NTT

RPKI 101: The use of RPKI to improve Internet routing

Job Snijders
IP Development Engineer

job@ntt.net

Agenda

- Overview of the global Internet routing system
- What challenges exist in the system?
- What is RPKI-based BGP Prefix Origin Validation?
- Collaboration with industry partners
- Study Resources
- Q & A

What is the Internet?

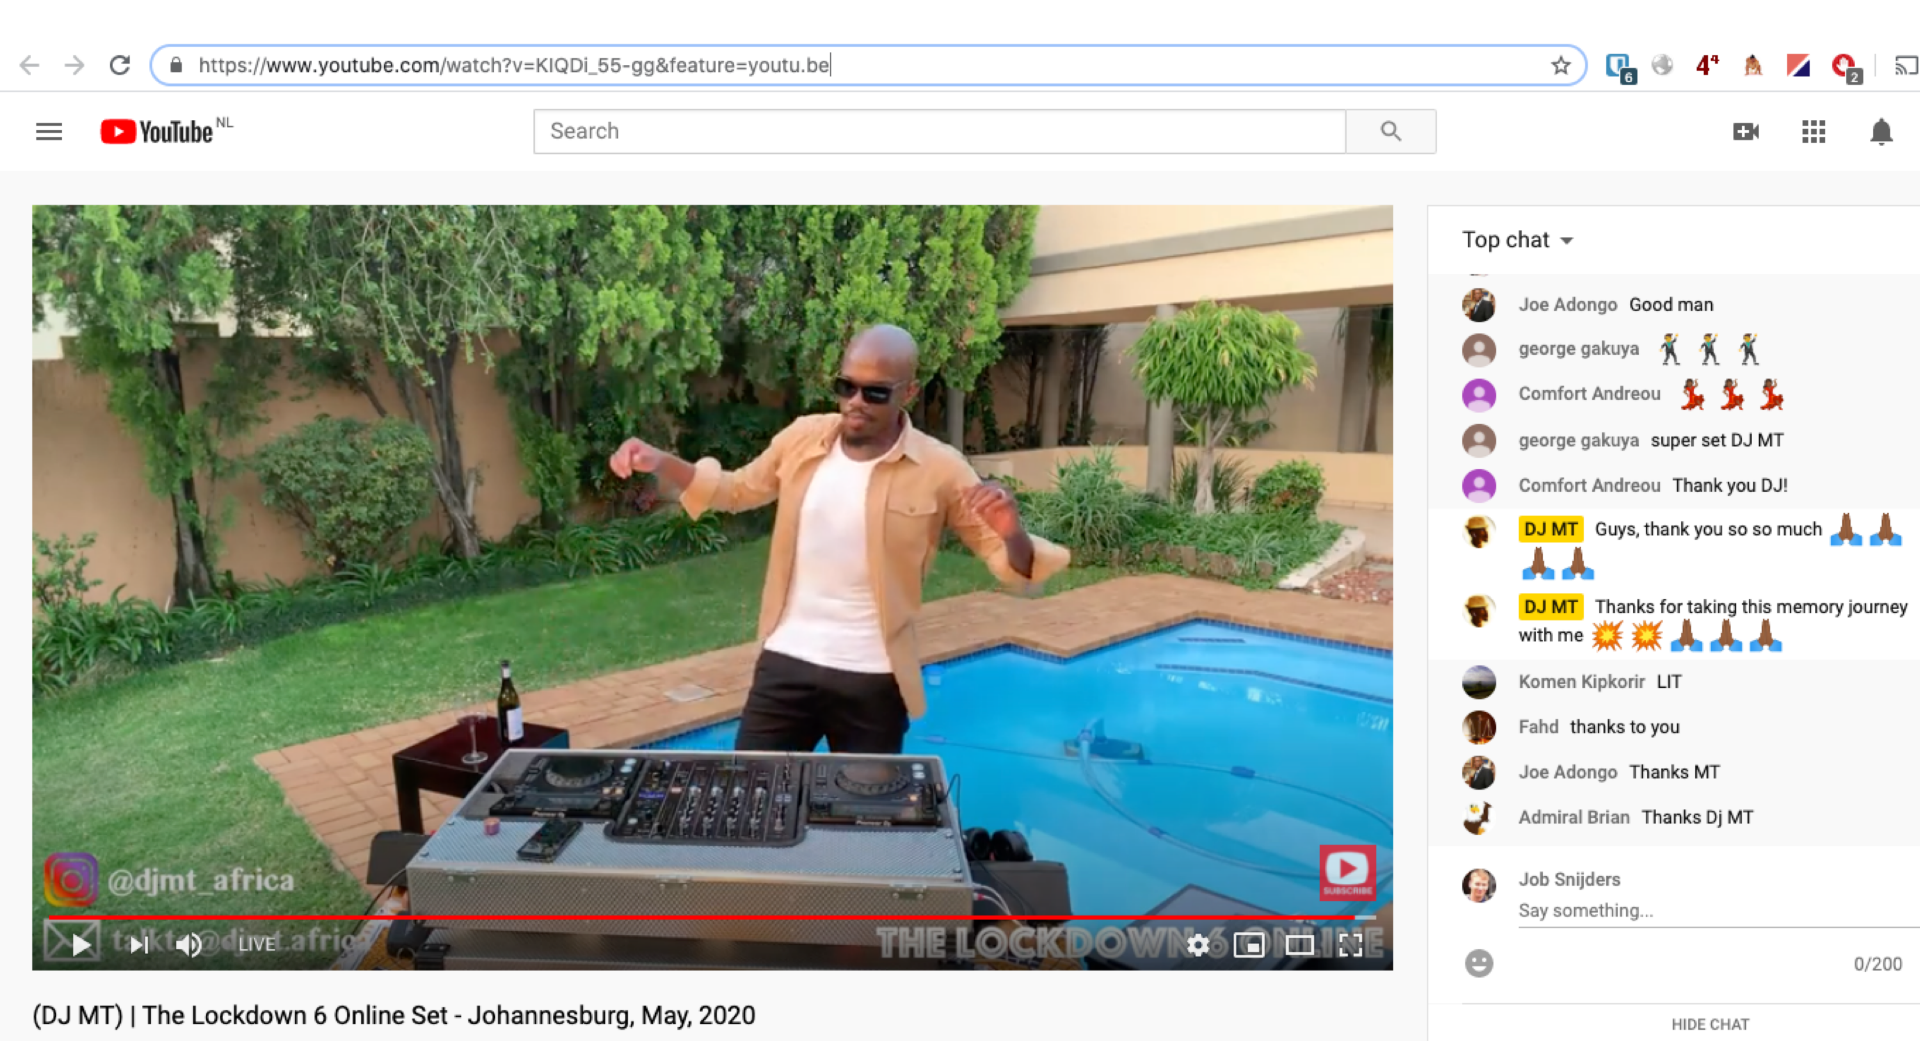

Joe Adongo Good man

george gakuya 🕺 🕺 🕺

Comfort Andreou 💃 💃 💃

george gakuya super set DJ MT

Comfort Andreou Thank you DJ!

