TALLINN UNIVERSITY OF TECHNOLOGY
Queen Mary University of London
Hochschule für Angewandte Wissenschaften Hamburg
Hamburg University of Applied Sciences
Freie Universität Berlin

RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem.

Or why Xmas online shopping might go wrong.

Matthias Wählisch, Robert Schmidt, Thomas C. Schmidt, Olaf Maennel, Steve Uhlig, Gareth Tyson

Starting Point

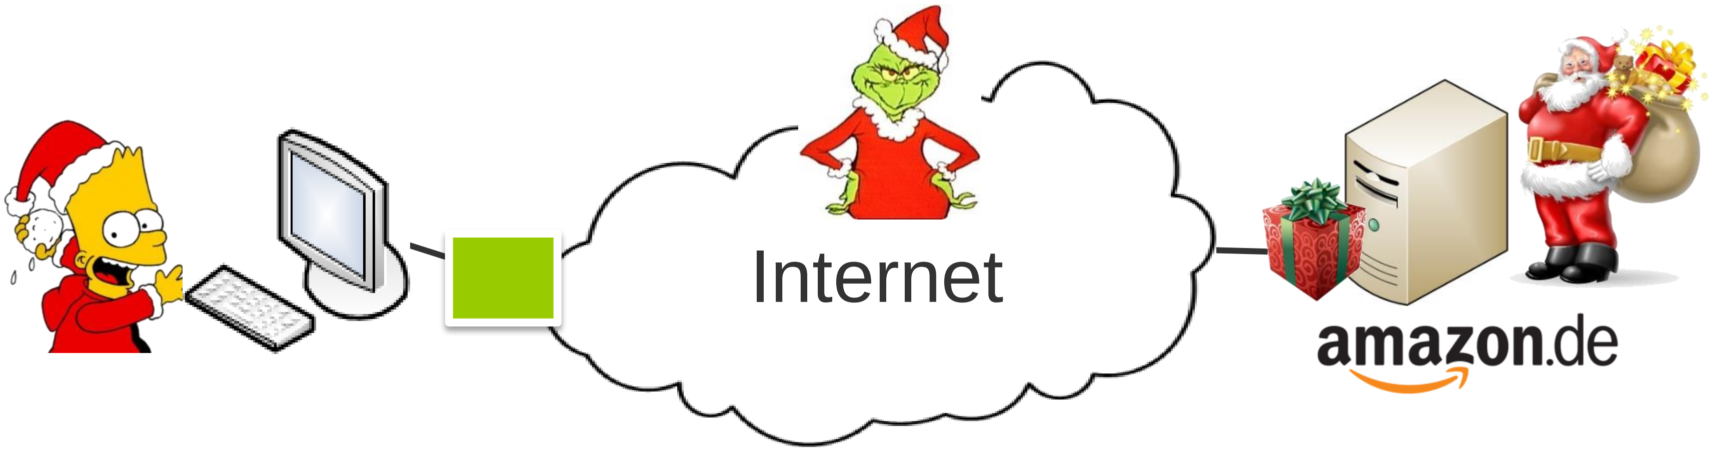

Recap: Internet in a Nutshell & Attacks

Autonomous Systems (ASes)

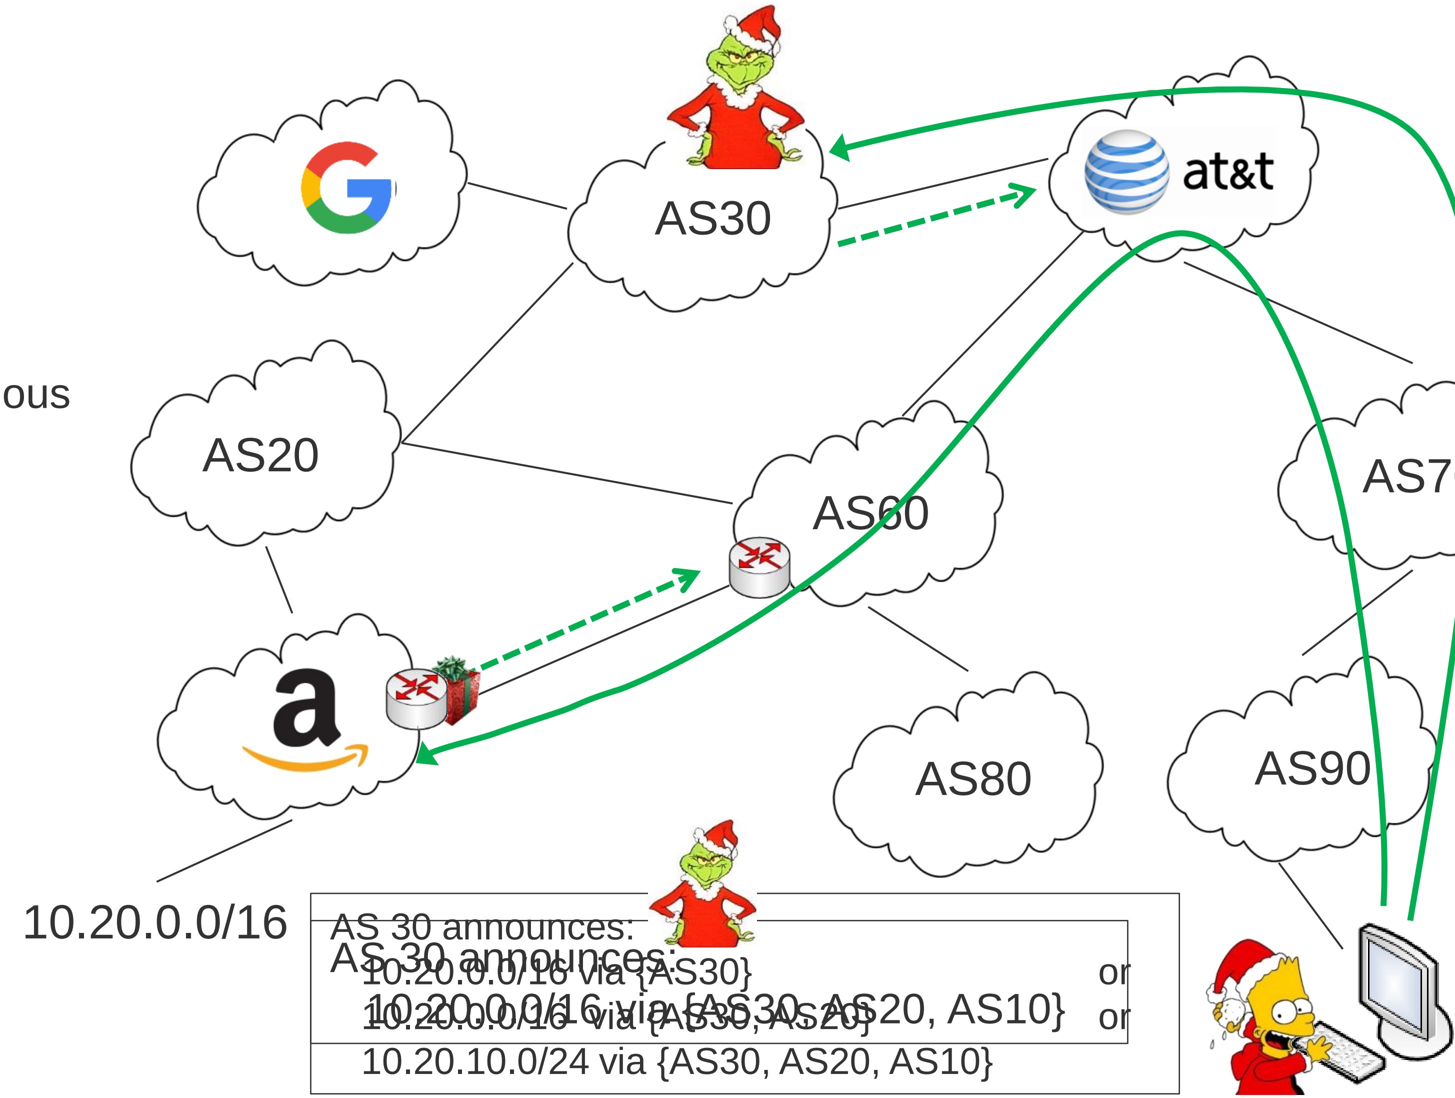

What is This Talk About?

