

# Routing Security Update Q2 2019

Job Snijders

NTT Communications / AS 2914

job@ntt.net

# What is it we are doing here?

- Facilitation of communication?
- Making money?
- Sharing a hallucination?

# Agenda

- Challenges
- Business impact of Origin Validation
- Software available
- Cleanup efforts
- Resources
- Deployment update

# The LARGEST challenge we have with RPKI OV

A few hundred people misconfigured their RPKI ROAs

This results in what in BGP are known as “RPKI Invalids”

But in business operations these are called “*false positives*”

Patrick Gilmore reminded me: “*Damn computer never does what I want it to do, it only does what I tell it to do*”

## **Reminder: if you misconfigure things, it hurts**

We as an industry have to consider that being tolerant to mistakes other people make, may jeopardize our own operation.

John Postel was right and wrong, he said “Be liberal in what you accept, and conservative in what you send”

It is now 2019 ... we have to rethink the harmful consequences

<https://tools.ietf.org/html/draft-iab-protocol-maintenance>

# The path towards Origin Validation deployment

It is quite simple.

**DEPLOY. NOW.**

RPKI based BGP Origin Validation,

With “Invalid == reject” routing policies

# **Every RPKI OV deployment contributed to less false positives, we have data to show**

**Because AT&T deployed, many networks fixed their ROAs**

**Because Cloudflare deployed, many networks fixed their ROAs**

**Because YYCIX, DE-CIX, others deployed; many networks fixed**

**Many organisations don't listen to nice requests, sometimes you need to introduce some discomfort before they are motivated to take action.**

## Study resources

<https://nusenu.github.io/RPKI-Observatory/unreachable-networks.html>

[https://labs.ripe.net/Members/nusenu\\_nusenu/the-rpki-observatory](https://labs.ripe.net/Members/nusenu_nusenu/the-rpki-observatory)


<https://medium.com/@nusenu/towards-cleaning-up-rpki-invalids-d69b03ab8a8c>

The industry has reduced the “false positives” by 50% in the last 6 months – KEEP PUSHING!!!!



# RPKI based traffic analysis with pmacct

pmacct: one slightly more complex use-case



## **pmacct's RPKI capabilities**

- RFC 6811 Origin Validation procedure is applied
- Mark traffic based on Validation Status, without deploying RPKI in your network
- This helps you understand the effects of rejecting “RPKI invalid” announcements
- Pmacct version 1.7.3


## Most importantly, pmacct recognises the 2 types

There are false positives which are:


- *Unrecoverable*, there is no alternative path
- *Implicitly repaired*, because there is a covering less-specific valid or unknown route.

There are from NTT's perspective no “*Unrecoverable*” important destinations, and honestly if we deploy OV, we are doing as they are asking us to do.

# A view from AS 2914 / NTT's global backbone



# Zooming in on a day – couldn't go smaller than a pixel




## Validator situation: very good

- NLNetlabs Routinator (rust, fast,)
- Cloudflare OctoRPKI / GoRTR (go, fast)
- OpenBSD rpki-client(1) (C, in private beta, most basic option)
- Dragon Research Labs RPKI Toolkit (Python + SQL)
- BBN's RPSTIR (C language)
- ~~RIPE NCC RPKI Validator version 3 (java, slowish, lots of features)~~

# OpenBSD's rpki-client(1)

- Started January 2019, almost done
- Runs as command line tool, not daemon
- Outputs all VRPs in OpenBGPD format
- Can be used to embed in carrier grade routers
- Can be used with GoRTR from Cloudflare
- Side effect: clean room implementation of rsync
  - (BSD license, “openrsync”)



# Using RPKI to clean up the IRR





# Applying Origin Validation to the IRR

- RPKI ROAs can be used for *BGP Origin Validation*
- But, what about applying the RFC 6811 “Origin Validation Procedure” to IRR data?
- Perhaps, we should consider unvalidated IRR data objects as if they are BGP announcements!

