Network Security and Measurement

- BGP Hijacking and RPKI -

Prof. Dr. Thomas Schmidt

http://inet.haw-hamburg.de | t.schmidt@haw-hamburg.de

Agenda

BGP Hijacking

Resource Public Key Infrastructure

Monitoring with the RTRlib

Measuring the RPKI

Steeling resources from the Internet

BGP HIJACKING

How can an Attacker Try to Hijack Your IP Prefix?

You
AS 123 announces IP prefix 10.20.0.0/16

Me
Receive a BGP update with
path 123, 10, 20, 567

Attacker

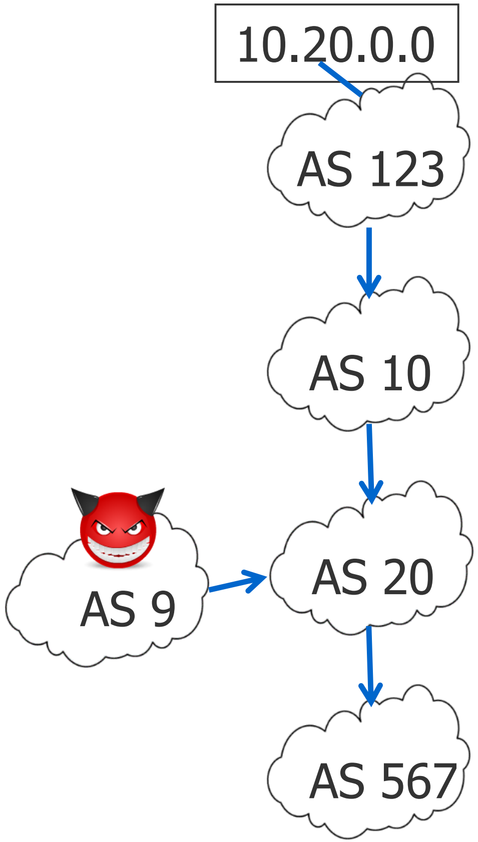

How can an Attacker Try to Hijack Your IP Prefix?

You

AS 123 announces IP prefix 10.20.0.0/16

Me

Receive a BGP update with path 123, 10, 20, 567

Receive a BGP update with path 9, 20

Attacker

Announces 10.20.0.0/16

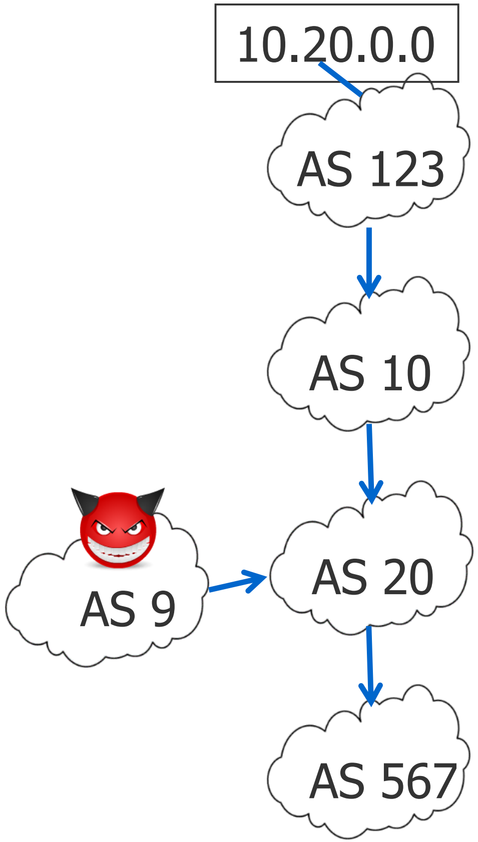

How can an Attacker Try to Hijack Your IP Prefix?

You

AS 123 announces IP prefix 10.20.0.0/16

Me

Receive a BGP update with path 123, 10, 20, 567

Receive a BGP update with path 9, 20

Receive a more specific prefix

Attacker

Announces 10.20.0.0/16

Announces 10.20.30.0/24

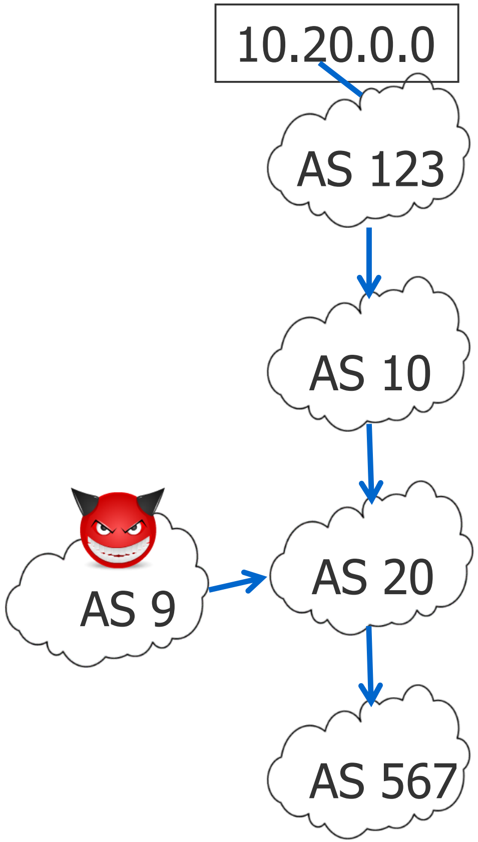

Hijacks in the Real World?

Prominent examples

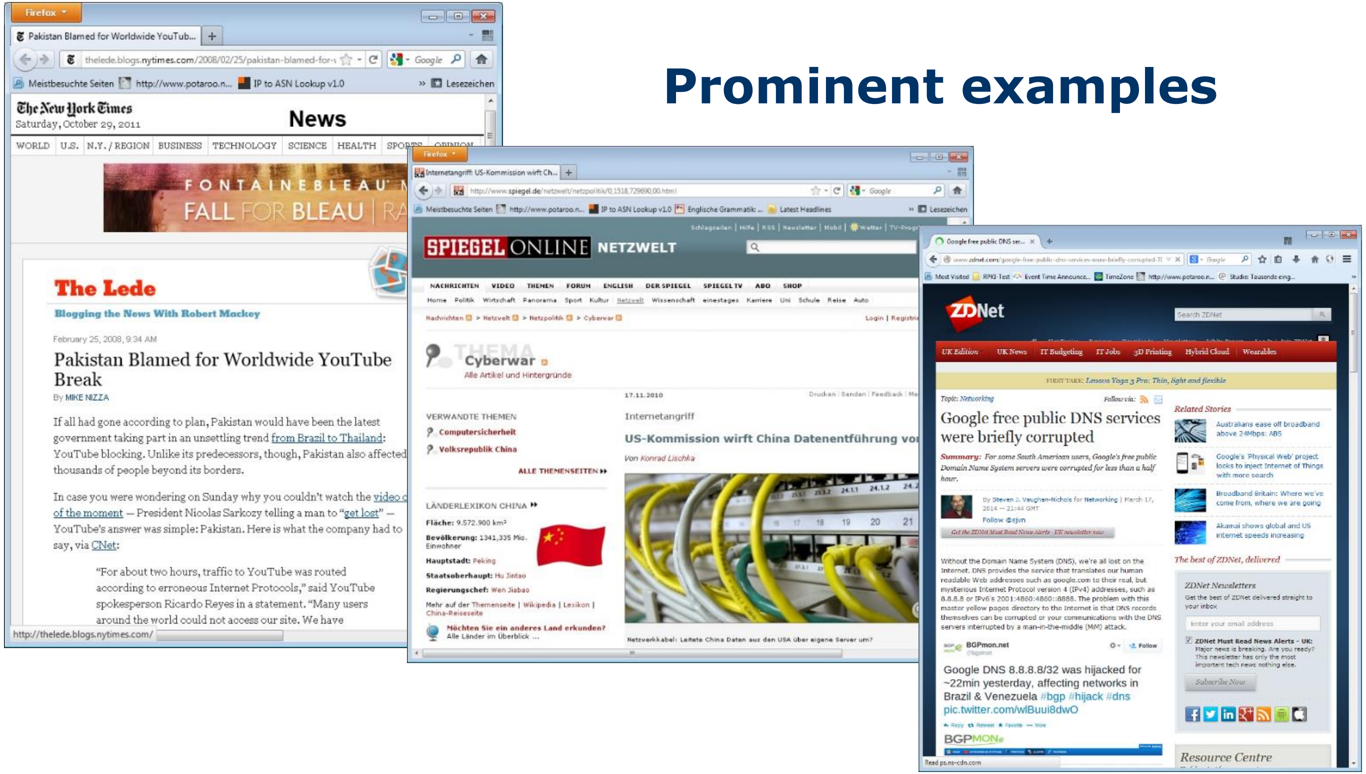

Hijacks in the Real World?