DJ MT Guys, thank you so so much 🙏 🙏 🙏 🙏

DJ MT Thanks for taking this memory journey with me 💥 💥 🙏 🙏 🙏

Komen Kipkorir LIT

Fahd thanks to you

Joe Adongo Thanks MT

Admiral Brian Thanks Dj MT

Job Snijders

(DJ MT) | The Lockdown 6 Online Set - Johannesburg, May, 2020

The BGP protocol connects our networks

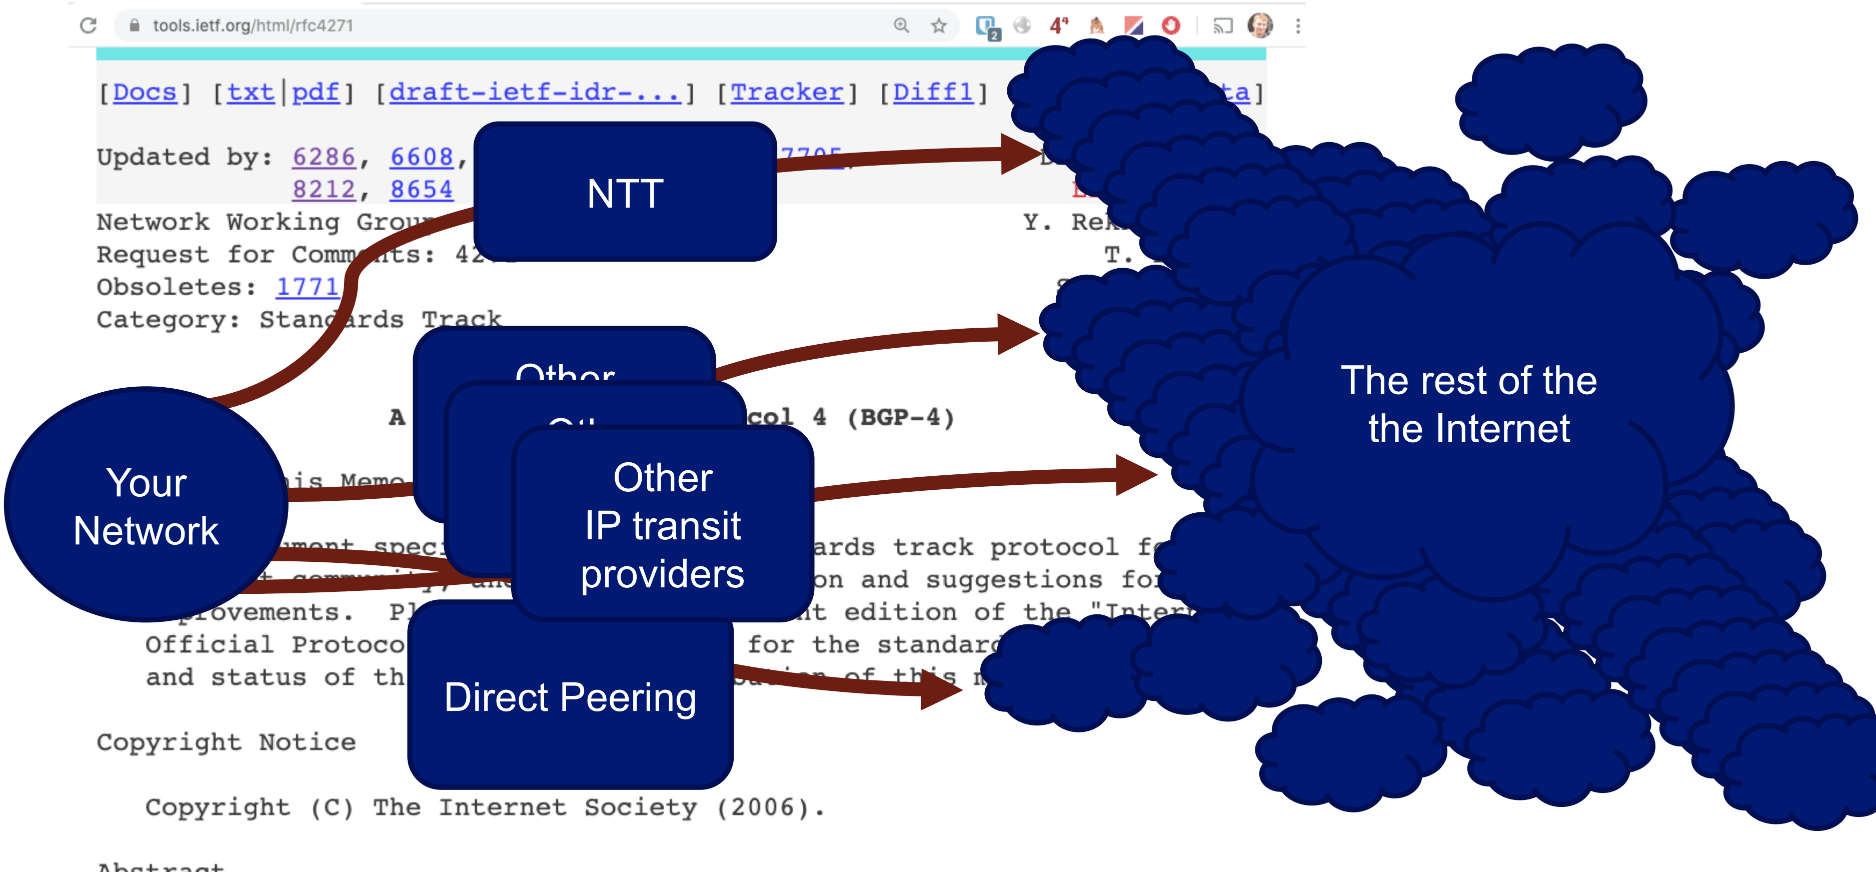

```
[Docs] [txt|pdf] [draft-ietf-idr-...] [Tracker] [Diff1]                ta]

Updated by: 6286, 6608,                7705,
            8212, 8654
Network Working Grou                                              Y. Re
Request for Comm     nts: 42                                        T.
Obsoletes: 1771
Category: Standards Track


                  A                   col 4 (BGP-4)

           is Memo

      ment speci                   ards track protocol f
      t communit   , an             on and suggestions fo
   provements.  Pl                  nt edition of the "Inter
   Official Protoco                 for the standard
   and status of th                                of this m

Copyright Notice

   Copyright (C) The Internet Society (2006).

Abstract
```

We share the Internet together

- The BGP Default-Free Zone is a shared resource, “pollution” in this shared routing system is problematic for everyone
- “Water conflicts” exist in the Internet:
  - Operator misconfigurations
  - BGP vendor software defects
  - Various types of malicious activity

Any problems upstream the “BGP river” can cause problems downstream!