How can you prevent your network from prefix hijacking?

How can you perform prefix origin validation?

What is the state of deployment of current countermeasures?

Why does the current web ecosystem challenges network security?

Why would you not deploy current security mechanisms in the backbone?

Agenda

1. Problem space
2. Proposed IETF solutions
3. Tools: Monitoring RPKI deployment
4. RPKI and the web ecosystem

RPKI

Problem

Original Design Choice (RFC 4271)

- BGP is based on trust between peers

Implications

- Any BGP speaker can claim to own an IP prefix
- Any BGP speaker can modify the AS path
- Receiver of a BGP update cannot verify the correctness of the data

Compromise

- Filtering
- Considering data of the Internet Routing Registry

⇒ This is not enough anymore!

Hijacks in the Real World?!

Prominent examples

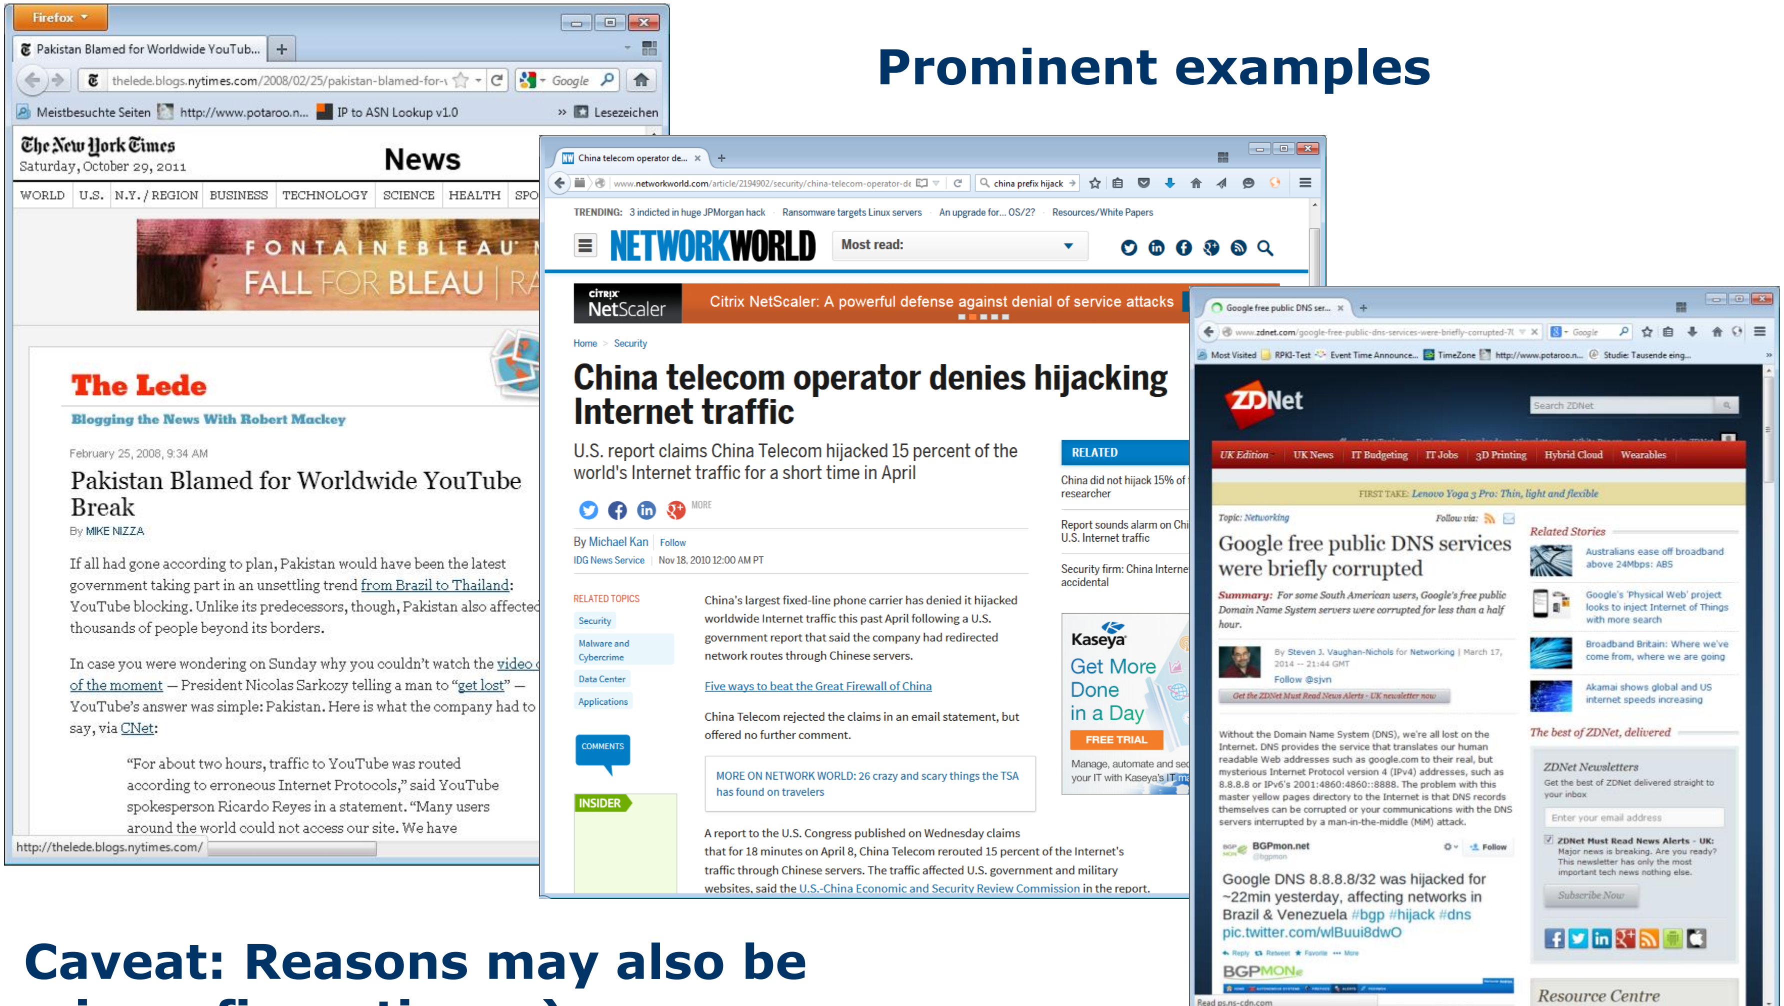

Caveat: Reasons may also be misconfiguration ;-)

Protection Concepts

1. **Prefix Origin Validation**
   - Mapping of IP prefixes and origin AS necessary
     - Including cryptographic proof
     - Prefix owner should be able to authenticate *Origin AS(es)*
   - BGP router compares BGP update with mapping
2. **Path Validation**
   - BGP path information are cryptographically secured
     - Paths will be signed

Challenges

- Cryptographic operations are complex
- Minimal additional load at routers

In the following we concentrate on 1.

