A Study of BGP Route Origin Registration and Validation

Measurements of RPKI and RouteViews

Daniele Iamartino
danieleiamartino@gmail.com
Cristel Pelsser
cristel@iij.ad.jp
Randy Bush
randy@psg.com

RIPE 69
(November 2014)

Recap

- RPKI deployed in 2012 in order to secure the Internet routing.
- **Route origin validation**: check if the *origin AS* of a BGP announcement is correct, using RPKI
  - Not completely *crypto-checked*, so can be violated, but should prevent vast majority of **accidental 'hijackings'** on the Internet today

Route Origin Validation

- ISP get a certificate signed by the CA of the RIR
- ISP sign a **ROA** (*Route Origin Authorization*) file and put on the RIR's RPKI repo
- Example ROA: (Prefix 10.0.0.0/16, AS42)
  - *Autonomous system* number 42 is authorized to announce prefix 10.0.0.0/16
  - When we receive a BGP announcement for 10.0.0.0/16, we check if the last AS on the AS_PATH is AS42.

Route Origin Validation: Maximum length

- If the ROA cover prefix 10.0.0.0/16, only that prefix can be announced.
- If we announce a longer prefix (ex: 10.0.1.0/24), even from the correct AS, the announcement will be invalid.

- Two ways to solve:
  - Create another ROA: 10.0.1.0/24, AS42
  - Set a ***maximum length*** in the ROA (ex: 10.0.0.0/16, maxlen: 24, AS42)
    - = "AS42 can announce prefix 10.0.0.0/16 or longer prefixes up to /24"
    - So 10.0.1.0/24 can be announced

Introduction

Questions:

- What is the **deployment** of RPKI?
- Are today's **BGP routes** valid against RPKI-based route origin validation?
- What happen if we filter invalid announcements today?

Steps

1. Look at the **ROA** (*Route Origin Authorization*) file publication on RPKI repos of all RIRs
2. Take **RIB dumps** from a BGP monitor and validate all route announcements

RPKI adoption on ROAs

Publication point	v4 host addresses covered by a ROA	v4 host addresses allocated by the RIR	% coverage
RIPE NCC	125,133,312	797,906,680	15.68%
ARIN	30,187,520	1,733,372,928	1.74%
LACNIC	19,089,408	189,833,472	10.05%
AfriNIC	2,814,464	119,534,080	2.35%
APNIC	744,960	872,194,816	0.08%
Total	177,969,664	3,712,814,976	**4.79%**

- **RIPE NCC** is leader in ROA registration
- Although **ARIN** has allocated most of the address space, it lags far behind most other RIRs in registrations
- Global IPv4 ROA coverage is 4.79%

Accepted ROAs

- We validate files in RPKI repos using the **rcynic** tool
- We have history of RPKI repositories since 2012
- So we validated all the history and plotted valid ROA files

Accepted ROAs (logscale)

