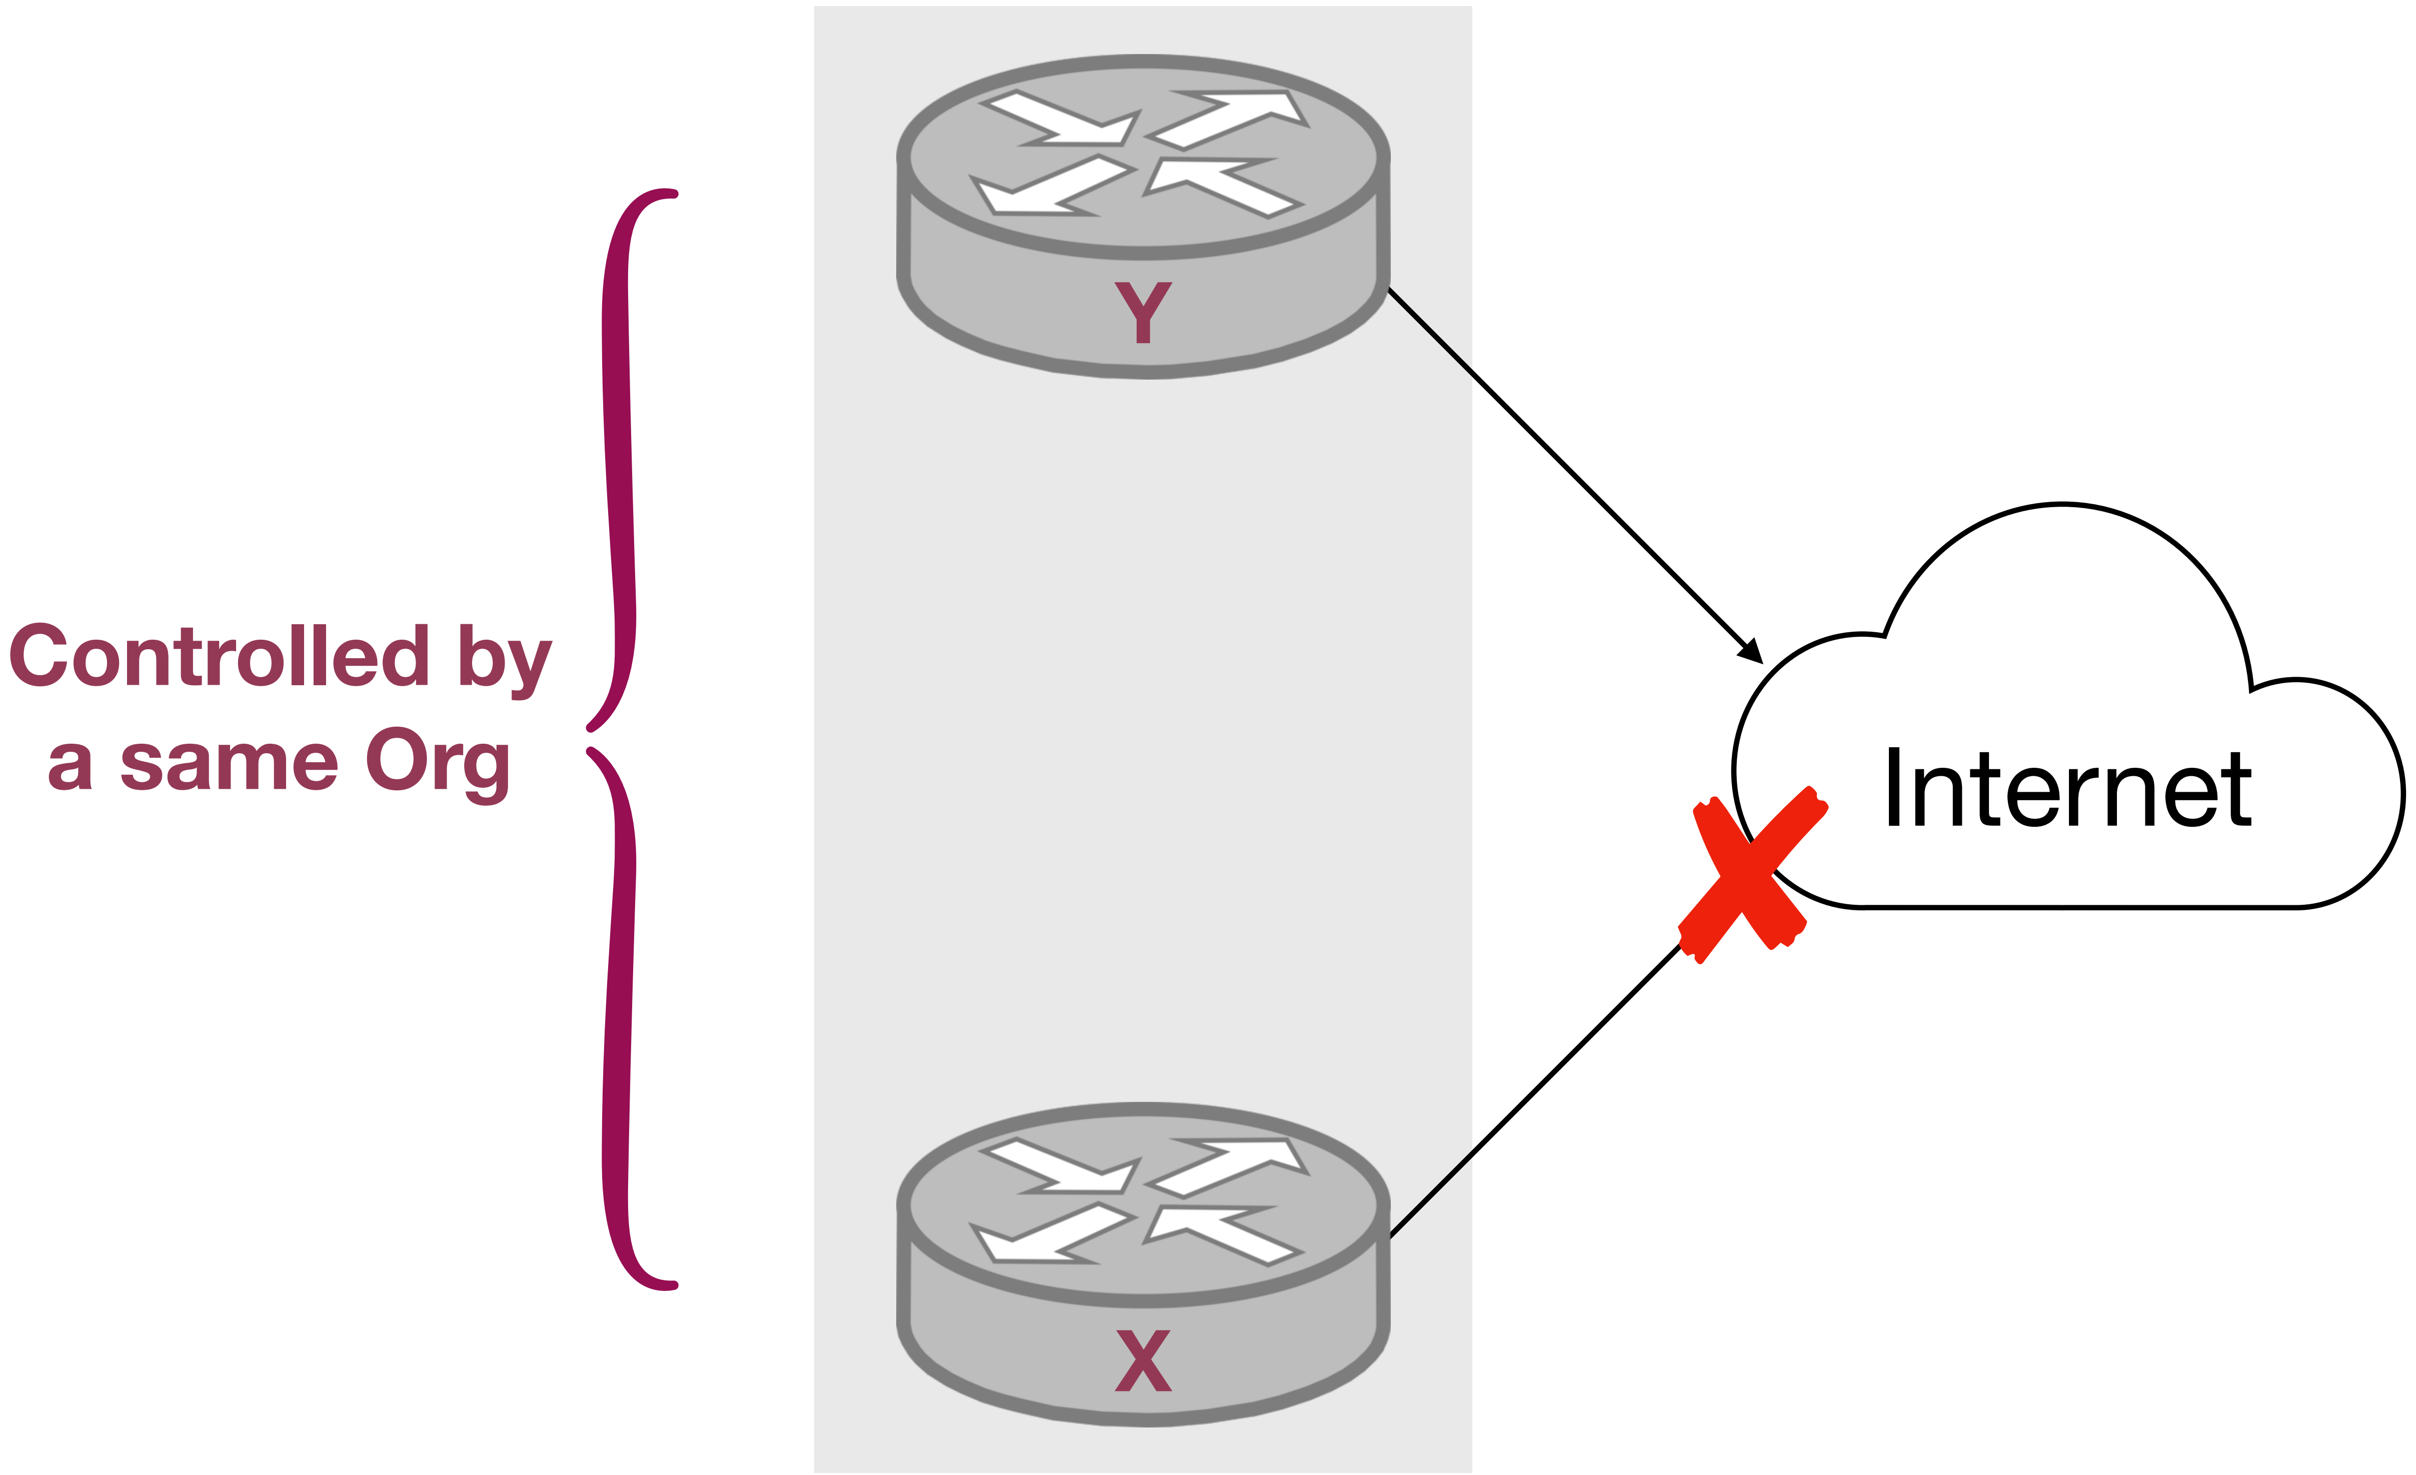

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

Impact of RPKI-Invalid: Disconnection

- We run active measurements, using RIPE Atlas probes to ping&traceroute all invalid prefixes.

Impact of RPKI-Invalid: Disconnection

- We run active measurements, using RIPE Atlas probes to ping&traceroute all invalid prefixes.

- We found only 3.1% of RPKI-Invalid will result to disconnection.

