Securing BGP - RPKI

ThaiNOG2018 - Bangkok

21 May 2018

Tashi Phuntsho (tashi@apnic.net)

Fat-finger/Hijacks/Leaks

- **Amazon (AS16509) Route53 hijack – April2018**
  - AS10279 (eNET) announced/originated more specifics (/24s) of Amazon Route53's prefix (205.251.192.0/21)
    - **205.251.192.0/24 ……. 205.251.199.0/24**
    - https://ip-ranges.amazonaws.com/ip-ranges.json
  - Its peers, like AS6939 (HE), shared these routes with 1000s of their own peers…
  - The motive?
    - During the period, DNS servers in the hijacked range only responded to queries for myetherwallet.com
    - Responded with addresses associated with AS41995/AS48693)

Fat-finger/Hijacks/Leaks

- **Route53 hijack (continued….)**
  - Resolvers querying any of the Route53 managed names, would actually be asking the authoritative servers controlled through the BGP hijack

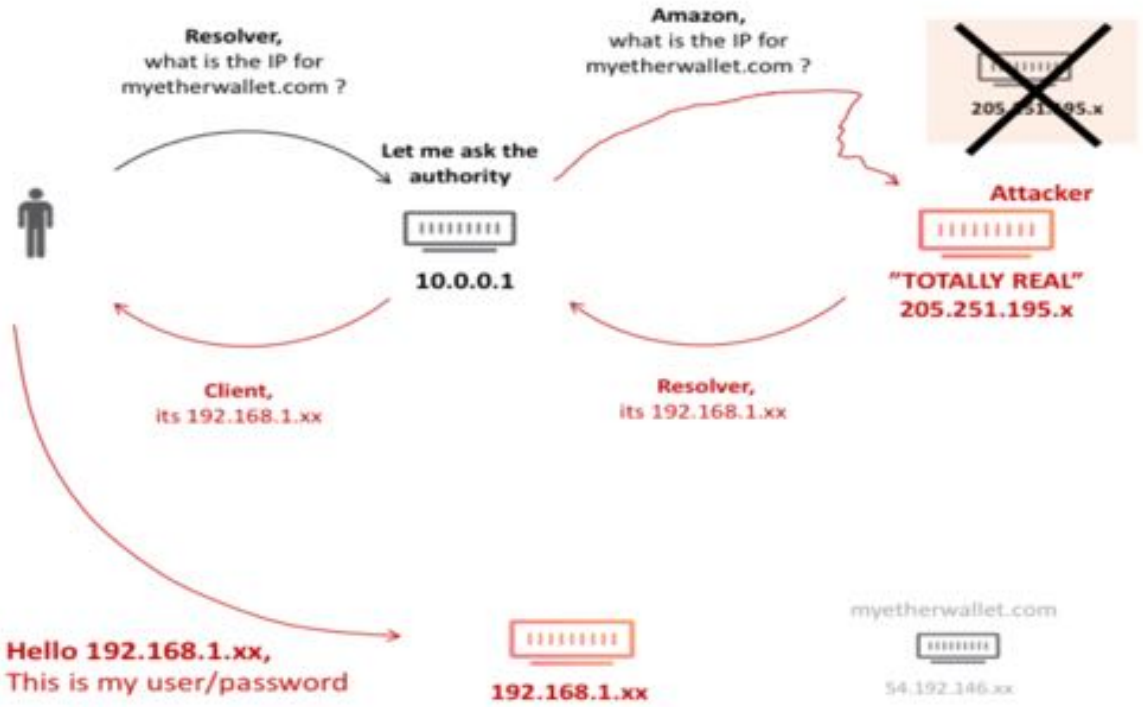

https://blog.cloudflare.com/bgp-leaks-and-crypto-currencies

Fat-finger/Hijacks/Leaks

- **Bharti (AS9498) originates 103.0.0.0/10**
  - Dec 2017 (~ 2 days)
  - No damage done – more than 8K specific routes!

- **Google brings down Internet in Japan**
  - Aug 2017 (~ 24 hours)
  - Google (AS15169) leaked *>130K* prefixes to Verizon (AS701) – in Chicago
    - Normally ~ 50 prefixes
    - ~25K of those were NTT OCN's (AS4713) more specifics
    - which was leaked onwards to KDDI and IIJ (and accepted)
  - Everyone who received the leaked more specifics, preferred the Verizon-Google path to reach NTT OCN!

Fat-finger/Hijacks/Leaks

- **Google leak (contd…)**

```
trace from Tokyo, Japan to Inuyama, Japan at 03:28 Aug 25, 2017
 1  *
 2  183.177.32.145    Equinix Asia Pacific                 Tokyo       Japan           0.249
 3  210.130.154.37    IIJ IPv4 BLOCK ( AS2497 )            Tokyo       Japan           0.618
 4  58.138.102.109    tky001bb11.IIJ.Net                   Tokyo       Japan           0.877
 5  58.138.88.86      sjc002bb12.IIJ.Net                   San Jose    United States   97.797
 6  152.179.48.117    TenGigE0-3-0-8.GW6.SJC7.ALTER.NET    San Jose    United States   97.869
 7  *
 8  152.179.105.110   google-gw.customer.alter.net         Chicago     United States   337.19
 9  108.170.243.197   Google Inc.                          Chicago     United States   246.325
10  *
11  209.85.241.43     Google Inc.                                      United States   256.188
12  72.14.238.38      Google Inc.                          Vancouver   Canada          247.849
13  209.85.245.110    Google Inc.                          Vancouver   Canada          249.291
14  *
15  108.170.242.138   Google Inc.                          Tokyo       Japan           246.267
16  211.0.193.21      OCN (AS4713) CIDR BLOCK 21           Tokyo       Japan           246.351
17  122.1.245.65      OCN (AS4713) CIDR BLOCK 81           Tokyo       Japan           246.426
18  *
19  153.149.218.10    OCN (AS4713) CIDR BLOCK 93           Ōsaka-shi   Japan           256.027
20  125.170.96.38     OCN (AS4713) CIDR BLOCK 77                       Japan           255.683
21  *
22  60.37.32.250      OCN (AS4713) CIDR BLOCK 70                       Japan           254.989
23  118.23.141.202    OCN (AS4713) CIDR BLOCK 86                       Japan           254.526
24  *
25  211.11.83.160     OCN (AS4713) CIDR BLOCK 23           Inuyama     Japan           256.212
```