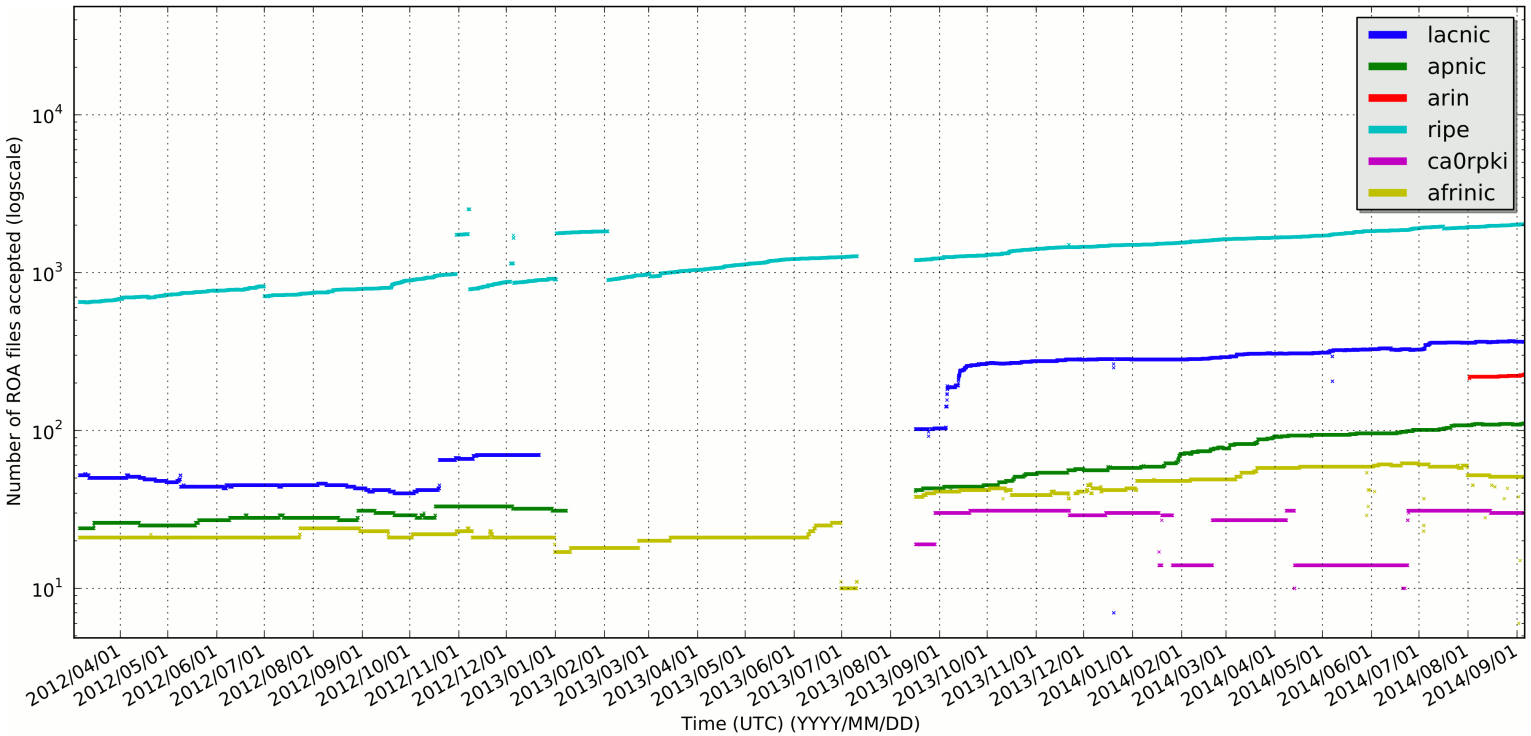

- **LACNIC** valid ROAs drops between Dec 2012 and Aug 2013
  - We believe this was expiration of their trust anchor.
- Aug 2013: Problem in our data collection
- **ARIN** data starts from Aug 2014 due to ARIN's legal barriers on data collection

BGP announcement origin-validation

- We want to validate real BGP announcements
- We have BGP announcement history for the same period as the RPKI repositories data
- How to validate?
  - One BGP **RIB** dump every 30 days since 2012
  - Search the **rcynic** dump just before that time, load all valid ROAs
  - For each announcement of the RIB, check if there is a valid covering ROA