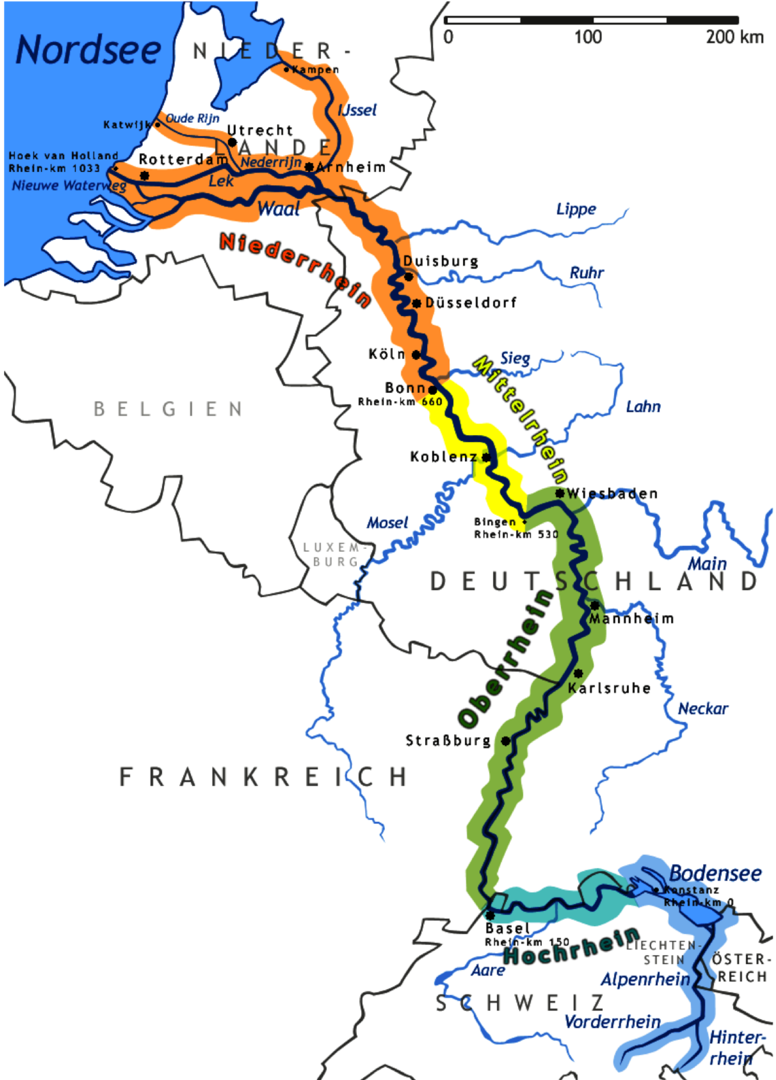

A simple problem scenario

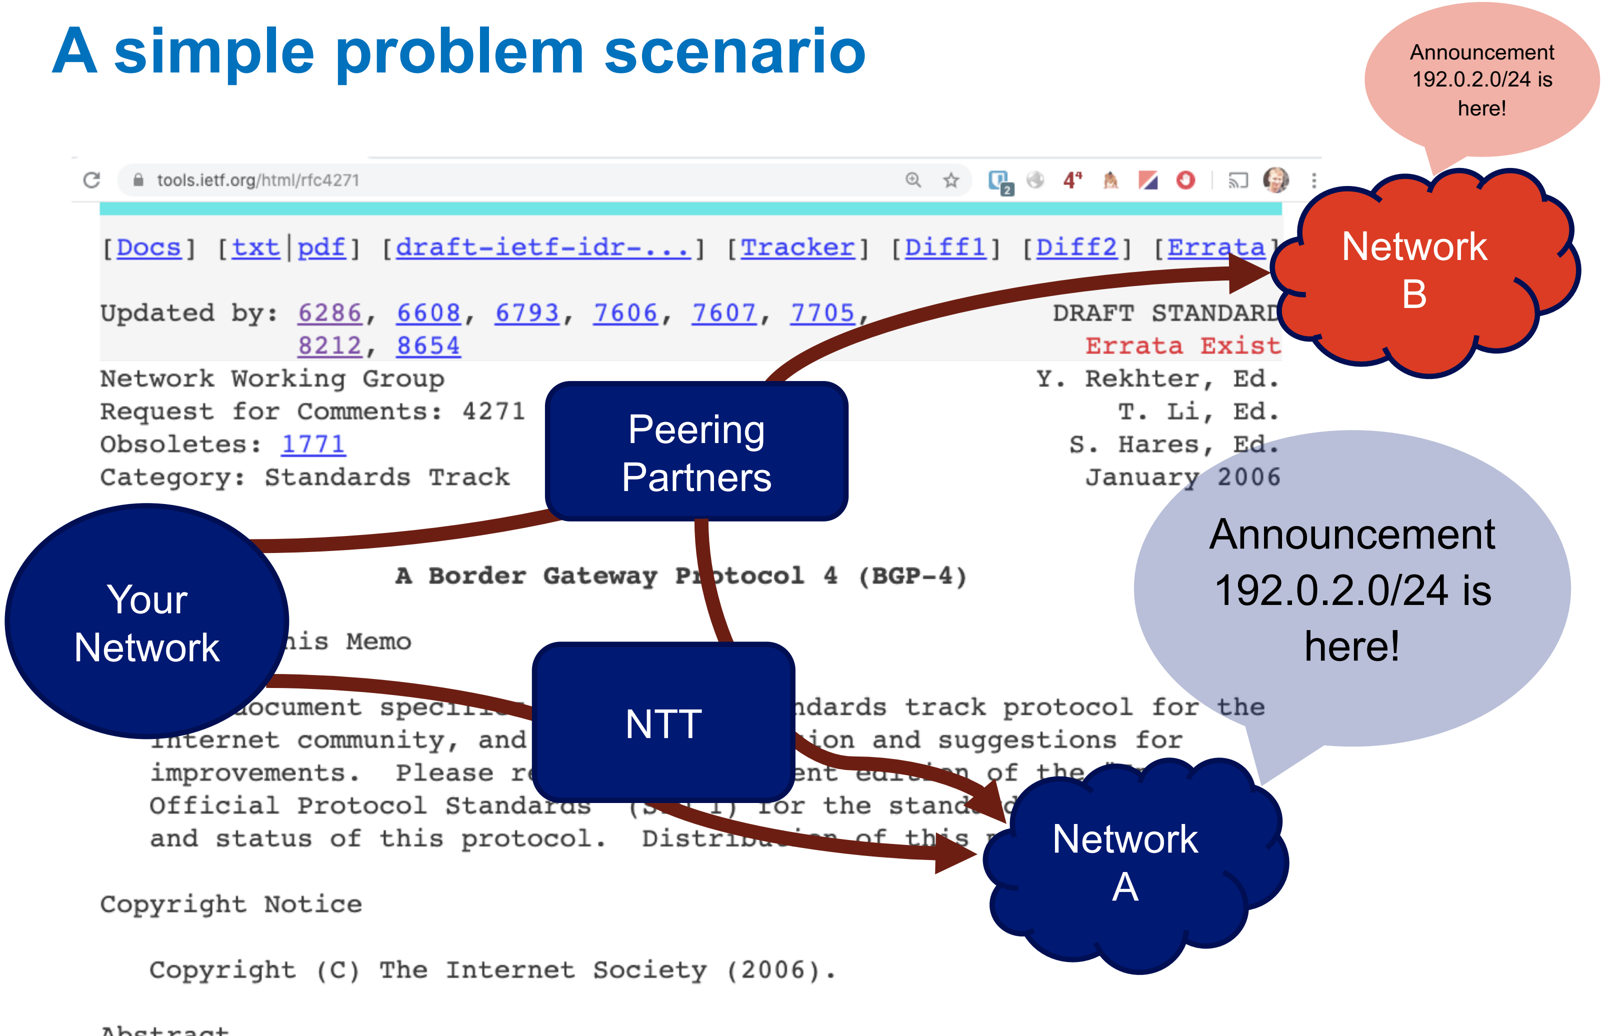

Announcement 192.0.2.0/24 is here!

Network B

Peering Partners

Your Network

NTT

Announcement 192.0.2.0/24 is here!

