



# Stalloris: RPKI Downgrade Attack

Tomas Hlavacek, Philipp Jeitner, Donika Mirdita, Haya Shulman and Michael Waidner

# Team

- Tomas Hlavacek
- Donika Mirdita
- Haya Shulman
- Michael Waidner


- **Philipp Jeitner**

- Network Security Researcher
- Just finished my PhD



Cybersecurity Analytics and Defences  
departement

- Network Security
- Routing and DNS Security



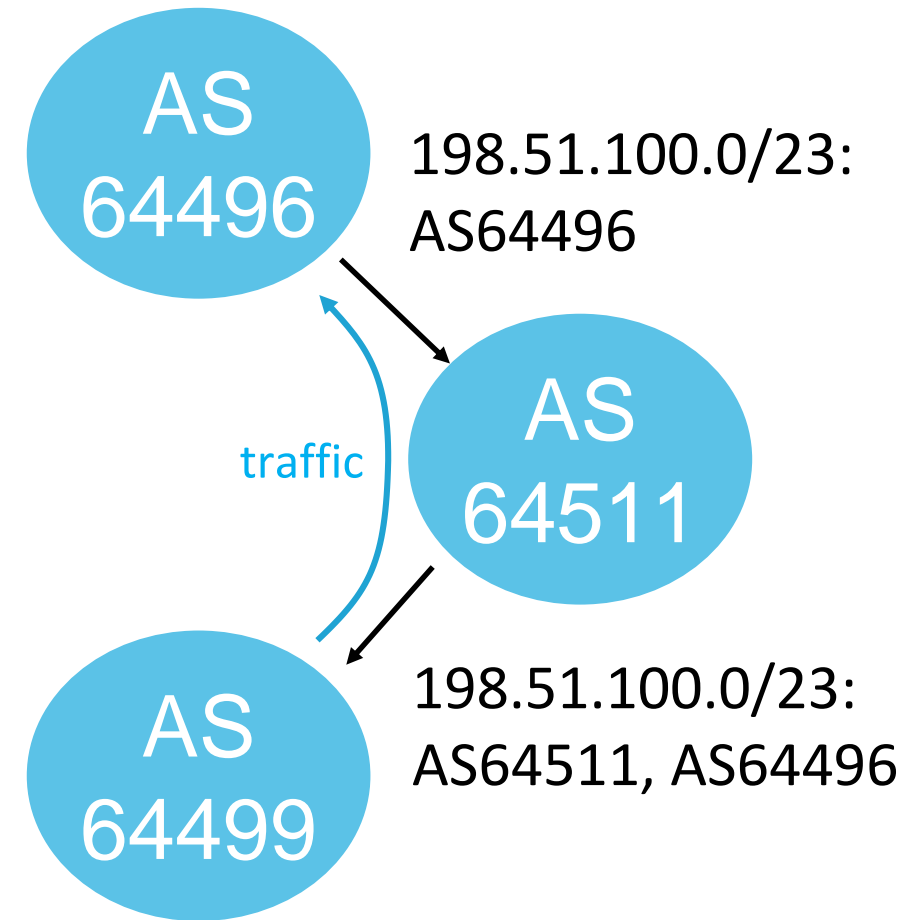
**ATHENE**  
National Research Center  
for Applied Cybersecurity