BGP announcement origin-validation

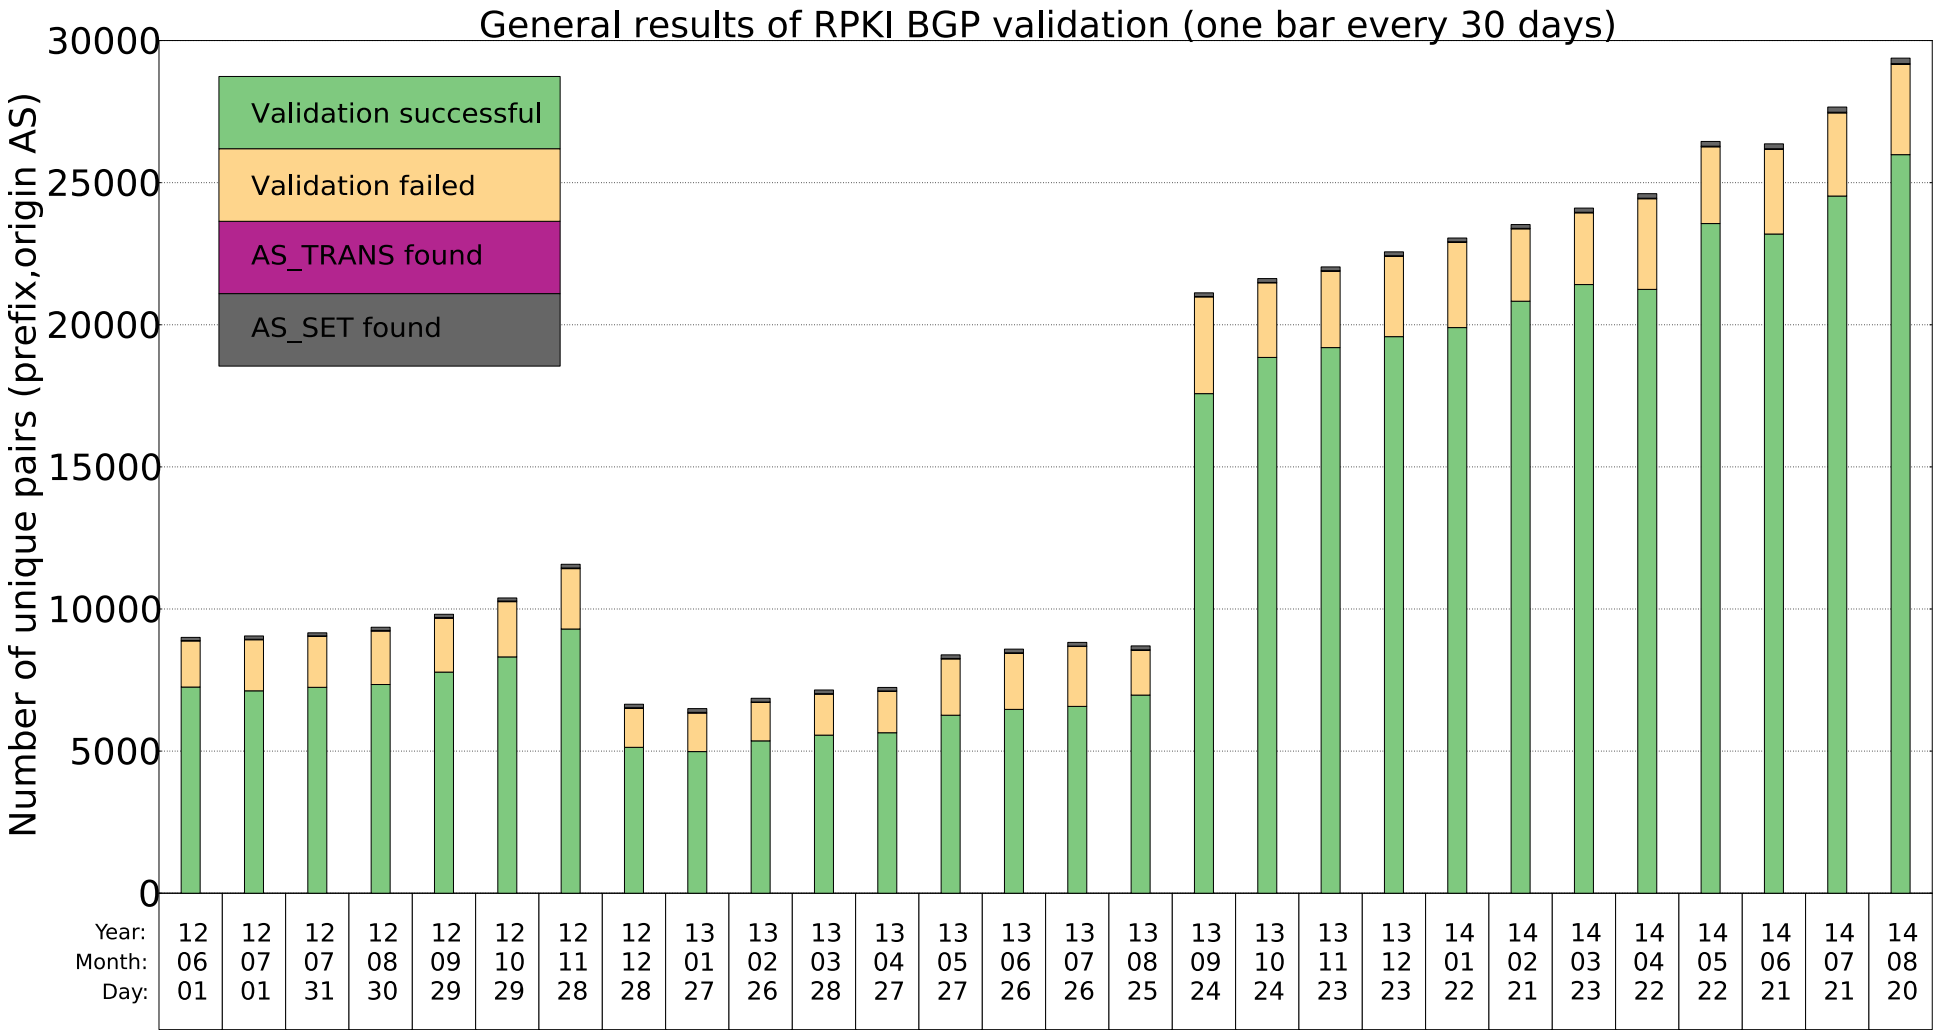

BGP announcement origin-validation

- We are not plotting "ROA not found" announcements (majority of them)
- Huge drop in the middle? LACNIC fault, as we saw before
- ~10% announcements are invalid
- It's more meaningful to look at validation of **prefixes**