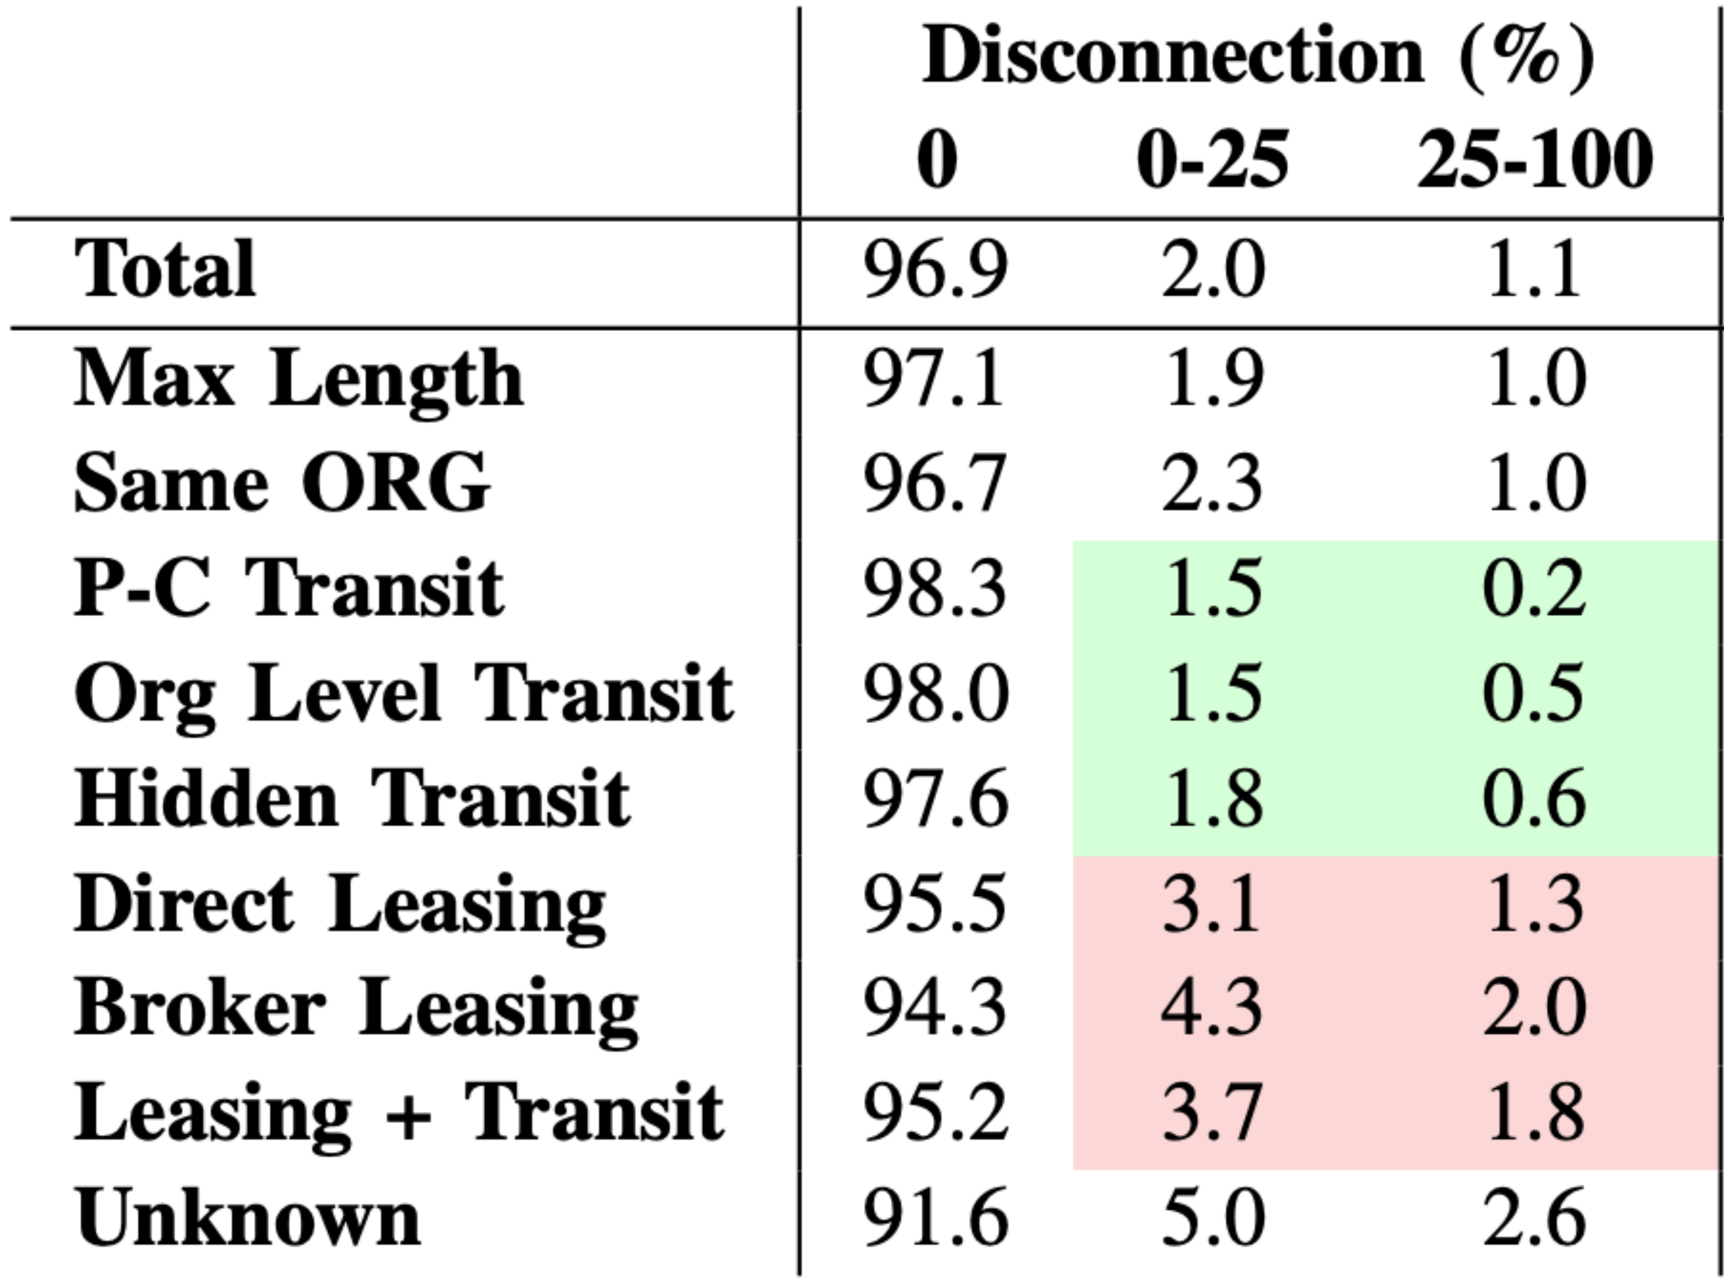

	Disconnection (%)		
	0	**0-25**	**25-100**
Total	96.9	2.0	1.1
Max Length	97.1	1.9	1.0
Same ORG	96.7	2.3	1.0
P-C Transit	98.3	1.5	0.2
Org Level Transit	98.0	1.5	0.5
Hidden Transit	97.6	1.8	0.6
Direct Leasing	95.5	3.1	1.3
Broker Leasing	94.3	4.3	2.0
Leasing + Transit	95.2	3.7	1.8
Unknown	91.6	5.0	2.6

Impact of RPKI-Invalid

- However, having alternative path doesn’t means no impact at all

Impact of RPKI-Invalid

- However, having alternative path doesn't means no impact at all.

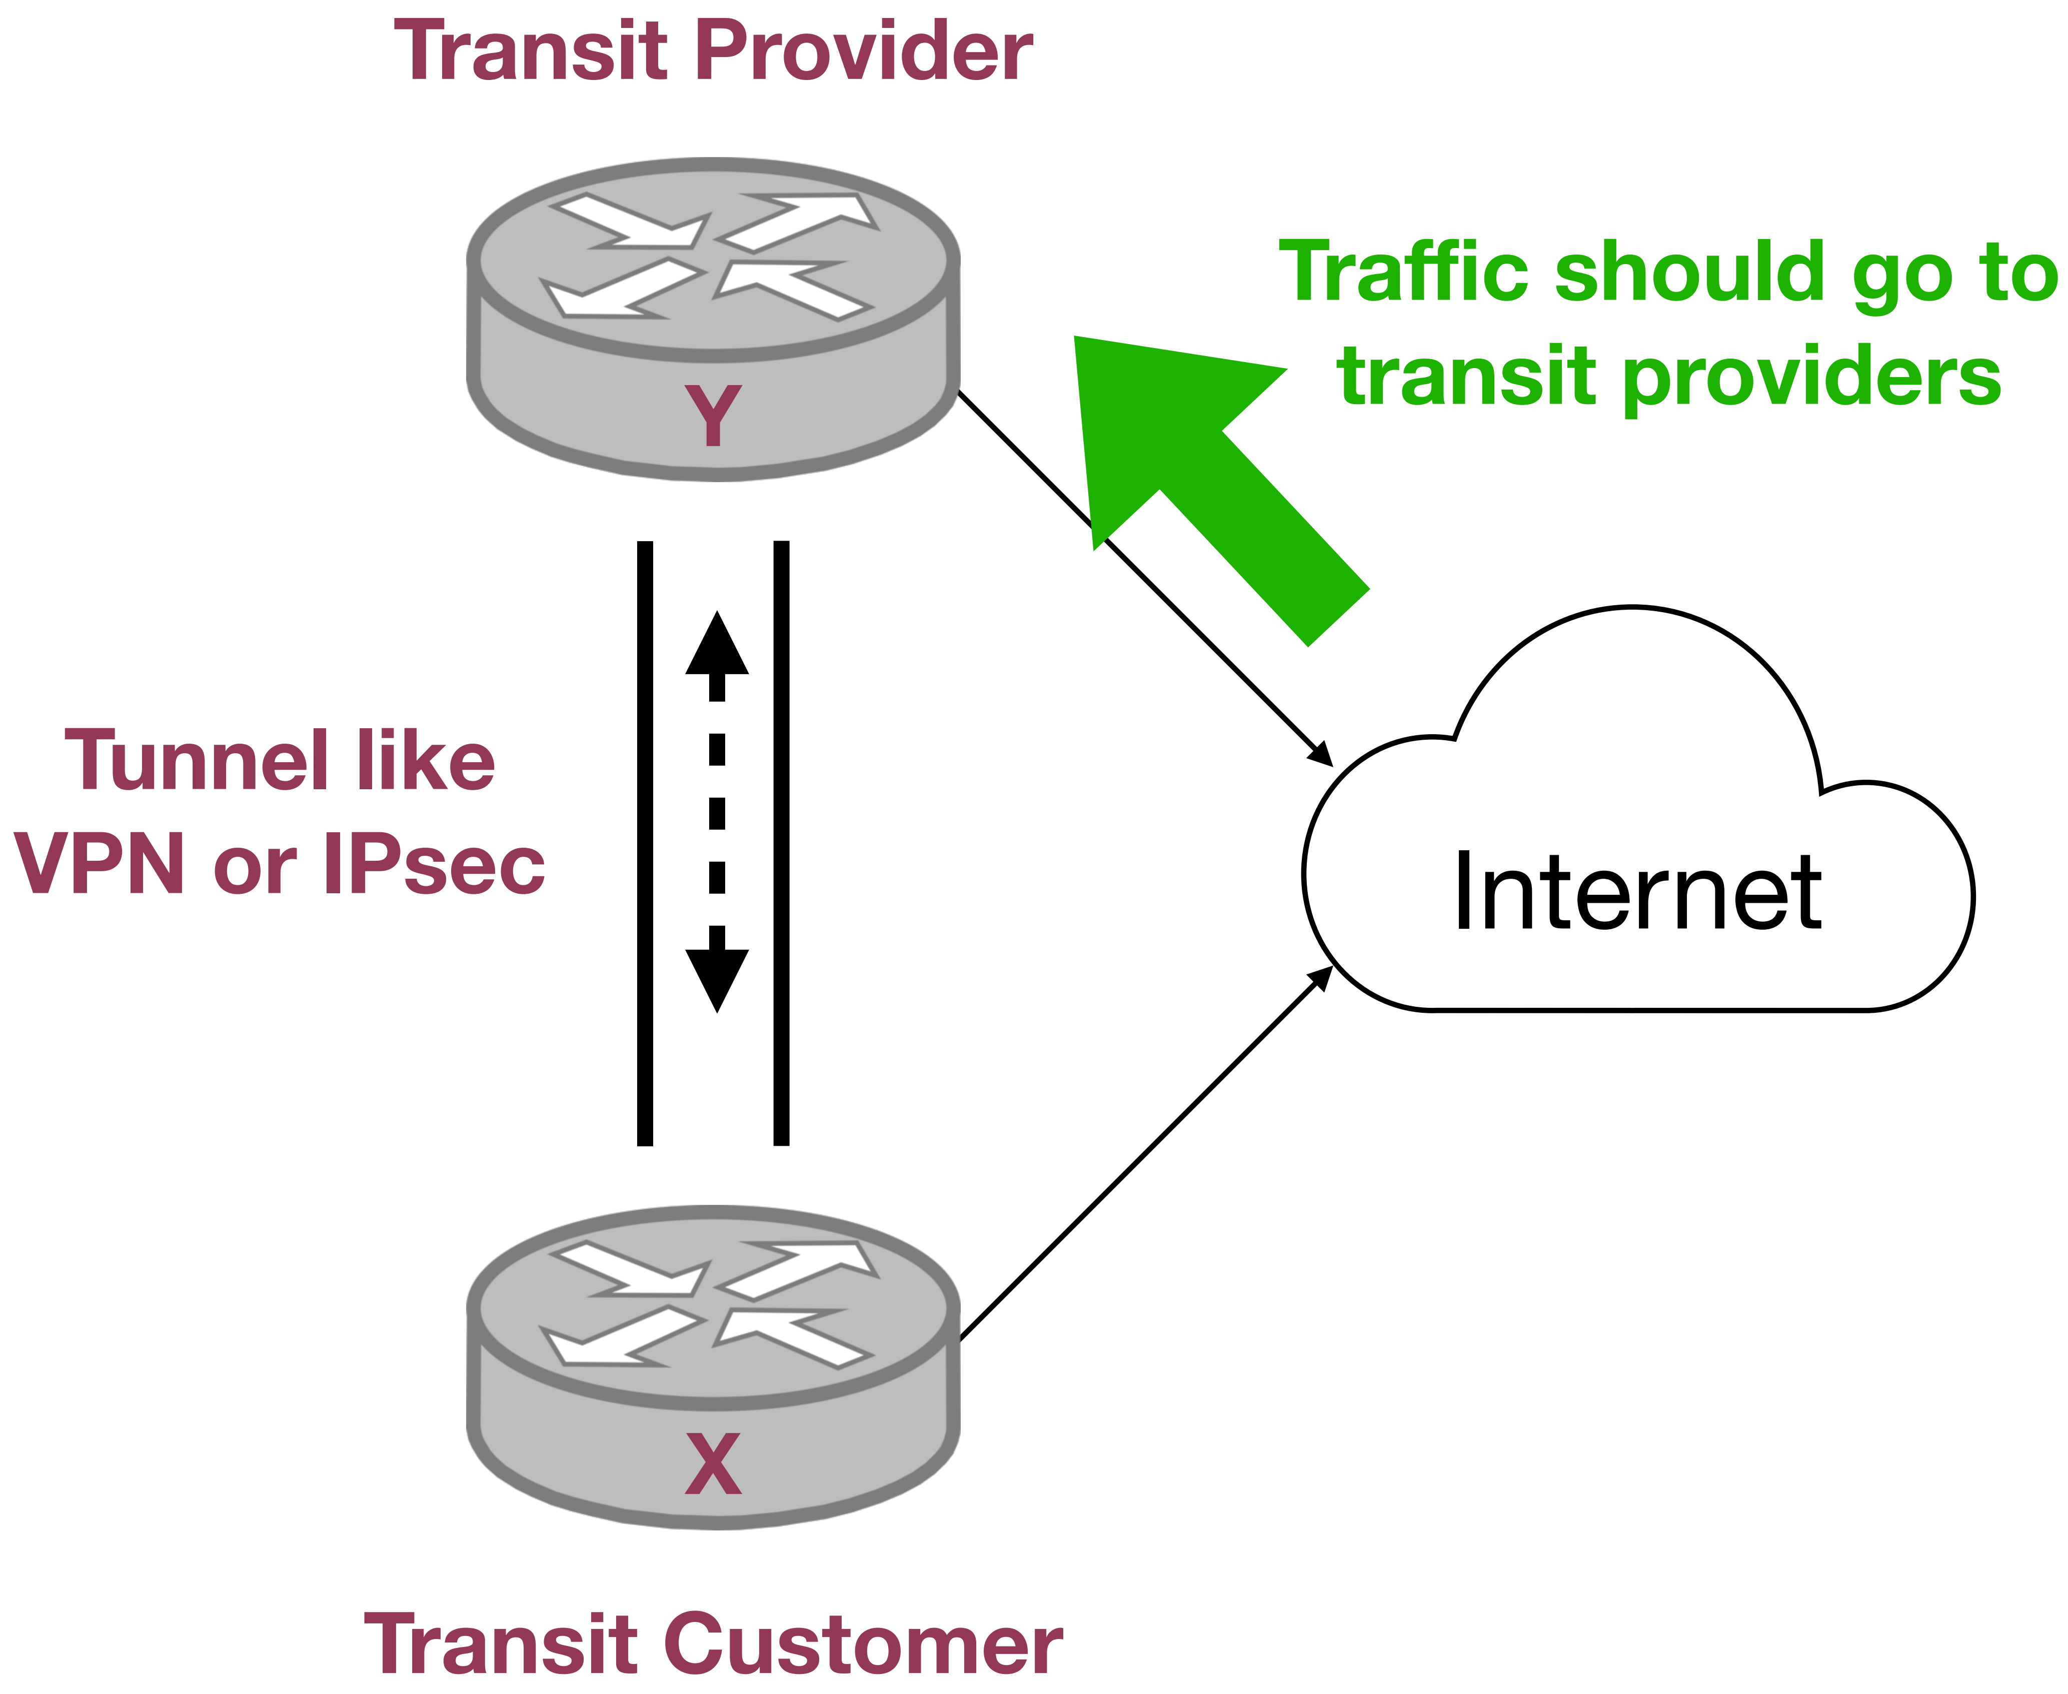

BGP Announcement:

129.21.0.0/16, AS X

129.21.0.0/20, AS Y

ROA:

129.21.0.0/16-20, AS X

Impact of RPKI-Invalid

- However, having alternative path doesn't means no impact at all.