Proposed Solution in the IETF

Resource Public Key Infrastructure (RPKI)

- System that allows to attest the usage of IP addresses and ASNs (i.e., Internet resources)
- RPKI includes cryptographically provable certificates

- Certificate hierarchy reflects IP-/AS-allocation in the Internet
  - Currently, each RIR creates a self-signed root certificate

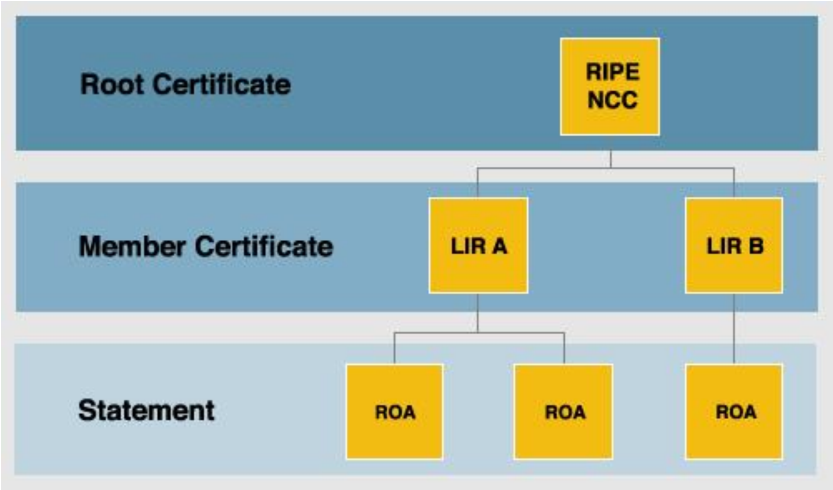

Source: RIPE

- Implementation of the RPKI started January 2011
- All RIRs participate

Routing Origination Authorization (ROA)

- Content of an ROA
  - Set of IP prefixes with minimal and maximal (optional) length
  - An AS number allowed to announce the prefixes
  - End-Entity-Certificate
- ROA will be signed with the certificate of the RPKI
- Note: Multiple ROAs per IP prefix possible

Example:

ROA

Valid from
01/10/2012 to
01/10/2013
+ E2E Cert

10.20.0.0/16-24 -> AS 123
80.90.0.0/16-16 -> AS 123

AS 123 is allowed to announce network range 10.20.0.0/16 to 10.20.0.0/24 and 80.90.0.0/16 from 1st Oct. 2012 until 1st Oct. 2013

Prefix Origin Verification & RPKI

Validation process consists of two steps

1. Validation of ROAs

- Performed at external cache

2. Validation of BGP updates

- Performed at BGP router
- No additional cryptographic operations necessary

How does the RPKI data comes to the BGP router?

Architecture Overview

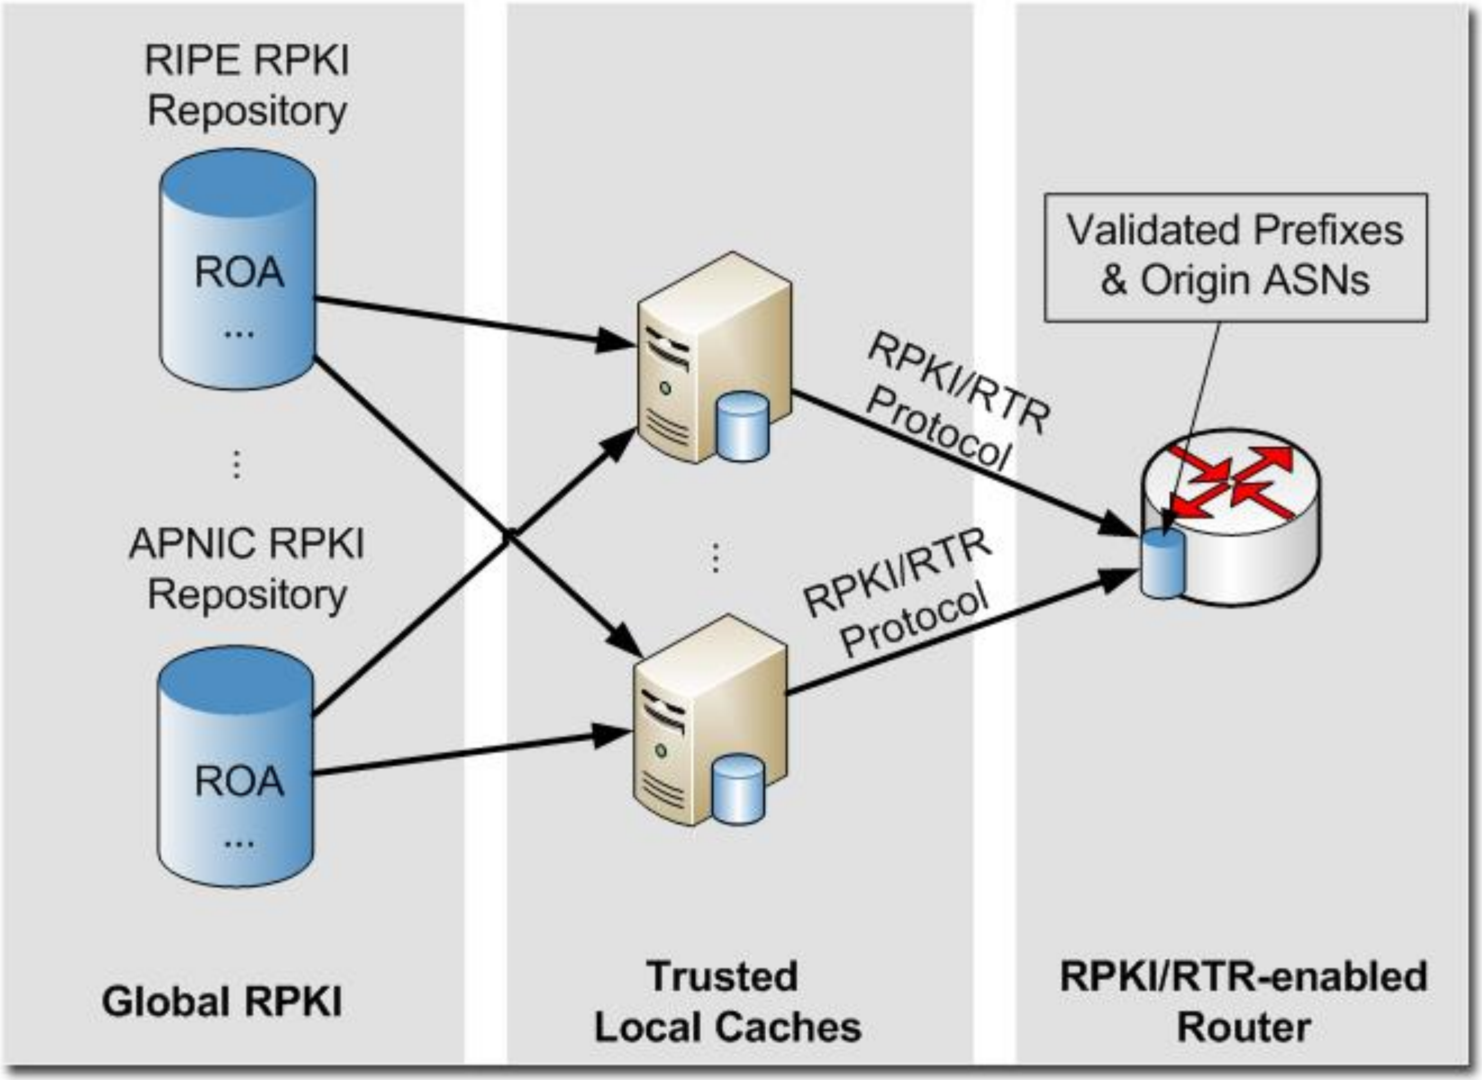

TOOLS

RTRlib [CSET@USENIX Security‘13]

General objective

- Open source implementation of the RPKI-RTR client protocol in C

Details

- Fetch validated prefixes + origin ASes from RPKI cache
- Keep the routers validation database in sync
- Provide an interface between local database and routing daemon to access validated objects
- Allow also for validation of BGP updates
- Conforms to relevant IETF RFCs/drafts