Valid/Invalid prefixes

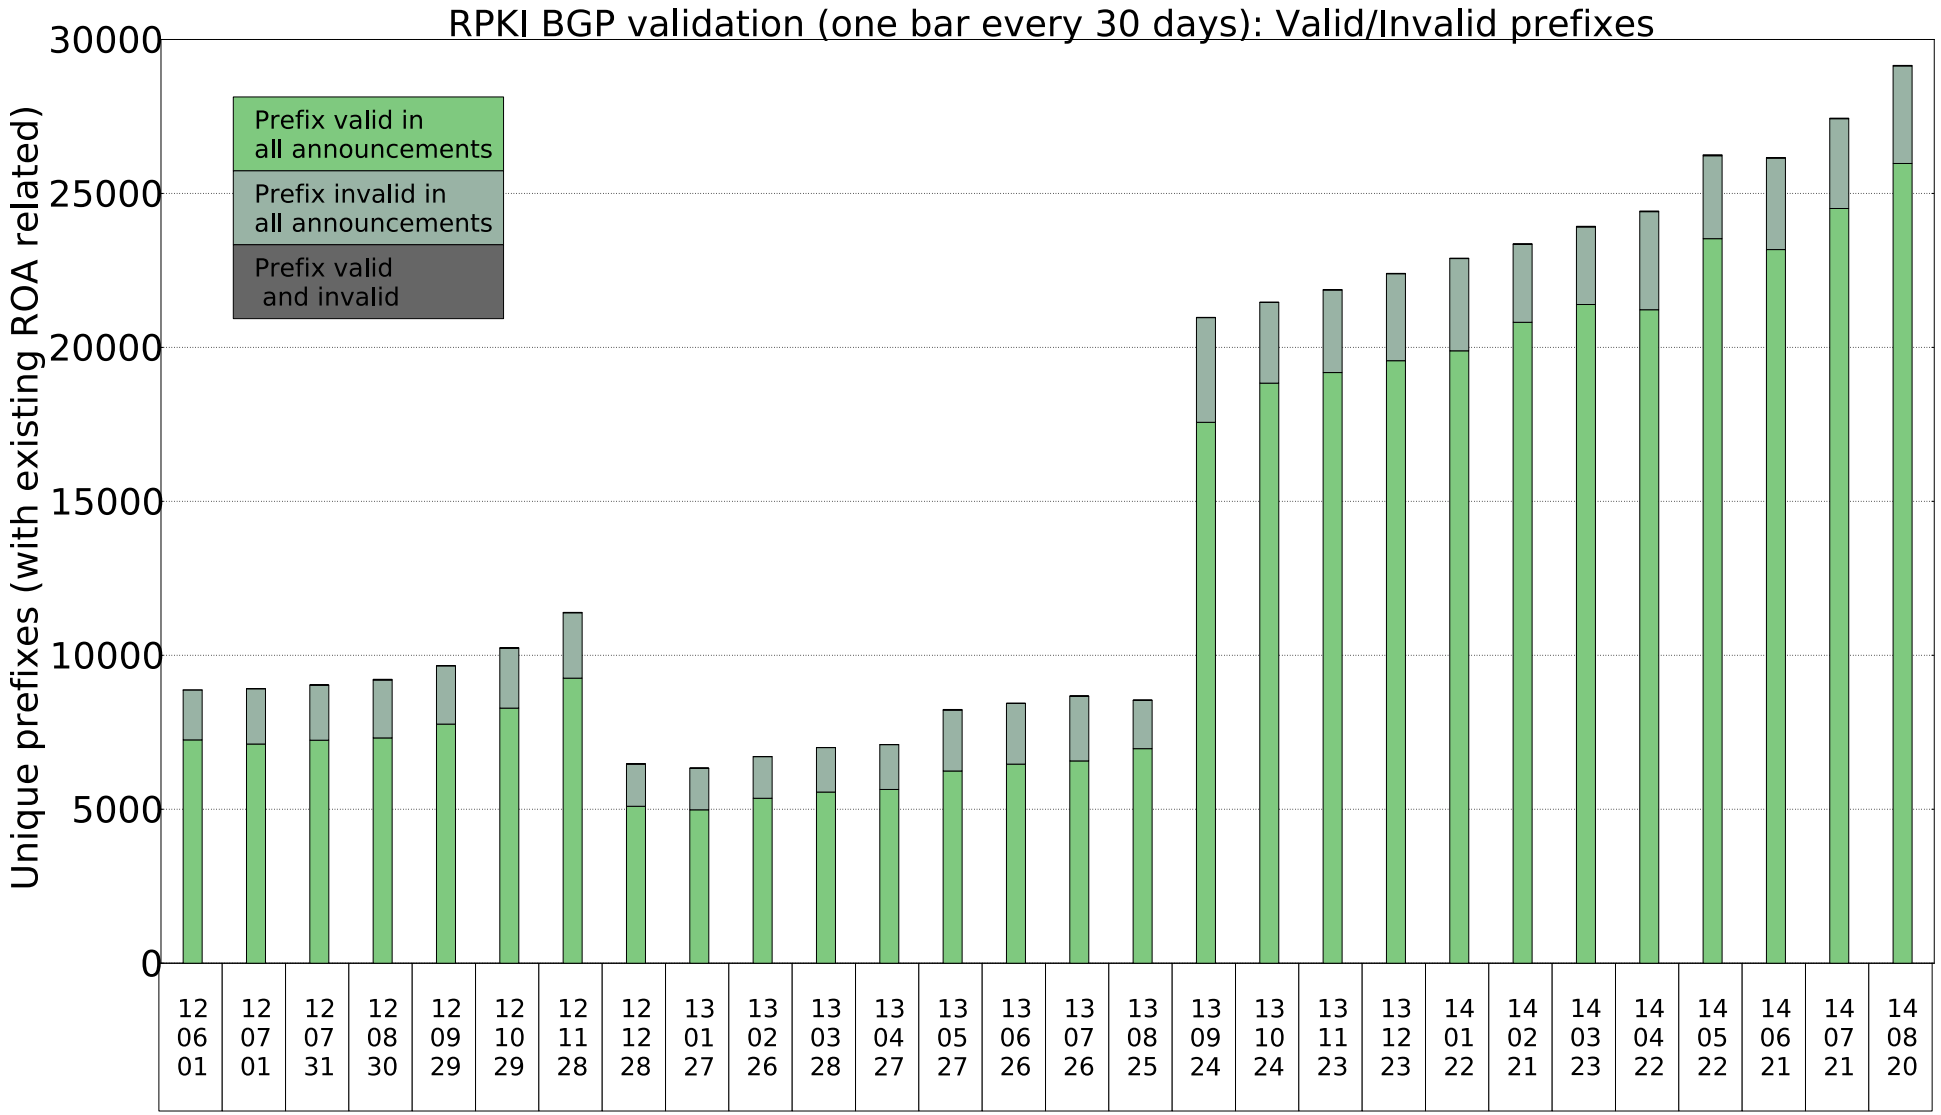

Valid/Invalid prefixes

- ~5% of global prefixes are RPKI-covered
- Even looking at prefixes only, we see 10% invalid prefixes
- Why invalid prefixes?
- Let's beak down reason of invalidity