Prominent examples

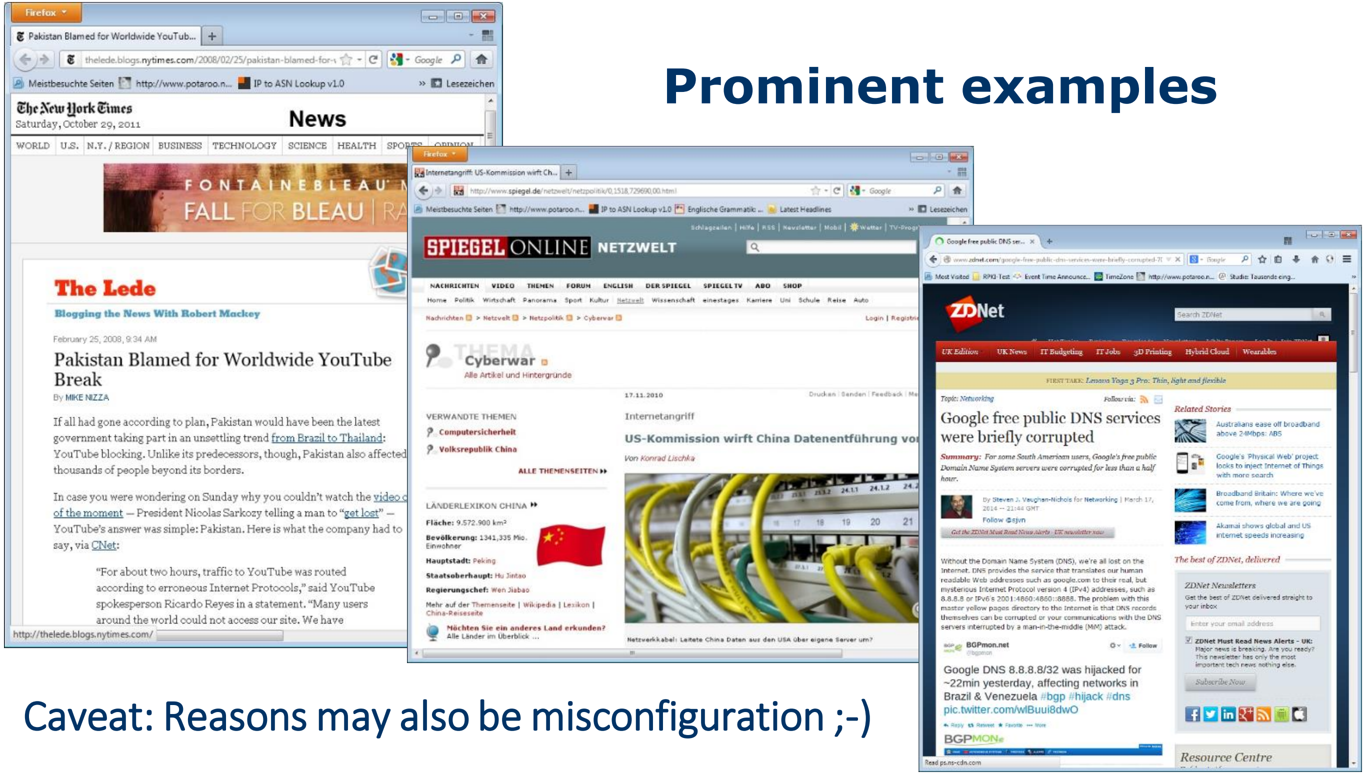

Caveat: Reasons may also be misconfiguration ;-)

Problem

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

$\Rightarrow$ This is not enough anymore!

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 hop-wise
   - BGP routers validate hops

Challenges

Cryptographic operations are complex

Minimize additional load at routers

Protection Concepts

RPKI: Resource Public Key Infrastructure
RFCs 6480, 6811

BGPsec: Secure BGP
RFC 8205

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 hop-wise
   - BGP routers validate hops

Challenges

Cryptographic operations are complex
Minimize additional load at routers

Validating the prefix origins

RPKI

Resource Public Key Infrastructure (RPKI)

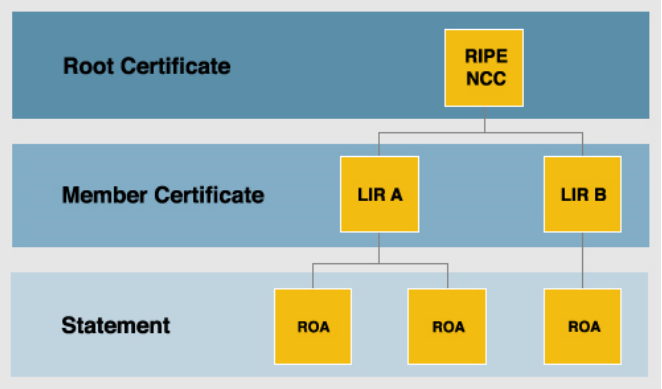

Source: RIPE

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

Implementation of the RPKI started January '11

All RIRs participate

Routing Origination Authorization (ROA)

Content of a 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 01/10/2020 to 01/10/2023 + 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. 2020 until 1st Oct. 2023

RPKI & ROA

All certificates including ROAs will be published in RPKI repositories

- Each RIR (including RIPE NCC ;) operates one
- ISPs can maintain their own repository
- Information of all repositories describe the overall picture

Check if AS is allowed to announce IP prefix
= check the corresponding ROA

- Corresponding ROA will be determined based on CIDR
- ROA needs cryptographic verification
- ROAs implements a positive attestation
  - If a ROA for a prefix exists, announcements of all origin ASes that are not included will be considered INVALID

Current Deployment:
IP prefixes in ROAs

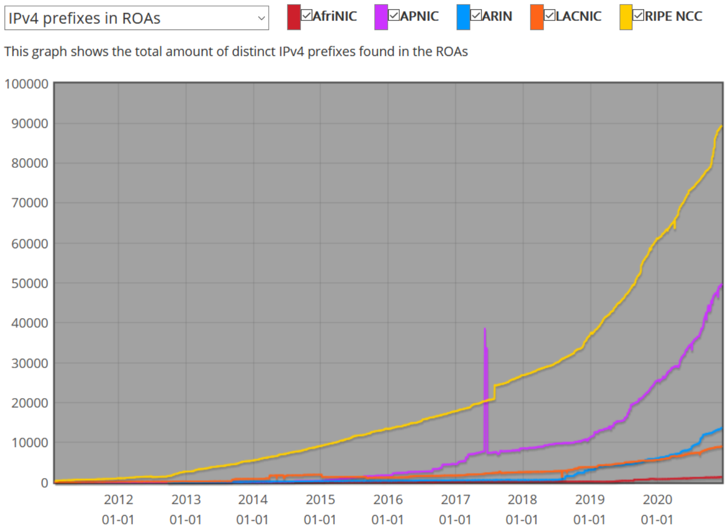

IPv4

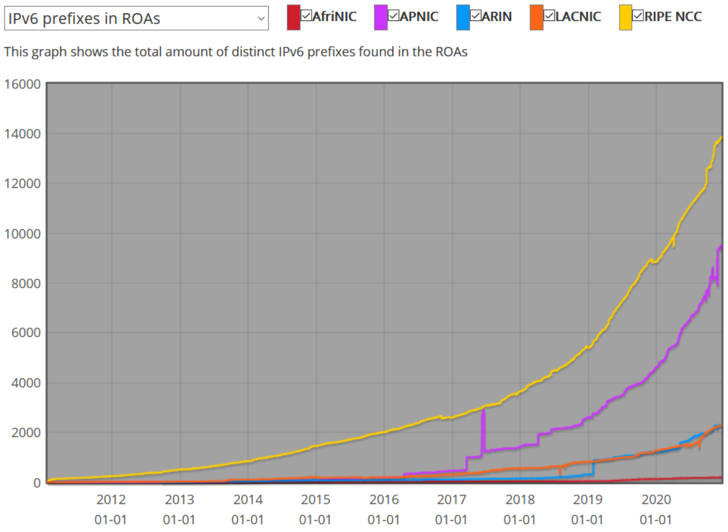

IPv6

http://certification-stats.ripe.net/

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

IETF “RPKI/RTR protocol” manages push of *valid* ROAs from cache to BGP router

- Implementations for Cisco and Juniper available
- Open Source BGP daemons on the way