Network A

tools.ietf.org/html/rfc4271

```
[Docs] [txt|pdf] [draft-ietf-idr-...] [Tracker] [Diff1] [Diff2] [Errata]

Updated by: 6286, 6608, 6793, 7606, 7607, 7705,            DRAFT STANDARD
            8212, 8654                                       Errata Exist
Network Working Group                                     Y. Rekhter, Ed.
Request for Comments: 4271                                     T. Li, Ed.
Obsoletes: 1771                                             S. Hares, Ed.
Category: Standards Track                                    January 2006


              A Border Gateway Protocol 4 (BGP-4)

   his Memo

   document speci          ndards track protocol for the
   Internet community, and          ion and suggestions for
   improvements.  Please re          ent edition of the "
   Official Protocol Standards  (S     ) for the stand
   and status of this protocol.  Distrib     of this

Copyright Notice

   Copyright (C) The Internet Society (2006).

Abstract
```

Our tool belt: BGP protection mechanisms

- Routing policies (via BGP communities) to enforce what was agreed upon between the two EBGP peers to be announced and propagated further into the routing system
- Maximum BGP Prefix Limits
- AS_PATH filters (http://peerlock.net/)
- IRR based BGP prefix-list filters to create “allowlists”
- Bogon prefix-filters applied as “blocklist”

…. and now also RPKI-based BGP Prefix Origin Validation!

RPKI-based BGP Origin Validation

The RPKI is a distributed database which can be cryptographically verified.

Through this database, Internet Number Resource holders (aka the owners of an IP Prefix) can publish their routing intentions: ROAs.

NTT then applies this validated information (in real-time!) to optimise the choices presented as input to the BGP best path selection process on the AS 2914 routers.

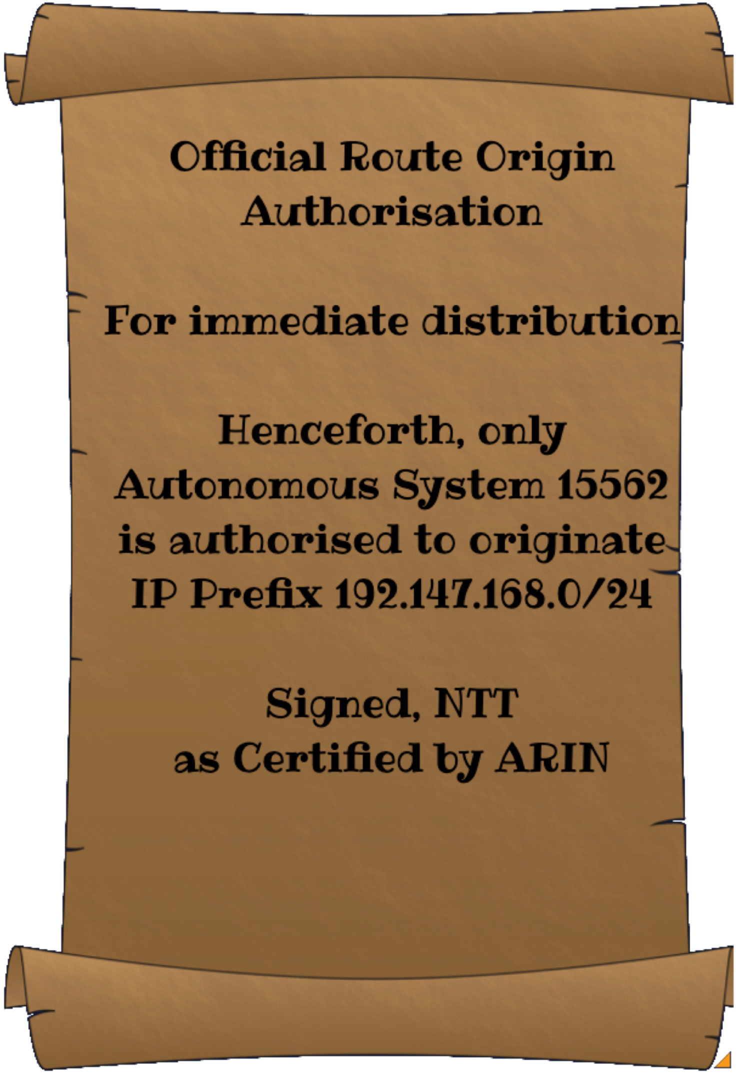

Official Route Origin
Authorisation

For immediate distribution

Henceforth, only
Autonomous System 15562
is authorised to originate
IP Prefix 192.147.168.0/24

Signed, NTT
as Certified by ARIN

Create RPKI ROAs via the Internet Registry

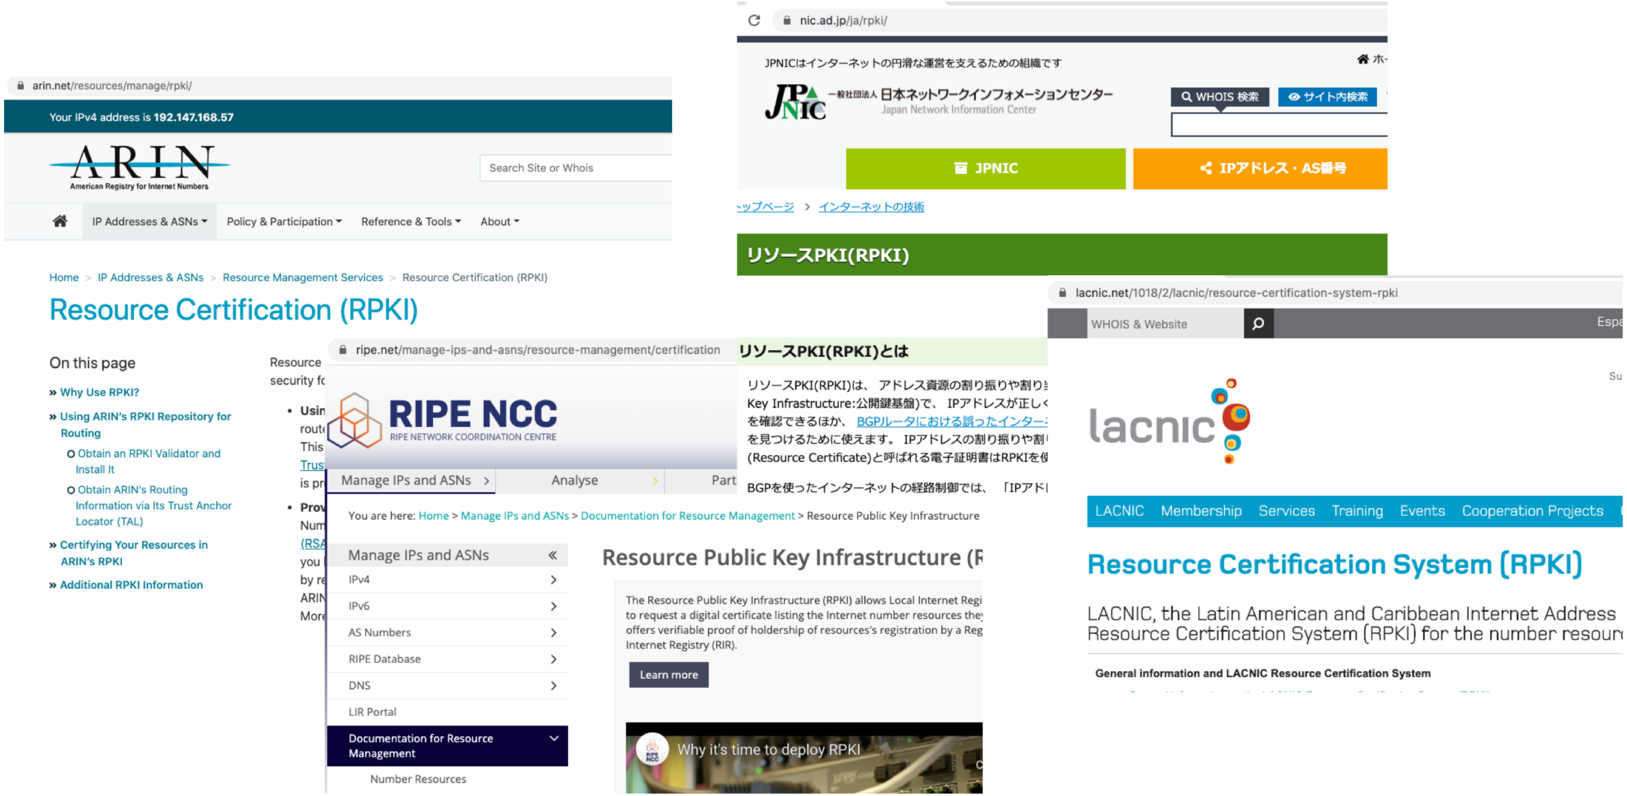

arin.net/resources/manage/rpki/

Your IPv4 address is **192.147.168.57**

ARIN American Registry for Internet Numbers

Search Site or Whois

IP Addresses & ASNs ▾ Policy & Participation ▾ Reference & Tools ▾ About ▾

Home > IP Addresses & ASNs > Resource Management Services > Resource Certification (RPKI)

