

RoVista: Measuring and Understanding the Route Origin Validation (ROV) in RPKI


Tijay Chung

(tijay@vt.edu, <https://tijay.github.io>)

Assistant Professor at Virginia Tech


Routing 101: Border Gateway Protocol (BGP)

- Each network resource owner announces its IP prefixes to the rest of routers, so that they can learn the path towards the owner.
- However, it has NONE of security consideration such as authorization



Resource PKI (Public Key Infrastructure)


- Public Key Infrastructure framework designed to secure Internet's routing structure; specifically BGP (developed starting in 2008)
- Currently more than 50% of IP spaces are verifiable with RPKI



Resource PKI (Public Key Infrastructure)

- Public Key Infrastructure framework designed to secure Internet's routing structure; specifically BGP (developed starting in 2008)
- Currently more than 50% of IP spaces are verifiable with RPKI

(Cryptographically verifiable)
Prefix-to-AS Mapping Database



Route Origin Authorization vs. Route Origin Validation



Route Origin Authorization vs. Route Origin Validation

Resource owner needs to create an assertion
(called ROA) and upload it to registry



Router

BGP announcement



Owner

AS 4385

129.21.0.0/16

Route Origin Authorization vs. Route Origin Validation



Router

BGP announcement



Owner

AS 4385

129.21.0.0/16

Router needs to download ROAs and
verify BGP announcements against them

Two questions

- How network operators use RPKI to “claim” their IP addresses?
- How network operators also use RPKI to “filter” invalid BGP announcements?

Two questions



Answering this question is “relatively”
straightforward

- How network operators use RPKI to “claim” their IP addresses?
- How network operators also use RPKI to “filter” invalid BGP announcements?

Two questions

- How network operators use RPKI to “claim” their IP addresses?
- How network operators also use RPKI to “filter” invalid BGP announcements?



This is not straightforward

Previous approaches (1)

<https://isbgpsafeyet.com/>

Is BGP **safe** yet? *No.*

Border Gateway Protocol (BGP) is the postal service of the Internet. It's responsible for looking at all of the available paths that data could travel and picking the best route.

Unfortunately, it isn't secure, and there have been some [major Internet disruptions](#) as a result. But fortunately there is a way to make it secure.

ISPs and other major Internet players (Sprint, Verizon, and others) would need to implement a certification system, called [RPKI](#).

[Test your ISP](#)

[Read FAQ](#)

Previous approaches (1)

<https://isbgpsafeyet.com/>

Is BGP **safe** yet? *No.*

Border Gateway Protocol (BGP) is the postal service of the Internet. It's responsible for looking at all of the available paths that data could travel and picking the best route.





Unfortunately, it isn't secure, and there have been some [major Internet disruptions](#) as a result. But fortunately there is a way to make it secure.

ISPs and other major Internet players (Sprint, Verizon, and others) would need to implement a certification system, called [RPKI](#).

[Test your ISP](#)

[Read FAQ](#)

valid.rpki.cloudflare.com

Announced By		
Origin AS	Announcement	Description
AS13335	104.16.0.0/12	 Cloudflare, Inc.
AS13335	104.18.32.0/19	  Cloudflare, Inc.
AS13335	104.18.32.0/20	  Cloudflare, Inc.
AS13335	104.18.47.0/24	  Cloudflare, Inc.

Previous approaches (1)

<https://isbgpsafeyet.com/>

Is BGP **safe** yet? *No.*

Border Gateway Protocol (BGP) is the postal service of the Internet. It's responsible for looking at all of the available paths that data could travel and picking the best route.

Unfortunately, it isn't secure, and there have been some [major Internet disruptions](#) as a result. But fortunately there is a way to make it secure.

ISPs and other major Internet players (Sprint, Verizon, and others) would need to implement a certification system, called [RPKI](#).

[Test your ISP](#)

[Read FAQ](#)

valid.rpki.cloudflare.com

Announced By		
Origin AS	Announcement	Description
AS13335	104.16.0.0/12	 Cloudflare, Inc.
AS13335	104.18.32.0/19	  Cloudflare, Inc.
AS13335	104.18.32.0/20	  Cloudflare, Inc.
AS13335	104.18.47.0/24	  Cloudflare, Inc.

invalid.rpki.cloudflare.com

Announced By		
Origin AS	Announcement	Description
AS13335	103.21.244.0/24	  Cloudflare, inc.


Previous approaches (2)



- Crowd-source based spreadsheet managed by network operators
- <http://rpki.exposed>


	May 4th 2020	Rejecting invalids	Rejecting invalids	Rejecting invalids		
Carrier	ASN	Transits	Peers	Customers	ROAs	Status
NTT	2914	n/a	yes	yes	done	done
GTT	3257	n/a	yes	yes	done	done
AT&T	7018	n/a	yes	no	in progress	in progress
Telia	1299	n/a	yes	yes	done	done
Workonline	37271	yes	yes	yes	done	done
Seacom	37100	yes	some	yes		done
KPN Eurorings	286	n/a // yes (*)	yes	yes	done	done
Freethought	41000	yes	yes	yes	done	done
Fusix	57866	yes	yes	yes	done	done
BIT	12859	yes	yes	yes	done	done
Tuxis	197731	yes	yes	yes	done	done
MaxiTEL (NL)	61349	yes	yes	yes	done	done
ColoClue	8283	yes	yes	no	done	done
Fiber Telecom	41327	yes	yes	yes	done	done
Sentia BV	8315	yes	yes	yes	done	done
Cadence Networks	47638	yes	yes	yes	done	done
Atom86	8455	yes	yes	yes	done	done
AMS-IX	6777	n/a	yes	n/a	done	done
NetNod	52005	n/a	yes	n/a		done

Previous approaches (3)

- Official blogpost, mailing list, and so on.

 **Orange Wholesale** 
@OrangeWholesale

NEW We're glad to announce that we have now fully completed the **#RPKI** implementation in our **#IPTransit** network **NEW** 
Is your **#telecom** business ready? Already client? You can check your status via RPKI Monitor on our Customer Portal
Learn more about **#AS5511**  oran.ge/39qZ1Xl



11:00 AM · Jun 27, 2022

AT&T/as7018 now drops invalid prefixes from peers

Jay Borkenhagen [jayb at braeburn.org](mailto:jayb@braeburn.org)
Mon Feb 11 14:53:45 UTC 2019

- Previous message (by thread): [BGP topological vs centralized route reflector](#)
- Next message (by thread): [AT&T/as7018 now drops invalid prefixes from peers](#)
- Messages sorted by: [\[date \]](#) [\[thread \]](#) [\[subject \]](#) [\[author \]](#)

FYI:

The AT&T/as7018 network is now dropping all RPKI-invalid route announcements that we receive from our peers.

We continue to accept invalid route announcements from our customers, at least for now. We are communicating with our customers whose invalid announcements we are propagating, informing them that these routes will be accepted by fewer and fewer networks over time.


Thanks to those of you who are publishing ROAs in the RPKI. We would also like to encourage other networks to join us in taking this step to improve the quality of routing information in the Internet.

Thanks!

Jay B.