Evaluation result of BGP update: VALID, INVALID, NOT_FOUND

- Combine the outcome with BGP policies

Architecture Overview

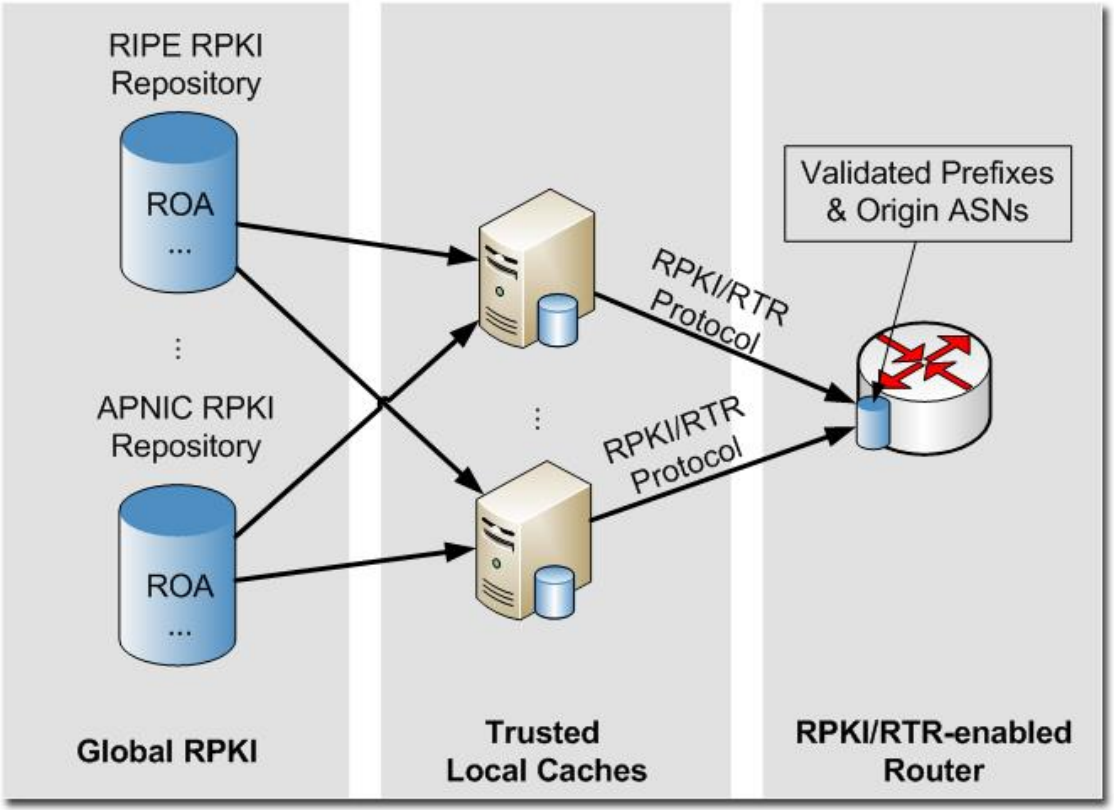

Validation Outcome

Validation of an ASN/Prefix pair against RPKI results in either

Valid

If at least one valid ROA exists that covers the announced prefix and matches the BGP origin AS, with max length less or larger than the BGP prefix length

Invalid

If no covering ROA matches the BGP origin AS or the announced prefix is more specific

Not Found

If no covering ROA exists

Validation Outcome - Examples

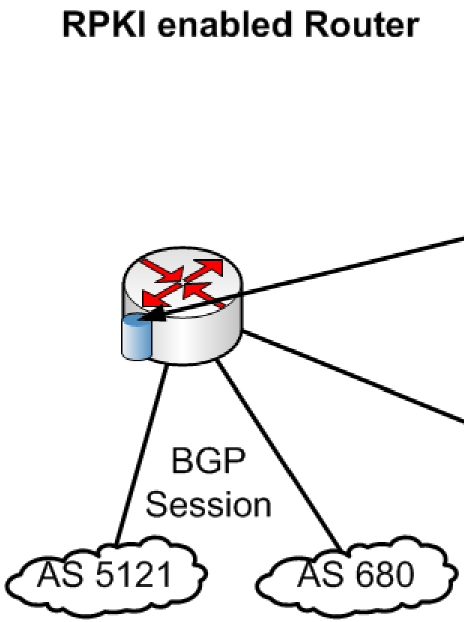

Validated Prefixes & Origin ASNs from RTR Cache

ASN	Prefix	MaxLen
100	10.20.0.0/16	24
200	50.60.70.0/24	32
200	50.60.70.0/24	28

Received BGP Updates and validation outcome

ASN	Prefix	Validation outcome
100	10.20.0.0/16	Valid
100	10.20.0.0/18	Valid
100	10.20.5.0/28	Invalid
200	10.20.0.0/16	Invalid
300	1.2.0.0/16	Not found

HAW HAMBURG

Validation Outcome - Examples

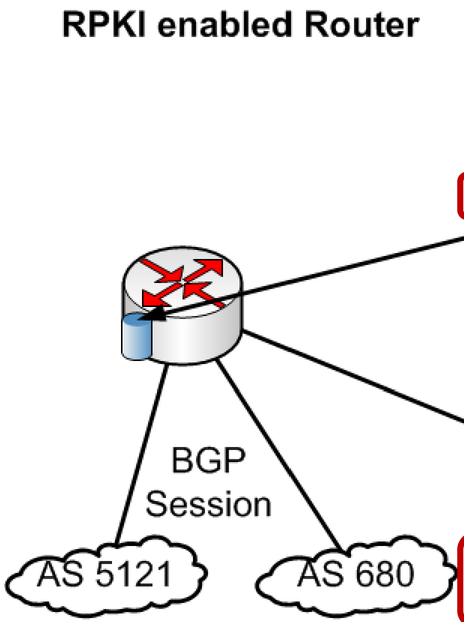

Validated Prefixes & Origin ASNs from RTR Cache

ASN	Prefix	MaxLen
100	10.20.0.0/16	24
200	50.60.70.0/24	32
200	50.60.70.0/24	28

Received BGP Updates and validation outcome

ASN	Prefix	Validation outcome
100	10.20.0.0/16	Valid
100	10.20.0.0/18	Valid
100	10.20.5.0/28	Invalid
200	10.20.0.0/16	Invalid
300	1.2.0.0/16	Not found

Validation Outcome - Examples

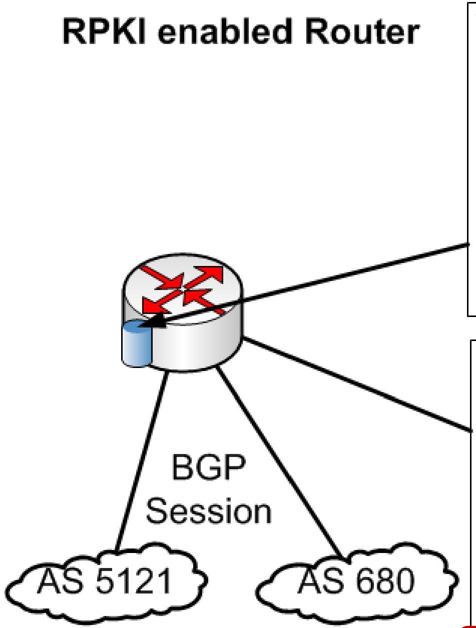

Validated Prefixes & Origin ASNs from RTR Cache

ASN	Prefix	MaxLen
100	10.20.0.0/16	24
200	50.60.70.0/24	32
200	50.60.70.0/24	28

Received BGP Updates and validation outcome

ASN	Prefix	Validation outcome
100	10.20.0.0/16	Valid
100	10.20.0.0/18	Valid
100	10.20.5.0/28	Invalid
200	10.20.0.0/16	Invalid
300	1.2.0.0/16	Not found

Validation Outcome - Examples

RPKI enabled Router

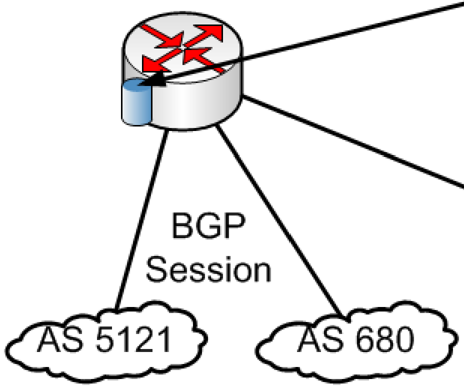

Validated Prefixes & Origin ASNs from RTR Cache

ASN	Prefix	MaxLen
100	10.20.0.0/16	24
200	50.60.70.0/24	32
200	50.60.70.0/24	28

Received BGP Updates and validation outcome

ASN	Prefix	Validation outcome
100	10.20.0.0/16	Valid
100	10.20.0.0/18	Valid
100	10.20.5.0/28	Invalid
200	10.20.0.0/16	Invalid
300	1.2.0.0/16	Not found