**Fraunhofer**  
SIT

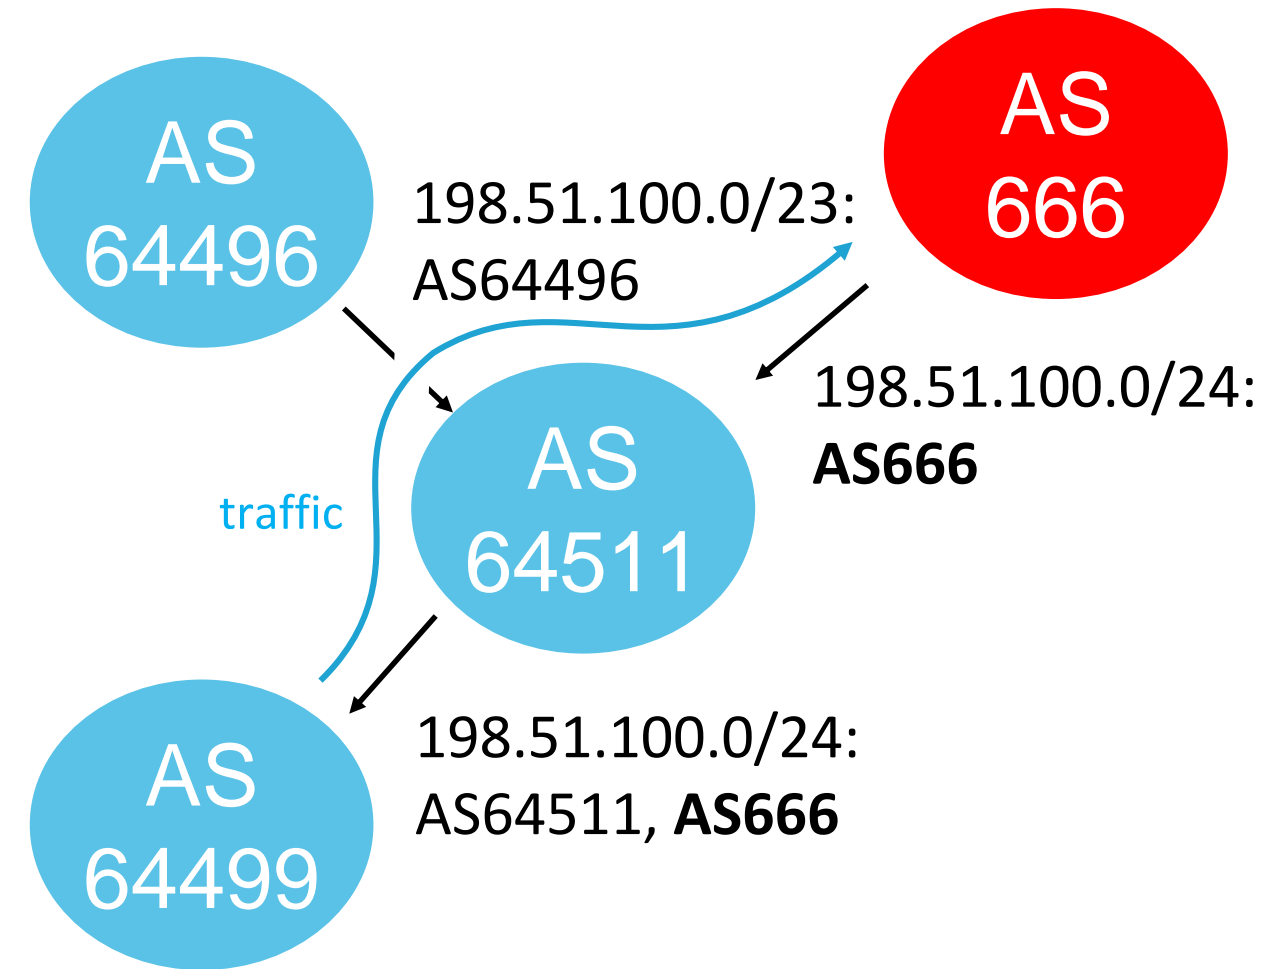
- BGP & BGP Security
- Ressource Public Key Infrastructure (RPKI)
- Downgrade attack against RPKI
- Feasability
- Mitigations

- Routing system of the Internet
- Networks (ASes) announce the IP prefixes they have
- Neighbors forward these announcements
- Everyone knows where to send traffic



# BGP Hijacks

- No built-in security
- Just announce a prefix you do not own, be MitM, profit?