Applications

- Extending BGP daemons Quagga and BIRD
- Integration into CAIDA BGPstream
- +++

Memory Consumption

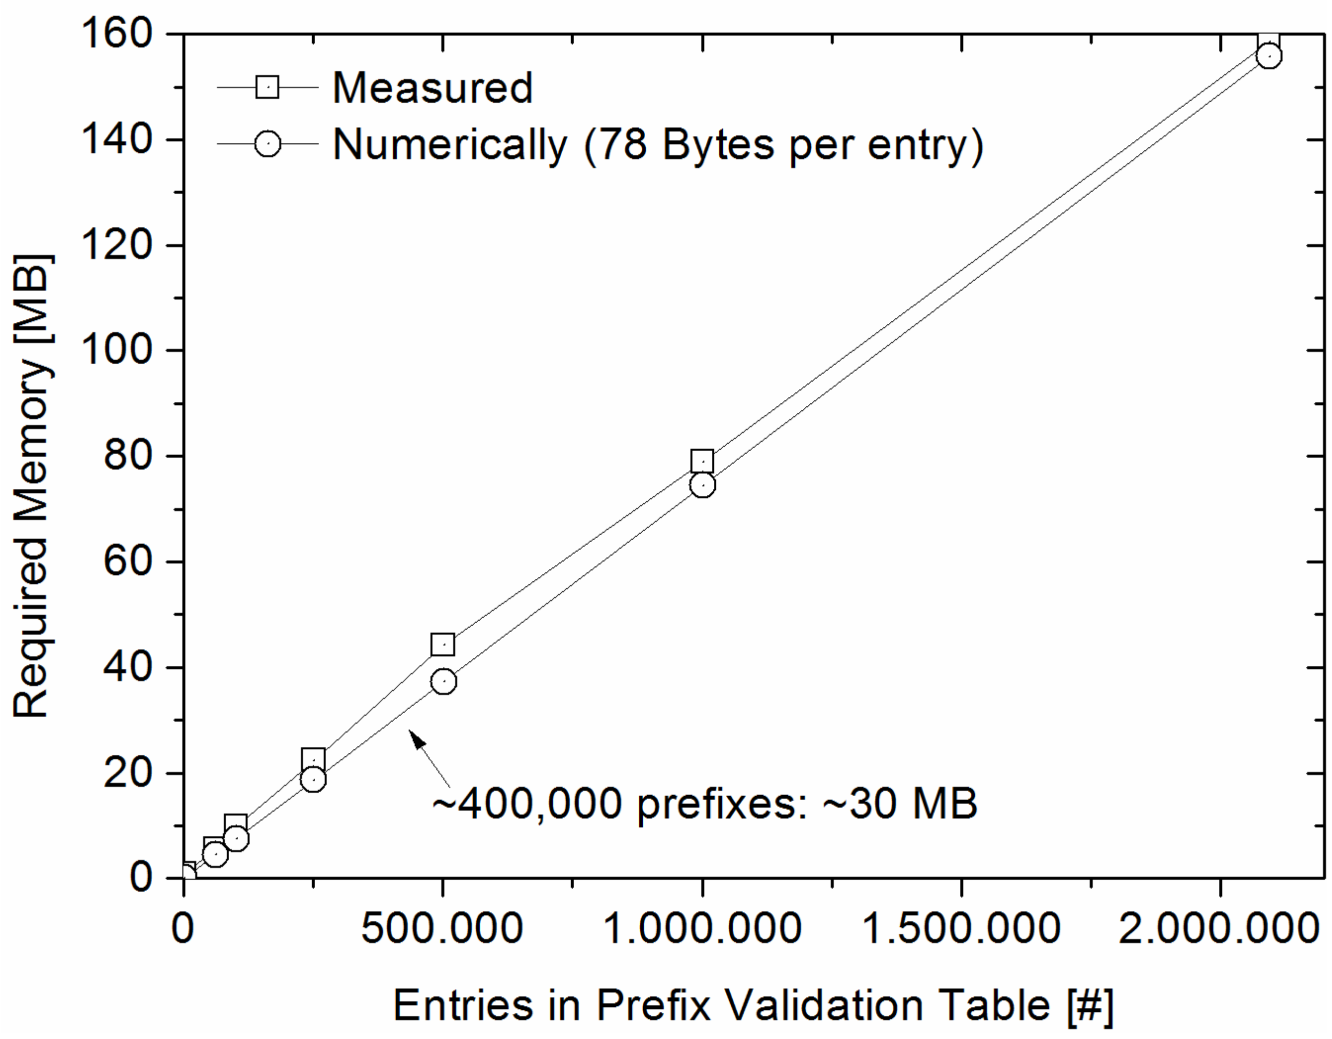

Delay While Loading ROA Data

Motivation: Router bootstrapping, Cache-Server-Reset

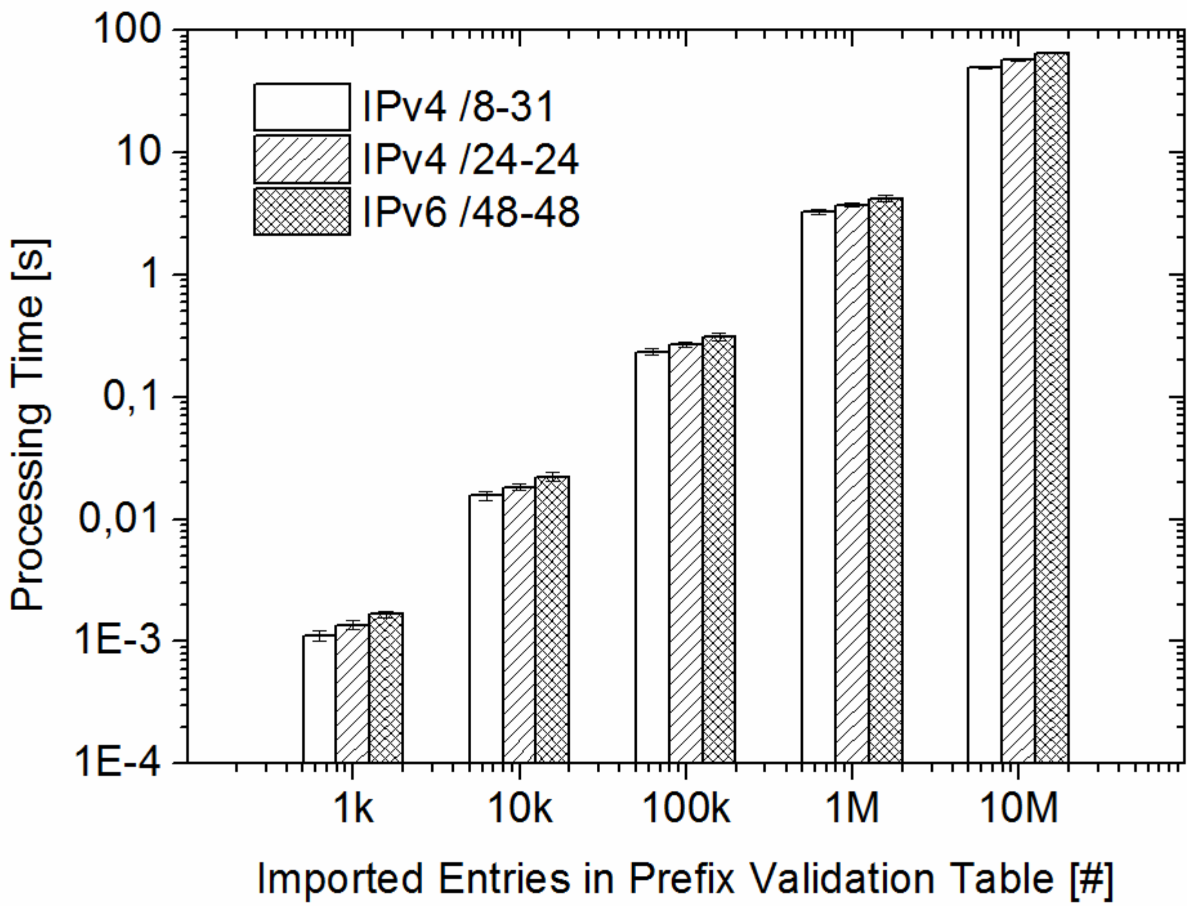

RPKI MIRO [Demo@SIGCOMM’15]

- Open source tool to monitor and explore RPKI repositories
- Modular architecture
  - Validator
  - Statistics
  - Browser
- Typical users
  - RIRs / CAs
  - Providers
  - Researchers
  - …
- https://github.com/rpki-miro
- http://rpki-browser.realmv6.org/