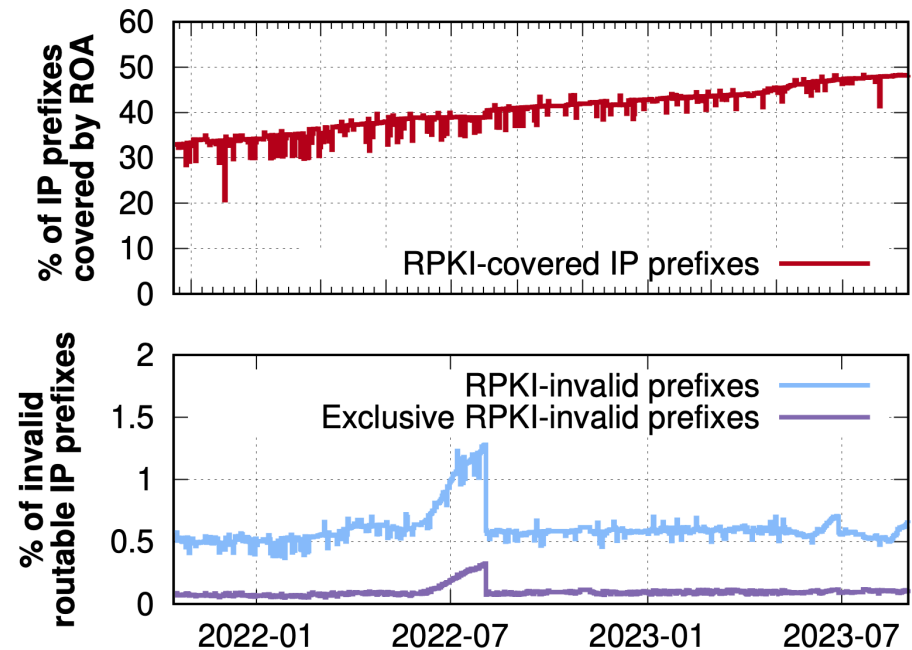
RoVista: Measuring and Understanding the ROV Status at Scale

- In-the-wild invalid prefixes
 - Due to misconfigurations or attacks, 0.5% of RPKI-covered BGP announcements are actually RPKI-invalid
- What if we can measure whether an AS can reach these RPKI-invalid prefixes?



RoVista: Measuring and Understanding the ROV Status at Scale

- In-the-wild invalid prefixes
 - Due to misconfigurations or attacks, 0.5% of RPKI-covered BGP announcements are actually RPKI-invalid
- What if we can measure whether an AS can reach these RPKI-invalid prefixes?




RoVista: Measuring and Understanding the ROV Status at Scale

- IP-ID Side-channel technique, which allows to infer the connectivity between two hosts (e.g., whether one host can receive a packet from other host)
- Preliminaries
 - TCP three-way handshake
 - IP-ID
 - IP Source Spoofing