Validation Outcome - Examples

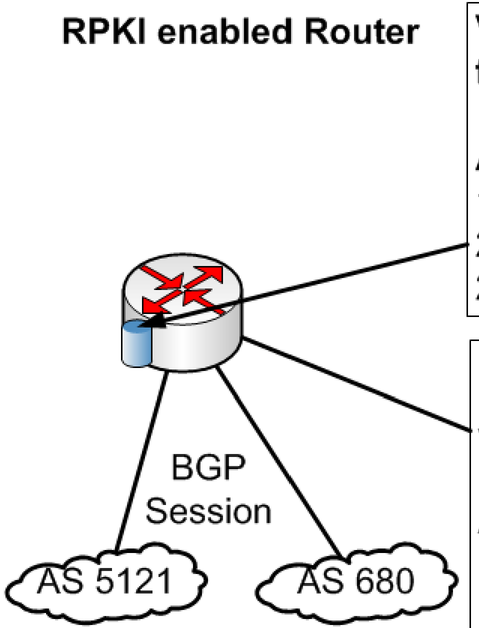

Validated Prefixes & Origin ASNs from RTR Cache

ASN	Prefix	MaxLen
100	10.20.0.0/16	24
200	50.60.70.0/24	32
200	50.60.70.0/24	28

Received BGP Updates and validation outcome

ASN	Prefix	Validation outcome
100	10.20.0.0/16	Valid
100	10.20.0.0/18	Valid
100	10.20.5.0/28	Invalid
200	10.20.0.0/16	Invalid
300	1.2.0.0/16	Not found

Zero-day Measurements: Valide vs. Invalide BGP Updates

Januar 2012

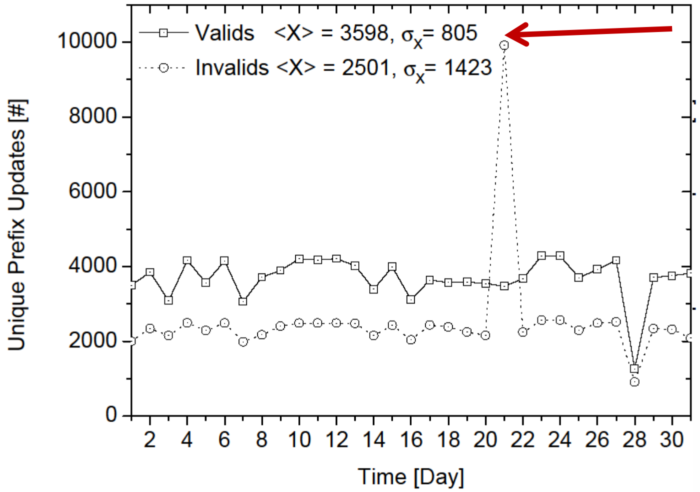

Mai 2012

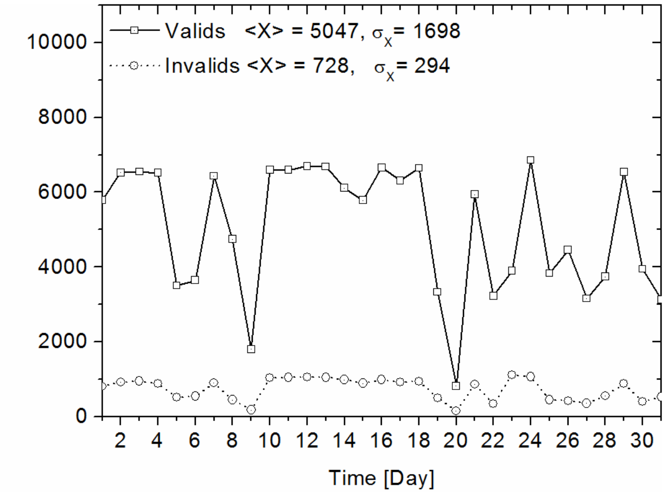

Number of invalids decreases over time

Zero-day Measurements: Valide vs. Invalide BGP Updates

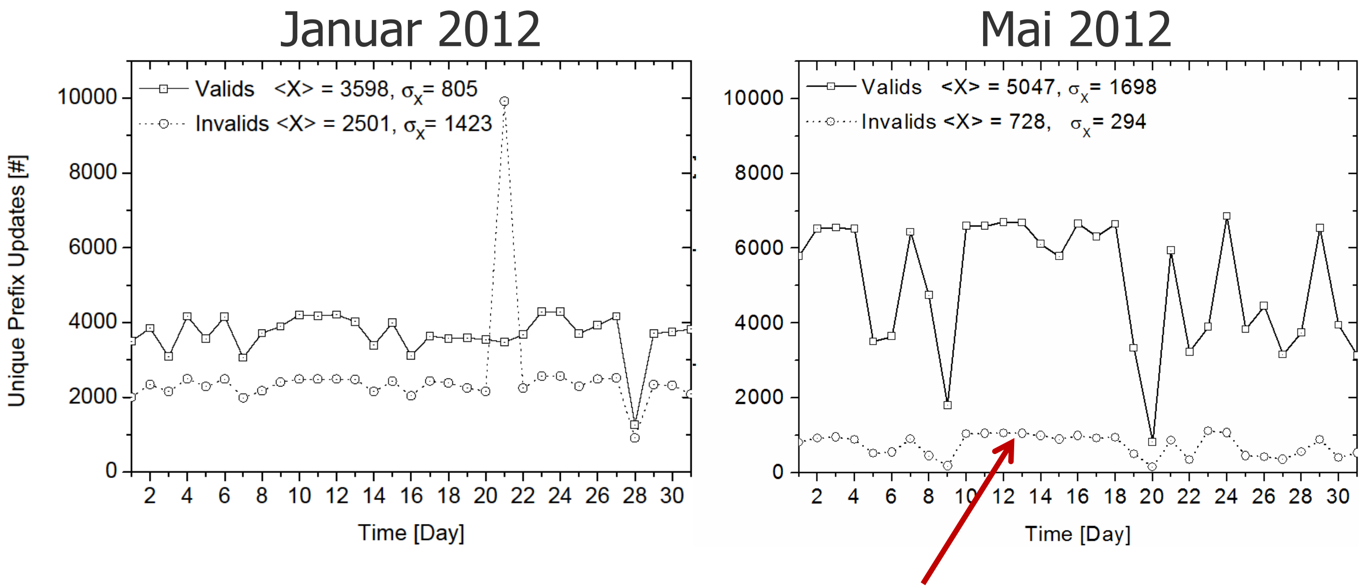

Number of invalids decreases over time

Are these updates really hijacks??

Some Common Pitfalls - Examples

Case 1: Missing Customer (or Sibling) Legitimation

ROA created: 12.0.0.0/8-9 -> AS 7018

AS 27487 announces 12.0.19.0/24

AS 2386 announces 12.1.216.0/24

⇒ Consider sub-allocations, start most specific

Both announcements are invalid if no ROAs exists

Case 2: (De-)Aggregation

ROA created: 78.192.0.0/10-10 -> AS 12322

Usual announcement: 78.192.0.0/10

For 30 minutes: 78.192.10.0/24 …

⇒ Configure the max ROA prefix length explicitly

Common Pitfalls – Overview (1)

Valid origin, announced prefix is more specific

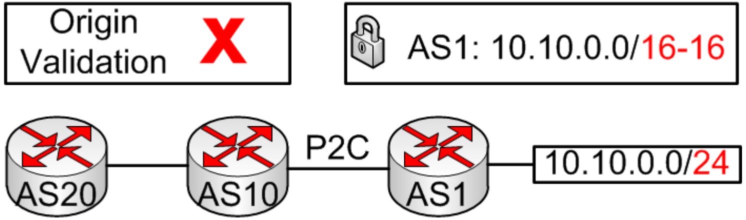

Provider does not consider customers

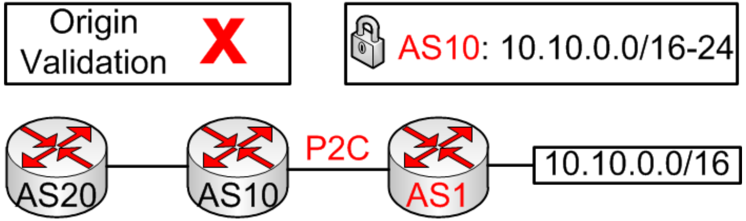

Common Pitfalls – Overview (2)

Additional AS of a company is not authorized

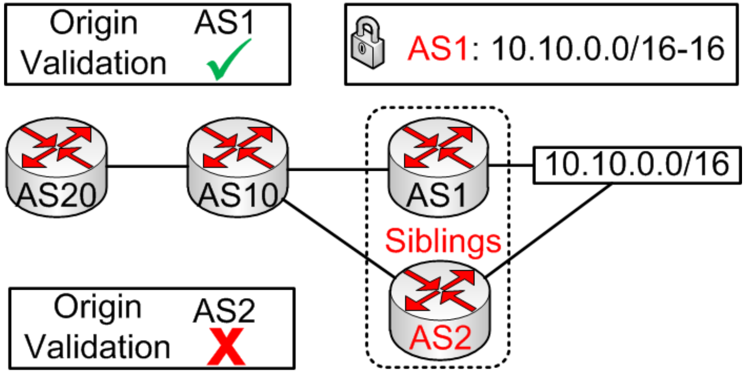

Monitoring with the RPKI Router Part

RTRLIB

What is the RTRlib?