- Systematic approach to BGP Security

- Certificates: Address block -> ASN

- Called Route Origin Authorization (ROA)

- Root of Trust: RIRs

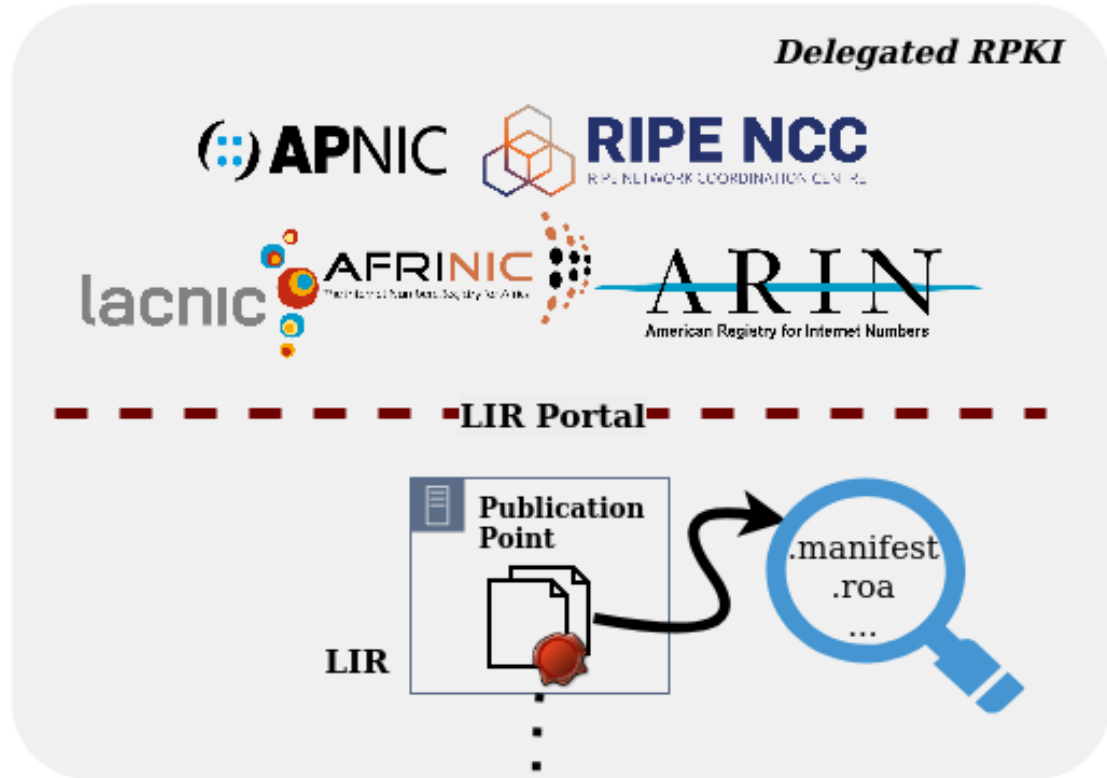
- Because RIRs allocate address blocks!


- Kind of like getting your TLS cert from the registry

```
{  
  "asn": "AS64496",  
  "prefix": "198.51.100.0/24",  
  "maxLength": 24,  
  "ta": "RIPE",  
}
```




# Route Origin Validation



- Now everyone has ROAs
  - (actually only 34.2%)
- How to check them?
  - Put them into BGP? 
  - Third-party system!