IP-ID Side-Channel

Preliminaries (1): TCP Handshake




IP-ID Side-Channel

Preliminaries (1): TCP Handshake




IP-ID Side-Channel

Preliminaries (1): TCP Handshake



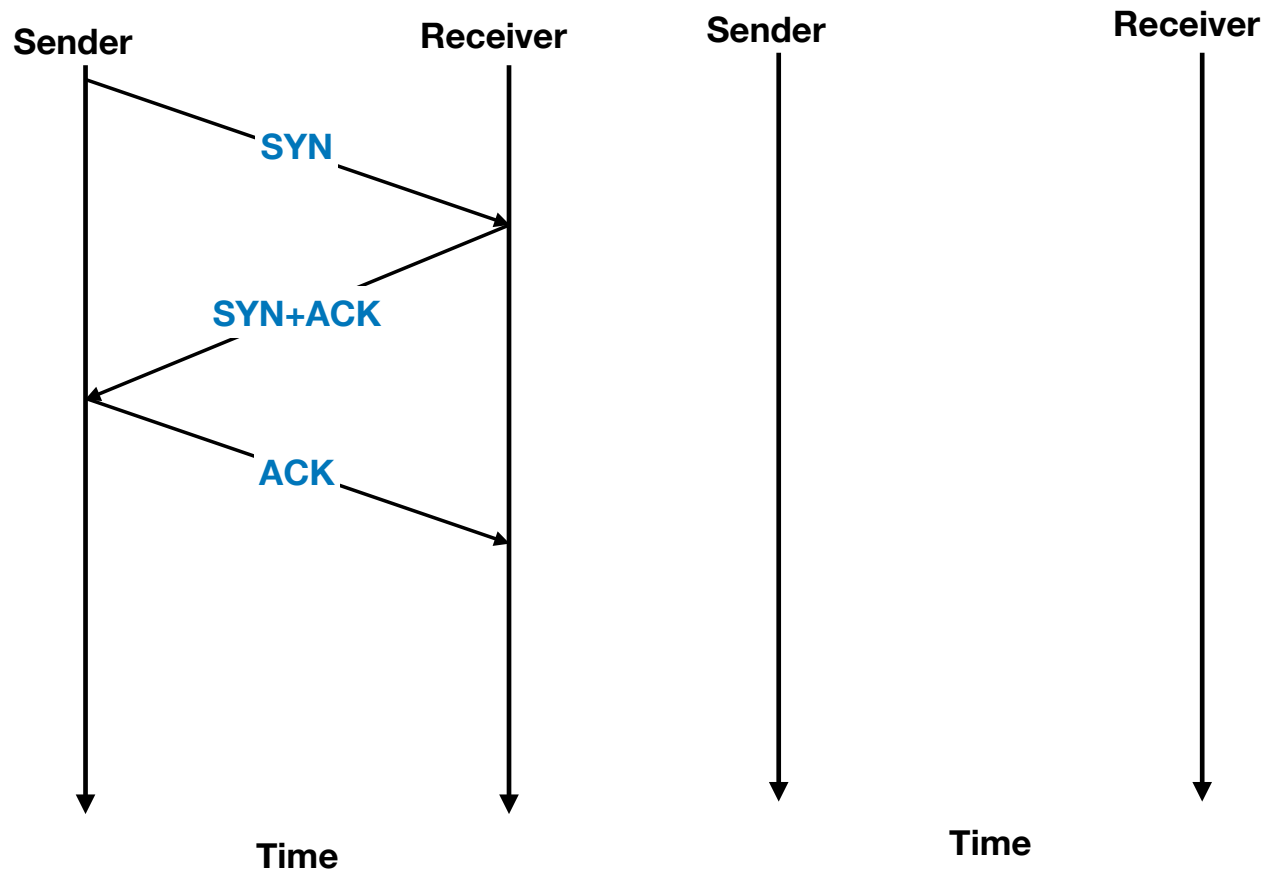
IP-ID Side-Channel

Preliminaries (1): TCP Handshake




IP-ID Side-Channel

Preliminaries (1): TCP Handshake



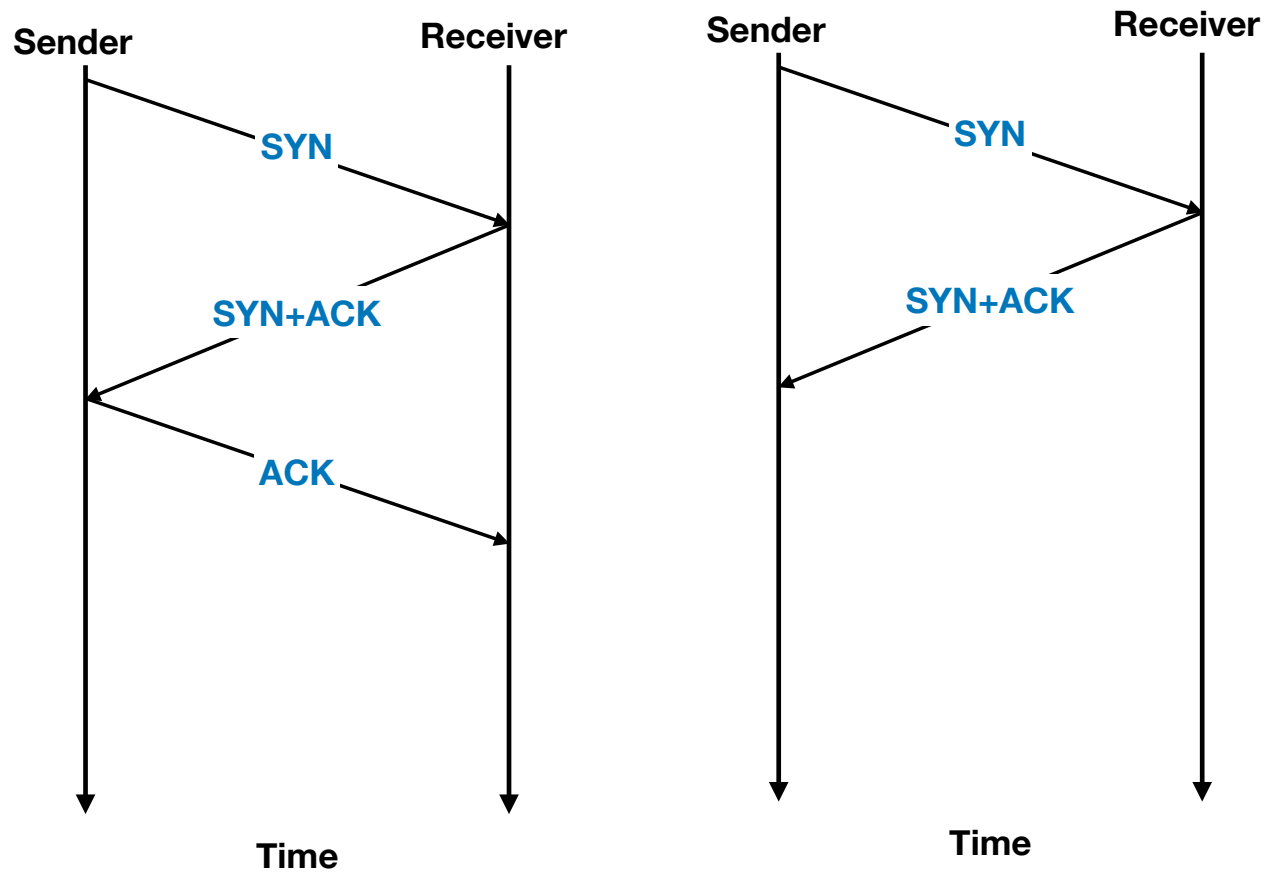
IP-ID Side-Channel

Preliminaries (1): TCP Handshake



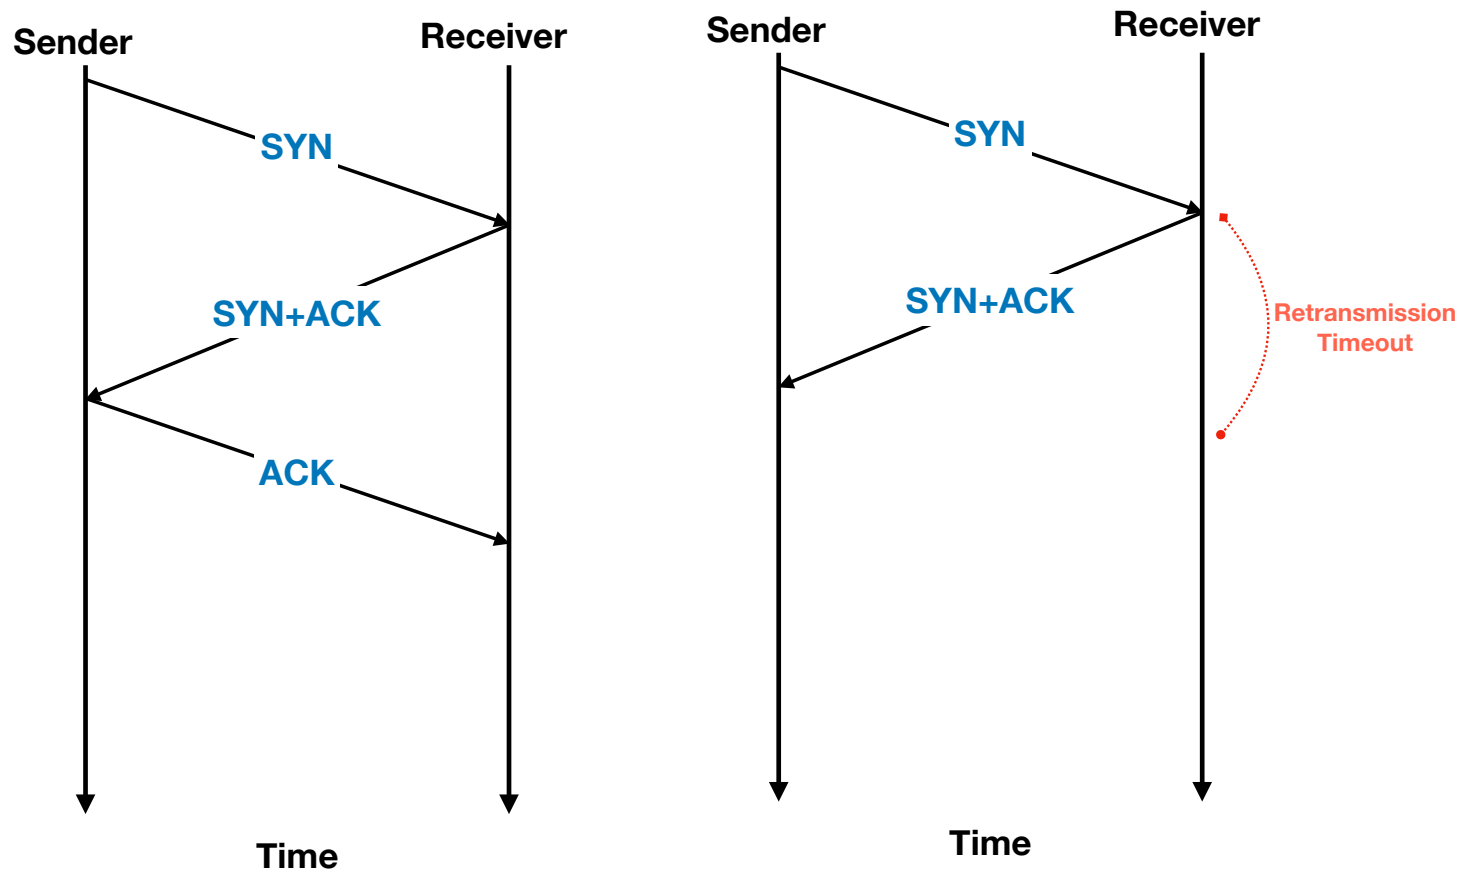
IP-ID Side-Channel

Preliminaries (1): TCP Handshake




IP-ID Side-Channel

Preliminaries (1): TCP Handshake




IP-ID Side-Channel

Preliminaries (1): TCP Handshake



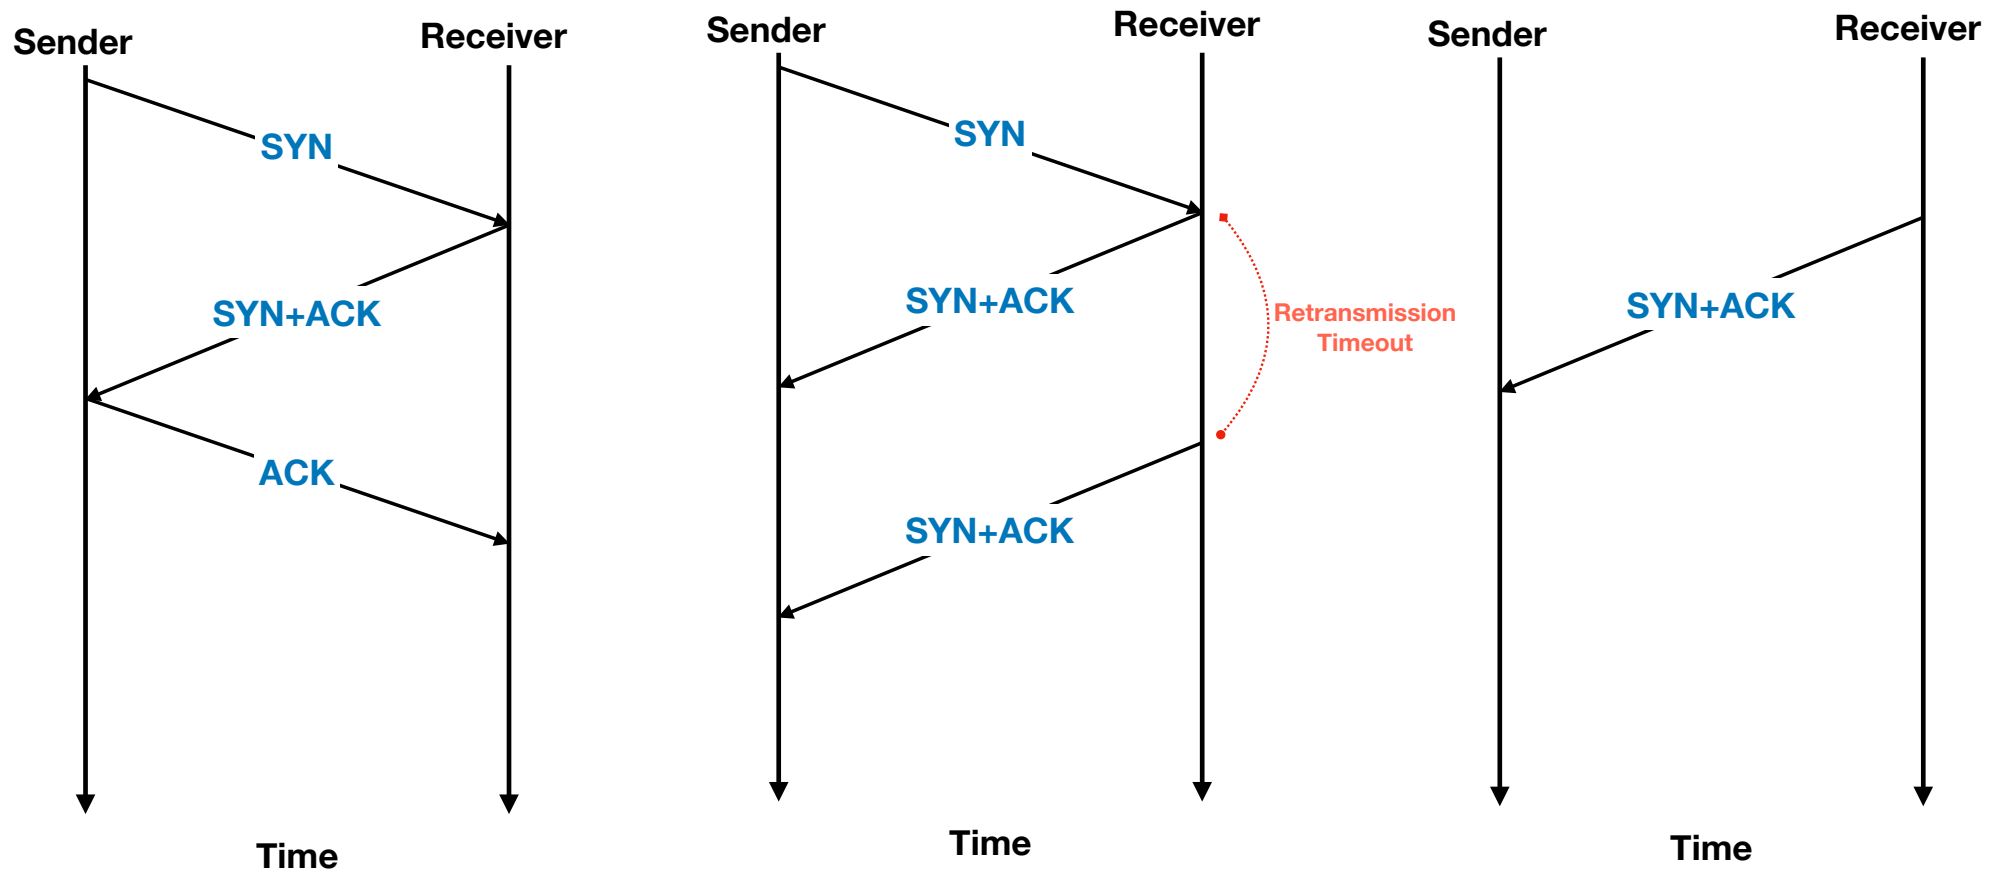
IP-ID Side-Channel

Preliminaries (1): TCP Handshake




IP-ID Side-Channel

Preliminaries (1): TCP Handshake