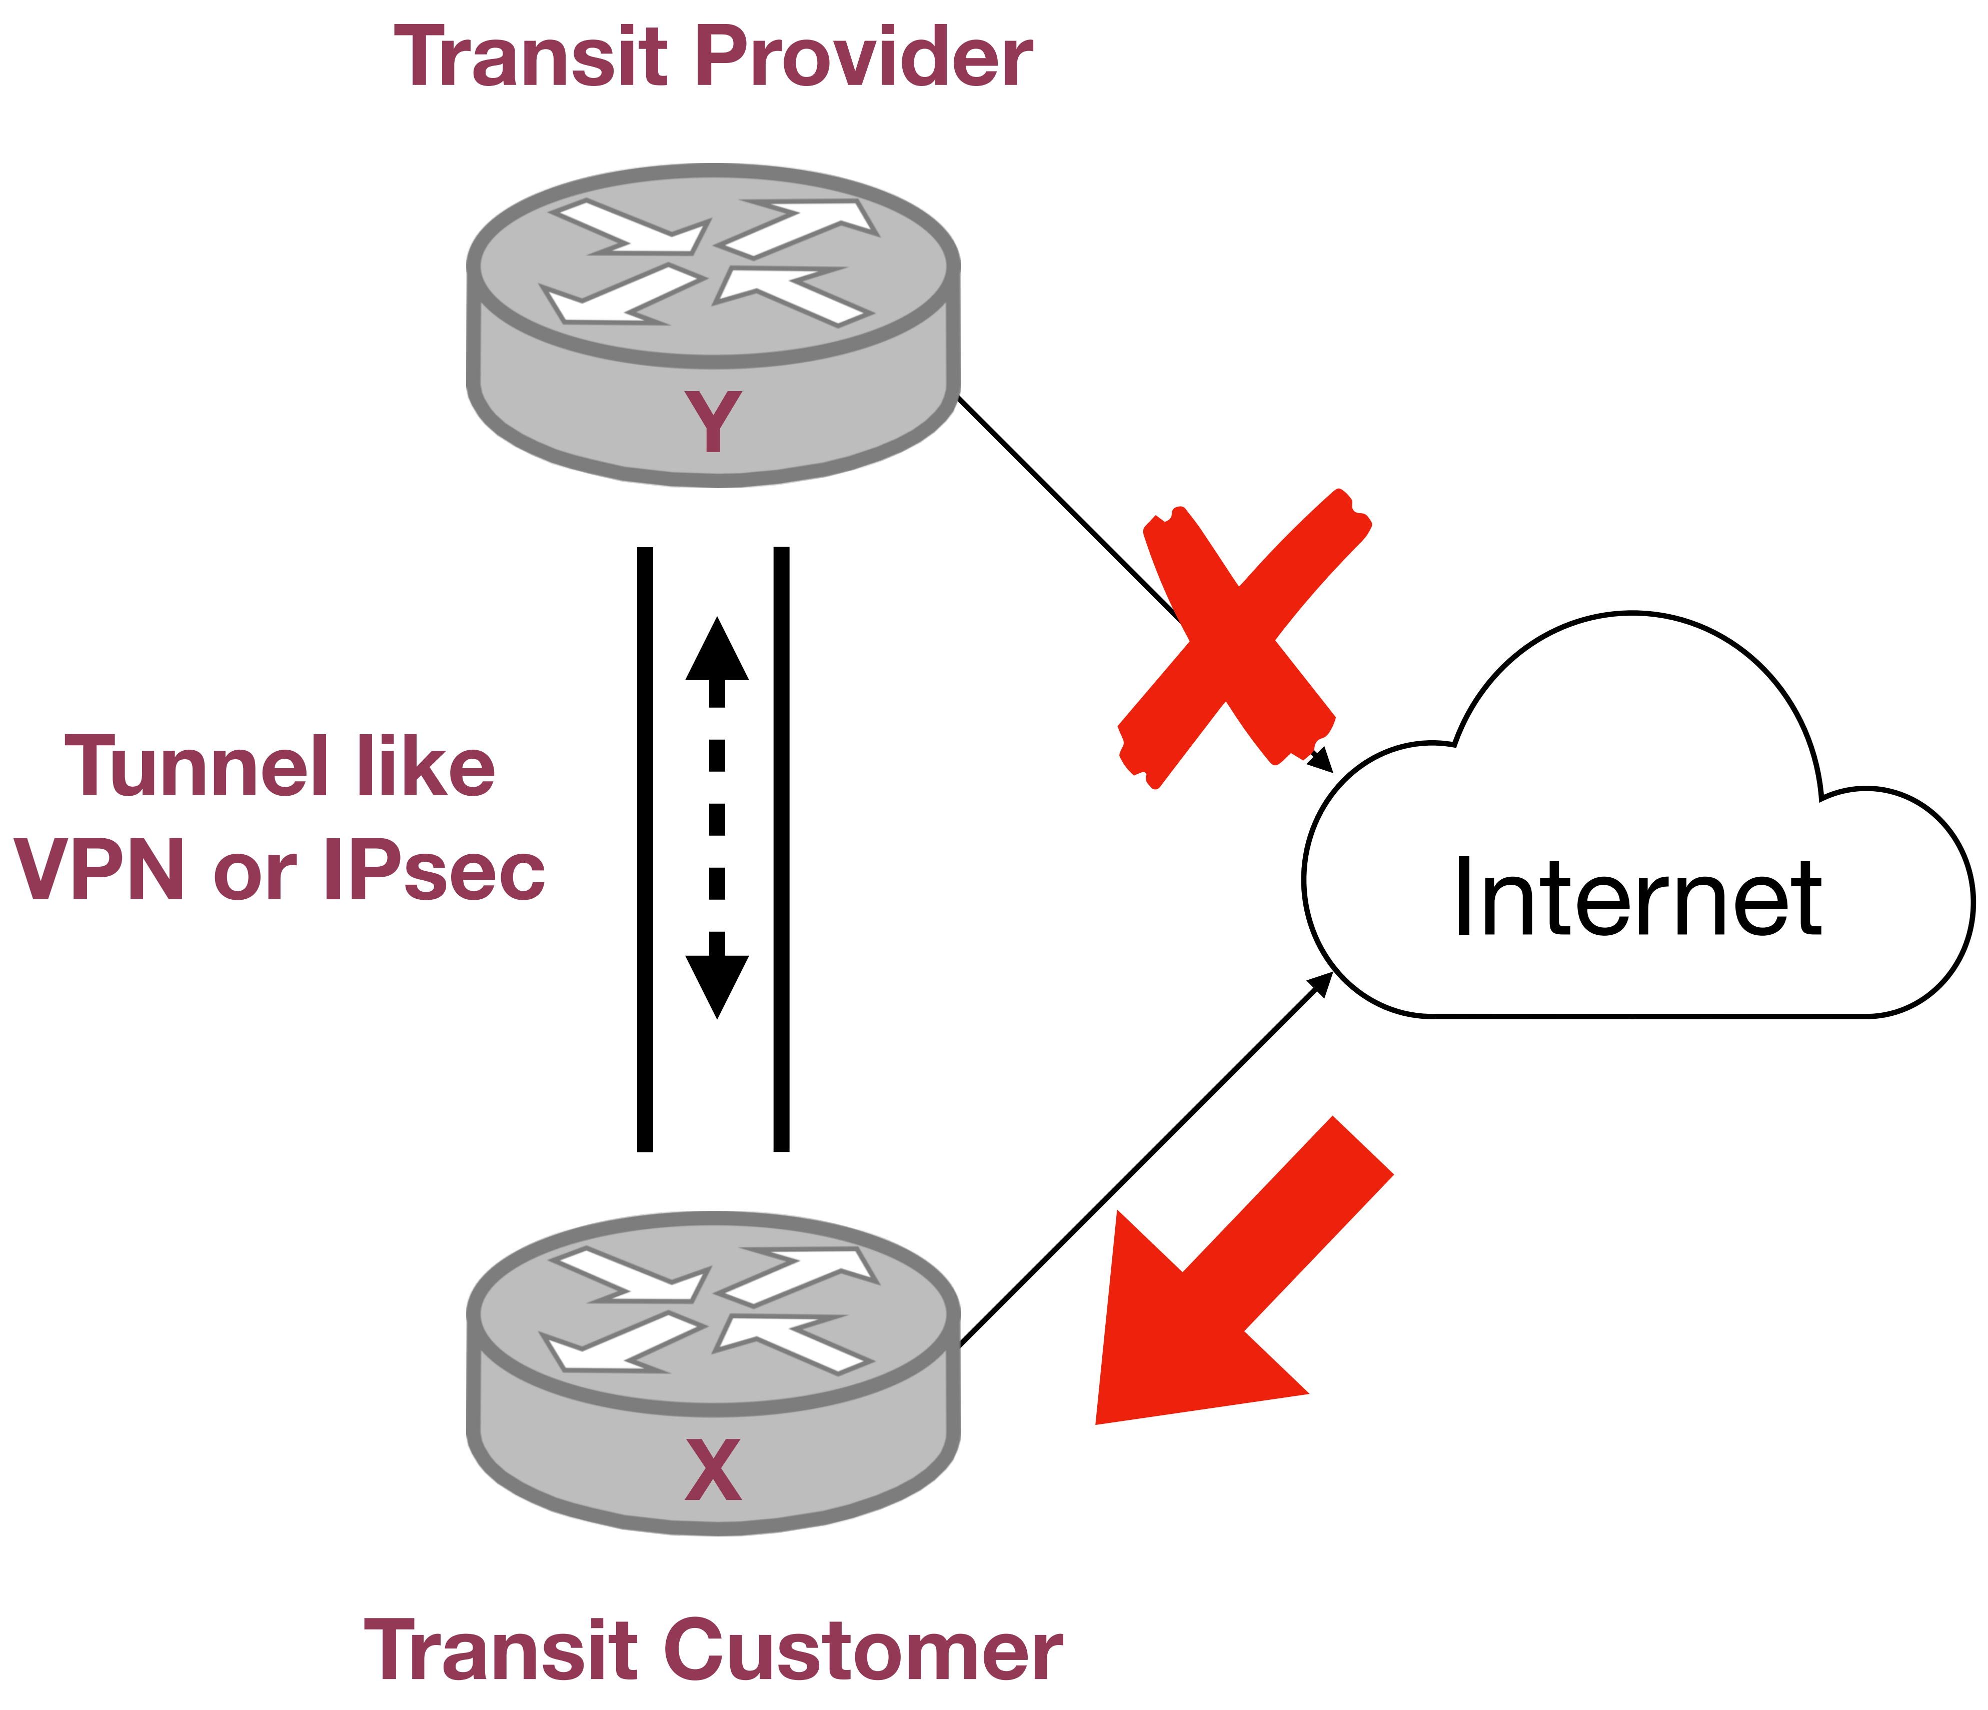

Transit Provider

Y

Tunnel like
VPN or IPsec

Internet

X

Transit Customer

BGP Announcement:

129.21.0.0/16, AS X

129.21.0.0/20, AS Y ✗

ROA:

129.21.0.0/16-20, AS X

Impact of RPKI-Invalid: Path Divergence

- However, having alternative path doesn’t means no impact at all.