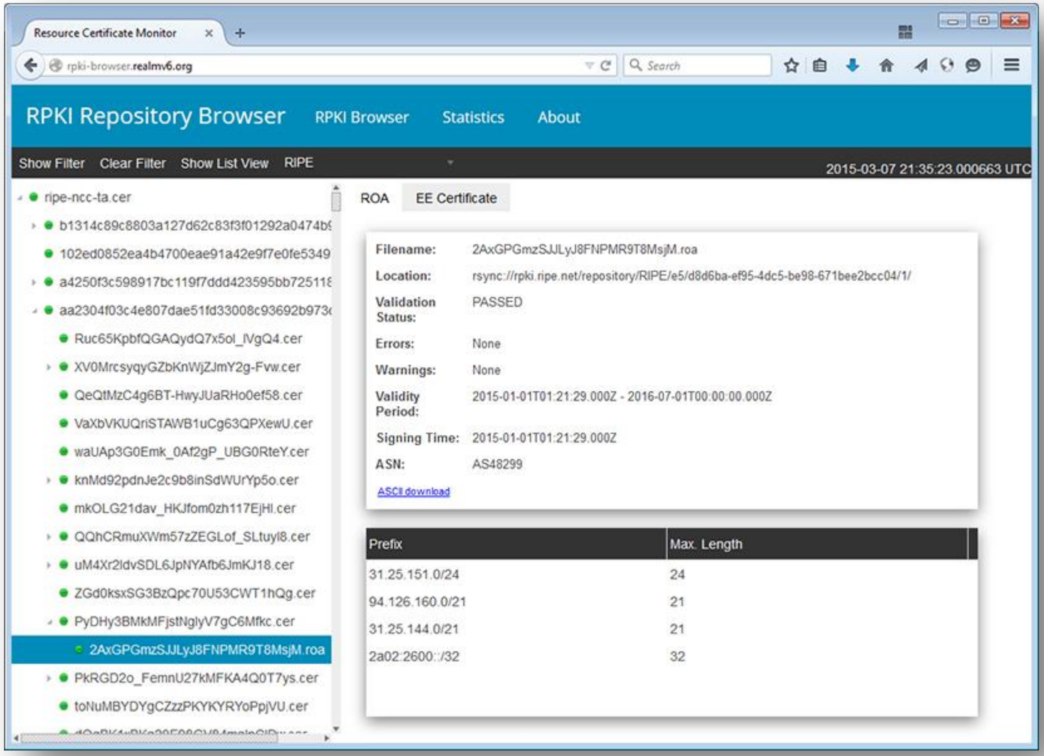

RiPKI:
RPKI & THE WEB ECOSYSTEM
[HOTNETS’15]

Motivation

Exclusive protection by TLS is insufficient!

1. Compromised trusted CAs
    - DANE rarely deployed
2. Forged certificates
    - DANE rarely deployed
    - Extended Validation rarely deployed [IMC’11]
    - Leveraged by prefix hijacking [Black Hat’15]
3. Blackholing
    - Implemented by prefix hijacking

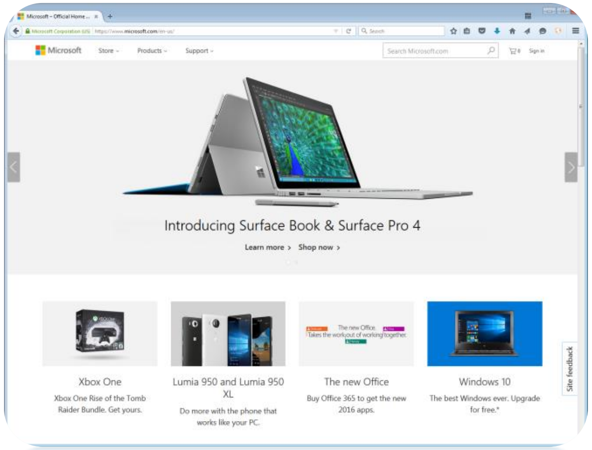

Attacker Model (in the Web Ecosystem)

Freie Universität Berlin

- Attacker is able to manipulate Internet routing
- Drop or forward redirected traffic to web server

Advantages compared to common DDoS attacks in the web

- DDoS and data manipulation are possible
- Attack does not need to affect all clients
- Web server is not aware of attack

Objectives

Freie Universität Berlin

Empirically explore the relationship between web hosting infrastructure and RPKI deployment (ROA creation).

Which web servers are secured by the RPKI?

Web Ecosystem

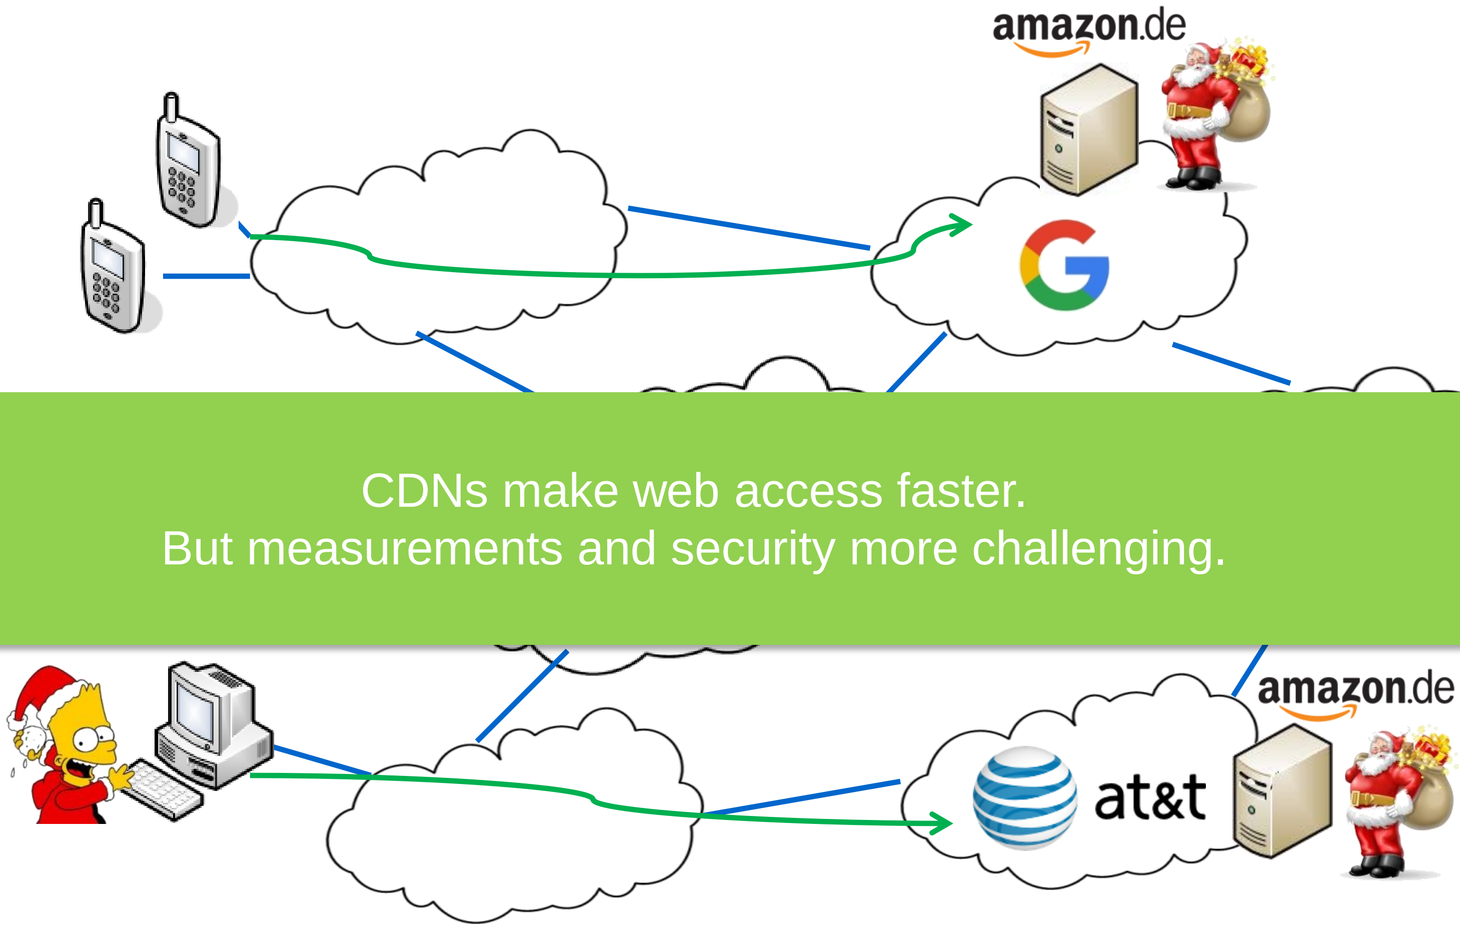

CDNs make web access faster.
But measurements and security more challenging.