IP-ID Side-Channel

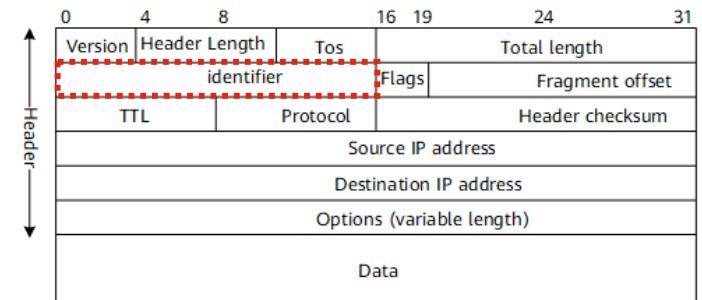
Preliminaries (1): TCP Handshake



IP-ID Side-Channel Preliminaries (2): IP-ID

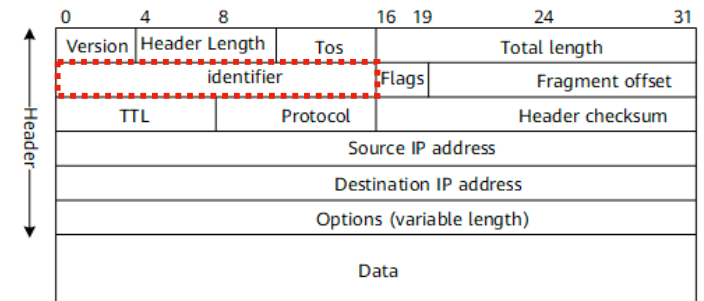
IP-ID Side-Channel Preliminaries (2): IP-ID

- IP ID was first introduced by RFC 791
- originally designed to assist packet fragmentation and reassembly by assigning an unique identifier for each packet



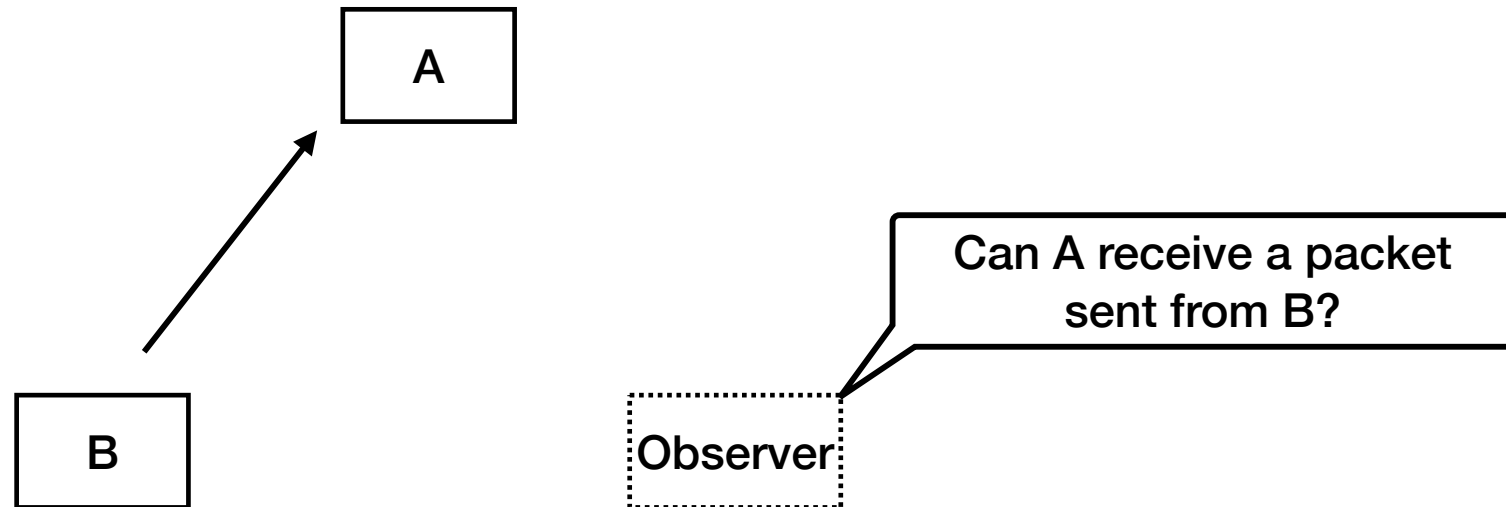
IP-ID Side-Channel Preliminaries (2): IP-ID

- IP ID was first introduced by RFC 791
 - originally designed to assist packet fragmentation and reassembly by assigning an unique identifier for each packet
- How to assign IPID?
 - **Global counter**
 - increments the IP-ID by 1 unit whenever it sends a new packet regardless of the destination IP address
 - ...