General objective

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

It’s open-source: http://rpki.realmv6.org

Applications

Extension of BGP daemons

- Now part of FRR, (Quagga), BIRD (code-wise), and commercial products

Monitoring of the RPKI deployment

- Integrate the library in your Python/Perl … scripts
- Particularly suitable for real-time monitoring

Testing purposes

- Evaluate performance of your RPKI/RTR cache server
- Play around with BGP update validation

Monitoring Scenario (Example)

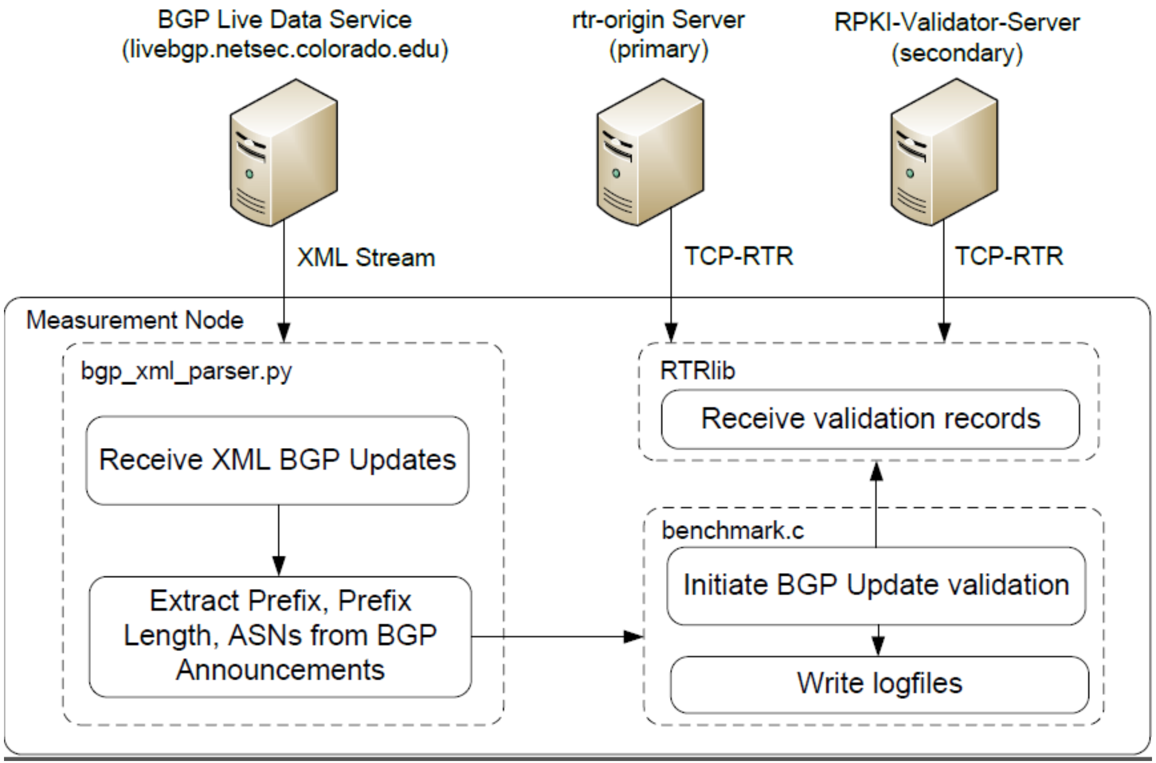

Going wild

MEASURING THE RPKI

Which web servers are secured by the RPKI?

Empirically explore the relationship between web hosting infrastructure and RPKI deployment.

[HotNets `15]