ROA:  
AS64496 owns  
198.51.100.0/24



# RPKI works!



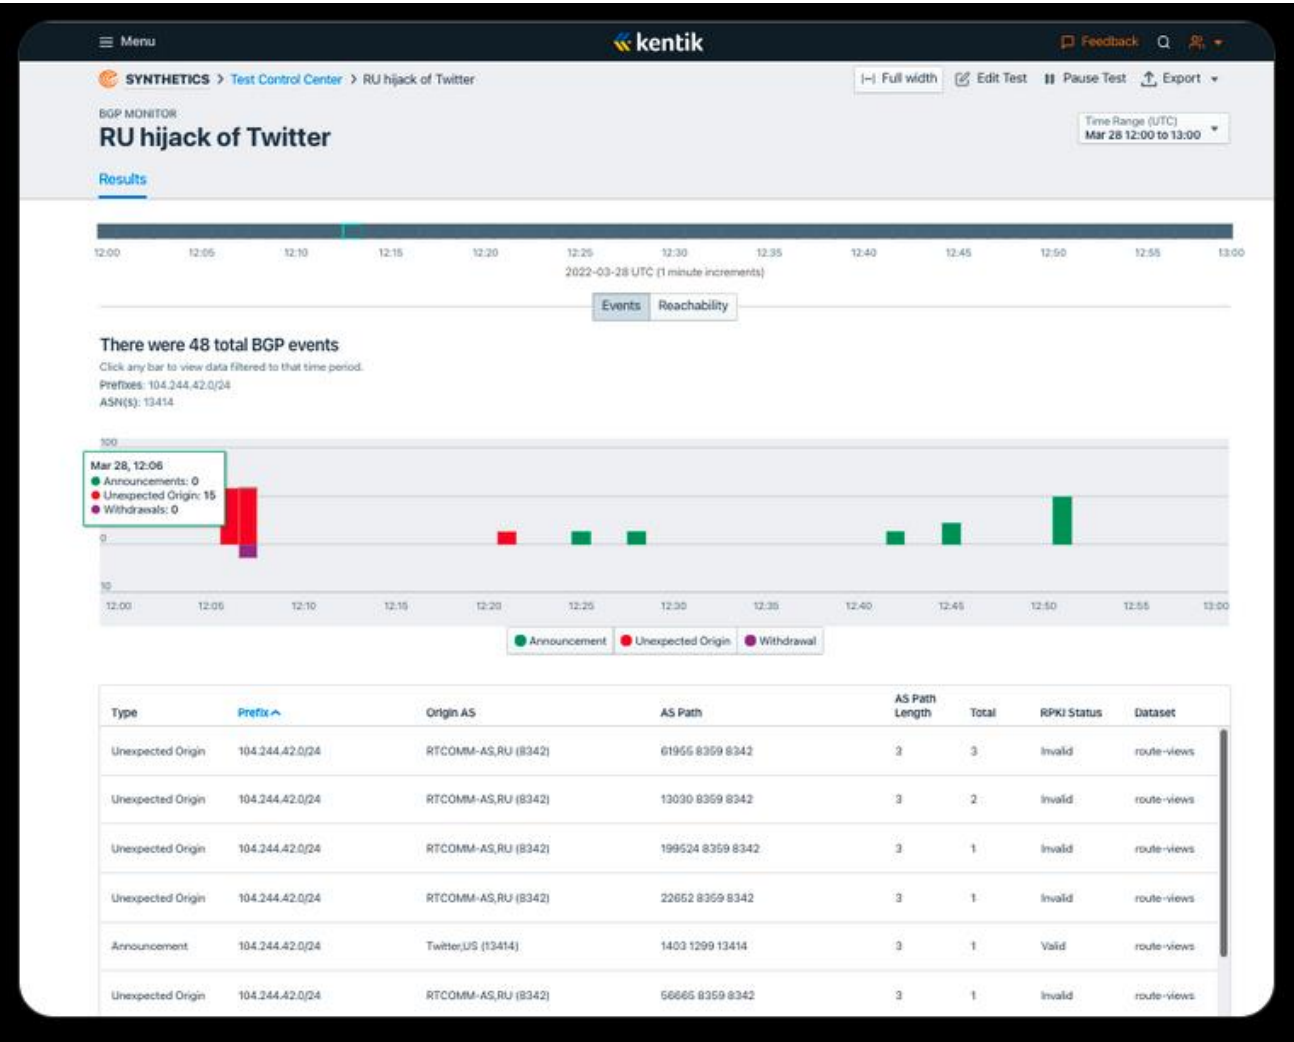
**Doug Madory**  
@DougMadory

From 12:05-12:50 UTC, RU telecom RTComm (AS8342) hijacked a prefix (104.244.42.0/24) belonging to Twitter.

The hijack didn't propagate far due to a RPKI ROA which asserted AS13414 was the rightful origin.

This is the same prefix hijacked during the coup in Myanmar last year.

5:29 nachm. · 28. März 2022 · Twitter Web App






# Attacking RPKI

- Integrity?
  - Create Malicious ROA? Breaking crypto is hard.
  - Fool the CAs? CAs are run by RIRs.
- Availability!
  - RPKI is a third party system to BGP
  - What if RPKI stops working?