IP-ID Side-Channel

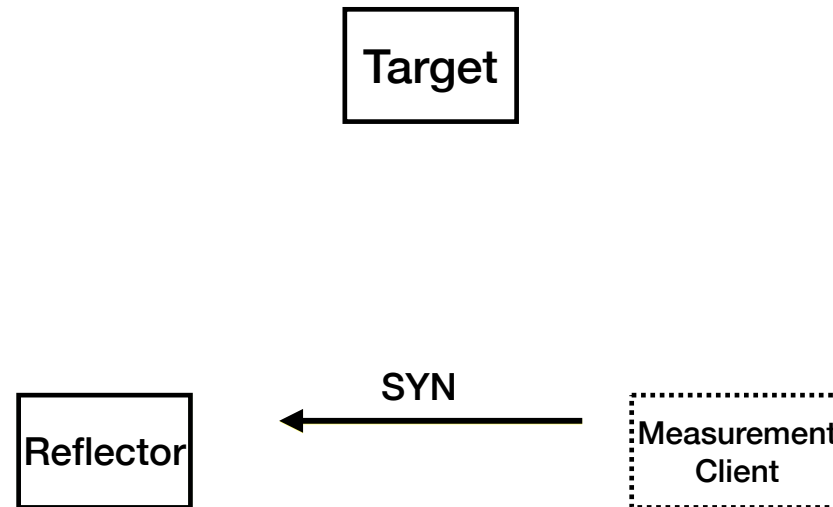
“Can we measure the connectivity of two **remote** end hosts?”