Reason for invalidity of prefixes

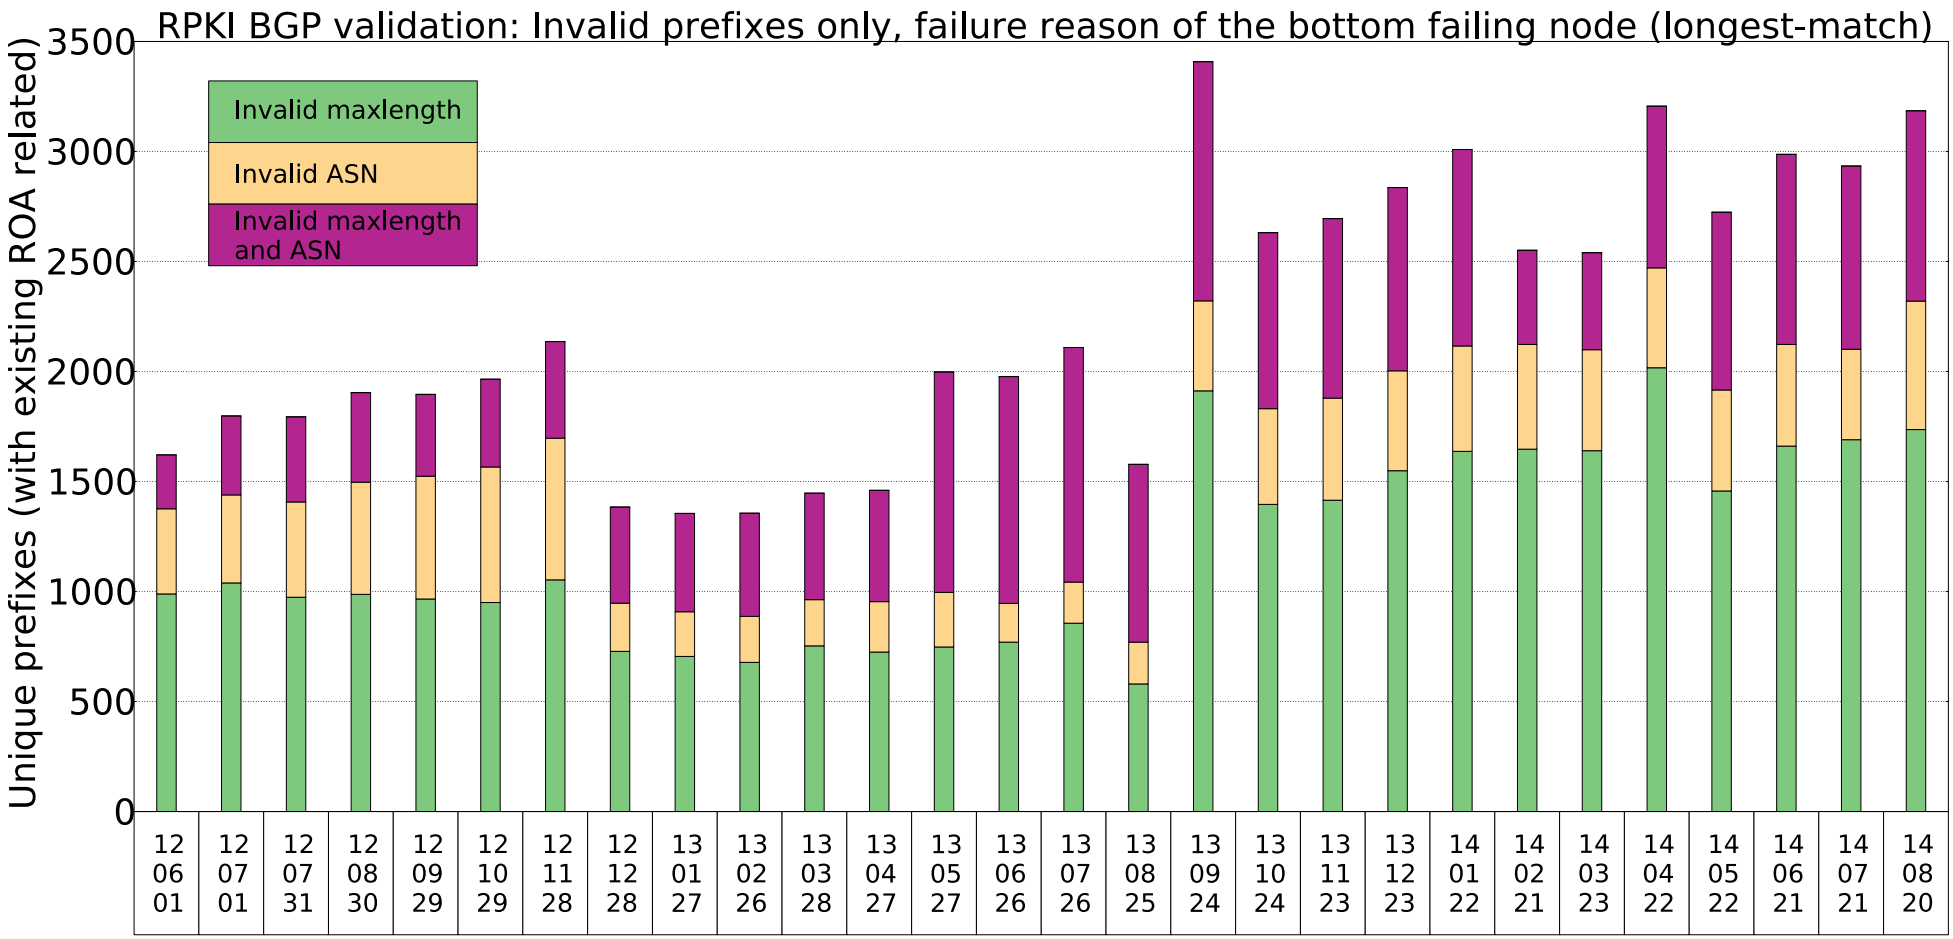

Reason for invalidity of prefixes

- Most of the problems: ***maxlength*** error
  - The origin AS is correct, ROA exists, but the announced prefix is longer
  - People registering ROA should be careful!
- What about **coverage**?:
  - Let's say we **drop** invalid prefixes that we receive. Do we lose connectivity?
  - An invalid prefix could be covered by another valid or "ROA not found" prefix
  - For example: announcement of 10.0.2.0/24 is invalid, but also 10.0.0.0/16 is announced and valid. The invalid prefix is covered by a valid.