# An example

```
route:          129.250.15.0/24
origin:         AS60068
descr:         AS60068 route object
descr:         this is a test of hijack possibilities
                with current state of RIPE/RADB security
                setup - this records covers IP address used for
                rr.ntt.net service
descr:         please note this is just a demonstrative object,
                with no real harmful intention
mnt-by:         DATACAMP-MNT
created:        2018-02-10T16:57:07Z
last-modified: 2018-09-04T19:07:32Z
source:        RIPE-NONAUTH
```

# The previous slide is in conflict with this ROA!

```
$ whois -h whois.bgpmon.net 129.250.15.0/24
% This is the BGPmon.net whois Service
% You can use this whois gateway to retrieve information
% about an IP address or prefix
% We support both IPv4 and IPv6 address.
%
% For more information visit:
% https://portal.bgpmon.net/bgpmonapi.php
Prefix: 129.250.0.0/16
Prefix description: NTT Communications backbone
Country code: US
Origin AS: 2914
Origin AS Name: NTT-COMMUNICATIONS-2914 - NTT America, Inc., US
RPKI status: ROA validation successful
First seen: 2019-02-23
Last seen: 2019-05-22
Seen by #peers: 71
```

# One effort: RIPE-NONAUTH IRR cleanup

Formal proposal: Apply the *Origin Validation* procedure to IRR objects in the RIPE-NONAUTH IRR database.

Helps remove wrong LACNIC, APNIC, ARIN, AFRINIC route registrations from RIPE-NONAUTH

Changes:

- 7 day hold period
- Notifications should be send to the IRR route object holder (if we can).

<https://www.ripe.net/participate/policies/proposals/2018-06>

Test tool: <https://github.com/job/ripe-proposal-2018-06>

## Another effort: IRRd version 4!



<https://github.com/irrdnet/irrd4>

## Another effort: IRRd version 4!

- IRRd version 3 is an organically grown, 20 year old code base, mostly in C, perl, ineffective database backend
- Reliability issues with irrd2 and irrd3 (have to restart often)
- Absolutely critical to NTT's daily operations, all NTT's prefix-filters are generated with this software

Funded by NTT Communications, developed by [Dashcare](#)

# Quick overview of the size of the old codebase

```
job@vurt irrd$ cloc .  
    189 text files.  
    185 unique files.  
    28 files ignored.
```

```
github.com/AlDanial/cloc v 1.74  T=2.25 s (71.9 files/s, 36938.2 lines/s)
```

Language	files	blank	comment	code
C	92	6645	9205	33967
Perl	10	812	877	12451
Bourne Shell	4	993	1308	9687
C/C++ Header	35	722	549	3608
yacc	1	326	111	1453
make	20	168	63	313
SUM:	162	9666	12113	<b>61479</b>

# IRRD version 4

Just ~ 10,000 lines of python



## Benefits of IRRd version 4

- Single modern architecture with extension options
- Code base is well documented, consistent, maintainable
- Extensive regression & integration testing
- QA checks compared to `rr.ntt.net` to ensure smooth transition
- BSD 2-Clause License

The next version of IRRd will do Origin Validation on IRR objects.

# Friends wrote a book, have a look

← → ↻ 🔒 <https://www.amazon.com/Day-One-Deploying-Routing-Security-ebook/dp/B07NDNBWB8/re...> ☆

**amazon** prime

All ▾ day one juniper security 🔍

📍 Deliver to Job **Amsterdam 1065 SZ** Departments ▾ Browsing History ▾ Job's Amazon.com Today's Deals EN 🌐


Buy a Kindle Kindle eBooks Kindle Unlimited Prime Reading Best Sellers & More Kindle Book Deals Free Reading Apps

Kindle Store > Kindle eBooks > Computers & Technology

**Look inside** ↴

**juniper** | Engineering Simplicity

DAY ONE: DEPLOYING BGP ROUTING SECURITY



Secure, field-tested, device and protocol configurations for running Junos® OS routers in the BGP default-free zone.

By Melchior Aelmans & Niels Rajjer

## Day One: Deploying BGP Routing Security Kindle Edition

by [Melchior Aelmans](#) (Author), [Niels Rajjer](#) (Author)

[Be the first to review this item](#)

> [See all formats and editions](#)