IP-ID Side-Channel Basic Idea

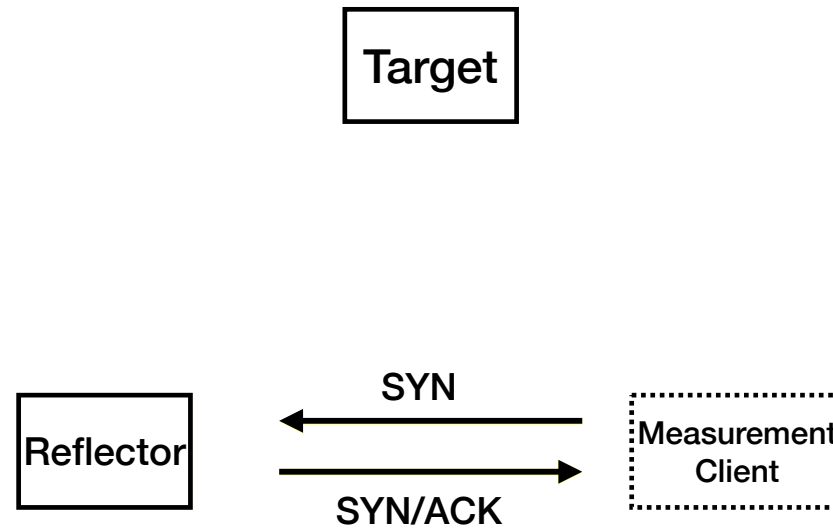


IP-ID Side-Channel Basic Idea



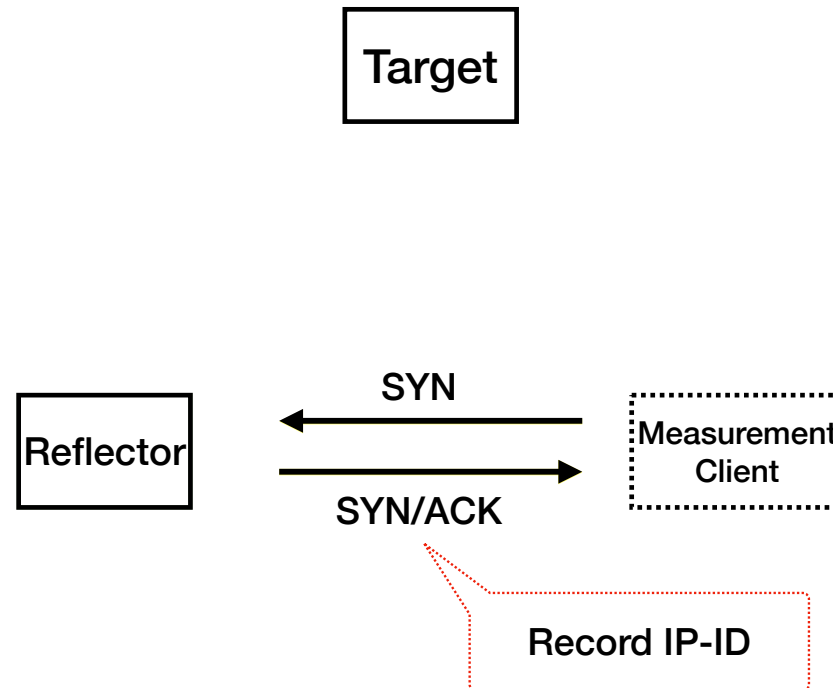
Time	IP-ID

IP-ID Side-Channel Basic Idea




Time	IP-ID

IP-ID Side-Channel Basic Idea




Time	IP-ID

IP-ID Side-Channel Basic Idea



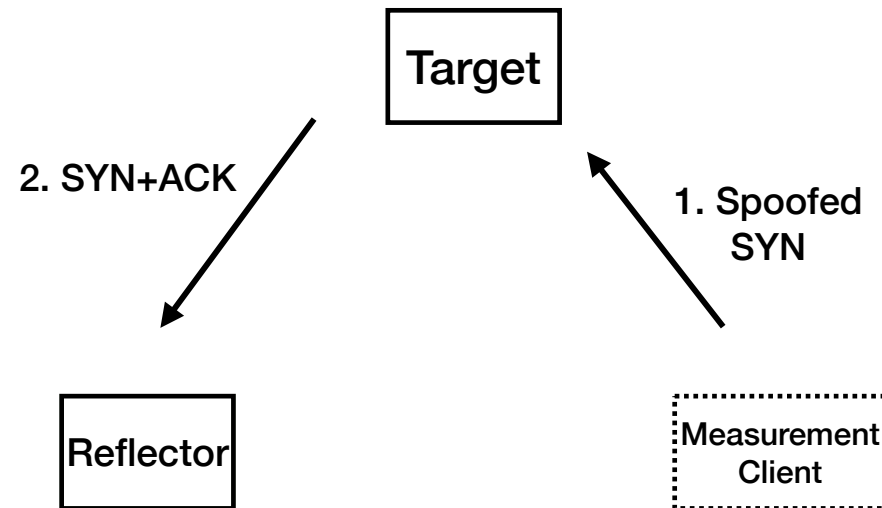
Time	IP-ID
t_0	1

IP-ID Side-Channel Basic Idea



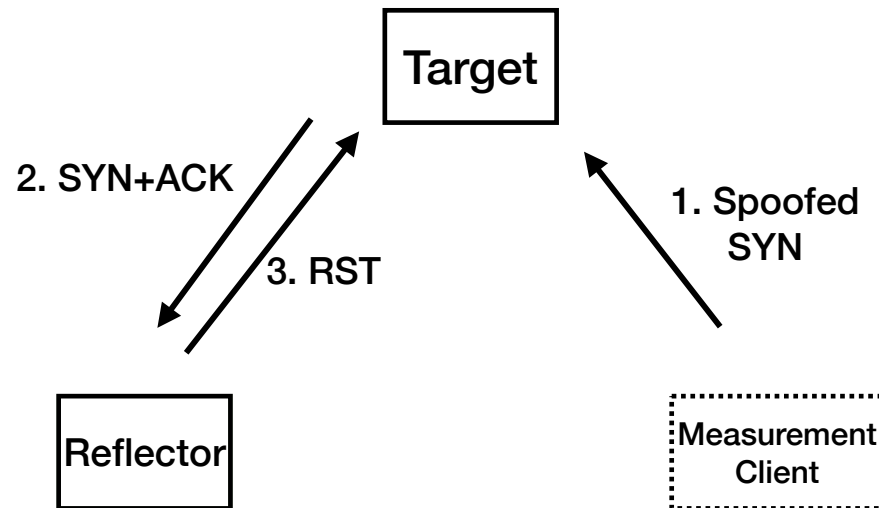
Time	IP-ID
t_0	1

IP-ID Side-Channel Basic Idea




Time	IP-ID
t_0	1

IP-ID Side-Channel Basic Idea




Time	IP-ID
t_0	1

IP-ID Side-Channel Basic Idea



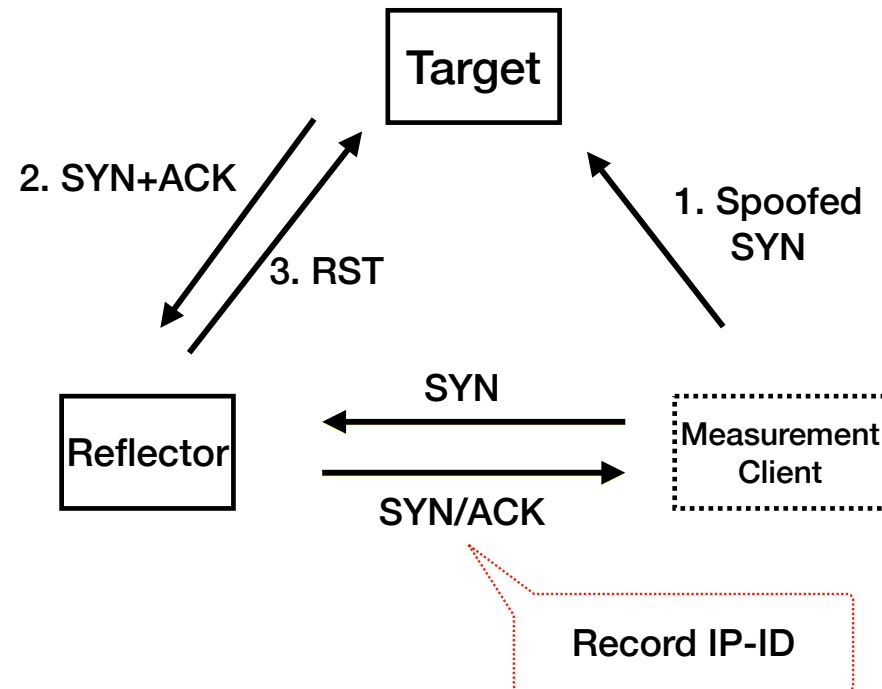
Time	IP-ID
t_0	1

IP-ID Side-Channel Basic Idea




Time	IP-ID
t_0	1

IP-ID Side-Channel Basic Idea



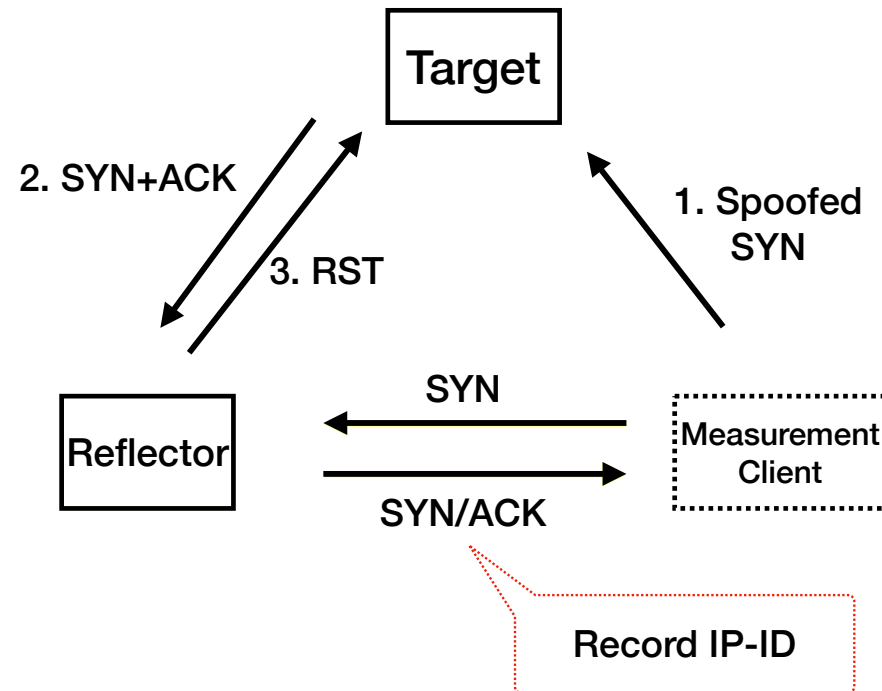
Time	IP-ID
t_0	1
t_1	3

IP-ID Side-Channel Basic Idea




Time	IP-ID
t_0	1
t_1	3

IP-ID Side-Channel Basic Idea



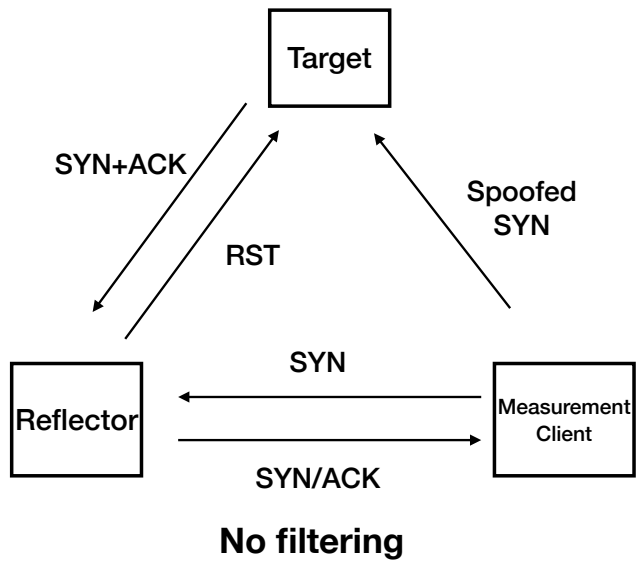
Time	IP-ID
t_0	1
t_1	3

IP-ID Side-Channel Basic Idea

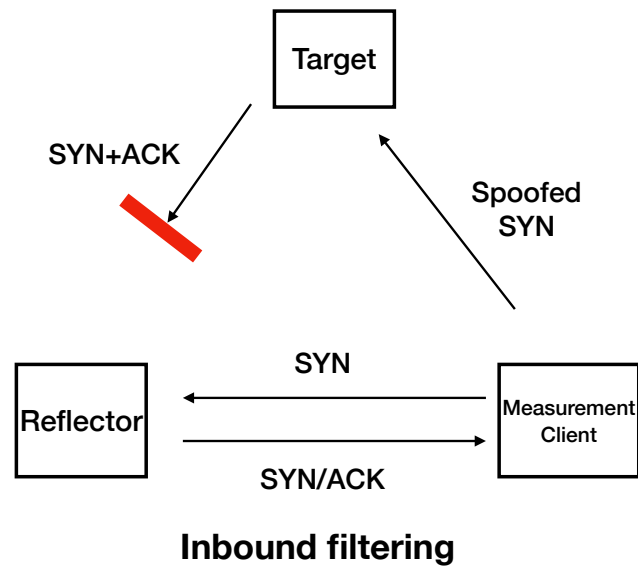
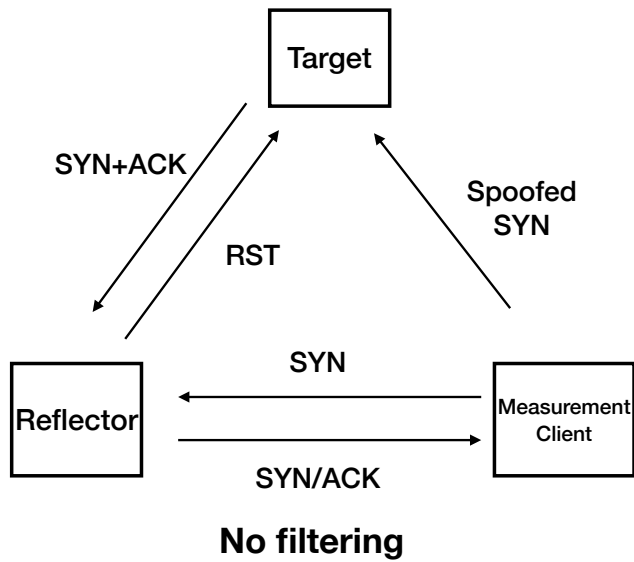


Time	IP-ID
t_0	1
t_1	3
t_2	4




IP-ID Side-Channel Three Cases






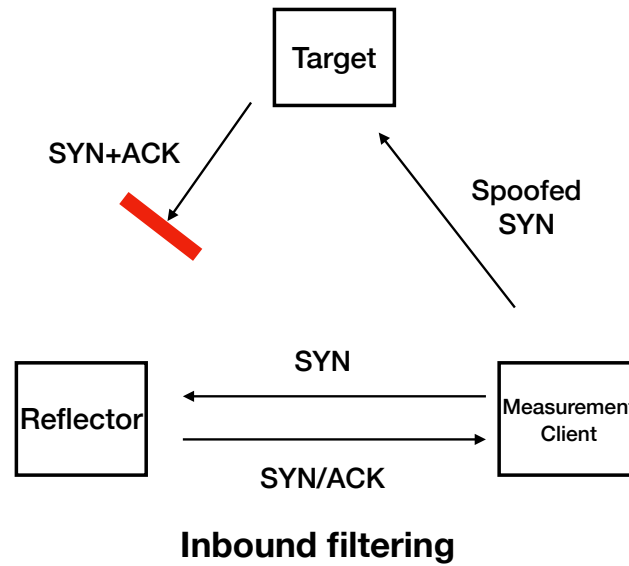


IP-ID Side-Channel Three Cases








IP-ID Side-Channel Three Cases









IP-ID Side-Channel Three Cases









IP-ID Side-Channel Three Cases








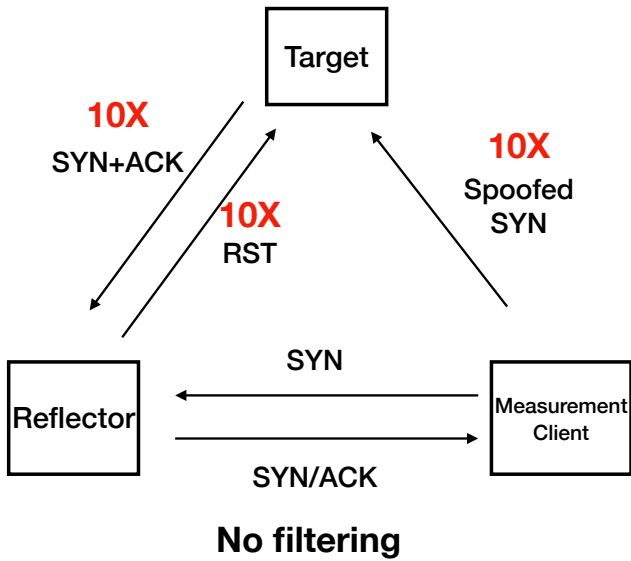
IP-ID Side-Channel Three Cases