Taking coverage into account

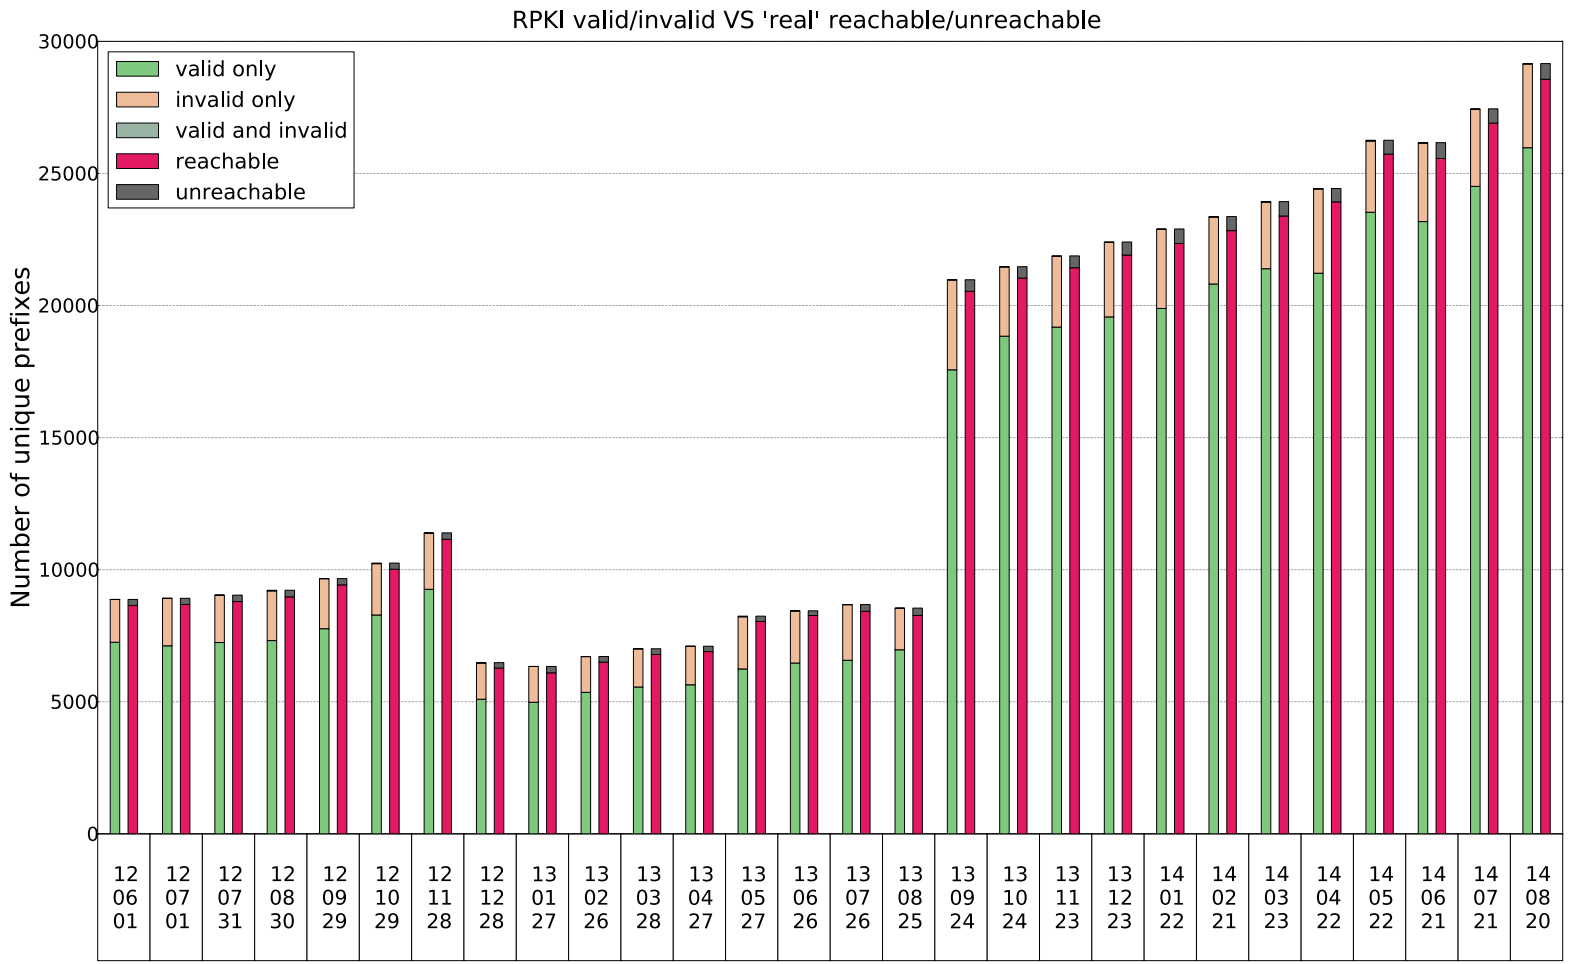

Around 80% of **invalid** prefixes are in fact **reachable**. They are "rescued" by another valid or a "ROA not found" covering prefix.

What is the most common error?

- When we see an announcement coming from the wrong origin AS, in **72%** of the cases we can find the correct AS in one of the AS paths of that prefix.
- Reason of this:
  - **ISP** with AS42 register a ROA for its **10.0.0.0/16**,AS42
  - **AS666, customer of ISP** do not register any ROA and announce **10.0.2.0/24**, AS666
  - We receive an announcement: 10.0.2.0/24 with AS_PATH: 100 200 **42** 666
  - The announcement of the customer is invalid because of wrong origin AS and maxlen, but the **correct AS** (of the ISP) **is on the AS_PATH**

What is the most common error?