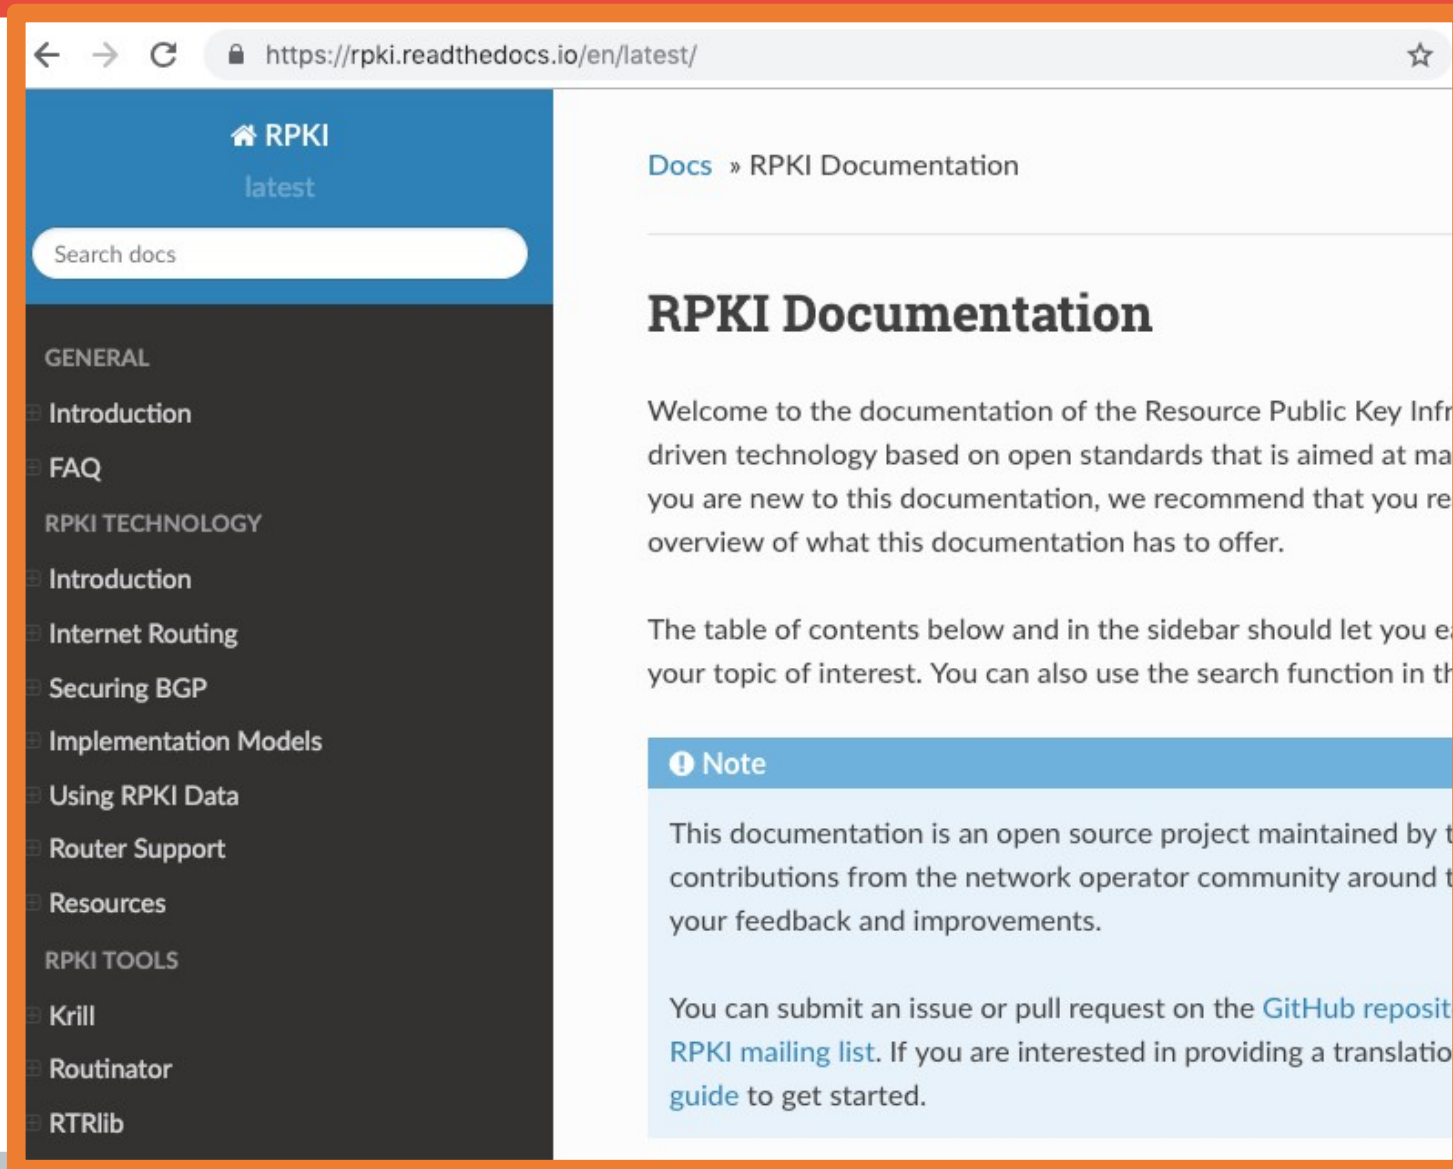
**Kindle**  
**\$1.12**

[Read with Our Free App](#)

Secure, field-tested, device and protocol configurations for running Junos® OS routers in the BGP default-free zone.