IP-ID Side-Channel Three Cases




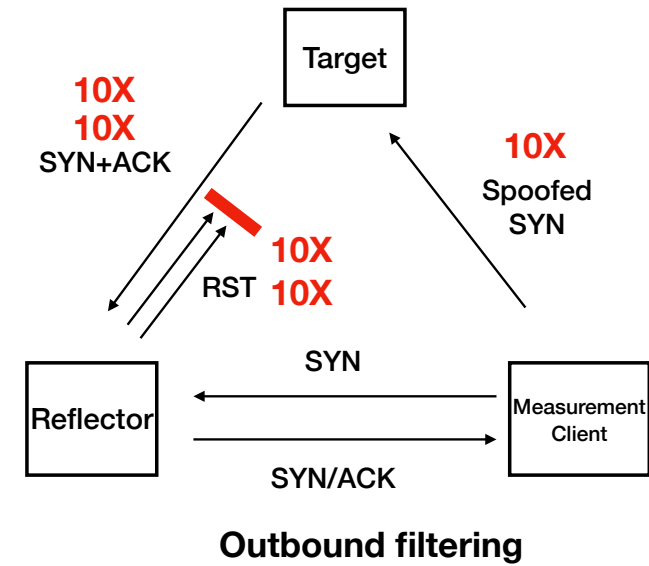




IP-ID Side-Channel Three Cases





IP-ID Side-Channel Three Cases



IP-ID Side-Channel Three Cases




ROV Detection




ROV Detection

Announcing RPKI-invalid prefix




Outbound filtering




- Let's apply IP-ID side-channel to "detecting ROV policy"
- When we find a host of which IP address is announced through **RPKI-invalid prefix**, we define them as "targets"

ROV Detection

Announcing RPKI-invalid prefix



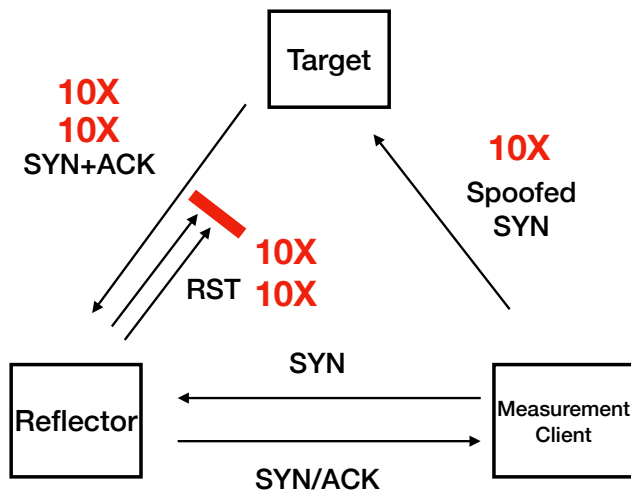
Outbound filtering




- Let's apply IP-ID side-channel to "detecting ROV policy"
- When we find a host of which IP address is announced through **RPKI-invalid prefix**, we define them as "targets"
- If a reflector can't send packets to the target, it may indicate that such RPKI-invalid prefixes are being filtered

ROV Detection

Announcing RPKI-invalid prefix



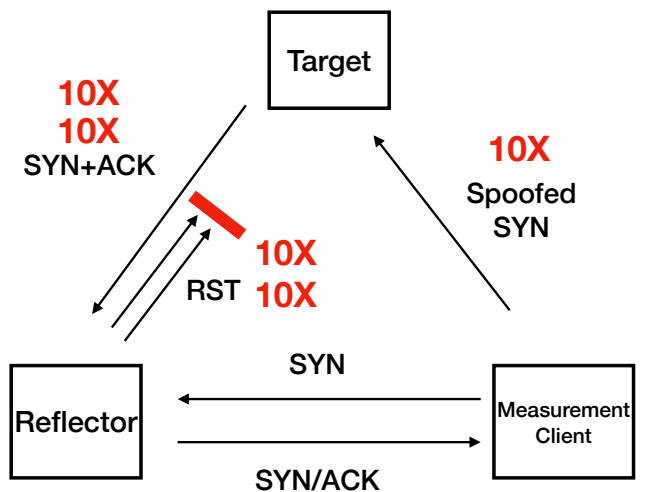
Outbound filtering




- Let's apply IP-ID side-channel to "detecting ROV policy"
- When we find a host of which IP address is announced through **RPKI-invalid prefix**, we define them as "targets"
- If a reflector can't send packets to the target, it may indicate that such RPKI-invalid prefixes are being filtered
- If we find many reflectors in the same AS that can't send packets to the target, it is highly likely due to their ROV policy