Resource Certification (RPKI)

On this page

- » Why Use RPKI?
- » Using ARIN's RPKI Repository for Routing
  - ○ Obtain an RPKI Validator and Install It
  - ○ Obtain ARIN's Routing Information via Its Trust Anchor Locator (TAL)
- » Certifying Your Resources in ARIN's RPKI
- » Additional RPKI Information

Resource
security fo

- Usin
  route
  This
  Trus
  is pr
- Prov
  Num
  (RSA
  you
  by re
  ARIN
  More

nic.ad.jp/ja/rpki/

JPNICはインターネットの円滑な運営を支えるための組織です

ホ-

一般社団法人 日本ネットワークインフォメーションセンター
Japan Network Information Center

WHOIS 検索 サイト内検索

JPNIC IPアドレス・AS番号

ップページ > インターネットの技術

リソースPKI(RPKI)

リソースPKI(RPKI)とは

リソースPKI(RPKI)は、 アドレス資源の割り振りや割り
Key Infrastructure:公開鍵基盤)で、 IPアドレスが正しく
を確認できるほか、 BGPルータにおける誤ったインター
を見つけるために使えます。 IPアドレスの割り振りや割
(Resource Certificate)と呼ばれる電子証明書はRPKIを使

BGPを使ったインターネットの経路制御では、 「IPアド

ripe.net/manage-ips-and-asns/resource-management/certification

RIPE NCC RIPE NETWORK COORDINATION CENTRE

Manage IPs and ASNs > Analyse > Part

You are here: Home > Manage IPs and ASNs > Documentation for Resource Management > Resource Public Key Infrastructure

Manage IPs and ASNs «

- IPv4 >
- IPv6 >
- AS Numbers >
- RIPE Database >
- DNS >
- LIR Portal
- Documentation for Resource Management
  - Number Resources

Resource Public Key Infrastructure (R

The Resource Public Key Infrastructure (RPKI) allows Local Internet Regi
to request a digital certificate listing the Internet number resources they
offers verifiable proof of holdership of resources's registration by a Reg
Internet Registry (RIR).

Learn more

Why it's time to deploy RPKI

lacnic.net/1018/2/lacnic/resource-certification-system-rpki

WHOIS & Website

Espa

Su

lacnic

LACNIC Membership Services Training Events Cooperation Projects

Resource Certification System (RPKI)

LACNIC, the Latin American and Caribbean Internet Address
Resource Certification System (RPKI) for the number resour

General information and LACNIC Resource Certification System

The other technical components

- The RIRs (ARIN, RIPE, APNIC, AFRINIC, LACNIC) operate the top level Certificate Authorities
- Organisations pull all published RPKI information from the Internet
- The RPKI Cache Validators construct the RPKI cache
- RPKI-to-Router (RTR) servers transport the Validated ROA Payloads (VRPs) to the EBGP routers

AFRINIC
APNIC
ARIN
lacnic
RIPE NCC
NTT
Telia
AT&T
gtt
CLOUDFLARE
rpki-client
ROUTINATOR
validador FORT
CISCO
BIRD
NOKIA
JUNIPER NETWORKS
GO AHEAD. EXPLORE. OpenBGPD

The process to deploy RPKI Origin Validation

- Organise engineering resources:
  - Test & deploy servers that will run the RPKI cache
  - Monitoring (connected RTR clients, number of VRPs, etc)
  - Design routing policies with “RPKI invalid == reject” in mind
  - Figure out where you can and cannot enable RPKI in the network
  - Read and write documentation about the changes
- Provide training to all relevant staff:
  - How to debug network issues now with RPKI in mind
  - What is RPKI? (questions will come up in NOC, operations, sales & marketing)

Then pick a date….. and do it!