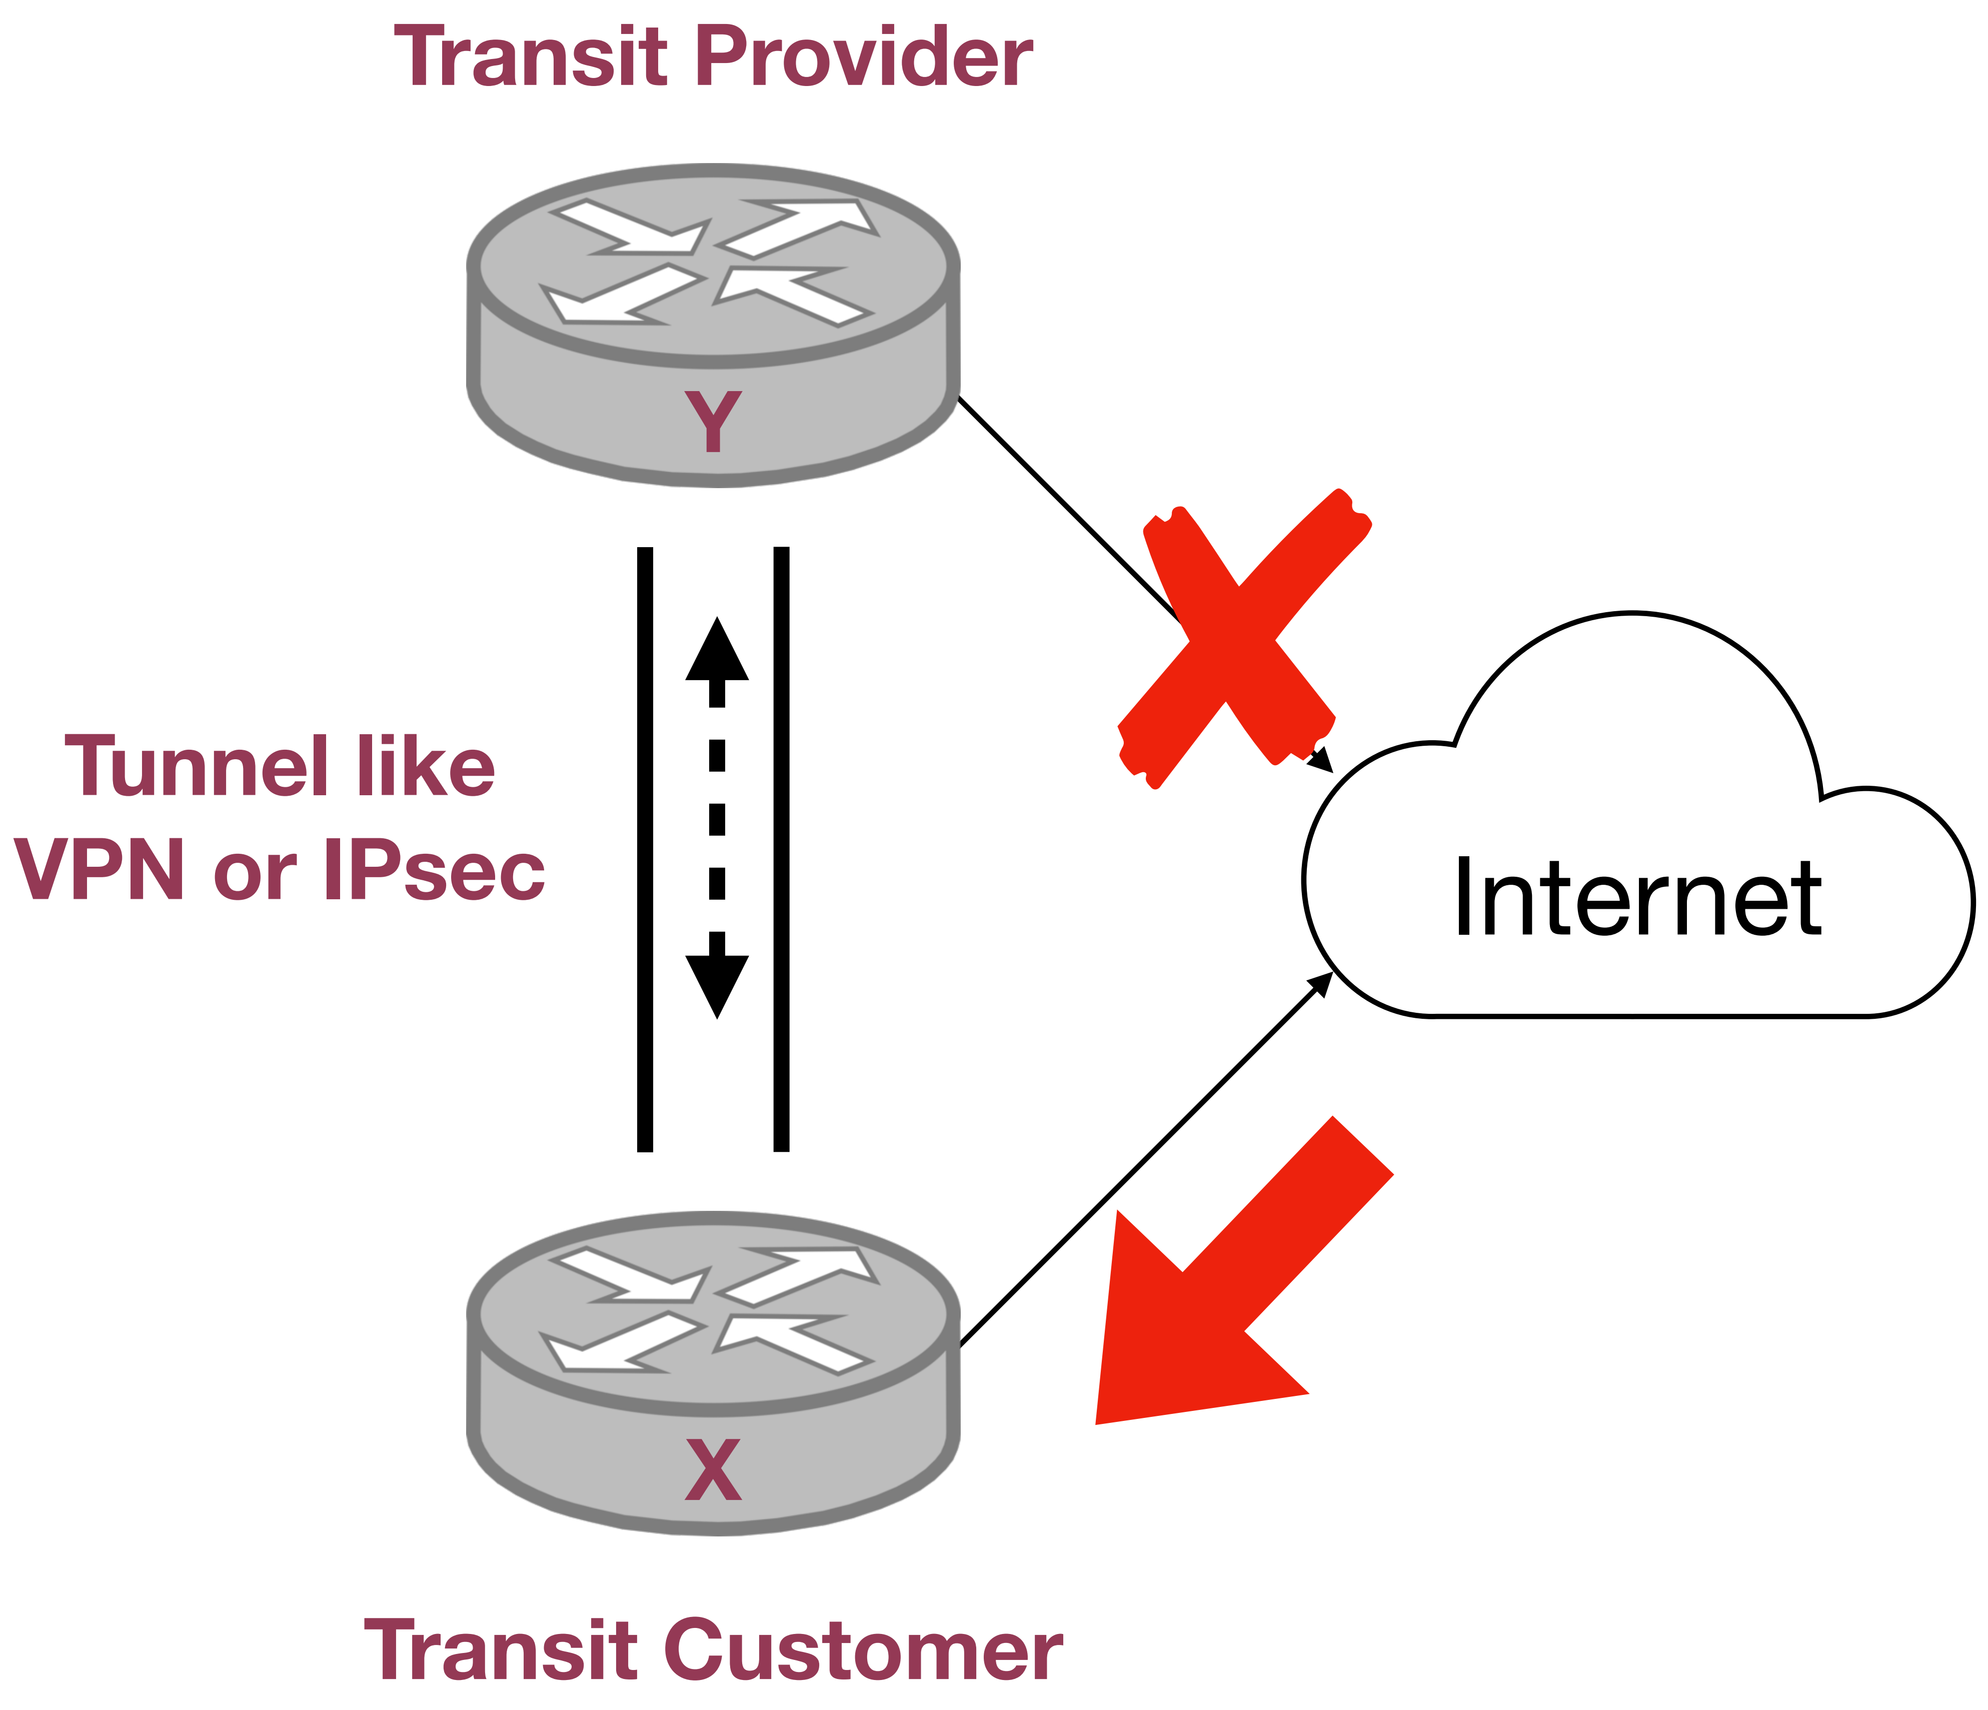

BGP Announcement:

129.21.0.0/16, AS X

129.21.0.0/20, AS Y ✗

ROA:

129.21.0.0/16-20, AS X

Impact of RPKI-Invalid: Path Divergence

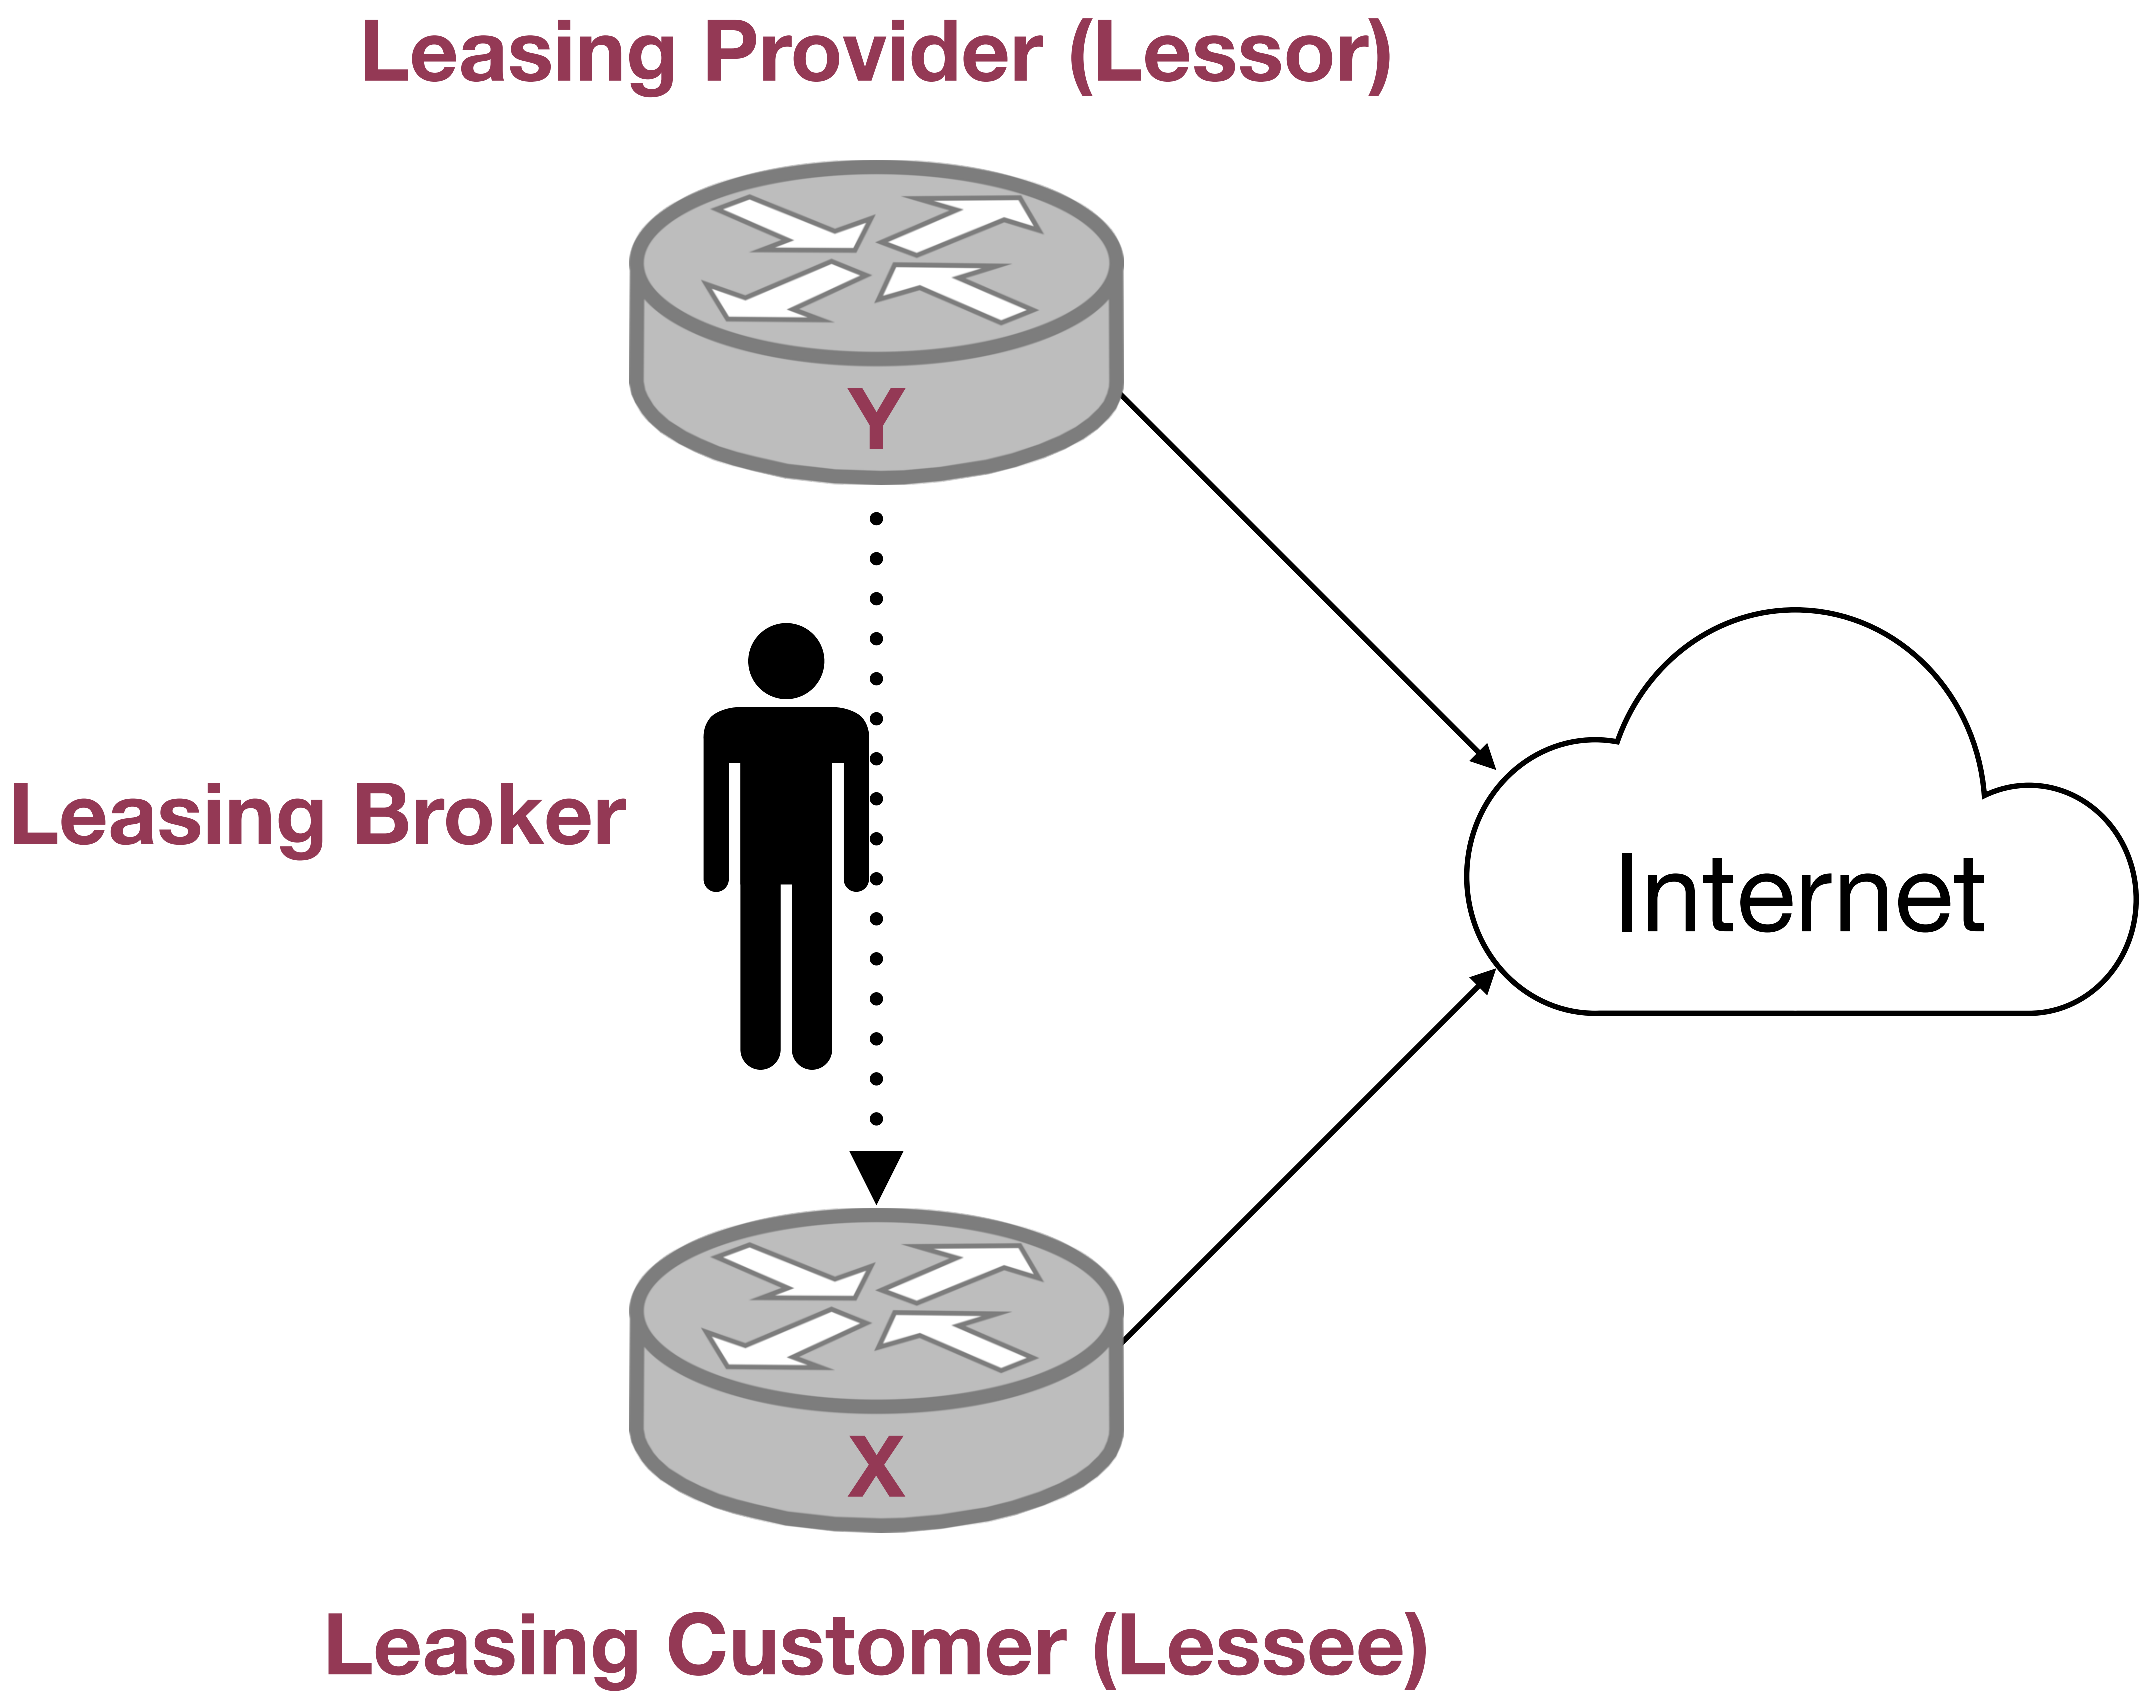

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

Impact of RPKI-Invalid: Path Divergence

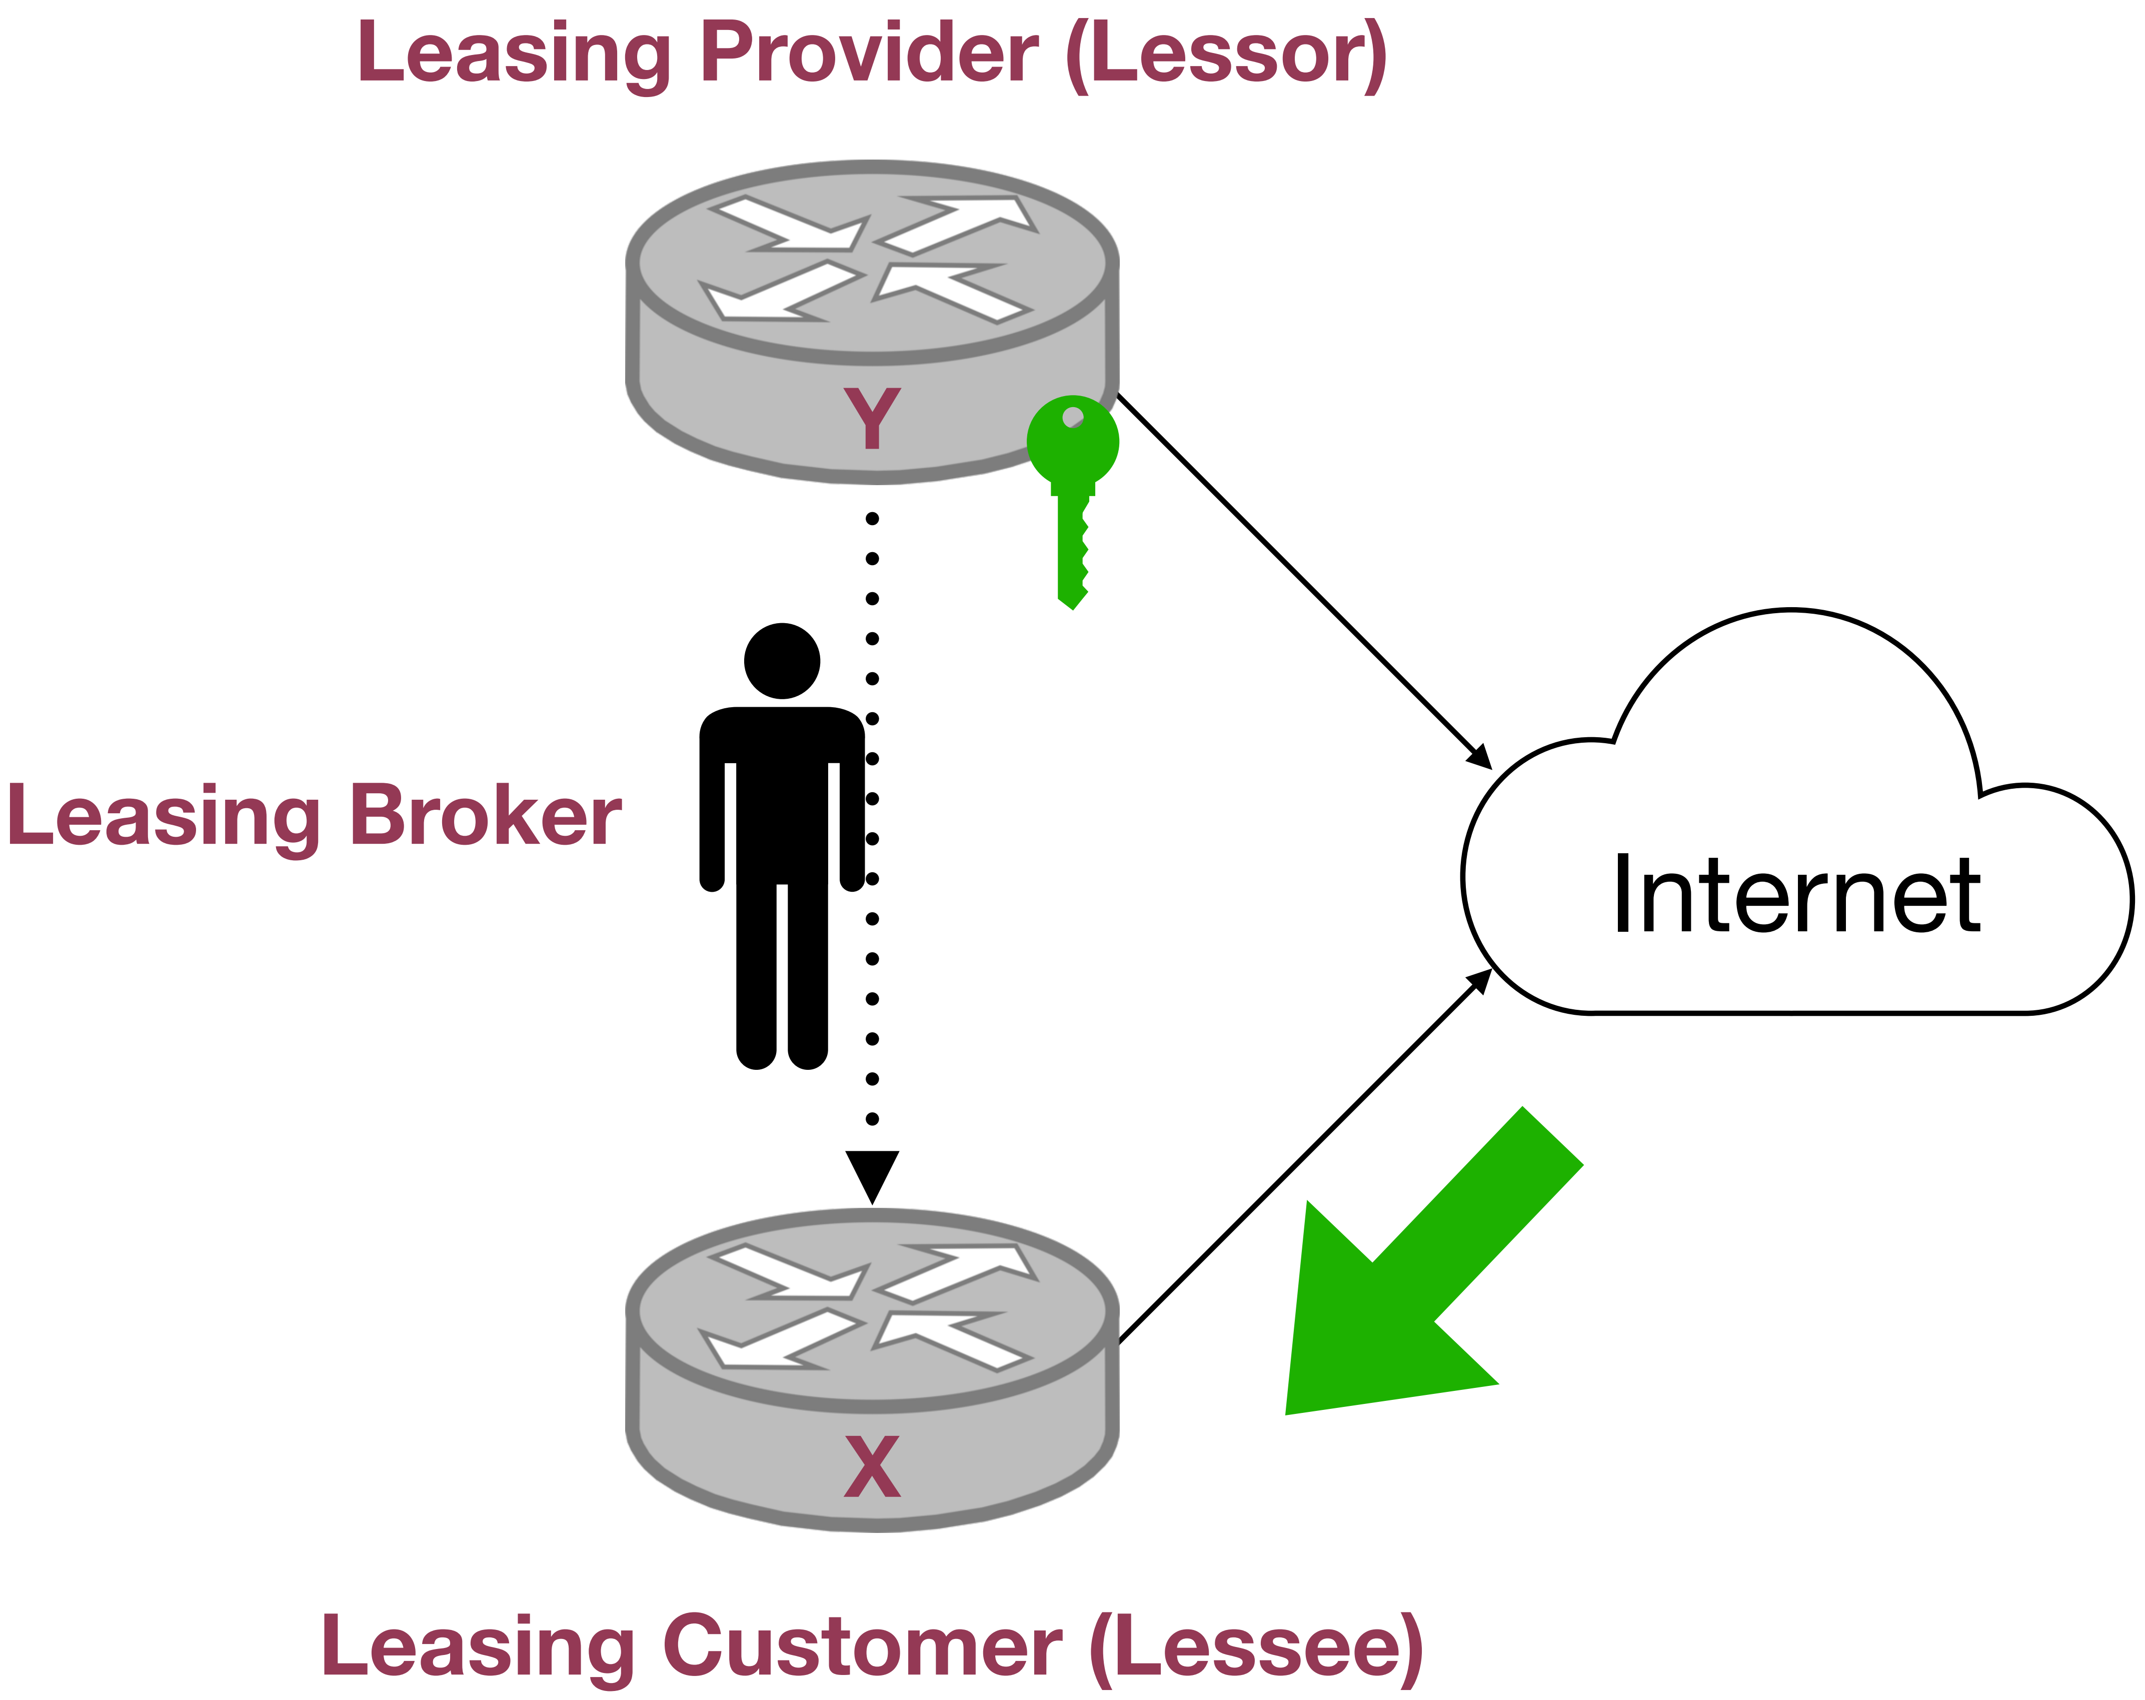