This book is intended for network administrators running Junos OS routers in the BGP default-free zone. It provides field-tested device and protocol configuration

# NLNetlabs made a website: [rpki.readthedocs.io](https://rpki.readthedocs.io)



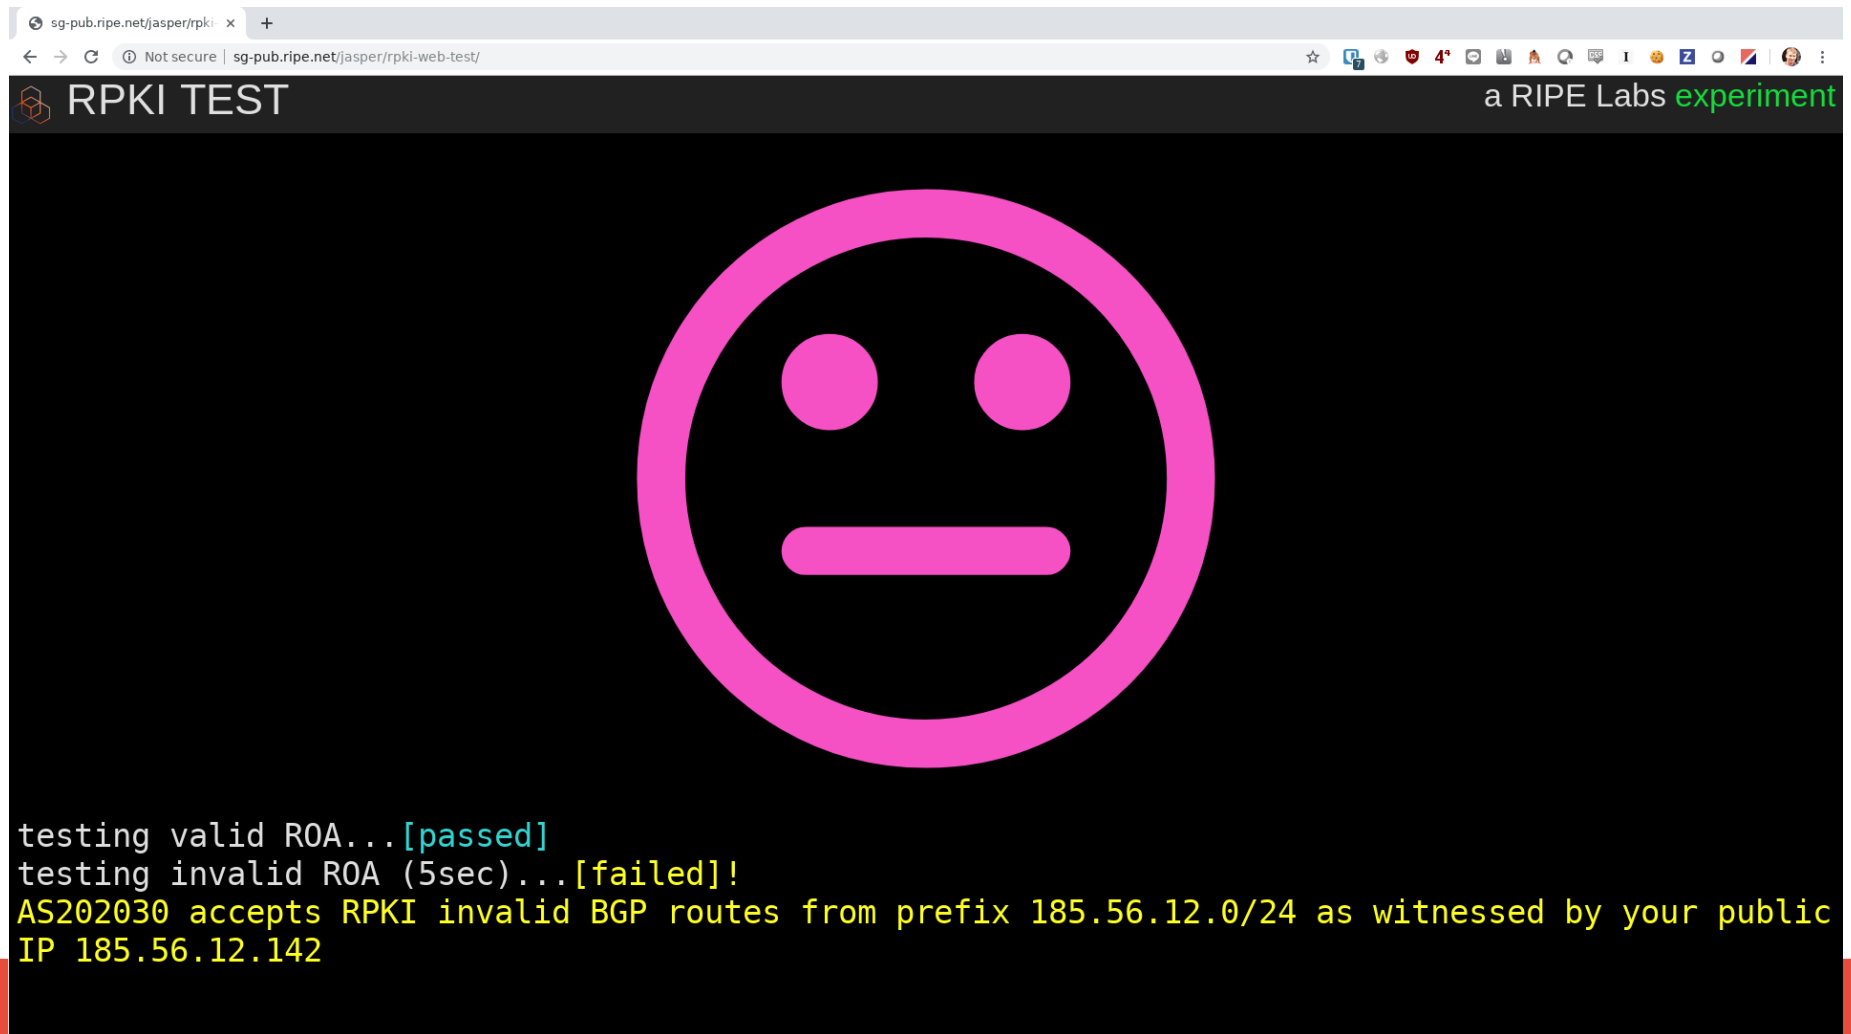
# RIPE Labs RPKI checker tool

<https://www.ripe.net/s/rpki-test>




# RIPE Labs RPKI checker tool

<https://www.ripe.net/s/rpki-test>



# Deployment update

- Cloudflare
- YYCIX


**YYCIX**  
Coreix Internet Exchange  
@yycix


Following

Good news! On Oct 7, the [#YYCIX](#) route servers started filtering prefixes which are RPKI ROA invalid. We are among the leaders in performing this validation -- probably the first [#IXP](#) in North America!

8:58 PM - 7 Oct 2018


11 Retweets 18 Likes



**Jerome Fleury**  
@Jerome\_UZ

Following

As of today, 75% of the [@cloudflare](#) PoPs (116/155) have RPKI strict validation enabled on all peering sessions. That's about 17,000 RPKI enabled peerings. Great work from [@lpoinsig](#)!



RPKI and BGP: our path to securing Internet Routing

# RPKI Deployment

- AT&T rejects invalids on peering sessions
- Nordunet rejects invalids on all EBGP sessions
- KPN / AS 286 rejects invalids on customer sessions
- Seacomm & Workonline drop invalids per April 2019
- INEX, AMS-IX, DE-CIX, France-IX, Netnod
- MSK-IX
- XS4ALL
- **THE RIPE MEETING NETWORK!!!**
- IX.br (.... soon? :-)
- **You.... ?**



AFRAID OF  
CHANGE?  
LEAVE IT  
HERE!

R PIZZA?  
GET TO SHARE ON  
ADVISOR!





# Question everything!

Feel free to ask questions, ask for clarifications

If you don't want to use the microphone, please email me

[job@ntt.net](mailto:job@ntt.net)

(I am happy to help competitors too)

*Network Engineers Without Borders!*