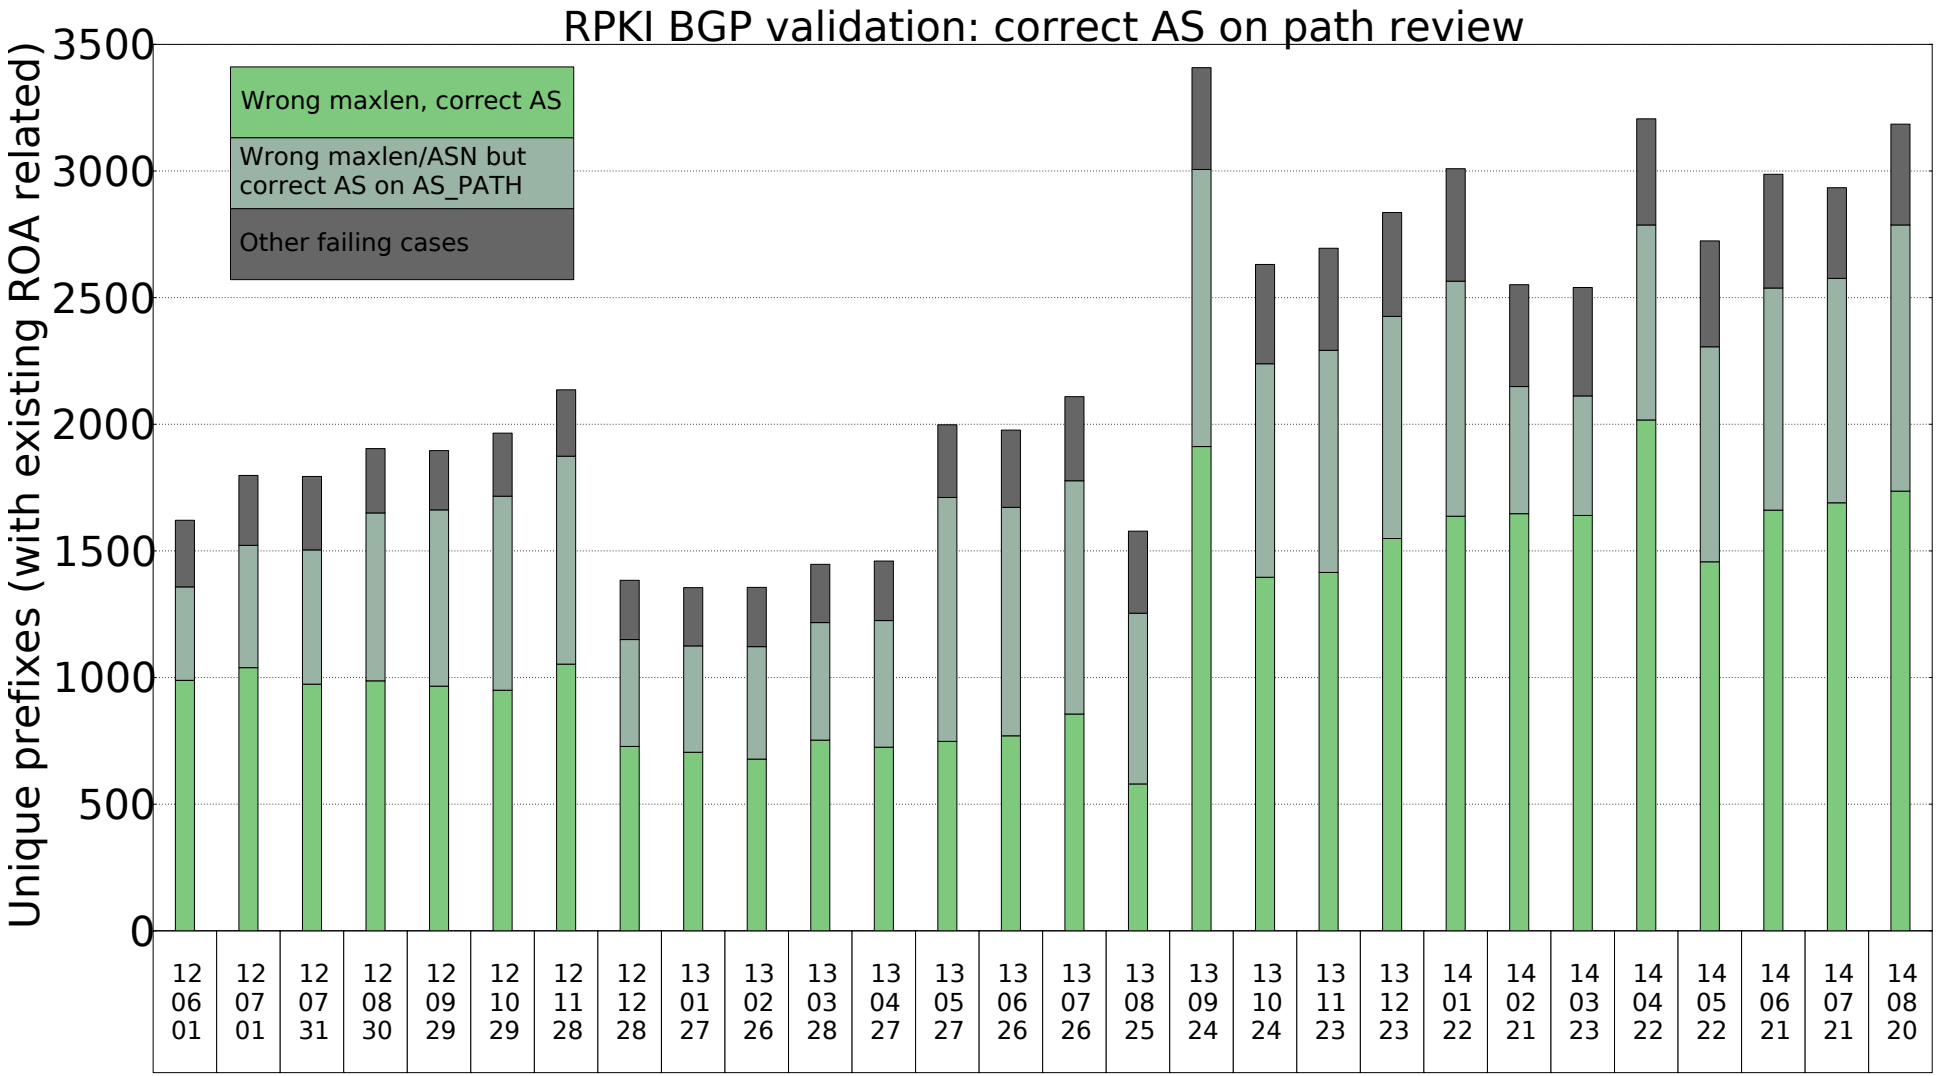

Measure on real traffic

- RPKI deployment is **about 5%**
- Is this 5% of prefixes where most of the Internet traffic is going?
- We measured the percentage of RPKI-covered traffic going through a big American research network for few days

Measure on a big research network

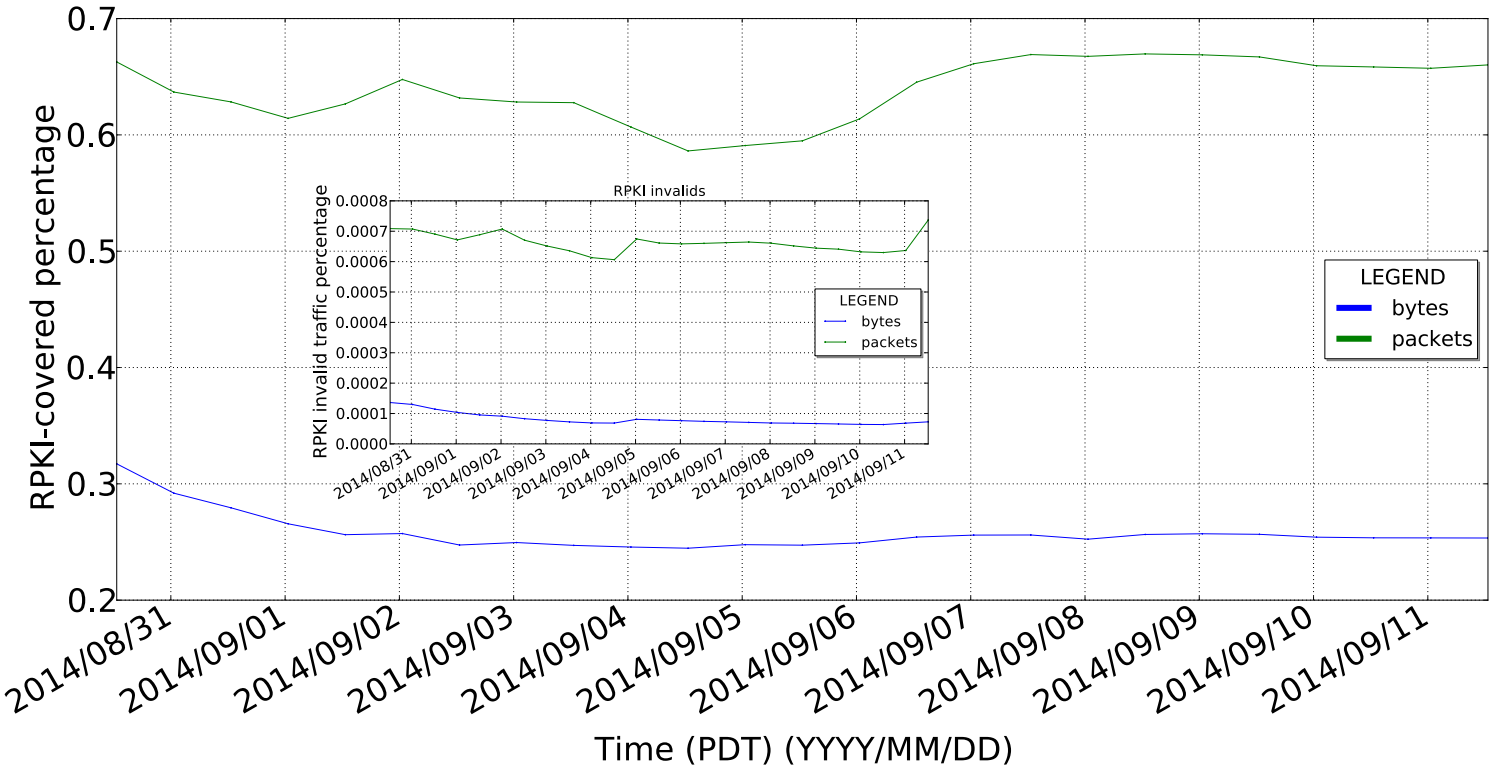

Only 0.3% of the bytes going though this network is RPKI-covered. So the 5% deployment is not an important part of the address space to this ISP

Help the Internet!

- Prefixes covered by RPKI are about 5%
  - RPKI deployment is good but still too **slow**.
- Help the Internet routing security is **easy**:
  - **Register your ROA** files on the RIR, and be sure to announce the same on BGP.
  - Start to deploy validation and filtering later

Help the Internet!

- The **top-ISP's ROA coverage problem** is very common, let's fix it!
  - Go to your customers announcing on BGP, tell them to register a ROA! (or register one for them)
- Lot of people misunderstood how to use "**maxlength**" in a ROA
  - Check that your announcements match what you registered!

FIN

Questions?