Web Ecosystem

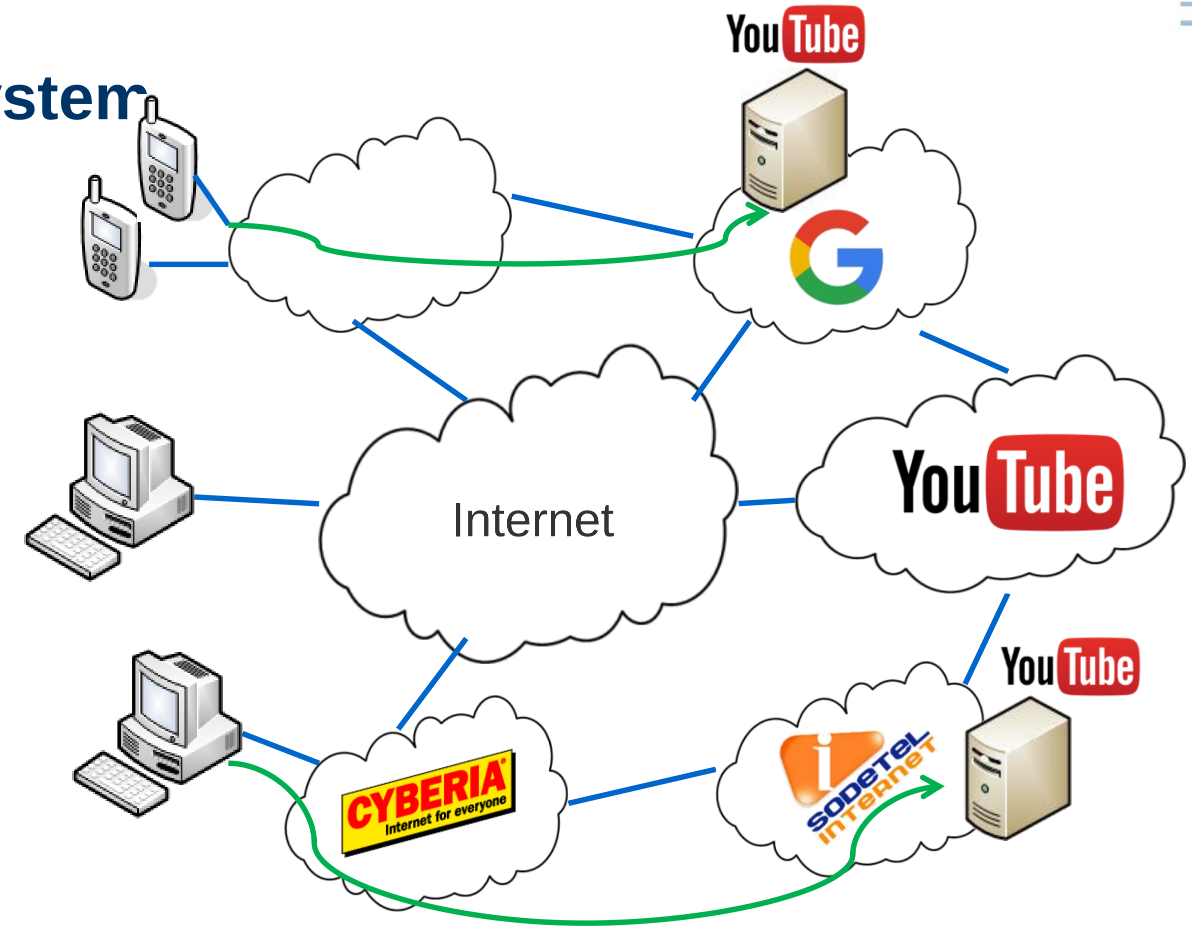

Web Ecosystem

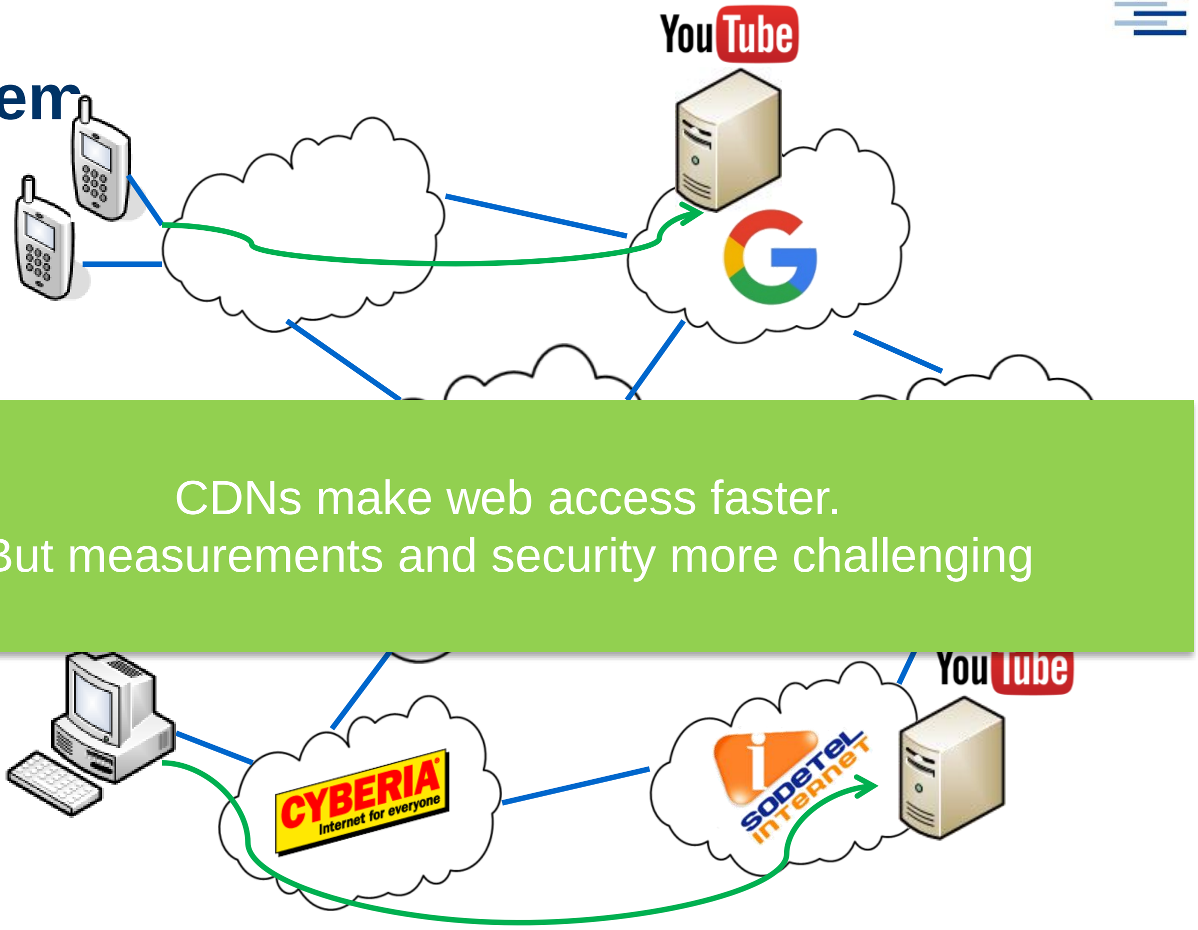

CDNs make web access faster.
But measurements and security more challenging

Measurement Methodology

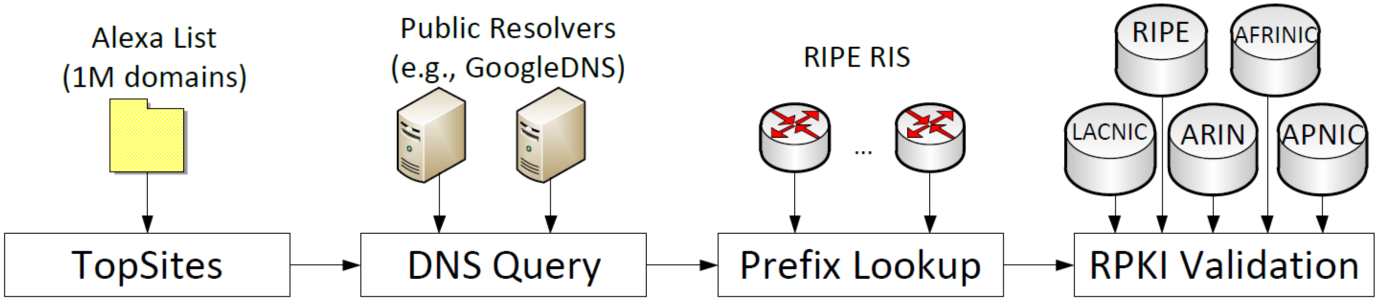

RPKI Validation Outcome for 1M Web Sites

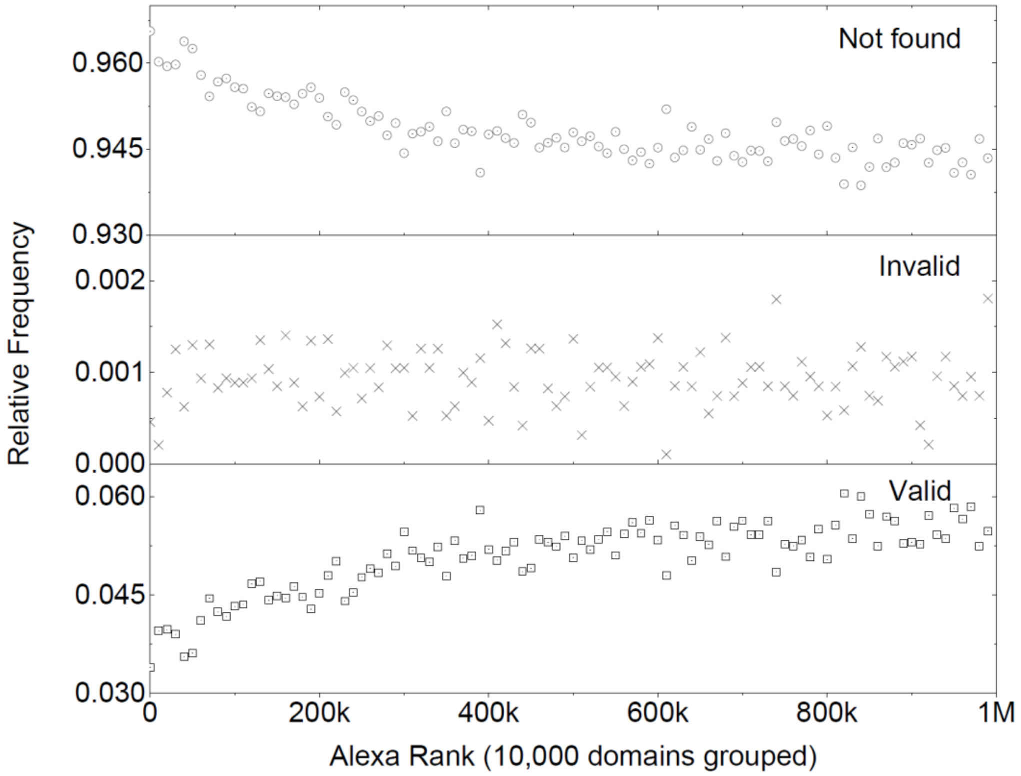

RPKI Validation Outcome for 1M Web Sites

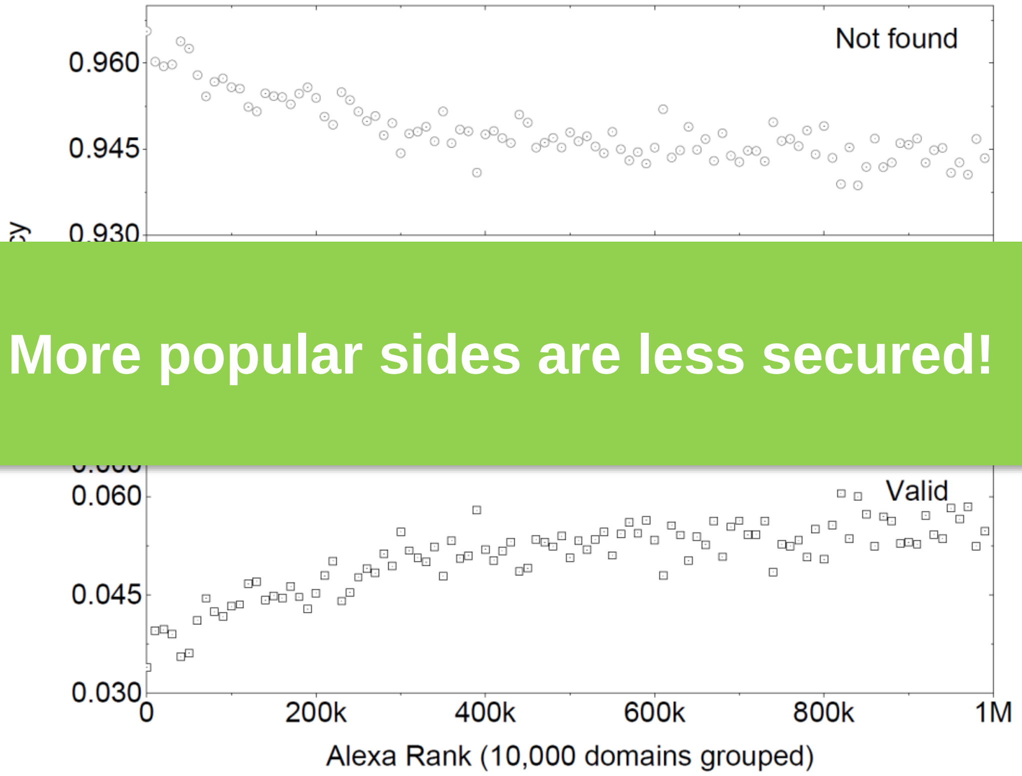

More popular sides are less secured!