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 **: Access to ROA**

Impact of RPKI-Invalid: Path Divergence

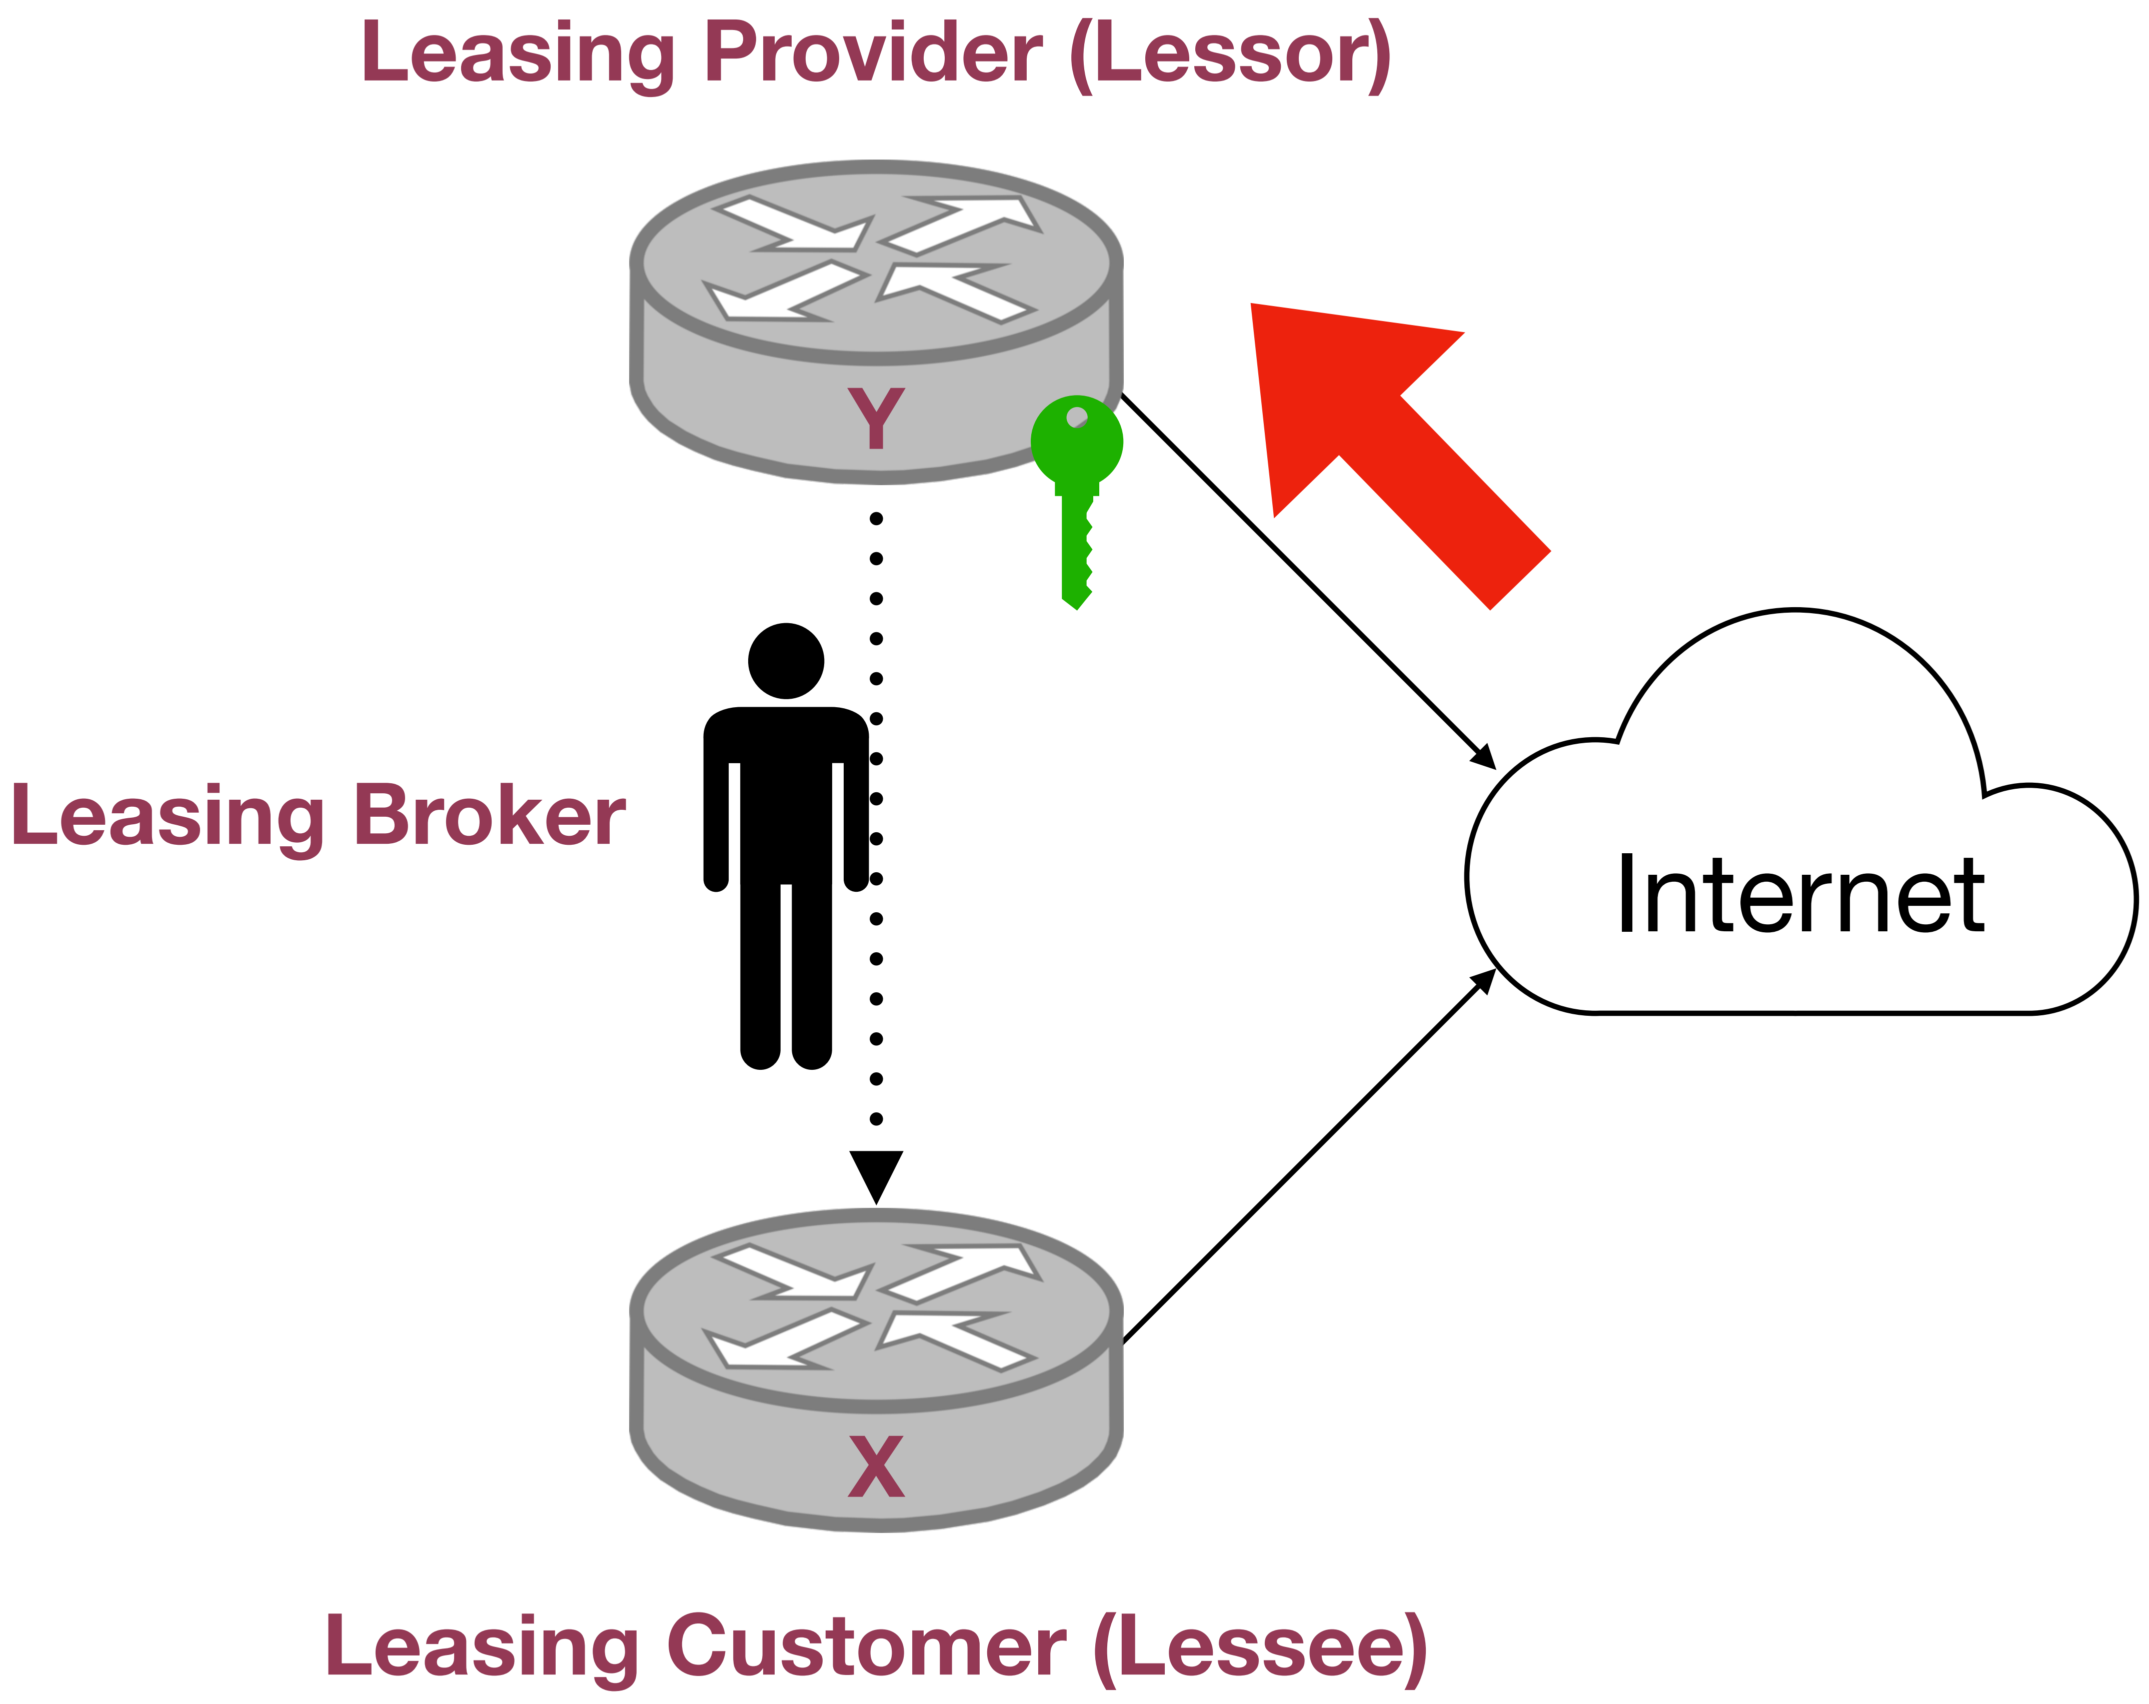

BGP Announcement:

129.21.0.0/16, AS Y

129.21.0.0/20, AS X ✗

ROA:

129.21.0.0/16-20, AS Y

🔑 : Access to ROA

Impact of RPKI-Invalid: Path Divergence

- Path divergence is more common (18.5%) compared to disconnection

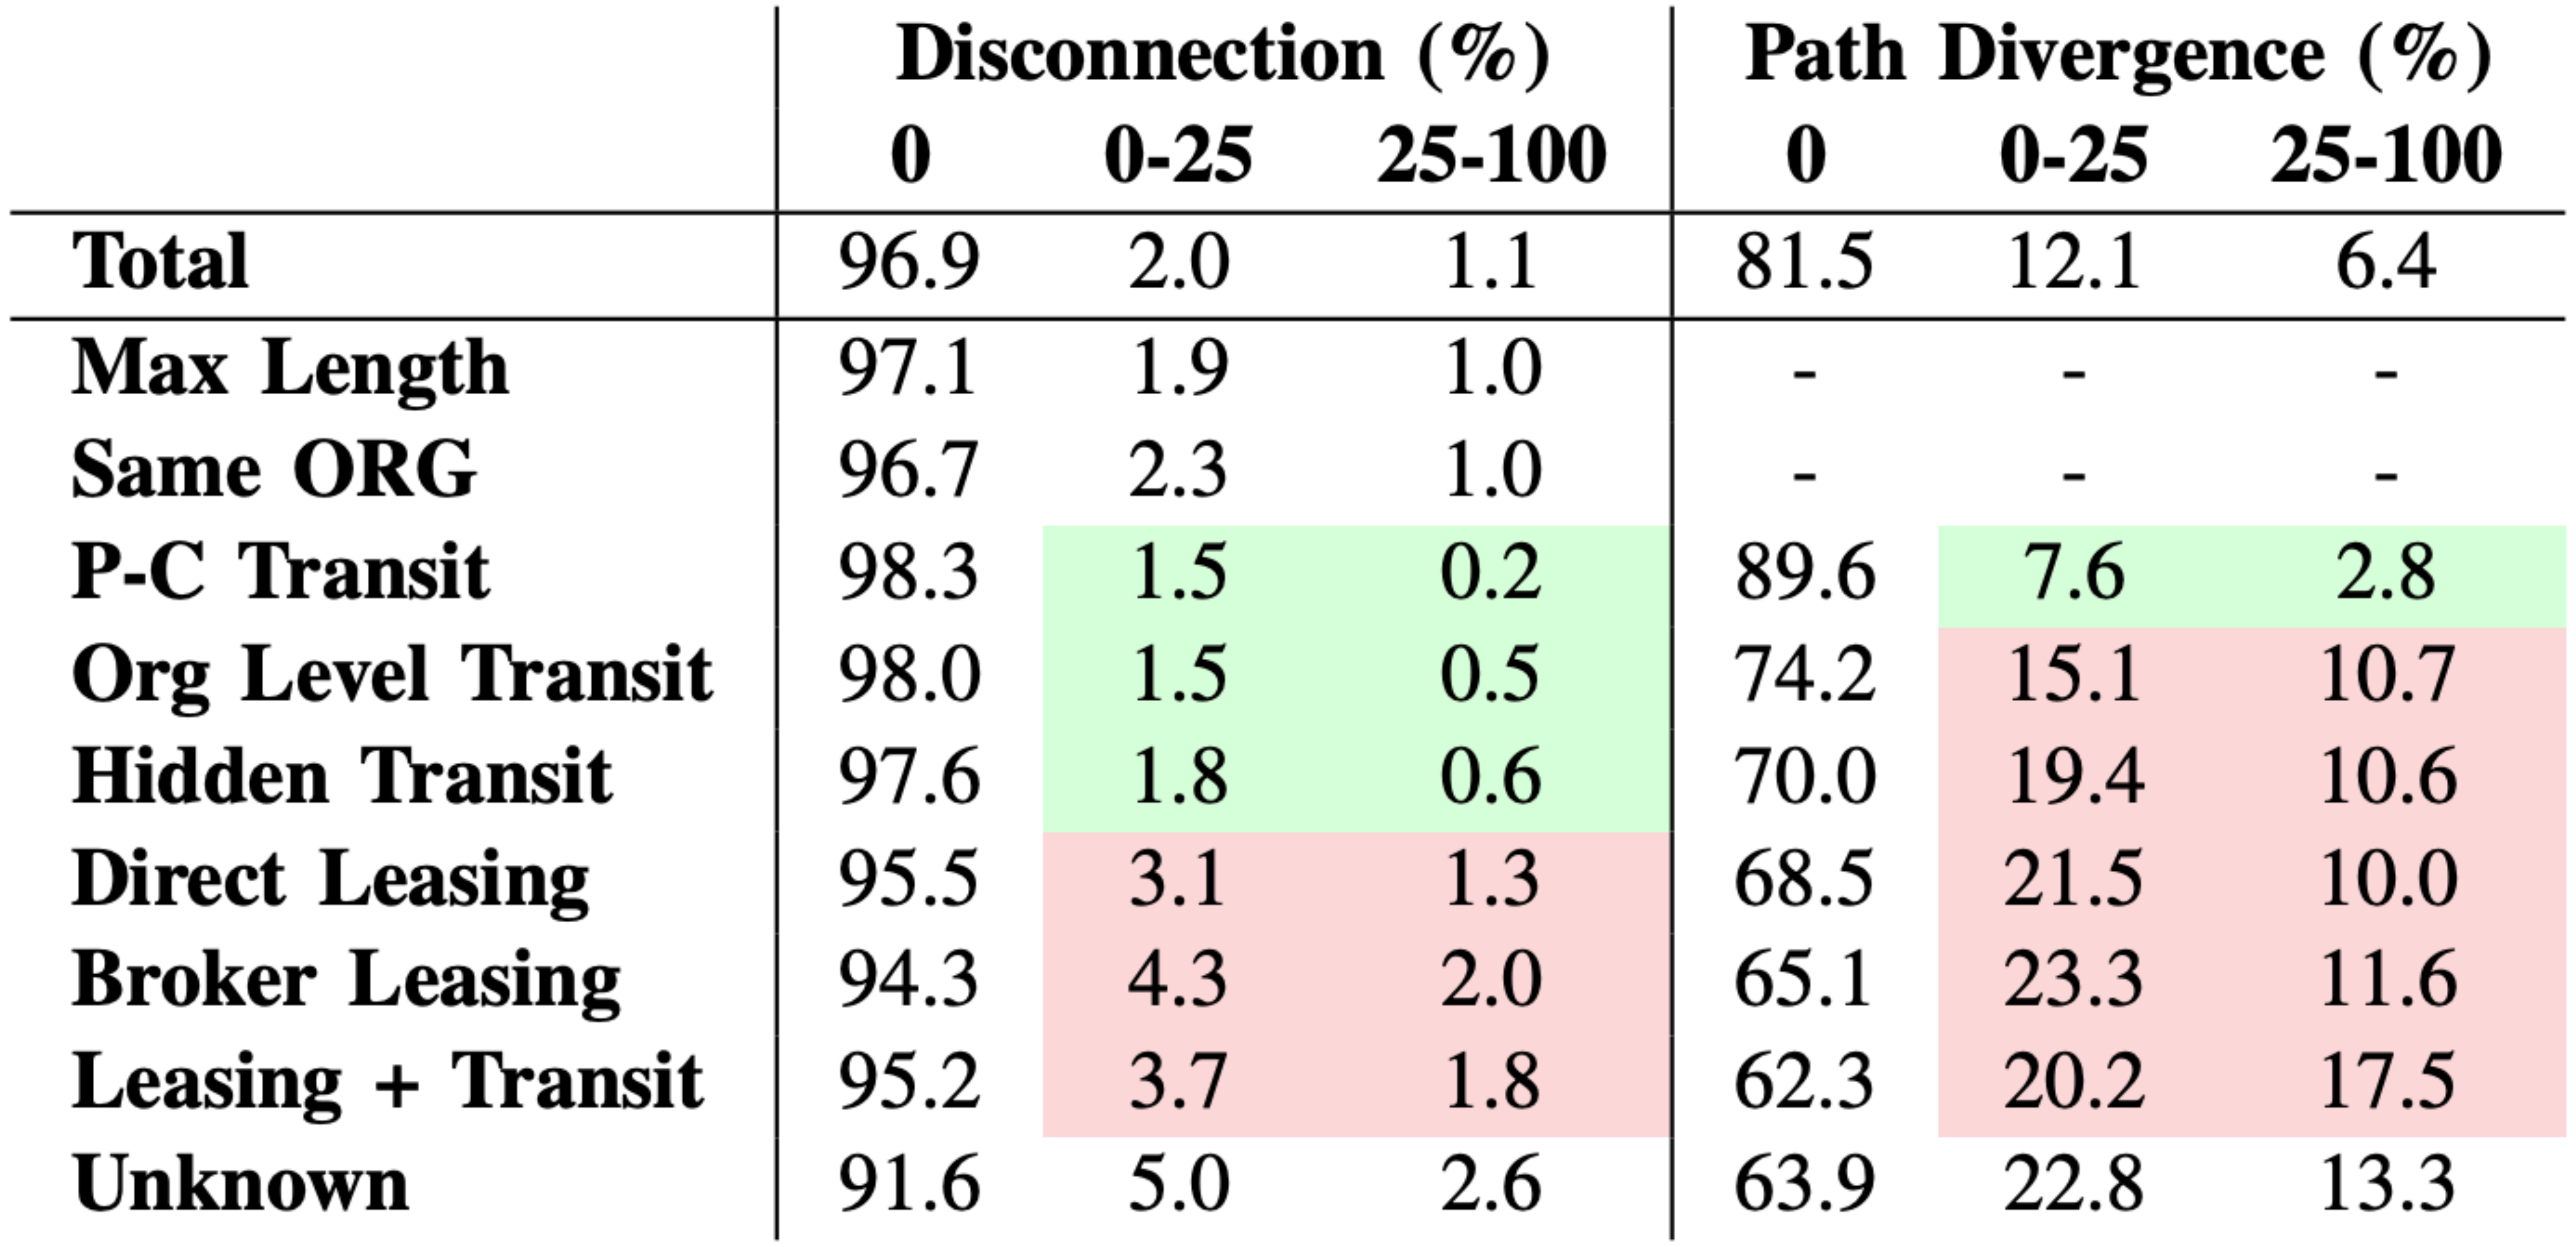

	Disconnection (%)			Path Divergence (%)		
	0	0-25	25-100	0	0-25	25-100
Total	96.9	2.0	1.1	81.5	12.1	6.4
Max Length	97.1	1.9	1.0	-	-	-
Same ORG	96.7	2.3	1.0	-	-	-
P-C Transit	98.3	1.5	0.2	89.6	7.6	2.8
Org Level Transit	98.0	1.5	0.5	74.2	15.1	10.7
Hidden Transit	97.6	1.8	0.6	70.0	19.4	10.6
Direct Leasing	95.5	3.1	1.3	68.5	21.5	10.0
Broker Leasing	94.3	4.3	2.0	65.1	23.3	11.6
Leasing + Transit	95.2	3.7	1.8	62.3	20.2	17.5
Unknown	91.6	5.0	2.6	63.9	22.8	13.3

Impact of RPKI-Invalid: Path Divergence

- Path divergence is more common (18.5%) compared to disconnection
- Leading to additional hops, impact performance and protection, and increases the risk of MITM