ROV Detection

Announcing RPKI-invalid prefix



Outbound filtering



- Let's apply IP-ID side-channel to “detecting ROV policy”
 - When we find a host of which IP address is announced through **RPKI-invalid prefix**, we define them as “targets”
- If a reflector can't send packets to the target, it may indicate that such RPKI-invalid prefixes are being filtered
- If we find many reflectors in the same AS that can't send packets to the target, it is highly likely due to their ROV policy
- ROV Score: the percentage of filtered RPKI-invalid prefixes on an AS

RoVista Measurement Results

Measurement Period	12/24/2021 ~ now
# of ASes	31K
# of countries	231

<https://rovista.netsecurelab.org/>

Cross-validation Comparison with the official sources

The list of ASes doing ROV

ISP	ASN	Source	ROV Ratio from RoVista
HEANet	1213	https://twitter.com/natural20/status/1366385420360155144	100%
Telstra	1221	https://lists.ausnog.net/pipermail/ausnog/2020-July/044367.html	100%
Sprint / T-Mobile	1239	https://www.sprint.net/policies/bgp-aggregation-and-filtering	100%
Telia	1299	https://www.teliacarrier.com/Our-Network/BGP-Routing/Routing-Security.html	100%
EBOX	1403	https://whois.arin.net/rest/asn/AS1403/pf?&s=AS1403	100%
IJ	2497	https://www.ij.ad.jp/en/dev/iir/pdf/iir_vol50_focus1_EN.pdf	100%
Belnet	2611	https://belnet.be/en/belnet-has-successfully-implemented-rpki	100%
NTT	2914	https://www.gin.ntt.net/support/policy/rr.cfm#RPKI	100%
TDC	3292	https://github.com/cloudflare/isbgpsafeyet.com/pull/523	100%
Swisscom	3303	https://twitter.com/swisscom_csirt/status/1300666695959244800	100%
Level3	3356	https://twitter.com/lumentechco/status/1374035675742412800	100%
Telstra	4637	https://www.zdnet.com/article/telstra-to-roll-out-rpki-routing-security-from-june-2020/	100%
Vocus	4826	https://blog.apnic.net/2021/05/13/vocus-rpki-implementation/	100%
Orange	5511	https://twitter.com/OrangeC/status/1541436188241891328	100%
Cyta	6866	https://blog.daknob.net/rpki-deployment-greece-feb-19/	100%
Hurricane Electric	6939	https://mailman.nanog.org/pipermail/nanog/2020-June/108277.html	100%
AT&T	7018	https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html	100%
Dhiraagu	7642	https://twitter.com/isseykun/status/1261758917467668481	0%
Comcast	7922	https://corporate.comcast.com/stories/improved-bgp-routing-security-adds-another-layer-of-protection-to-network	100%
ColoClue	8283	https://github.com/coloclue/kees	100%
Atom86	8455	https://www.linkedin.com/pulse/atom86-leveraging-rpki-make-internet-safer-place-ralph-dirkse/	100%
RETN	9002	https://twitter.com/RETNnet/status/1333735456408793089	92.5%
BIT	12859	https://www.bit.nl/news/2081/88/Registratie-van-RPKI-informatie-voor-een-veilige-routering-informatie-voor-een-veilige-routering	0%
Amazon	16509	https://aws.amazon.com/blogs/networking-and-content-delivery/how-aws-is-helping-to-secure-internet-routing/	100%
ASERGO	30736	https://twitter.com/asergogroup/status/1258377169526546432	100%
Jaguar	30781	https://twitter.com/JDescoux/status/1253344721201696768	100%
Seacom	37100	https://www.ripe.net/participate/mail/forum/routing-wg/PDZlMzAzMzhLWVhOTAtNzlxOC1lMzI0LTBjZjMyOGI1Y2NkM0BzZWZjb20ubXU+	100%
NAPAfrica	37195	https://www.napafrika.net/technical/rpki-handy-hints/	100%
Workonline	37271	https://as37271.fyi/routing-policy/	100%
Freethought	41000	https://twitter.com/freethoughtnet/status/1222841548771090432	100%
Fiber Telecom	41327	https://www.peeringdb.com/asn/41327	100%
HOPUS	44530	https://twitter.com/afenioux/status/1305430383345971201	100%
NAP.EC	52482	https://www.aeprovi.org.ec/es/implementacion-de-rpki-y-validacion-de-origen-bgp-en-ecuador	100%
Scaleway	54265	https://mailman.nanog.org/pipermail/nanog/2020-April/107295.html	100%
Terrahost	56655	https://twitter.com/TerraHost/status/1259311449073168384	100%
KAPSI	57692	https://twitter.com/atonkyra/status/1253609926221496322	100%
Fusix	57866	https://fusix.nl/deploying-rpki/	100%
Gigabit ApS	60876	https://mailman.nanog.org/pipermail/nanog/2020-April/107295.html	0%
Tuxis	197731	https://twitter.com/Tuxis_IE/status/1105060034873049091	100%


The list of ASes not-doing ROV

ISP	ASN	Source	ROV Ratio from RoVista
Deutsche Telekom	3320	https://twitter.com/deutsche Telekom/status/1252177058555473920	0%
Worldstream	49981	https://twitter.com/worldstream/status/1257670396461166593	0%

They had enabled ROV in early 2018, but they retracted ROV because of the Juniper router issue in 2018

Status Quo (1)


% of “fully protected” ASes



- The percentage of ASes with 100% ROV scores is increasing over time: 19%
- The ASes with 100% ROV scores don't necessarily indicate ROV “deployment”.

Status Quo (2)


% of “Partially Protected” ASes



- The percentage of ASes with higher than 0% ROV scores are also increasing: 79%


Case-study: How Quickly RoVista Detect ROV impact?

- During our measurement period, we find two ASes (Orange and KPN) officially announced their ROV “deployment”




1. Orange announced on June 27th, 2022 and RoVista detects the spike on June 6th, 2022
2. KPN announced on March 16th, 2022 and RoVista detects the spike on March 14th, 2022.

Case-Study: KPN: Collateral Benefits of ROV



Case-Study: KPN: Collateral Benefits of ROV



In case of Orange, the scores of all of their 20 customers that we measure jump to 100% simultaneously

Limitation

- RoVista **cannot** measure the ROV protection score of IXPs since it is infeasible for us to find measurement nodes in IXPs.
- RoVista relies on hosts announcing RPKI-invalid prefixes using the public BGP collectors, thus may have a limited coverage.
- **ROV protection score does not directly indicate the ROV “Deployment” status of the AS** — thus 100% ROV score does not necessarily mean that the AS has deployed ROV (it may be due to their providers)

Summary for RoVista

- RoVista is a data-plane based methodology to measure the ROV status of network operators by using (1) in-the-wild RPKI-invalid prefixes and (2) IP-ID Side-channel technique.
- We are releasing our results at <https://rovista.netsecurelab.org/> with APIs: please find your AS and contact us if discrepancies are found.
- The paper was published at Internet Measurement Conference (IMC'23)
- If you're a network operator, not participated in the survey yet, please help us: <https://www.surveymonkey.com/r/MANRSROVAdoptionSurvey>



Thank you

- This research has been generously supported by NSF, Comcast Innovation Fund, and Google.