What it can look like when you enable RPKI ROV

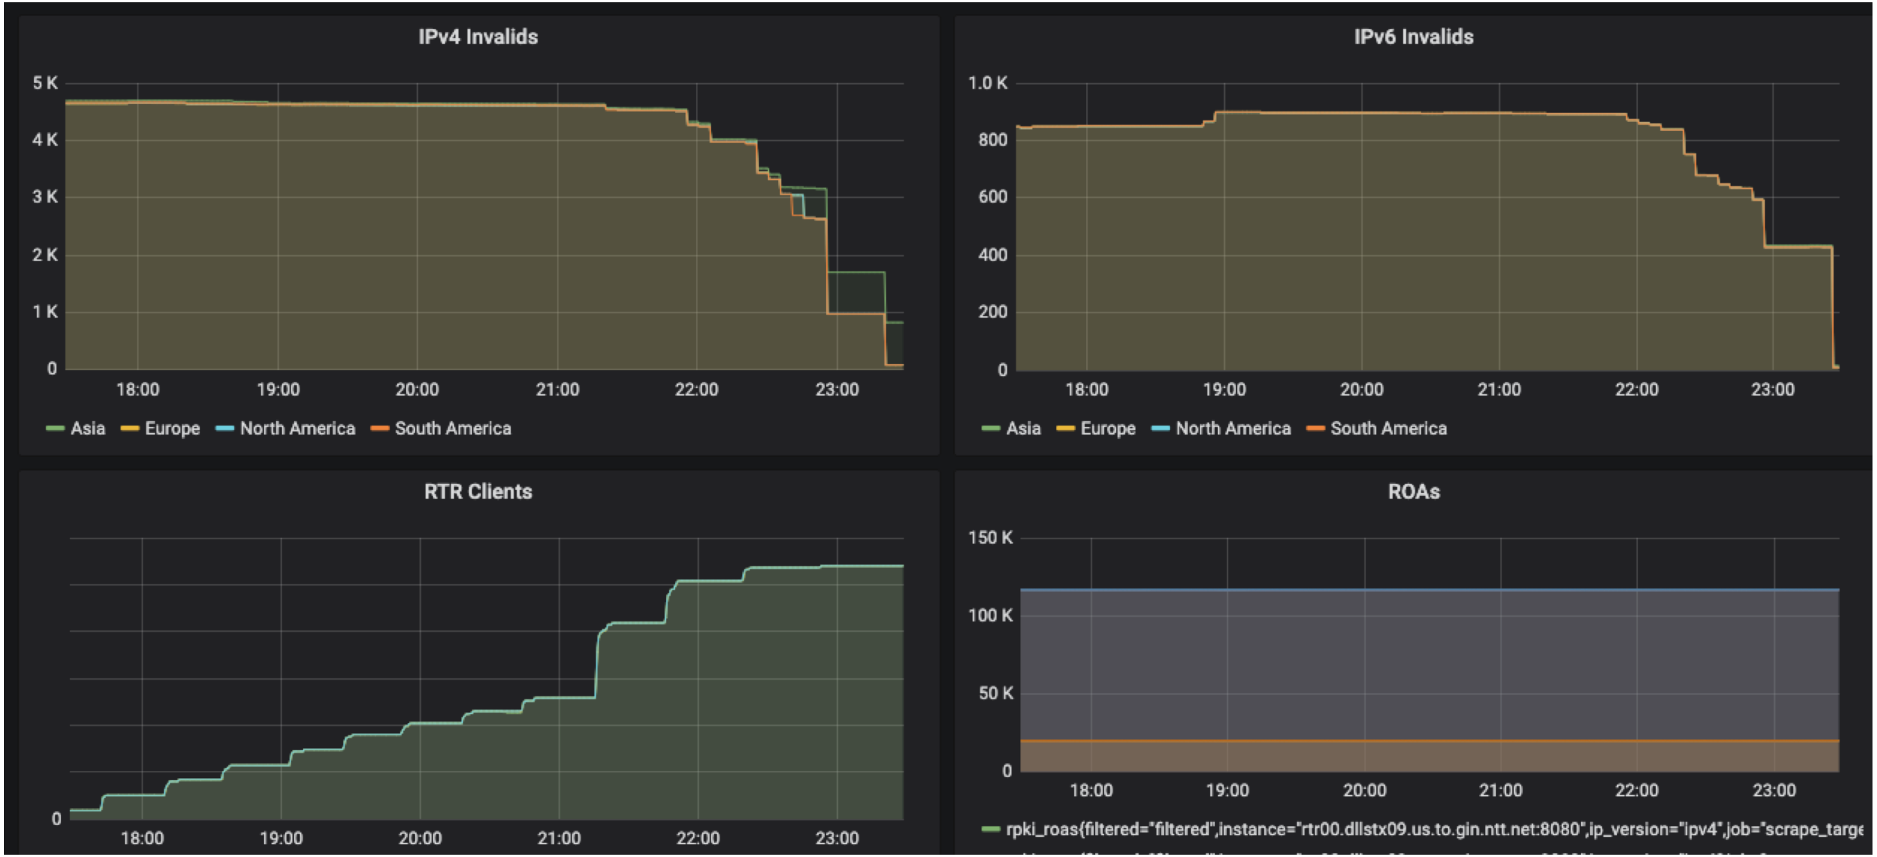

And what it looked like later on

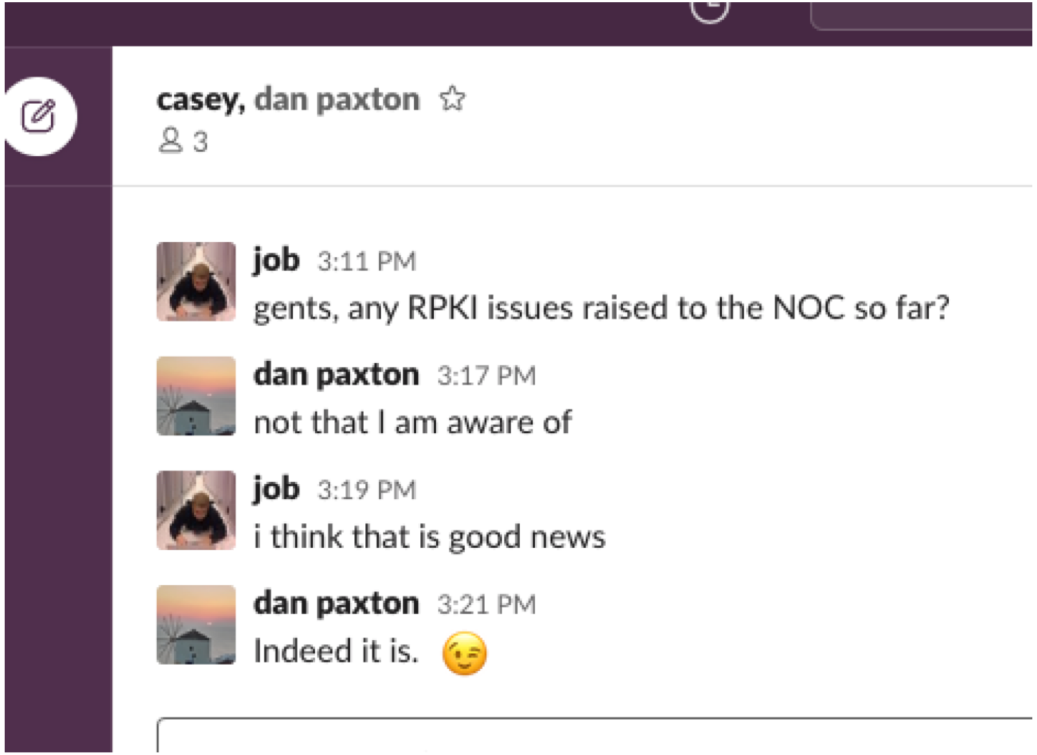

casey, dan paxton ☆

3

job 3:11 PM

gents, any RPKI issues raised to the NOC so far?

dan paxton 3:17 PM

not that I am aware of

job 3:19 PM

i think that is good news

dan paxton 3:21 PM

Indeed it is. 😉

Our deployment experience

- Started out with setting up RPKI caches and RTR servers (3 on 3 different continents) based on OpenBSD rpki-client and GoRTR.
- Extensive lab testing to test correct functioning of all software pieces
- Analysed potential impact of enabling RPKI on NTT's global IP traffic profile using pmacct's RPKI integration.
- Identified which customers who might be impacted by the change (very few), send out notification emails to those.
- Found a few (mostly cosmetic) software defects in vendor code, and identified a list of devices on the network that do not support RPKI.
- Deployment in production environment was done through a single flag day. RTR sessions brought up and policy immediately updated.