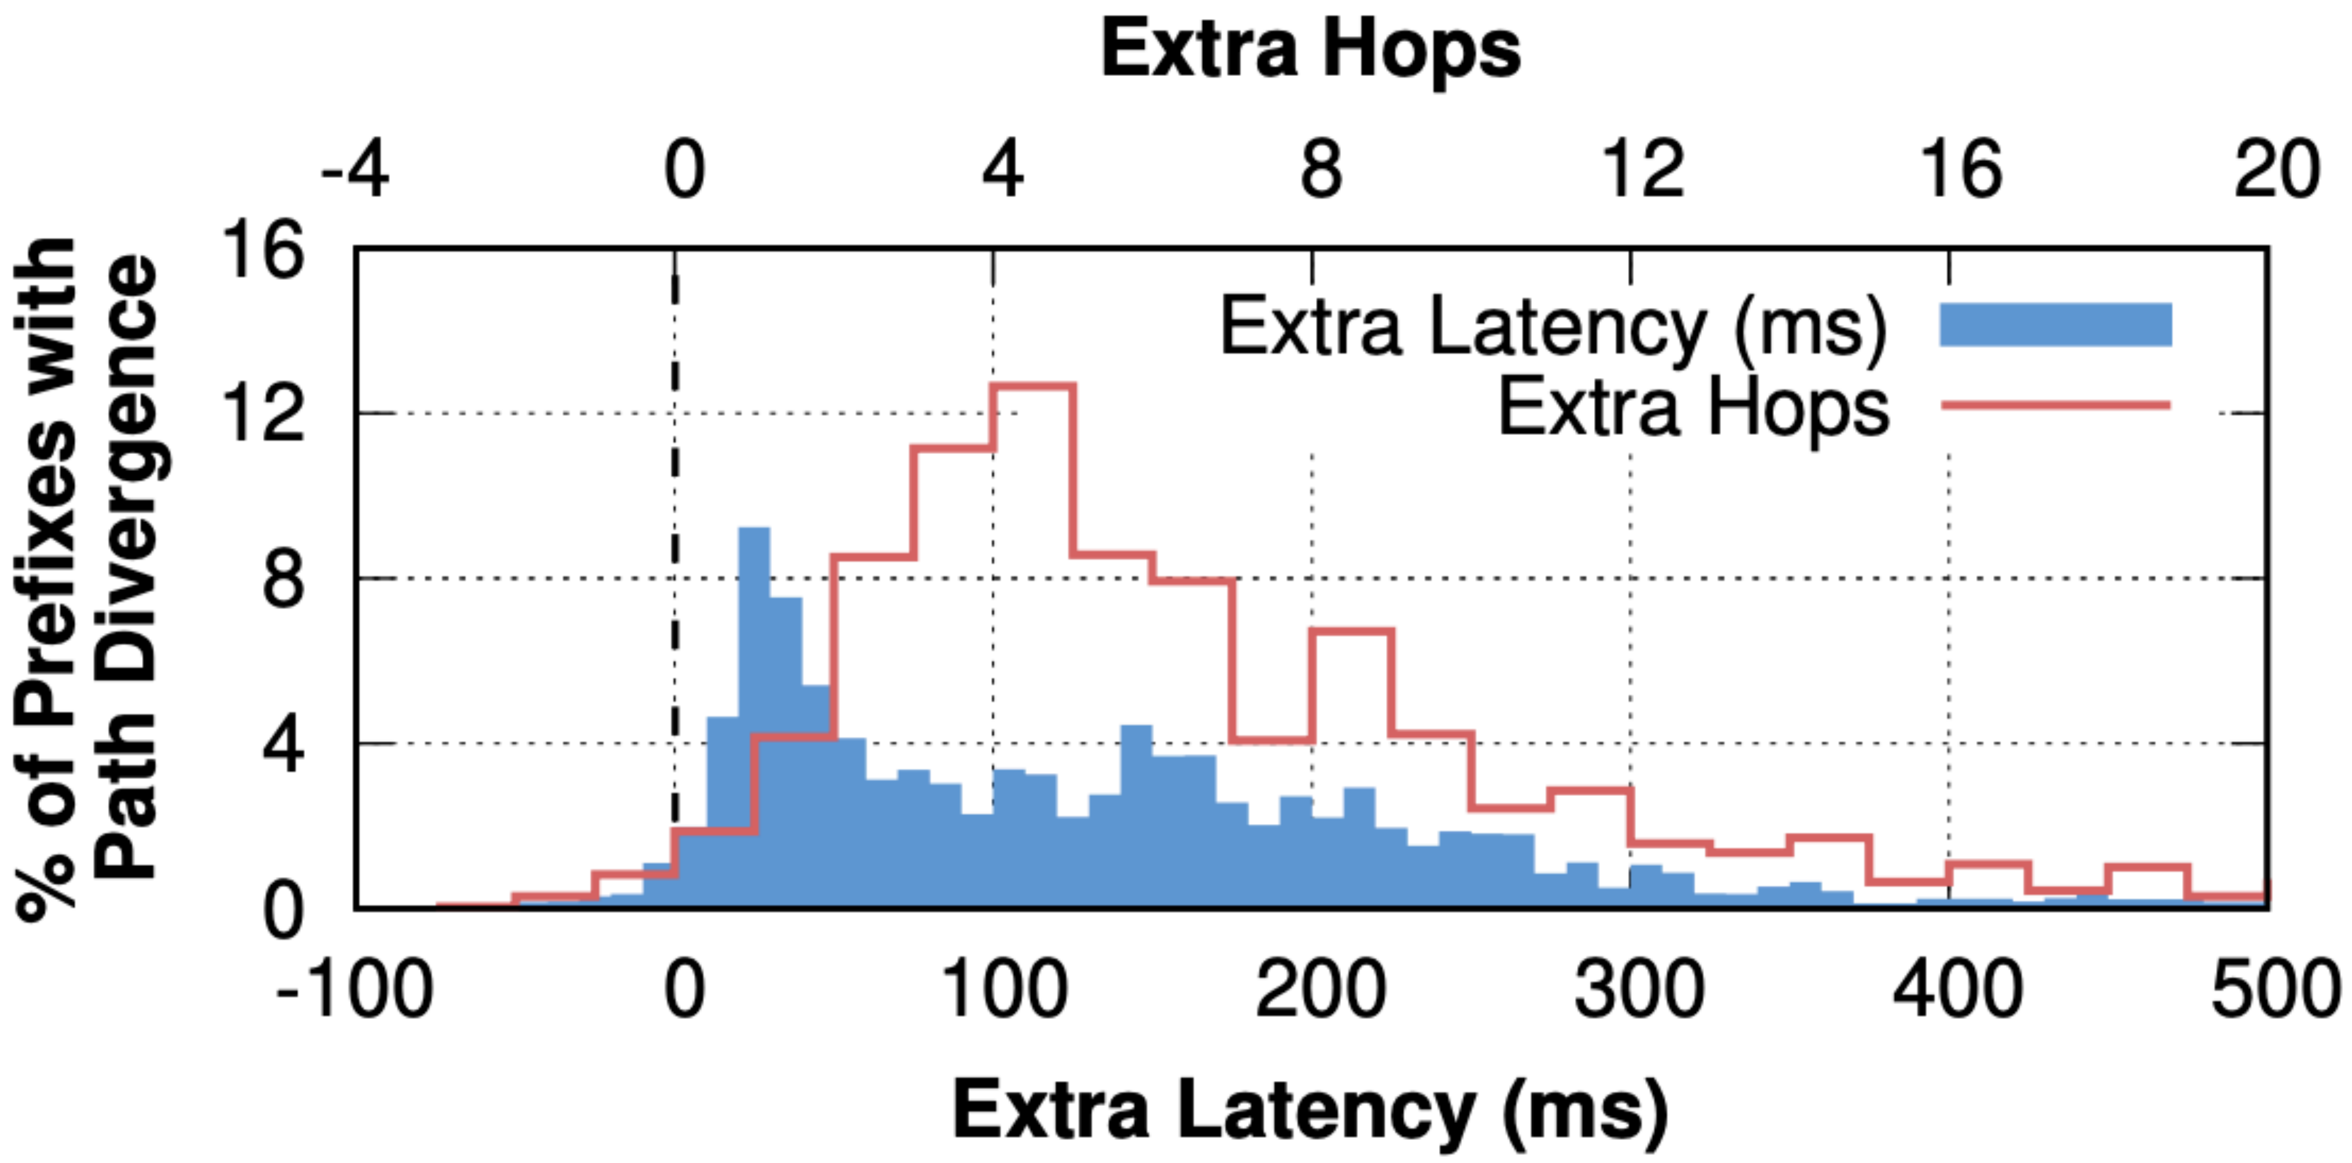

Fixing of RPKI-Invalid

- RPKI-Invalid prefixes are fixed quickly only if they result in disconnection

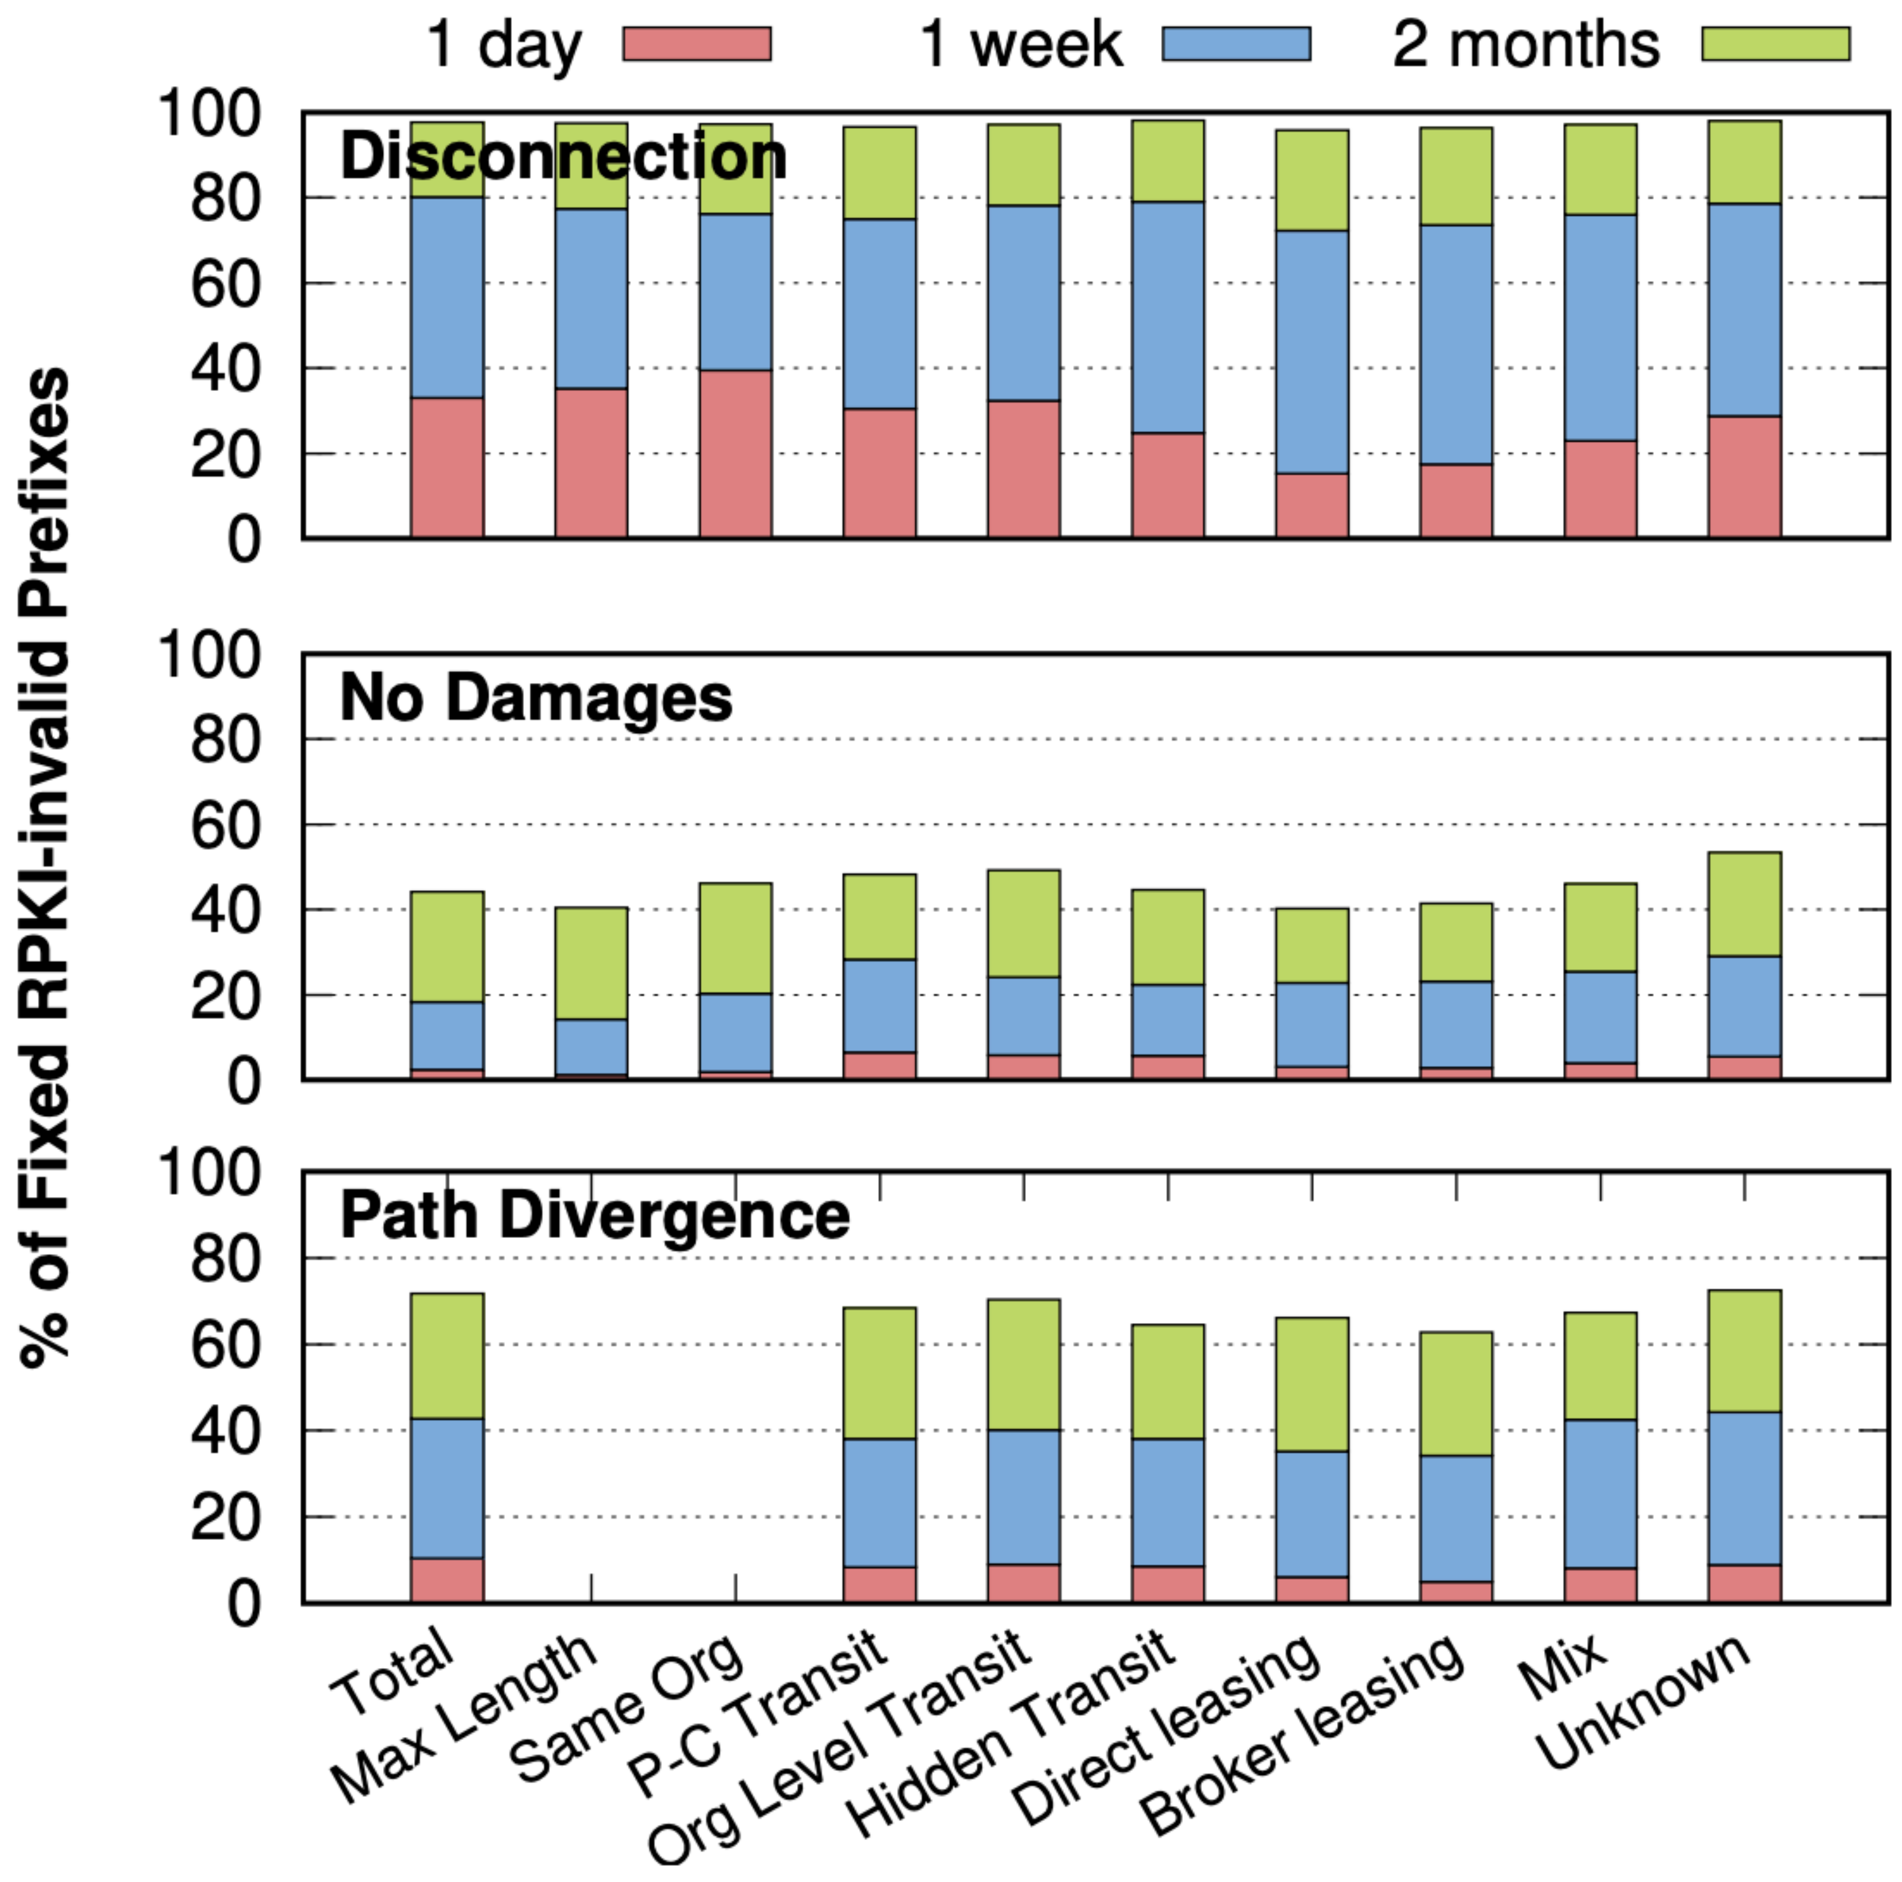

Summary

- Why RPKI-Invalid happens: Misconfiguration is the major cause of RPKI-Invalid (**96.9%**), and the complexity of **transit and leasing business** making them hard to detect and mitigate.
- Impact of RPKI-Invalid: RPKI-Invalid could result in not only disconnection, but also path divergence, which is harder to detect but still **impacting security and performance**.

Summary

- Why RPKI-Invalid happens: Misconfiguration is the major cause of RPKI-Invalid (**96.9%**), and the complexity of **transit and leasing business** making them hard to detect and mitigate.

- Impact of RPKI-Invalid: RPKI-Invalid could result in not only disconnection, but also path divergence, which is harder to detect but still **impacting security and performance**.

- We need better management pipeline, better registration record delegation, and monitoring/automated tools to help ASes prevent misconfigurations

Summary

- Why RPKI-Invalid happens: Misconfiguration is the major cause of RPKI-Invalid (**96.9%**), and the complexity of **transit and leasing business** making them hard to detect and mitigate.

- Impact of RPKI-Invalid: RPKI-Invalid could result in not only disconnection, but also path divergence, which is harder to detect but still **impacting security and performance**.

- We need better management pipeline, better registration record delegation, and monitoring/automated tools to help ASes prevent misconfigurations

- Interview/Survey for 16 ISPs and leasing brokers

Thanks for Listening!

Code & Dataset: https://roa-misconfig.netsecurelab.org/

Weitong Li[1], **Tao Wan**[2], **and Tijay Chung**[1]

[1]Virginia Tech, [2]CableLabs