Challenges

- DNS resolution results may depend on the location
- DNS resolution is time-consuming

⇒ We use stable, public ORDNS servers

- Embedded content

⇒ This study focuses on landing page

- Selecting domain names

⇒ Prefix www and w/o www

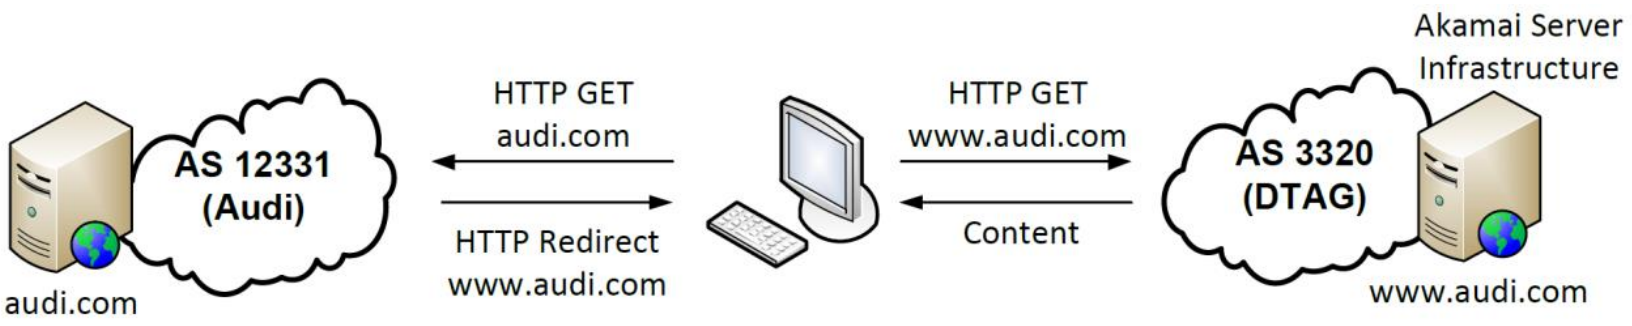

Overview: Measurement Methodology

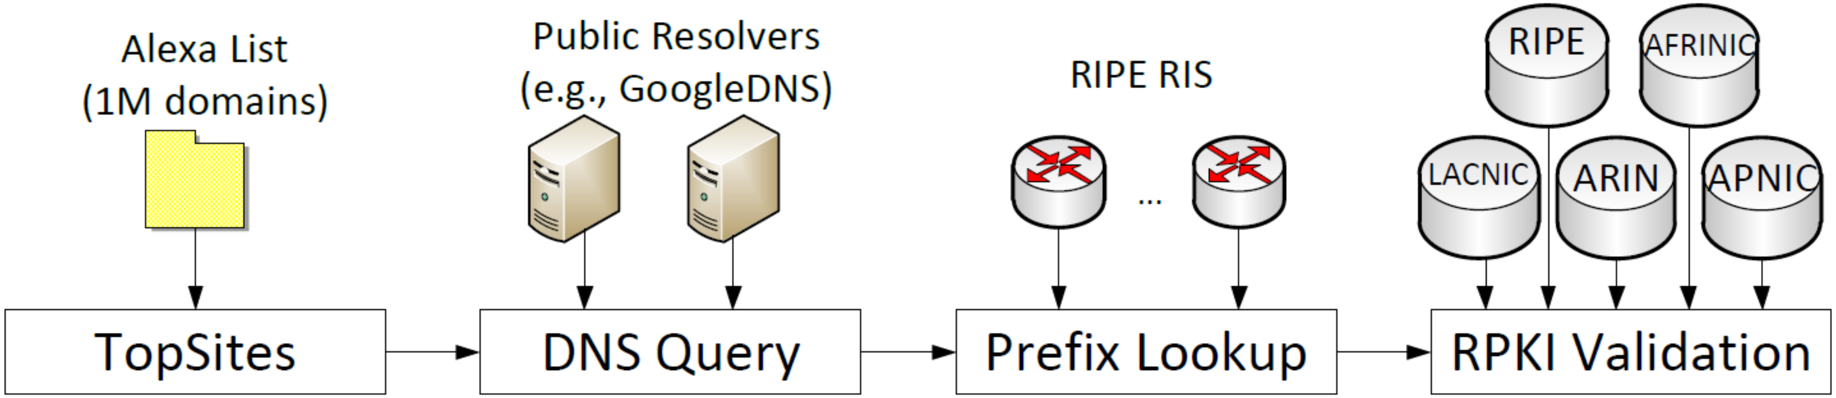

Side Result: Reducing Measurement Overhead?