# Making RPKI stop working

- Relying Parties (RPs) need to download ROAs from Publication Points (PPs)
  - If download fails, RPs will not have ROAs and assume RPKI has not been deployed
- Plan:
  - Break communication with PP
  - RPs cannot fetch information
  - RPKI turned off (RPKI state unknown)
  - Start BGP hijack



# RP cache and manifests

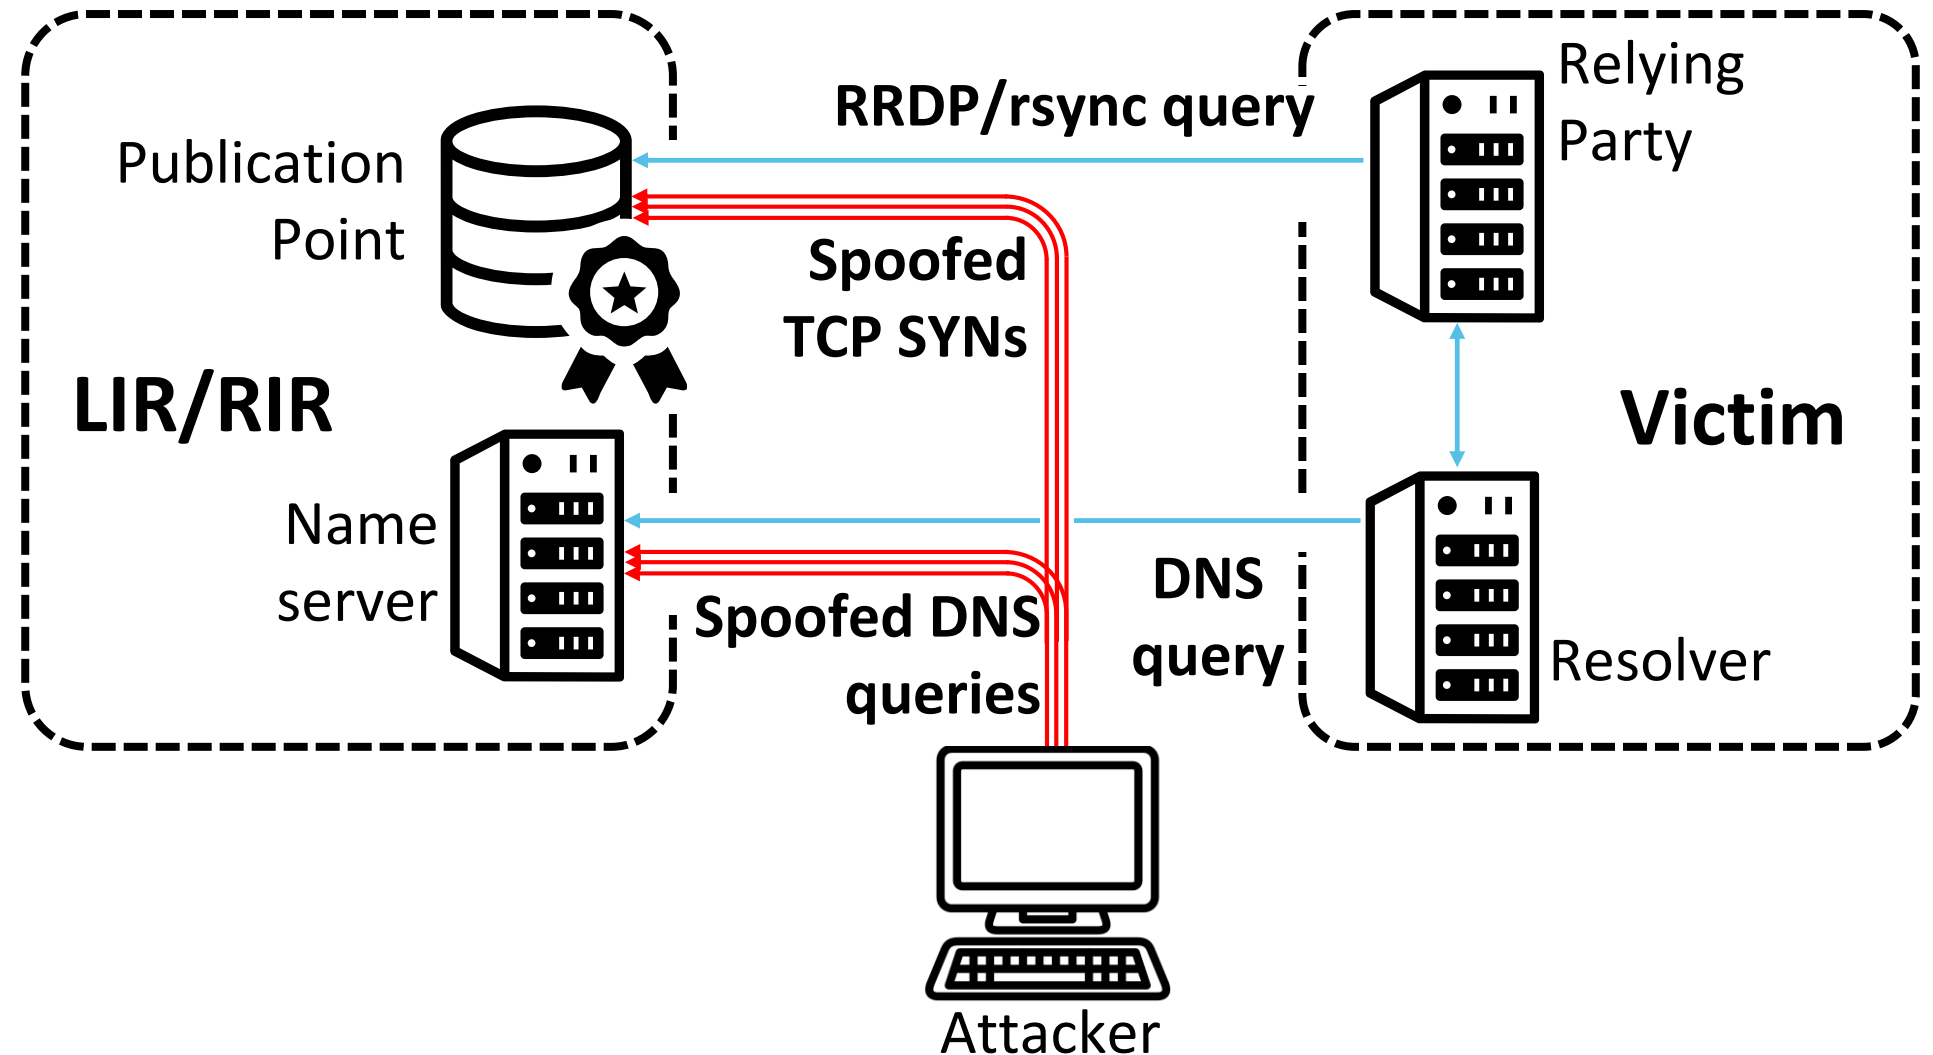
- RPs cache old data until expiry
  - ROAs expire pretty slowly (1 year)
- Manifests
  - Essentially a signed index
  - Designed to prevent replay attacks
  - ROAs not listed in manifest get removed
  - **Short expiry time! (1 day)**
    - Effectively only 6 hours of attack time because of deterministic re-generation