RPKI is an emergent industry trend

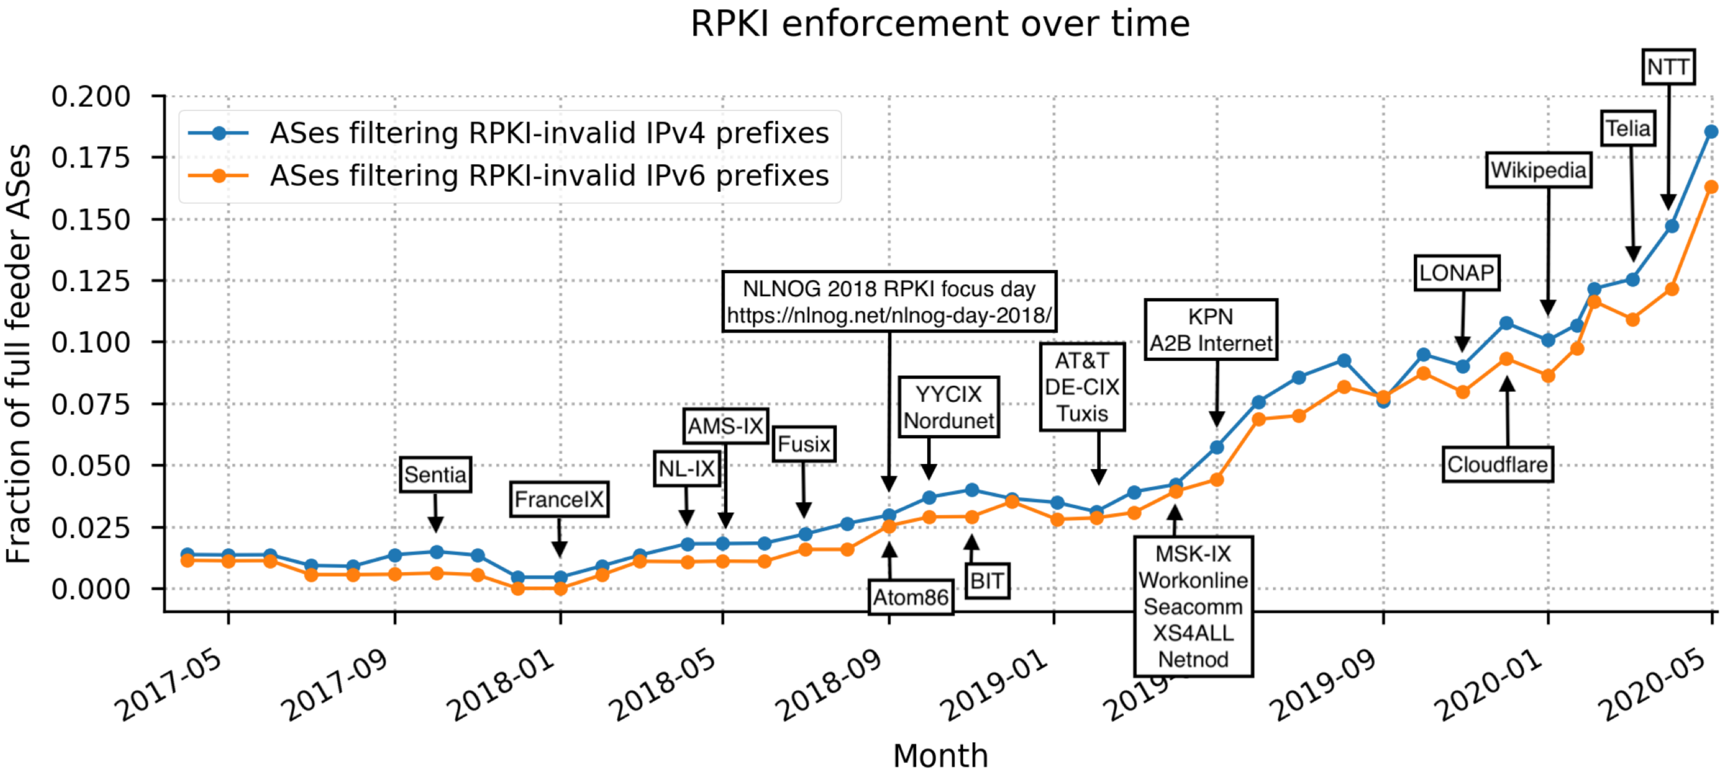

Source: https://twitter.com/JobSnijders/status/1256326712347881473

At present, 2598 Autonomous Systems in the BGP Default-Free Zone appear to apply Origin Validation (as measured from NTT’s perspective)

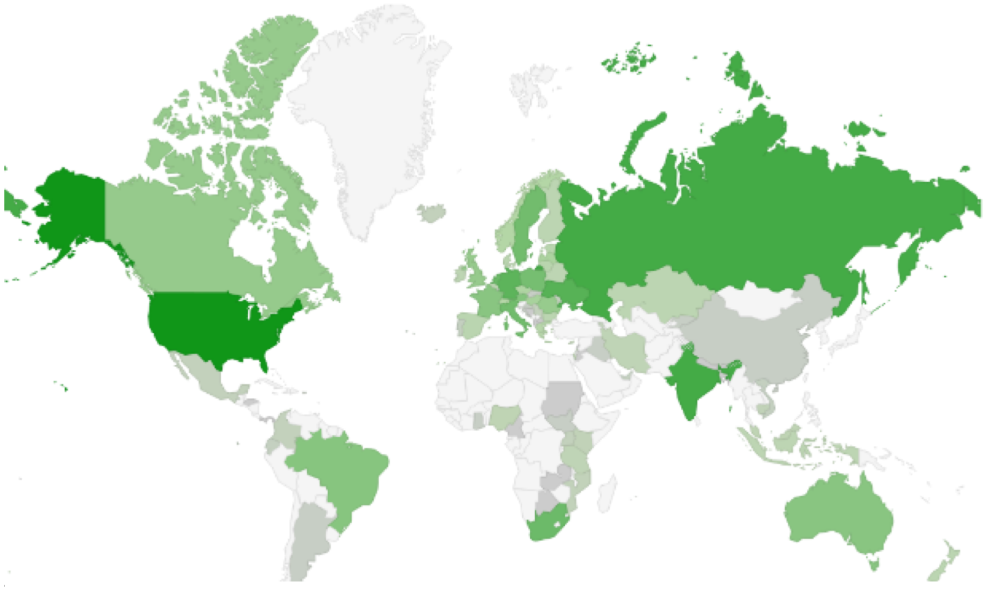

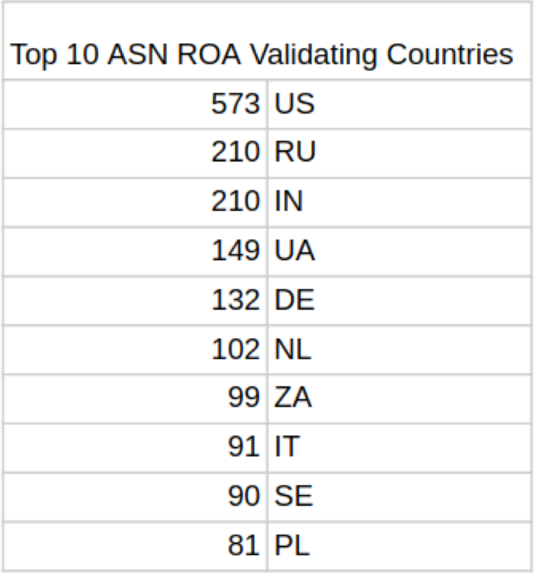

Top 10 ASN ROA Validating Countries	
573	US
210	RU
210	IN
149	UA
132	DE
102	NL
99	ZA
91	IT
90	SE
81	PL