Validation in Web-Browser

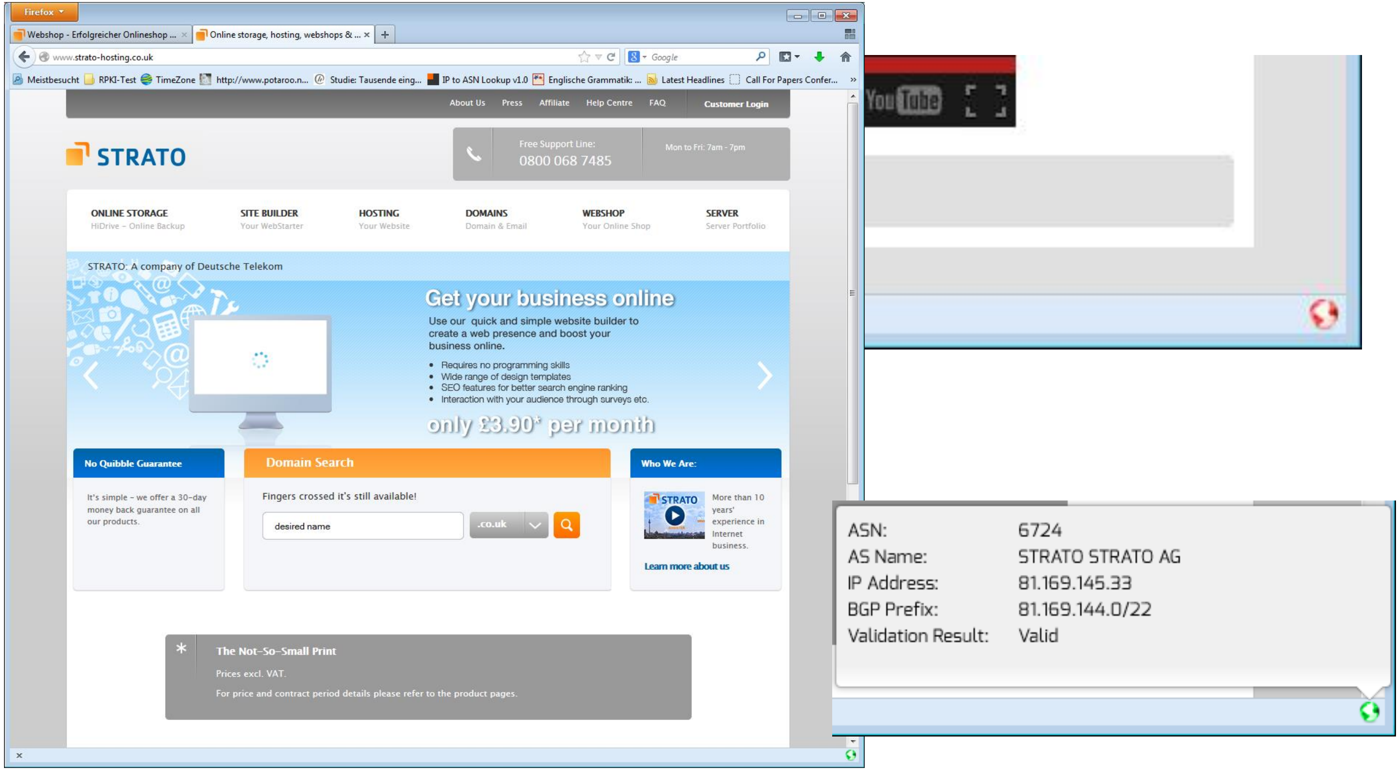

Study: ROA and ROV [SIGCOMM CCR ´18]

Route Origin Authorization (ROA)	Prefix owner authorizes AS to originate a set of prefixes
Route Origin Validation (ROV)	BGP router validates received routes using ROA information

Motivation & Research Problem

Goal: Which ASes use ROV-based filtering policies?

Assess impact of defense mechanisms

Track deployment over time

Create an incentive to deploy

Challenge: Private router configurations must be inferred

Controlled Experiments: Setup

Hand-crafted ROAs *and* BGP Updates

Goal: Find ASes that filter invalid routes

BGP	RPKI
Announce prefixes P_A (Anchor) and P_E (Experiment)	Issue ROAs for both prefixes
✓ Same RIR DB route object ✓ Same prefix length ✓ Announced at the same time ✓ Announced to same peers ✓ Announced from same origin AS	P_A announcement is always *valid*. Periodically change ROA for P_E : ➢ Flips announcement from *valid* to *invalid* to *valid* daily.