# Breaking communication

## Low-rate attack:

- Exploit rate-limiting on PP/NS
- Send spoofed requests
- Victim gets blocked
- After 6 hours: ROAs removed from cache due to expired manifest



# Rate-limiting in RPKI

## Tested rate-limiting in PPs

- DNS RRL & TCP Syn rate-limiting
- Typically implemented to prevent DoS

## Results

- 47% of PPs do it (limit < 10,000 pkt/s)
- Affects 60% of RPKI-protected IPv4 space
- 3% of IPv4 are protected by PPs with very low (<60 pkt/s) rate-limit

		total addresses	% of assigned address space	% of ROA-protected
has ROA	v4	64 * /8	34.2 %	100.0 %
	v6	322 * /24	40.0 %	100.0 %
vulnerable (all)	v4	39 * /8	20.4 %	59.6 %
	v6	122 * /24	15.2 %	37.9 %
vulnerable (low ratelimit)	v4	2 * /8	1.1 %	3.1 %
	v6	10 * /24	1.3 %	3.2 %

# Feasibility (1)

So 60% of Ipv4 can be attacked?

Example:

- Rate-limit is 1,000 pkt/s,  
attacker sends 10,000 pkt/s
- Connection success is ~ 10%
- But RPs will **retry**



# Feasibility (2)

(Scenario)	$n_{attempts}$	$t_{attack}$	$t_{sleep}$	$n_{retries}$
(1)	24	old manifest 6 hours	unbound (blocked) 900 s	unbound (blocked) 1
(2)	864	fresh manifest 1 day	routinator (normal) 600 s	bind9 / linux tcp 6
(3)	23040	long-valid manifest 2 days	RIPE NCC validator 120 s	unbound normal 16

## Simulation using different scenarios

- Feasible for low rate-limits (< 60 pkt/s), higher ones are challenging due to retries in 6 hours

# We have to try harder

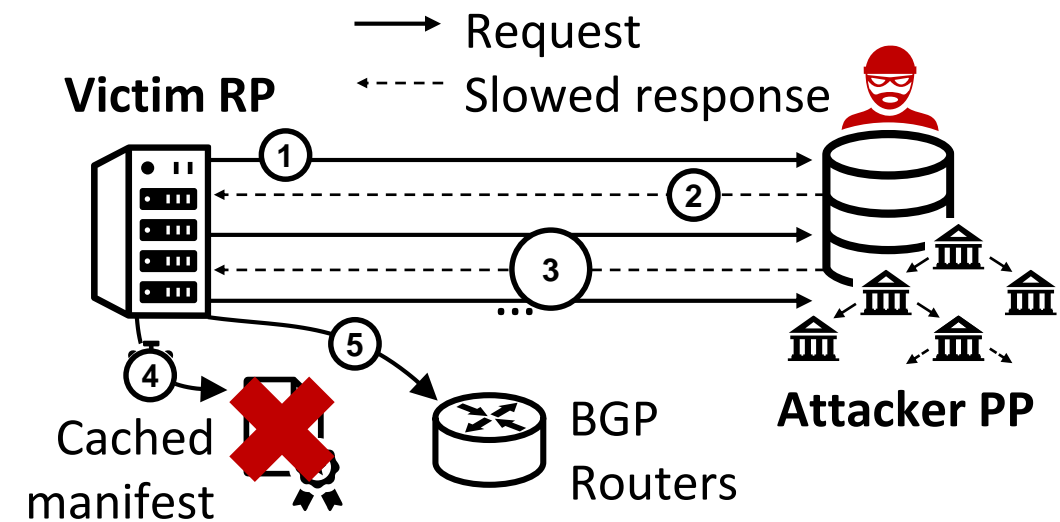
- RPKI allows *delegation*
  - LIRs can run their own Publication Point
  - **Attackers** can run their own Publication Point
  - and RPs have to contact **them**
  - *Can we exploit this to break the RP?* 🤔

## Stalloris

- Attacker becomes malicious Publication Point
- Sends responses as slow as possible
- Hinders RP from doing many retries

Simulation shows this makes attack feasible for high rate-limits and less-favorable scenarios

- Becoming a PP also helps time the attack with spoofed queries



# Wrapping up

- Third-party system allows attacks on availability
- Rate-limiting can be exploited to block legitimate requests from off-path
- Short manifest validity makes attacks feasible
- Attackers can become PPs and prevent RPs from doing their work

# Recommendations

- Publication points
  - Avoid low rate limits: Limiting to 60 pkt/s/IP is very easy to spoof
  - Longer manifest validities, e.g., 1 week
  - Randomize when manifests are re-generated
  - More robust deployment/redundancy
- Relying parties
  - Monitor connection failures
  - Limit processing time/PP and limit tree size under one PP

# Thank You!



## Stalloris: RPKI Downgrade Attack

Tomas Hlavacek, Philipp Jeitner, Donika Mirdita,  
Haya Shulman and Michael Waidner

## Contact: Philipp Jeitner

[philipp.jeitner@sit.fraunhofer.de](mailto:philipp.jeitner@sit.fraunhofer.de)