Source: Ben Cox, RIPE 80, Routing Working Group Session

~20% of IP space is covered by RPKI ROAs

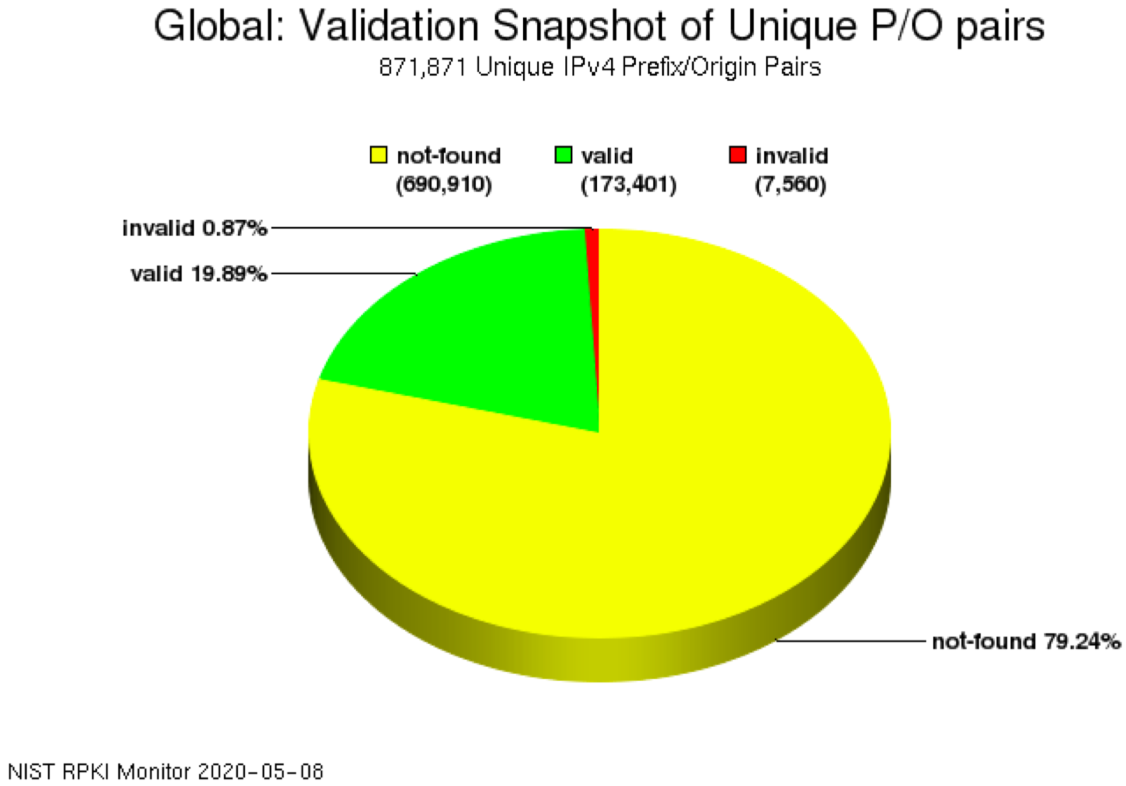
Global: Validation Snapshot of Unique P/O pairs

871,871 Unique IPv4 Prefix/Origin Pairs

NIST RPKI Monitor 2020-05-08

https://rpki-monitor.antd.nist.gov/

Expected fail-positions of RPKI and BGP for incremental deployment on the global Internet

- RPKI is an opportunistic security layer, applied on top of existing best practices related to inter-domain routing. Creation of ROAs activates the Origin Validation protection mechanism in NTT's EBGP policies. The cryptographic validation procedure as developed through open standards and open source efforts, will discard malformed, invalid or otherwise distrusted **RPKI objects**. This is a **fail-secure** feature.
- RPKI is only used to reject RPKI "Invalid" BGP announcements (RFC 6811). Only BGP route announcements with the RPKI "Not-Found" and "Valid" state are expected to propagate through AS 2914. This is a **cryptographically actuated** coalescing pipeline **filter** applied to BGP routing information.
- Should all RTR servers become unreachable from the **EBGP** router's perspective, our routing policy assigns any BGP announcement the "Not-Found" state. This means the route announcement will not be rejected because of RPKI. This is a **fail-safe** feature.
- The above arrangement provides the Internet with an **incremental deployment strategy**.
- Changes to RPKI ROAs are expected to propagate within the global system in about an hour.

RPKI Timing

Any changes an operator makes to their RPKI ROAs are expected to propagate through the RPKI supply chain into the global Internet routing system in about an hour.

Other applications of the RPKI in the IRR space

RIPE-731 RPKI based filter process applied to the “RIPE-NONAUTH” IRR. RPKI can now be used to identify stale (incorrect) IRR objects and remove those automatically.

IRRd

NTT’s open source IRRd v4.1.0-beta3, can apply a similar IRR clean-up mechanism to NTTCOM and rr.ntt.net’s IRR mirror instance. Release & deployment timeline expected to be in the second half of 2020.

Study resources

Excellent community maintained documentation with NLNetLabs as editor

https://rpki.readthedocs.io/

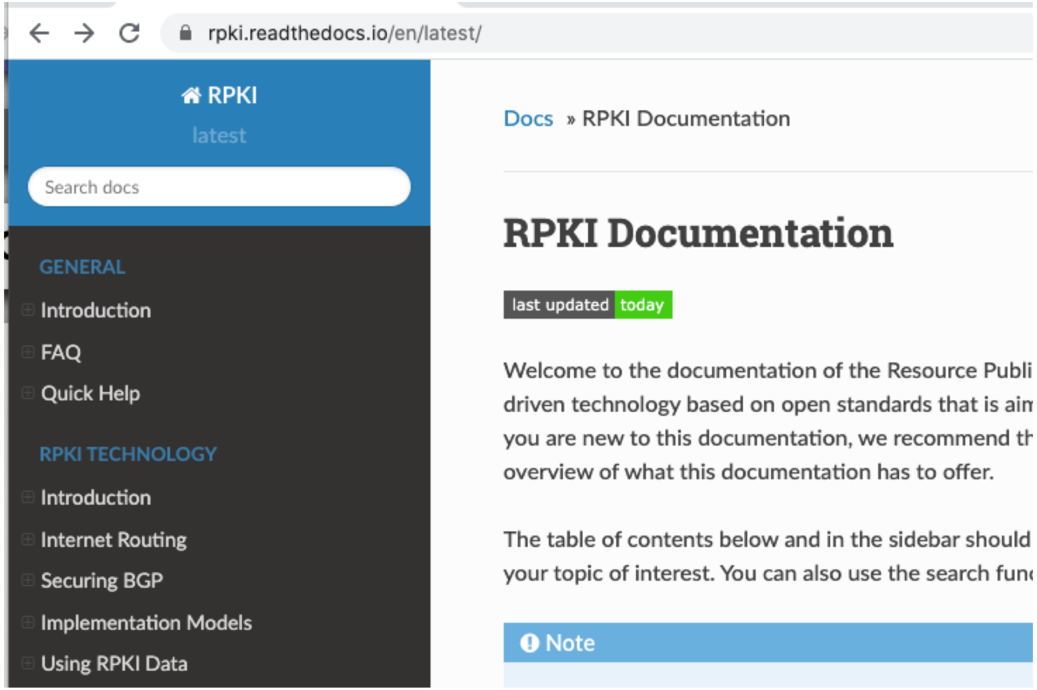

rpki.readthedocs.io/en/latest/

Docs » RPKI Documentation

RPKI Documentation

last updated today

Welcome to the documentation of the Resource Publi
driven technology based on open standards that is aim
you are new to this documentation, we recommend th
overview of what this documentation has to offer.

The table of contents below and in the sidebar should
your topic of interest. You can also use the search func

Note

IETF RFC Specifications

The RPKI architecture is documented in RFC 6480.

The RPKI specification is documented in a spread out series of RFCs:
RFC 6481, RFC 6482, RFC 6483, RFC 6484, RFC 6485, RFC 6486,
RFC 6487, RFC 6488, RFC 6489, RFC 6490, RFC 6491, RFC 6492,
RFC 6493, RFC 7935, RFC 7318, RFC 7330, RFC 8630, RFC 8481,
RFC 8416, RFC 8183, RFC 8182

Does RPKI resolve all Internet routing problems?

Short answer: No. There is no silver bullet.

However, there are multiple ongoing collaborative work projects in open standards bodies open source software, to bring further improvements to the Internet routing system.

IETF®

Request to all Internet network operators

- Create RPKI ROAs for Internet Number Resources
- Work to deploy RPKI-based BGP Origin Validation such that RPKI invalid route announcements are rejected on all EBGP sessions
(especially all transit, peering, and route server sessions)

Coordination platforms for RPKI and Internet Routing

- Network Information Centers such as the Internet Registries (both RIRs and NIRs)
- Regional Network Operator Groups (NOGs) and MANRS
- The RPKI mailing list at NLNetLabs: https://lists.nlnetlabs.nl/mailman/listinfo/rpki
- Hundreds of operators are connected to the #IX IRC channel on irc.terahertz.net

Together we do great things