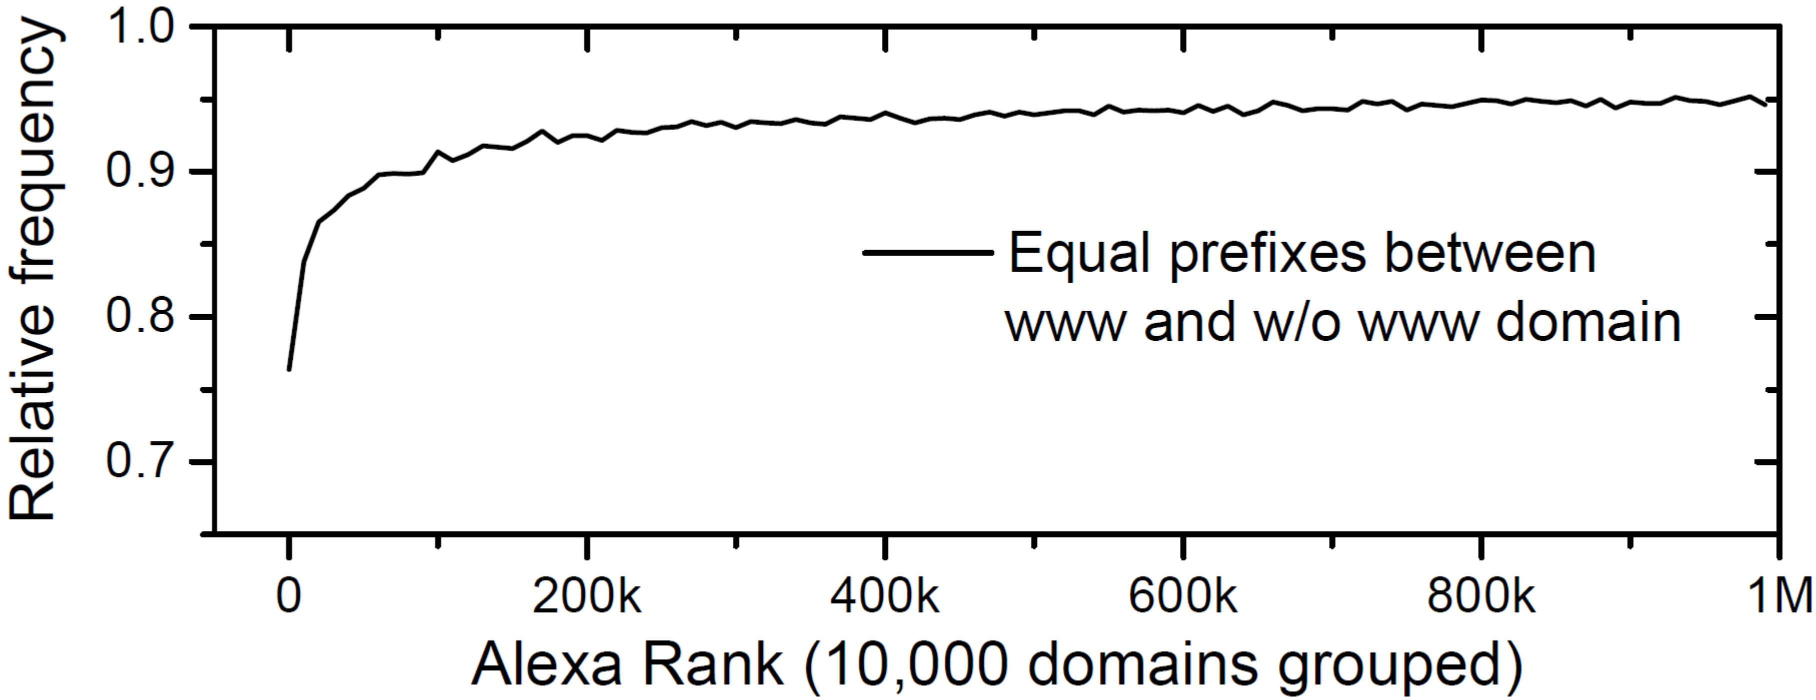

RPKI Validation Outcome for 1M Web Sites

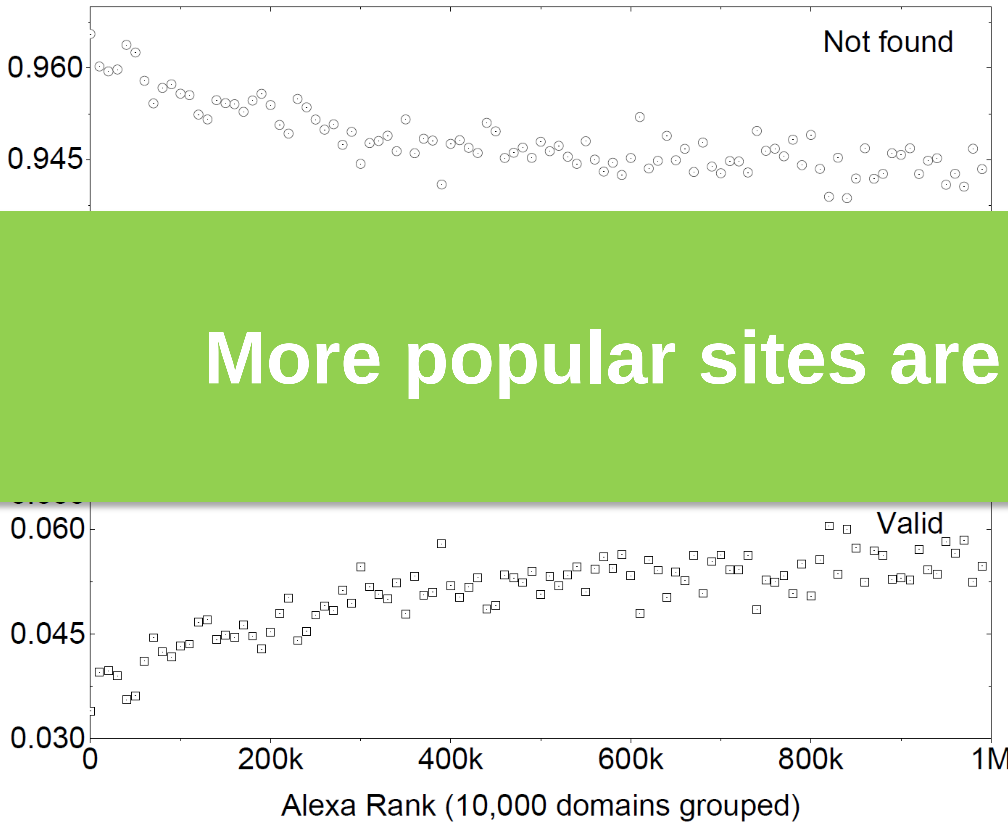

More popular sites are less secured!

130	booking.com	✓ (4/4)	✓ (2/2)

Popularity of CDNs Across Ranks

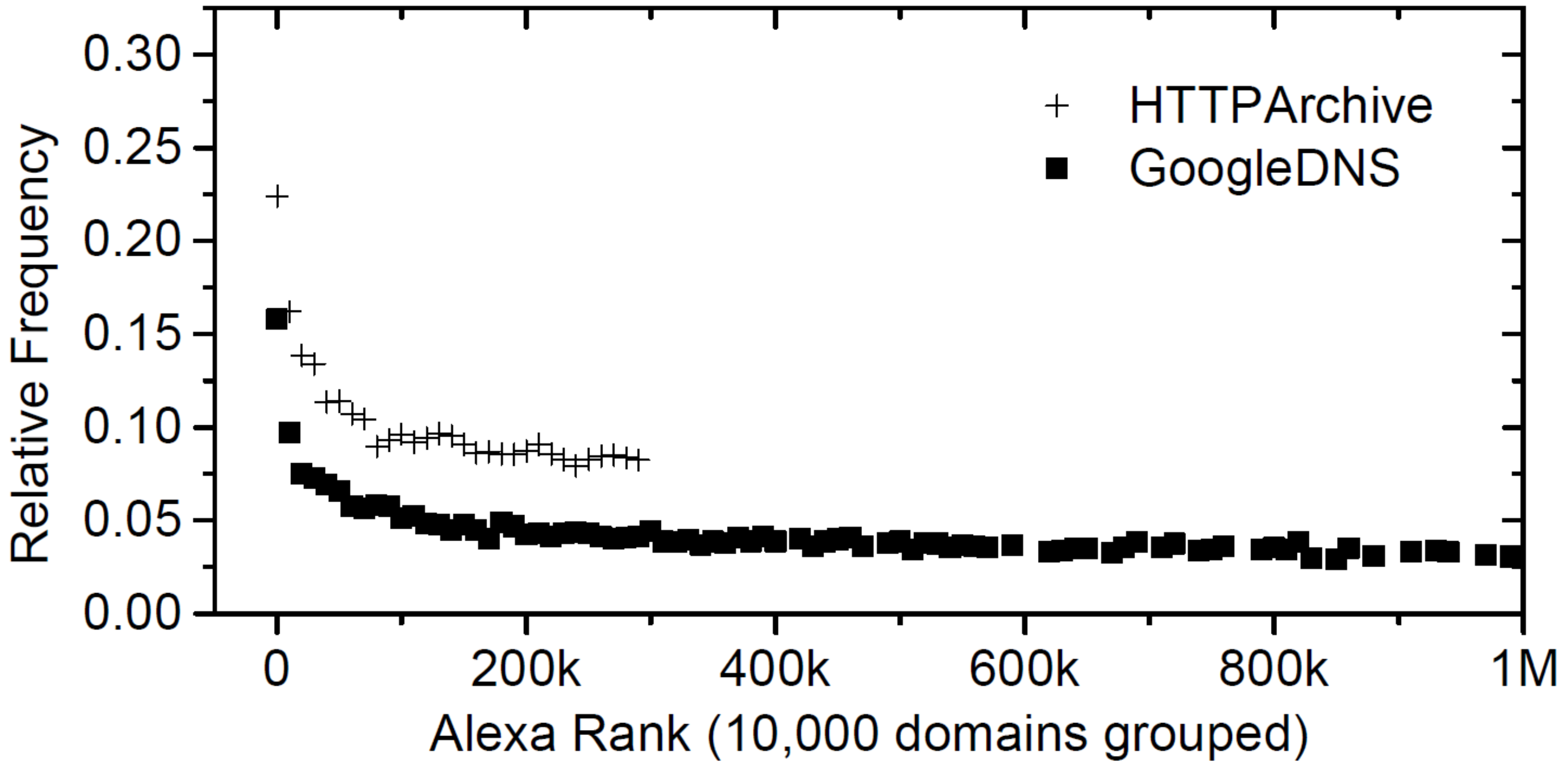

Do CDNs Push a Specific Rank?