Controlled Experiments

Initial Situation: Origin AS and vantage point AS peer directly

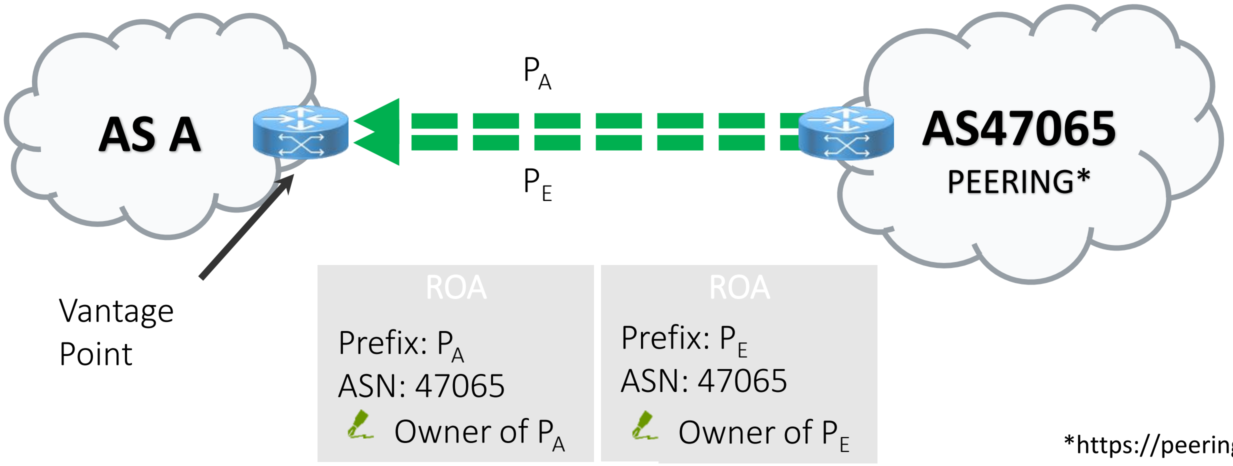

*https://peering.usc.edu/

Controlled Experiments

Initial Situation: Origin AS a[…]ntage[…]AS peer[…]irectly

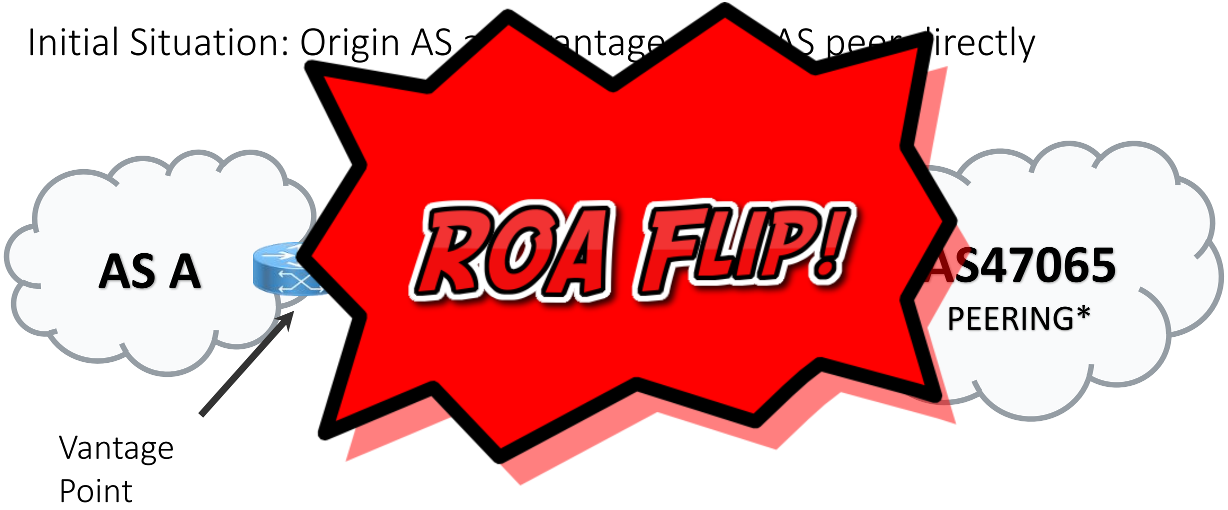

*https://peering.usc.edu/

Controlled Experiments

Observation: Vantage point exports no route for P_E

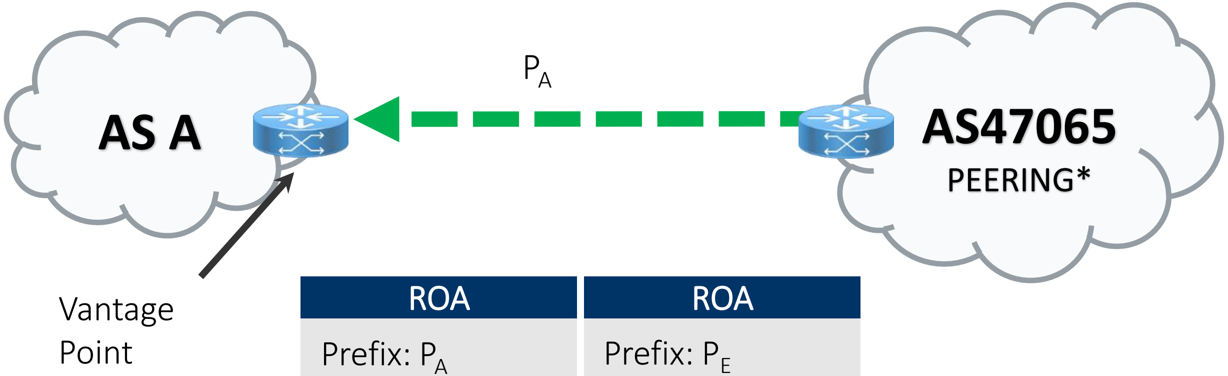

ROA	ROA
Prefix: P_A ASN: 47065 Owner of P_A	Prefix: P_E ASN: 51224 Owner of P_E

*https://peering.usc.edu/

Controlled Experiments

Observation 1: Vantage point exports no route for P_E

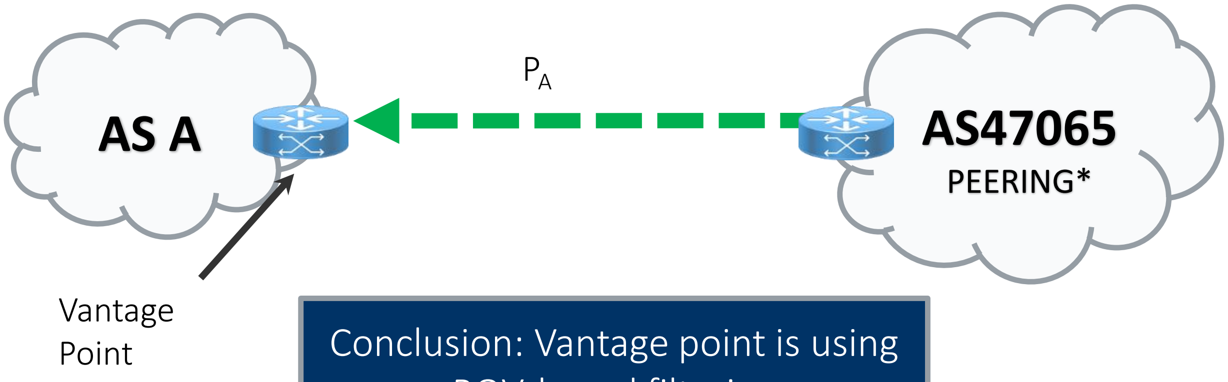

Conclusion: Vantage point is using ROV-based filtering

*https://peering.usc.edu/

Controlled Experiments Results

Before October 20^{th} 2017:

- (At least) Three ASes drop invalid routes

October 20^{th} 2017:

- AMS-IX Route Server changes ROV based filtering to ‘opt-out’
- 50+ ASes “drop” invalid routes

Full talk on [Youtube]

Literature

Andreas Reuter, Randy Bush, Italo Cunha, Ethan Katz-Bassett, Thomas C. Schmidt & Matthias Wählisch (2018).

Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *ACM SIGCOMM Computer Communication Review, 48*, 19-27.

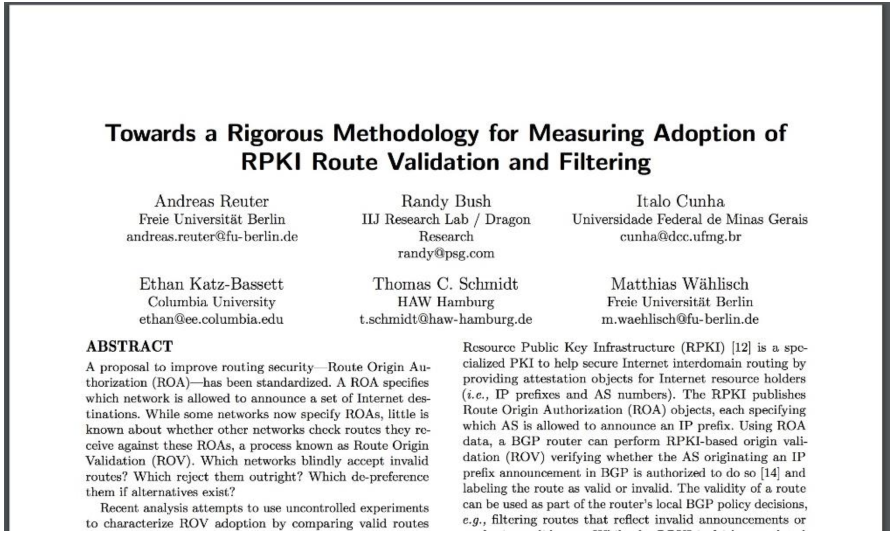

Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering

Andreas Reuter
Freie Universität Berlin
andreas.reuter@fu-berlin.de

Randy Bush
IIJ Research Lab / Dragon Research
randy@psg.com

Italo Cunha
Universidade Federal de Minas Gerais
cunha@dcc.ufmg.br

Ethan Katz-Bassett
Columbia University
ethan@ee.columbia.edu

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
m.waehlisch@fu-berlin.de

ABSTRACT

A proposal to improve routing security—Route Origin Authorization (ROA)—has been standardized. A ROA specifies which network is allowed to announce a set of Internet destinations. While some networks now specify ROAs, little is known about whether other networks check routes they receive against these ROAs, a process known as Route Origin Validation (ROV). Which networks blindly accept invalid routes? Which reject them outright? Which de-preference them if alternatives exist?

Recent analysis attempts to use uncontrolled experiments to characterize ROV adoption by comparing valid routes

Resource Public Key Infrastructure (RPKI) [12] is a specialized PKI to help secure Internet interdomain routing by providing attestation objects for Internet resource holders (*i.e.*, IP prefixes and AS numbers). The RPKI publishes Route Origin Authorization (ROA) objects, each specifying which AS is allowed to announce an IP prefix. Using ROA data, a BGP router can perform RPKI-based origin validation (ROV) verifying whether the AS originating an IP prefix announcement in BGP is authorized to do so [14] and labeling the route as valid or invalid. The validity of a route can be used as part of the router's local BGP policy decisions, *e.g.*, filtering routes that reflect invalid announcements or

ROV Deployment Monitor: rov.rpki.net

Literature

M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, G. Tyson (2015).

RIPPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem. *14th ACM Workshop on Hot Topics in Networks (HotNets)*.

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

Matthias Wählisch
Freie Universität Berlin
m.waehlisch@fu-berlin.de

Robert Schmidt
Freie Universität Berlin
rs.schmidt@fu-berlin.de

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Olaf Maennel
Tallinn U. of Technology
olaf.maennel@ttu.ee

Steve Uhlig
Queen Mary Univ. London
steve@eecs.qmul.ac.uk

Gareth Tyson
Queen Mary Univ. London
g.tyson@qmul.ac.uk

ABSTRACT

Web content delivery is one of the most important services on the Internet. Access to websites is typically secured via TLS. However, this security model does not account for prefix hijacking on the network layer, which may lead to traffic blackholing or transparent interception. Thus, to achieve comprehensive security and service availability, additional protective mechanisms are necessary such as the RPKI, a recently deployed Resource Public Key Infrastructure to prevent hijacking of traffic by networks. This paper argues two positions. First, that modern web hosting practices make route protection challenging due to the propensity to spread

Keywords

BGP, RPKI, secure inter-domain routing, deployment, hosting infrastructure, CDN

1. INTRODUCTION

Website security is a long pursued and rather esoteric goal. Traditionally, it has been approached from an end-to-end perspective (*e.g.* TLS), largely because this is easily within the sphere of control of any web provider. However, as evidenced by many prominent attacks, this is frequently insufficient. This is because various third party infrastructure dependencies