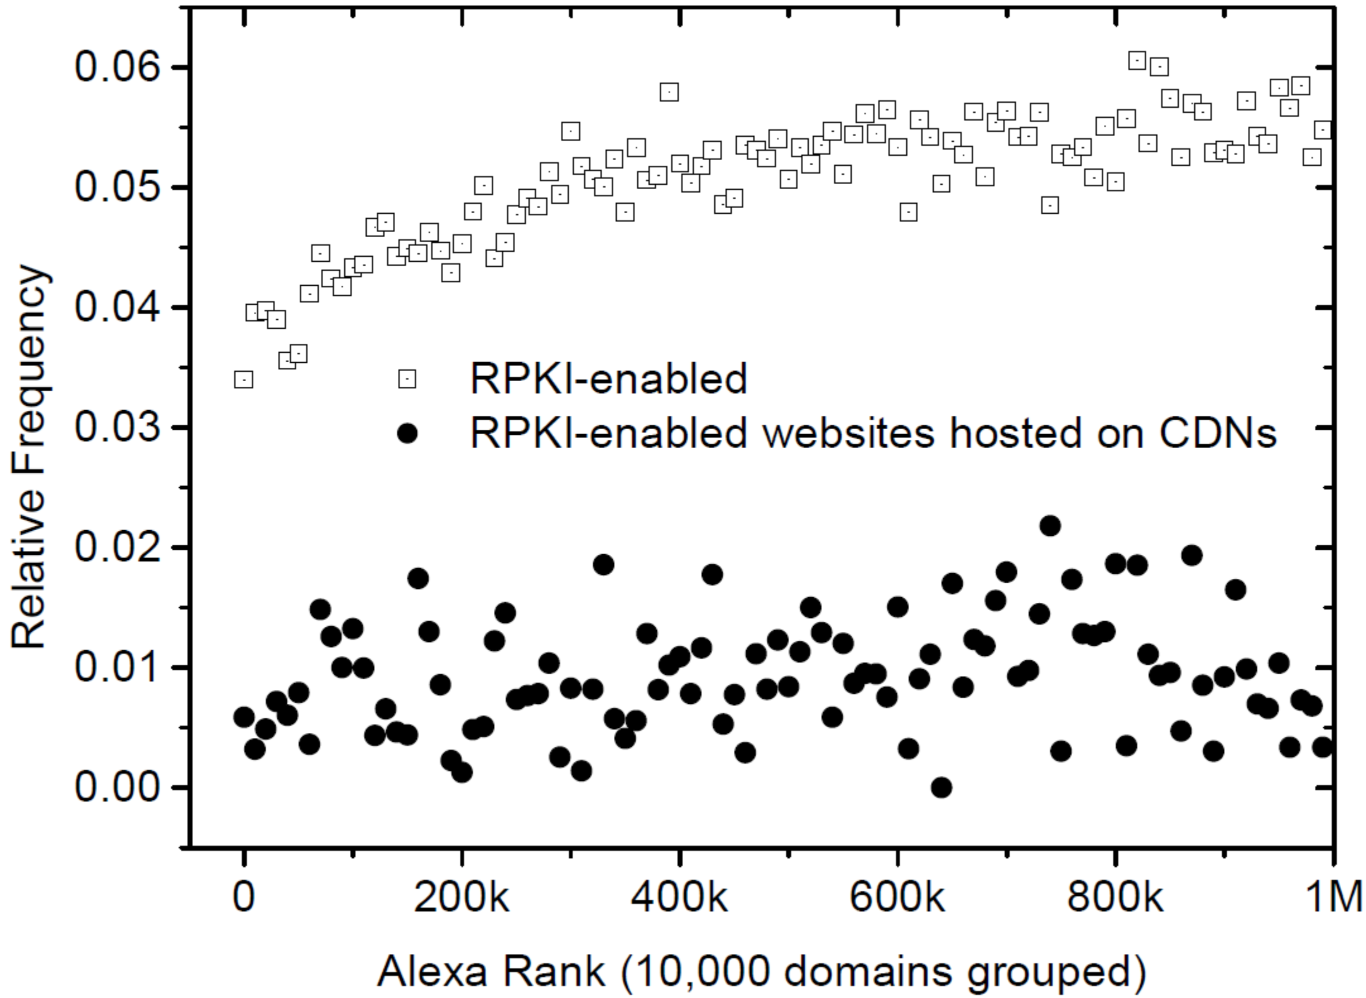

Reasons for not Deploying RPKI

- Political reasons
  - RIR are trust anchors
  - Local law may instruct RIR to revoke certificates
  - ROAs become invalid
  - Out of control of the operator

- Business reasons
  - RPKI implements a positive attestation model
  - ISPs have to add prefix-AS relation in advance
  - Might conflict with business policies

- Cost and complexity reasons

First Steps Towards Improved Browsing Experience

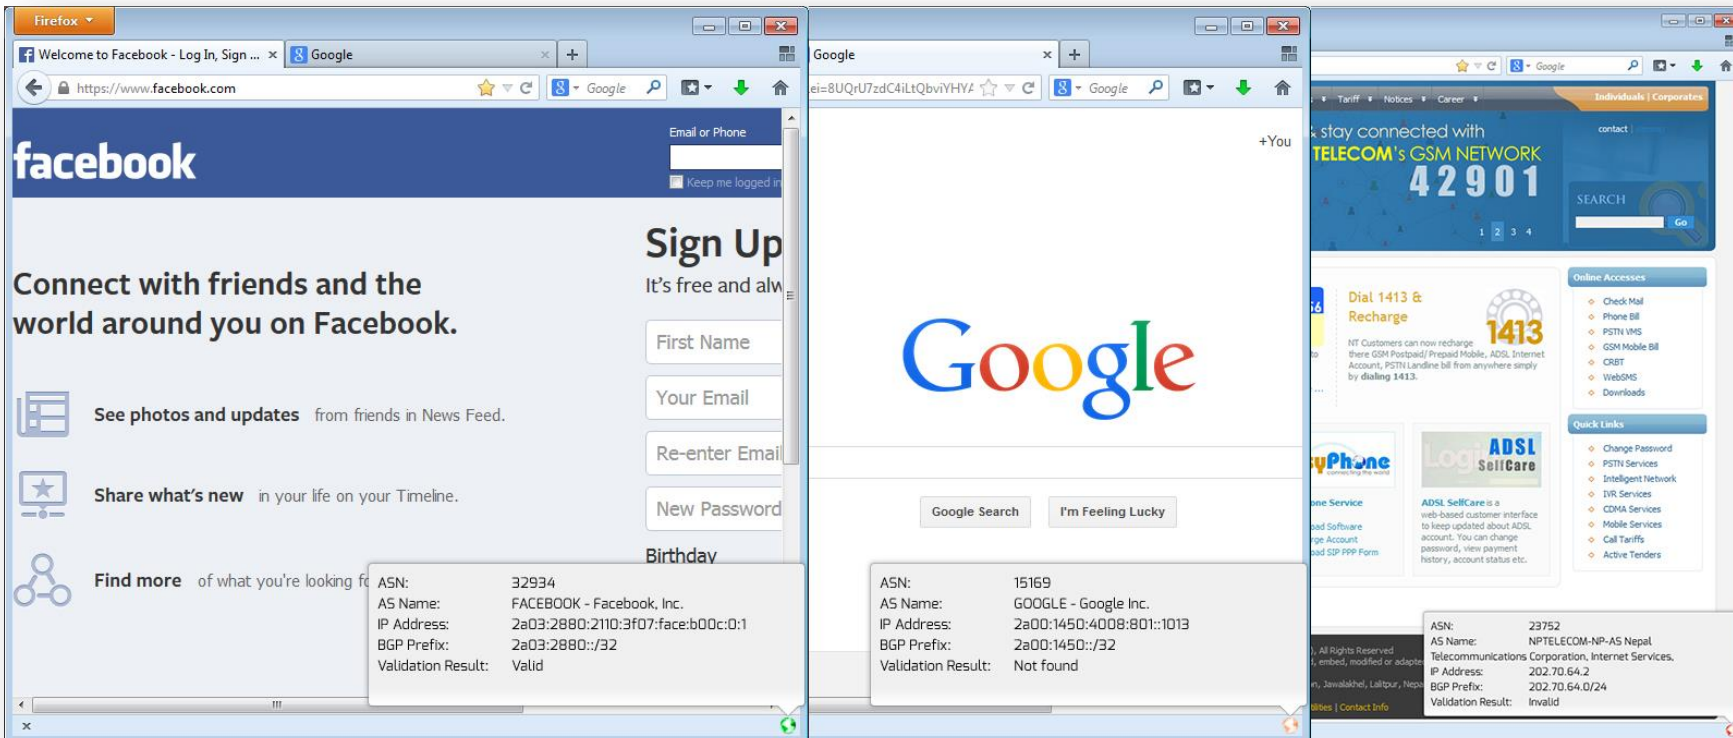

Conclusion

- RPKI is one building block in securing e2e communication
- CDNs are hesitant in deploying RPKI, popular sites are less secure
- CDN content benefits from RPKI deployment in 3rd party networks

Future research topics

- Improve web measurement methodology
  - Accelerate DNS measurements …
- Consider embedded content from external sites
- Improve securing web (content delivery) architecture
- Understand better *why* operators do not deploy security
  - Deployment comparison with DNSSEC