After leak (JP->JP)

```
trace from London, England to Nürnberg, Germany at 03:30 Aug 25, 2017
 1  *
 2  195.66.248.190    fe0-2.tr2.linx.net                   London      United Kingdom  0.327
 3  195.66.249.10     ge0-2-502.tr5.linx.net               London      United Kingdom  0.441
 4  195.66.249.13     ge0-2-501.tr4.linx.net               London      United Kingdom  0.477
 5  195.66.248.10     uunet-uk-transit.thn.linx.net        London      United Kingdom  0.507
 6  158.43.193.245    POS0-0.CR2.LND6.ALTER.NET            London      United Kingdom  0.497
 7  140.222.239.41    0.xe-0-0-0.IL1.NYC50.ALTER.NET       New York    United States   108.146
 8  146.188.4.197     xe-0-0-1.IL1.NYC41.ALTER.NET         New York    United States   75.719
 9  140.222.234.221   0.et-10-1-0.GW7.CHI13.ALTER.NET      Chicago     United States   94.793
10  152.179.105.110   google-gw.customer.alter.net         Chicago     United States   224.352
11  *
12  216.239.40.189    Google Inc.                          Northlake   United States   202.193
13  216.239.58.255    Google Inc.                                                      203.995
14  216.239.58.12     Google Inc.                                                      207.026
15  209.85.253.184    Google Inc.                          Luxembourg  Luxembourg      212.944
16  209.85.252.215    Google Inc.                                                      213.112
17  108.170.252.71    Google Inc.                                                      213.265
18  72.14.222.53      Google Inc.                                      Germany         212.861
19  188.111.165.169   Vodafone GmbH                                    Germany         227.077
20  178.7.138.112     Vodafone D2 GmbH                     Nürnberg    Germany         234.226
```

After leak (EU->JP)

```
trace from Tokyo, Japan to Inuyama, Japan at 04:44 Aug 24, 2017
 1  *
 2  202.177.203.50    xe-0-0-0.gw401.ty2.ap.equinix.com    Tokyo       Japan           0.717
 3  183.177.32.143    xe-1-1-1.gw402.ty1.ap.equinix.com    Tokyo       Japan           0.755
 4  143.90.232.25     25.143090232.odn.ne.jp               Tokyo       Japan           1.411
 5  143.90.161.73                                          Tokyo       Japan           2.757
 6  143.90.47.14      STOrs-01Te0-1-0-1.nw.odn.ad.jp       Tokyo       Japan           3.552
 7  210.252.167.230   230.210252167.odn.ne.jp              Tokyo       Japan           4.094
 8  *
 9  60.37.54.105      OCN (AS4713) CIDR BLOCK 70           Tokyo       Japan           4.088
10  125.170.97.85     OCN (AS4713) CIDR BLOCK 77                       Japan           4.017
11  125.170.97.74     OCN (AS4713) CIDR BLOCK 77           Ōsaka-shi   Japan           12.263
12  153.149.219.22    OCN (AS4713) CIDR BLOCK 93           Ōsaka-shi   Japan           12.362
13  153.146.148.18    OCN (AS4713) CIDR BLOCK 93           Tokyo       Japan           14.45
14  60.37.32.250      OCN (AS4713) CIDR BLOCK 70                       Japan           13.116
15  118.23.141.202    OCN (AS4713) CIDR BLOCK 86                       Japan           13.332
16  118.23.142.99     OCN (AS4713) CIDR BLOCK 86                       Japan           22.307
17  211.11.83.160     OCN (AS4713) CIDR BLOCK 23           Inuyama     Japan           15.672
```

Before leak (JP->JP)

https://dyn.com/blog/large-bgp-leak-by-google-disrupts-internet-in-japan/

Fat-finger/Hijacks/Leaks

- **Google (AS15169) services downed**
  - Nov 2012 (~ 30 minutes)
  - Moratel Id (AS23947) leaked Google prefixes to its upstream (AS3491)
    - **`AS path: … 3491 23947 15169`**

- **YouTube (AS36561) Incident**
  - Feb 2008 (down for ~ 2 hours)
  - PT (AS17557) announced 208.65.153.0/24 (208.65.152.0/22)
    - Propagated by AS3491 (PCCW)

How do we address these…

- **Filters!!!**
  - On both ends of a eBGP session
    - AS-PATH, prefix-list, max-prefix limit
  - only announce/originate your own prefix (and your downstream)
  - Only accept your peer's prefix (and their downstream)

Current practice

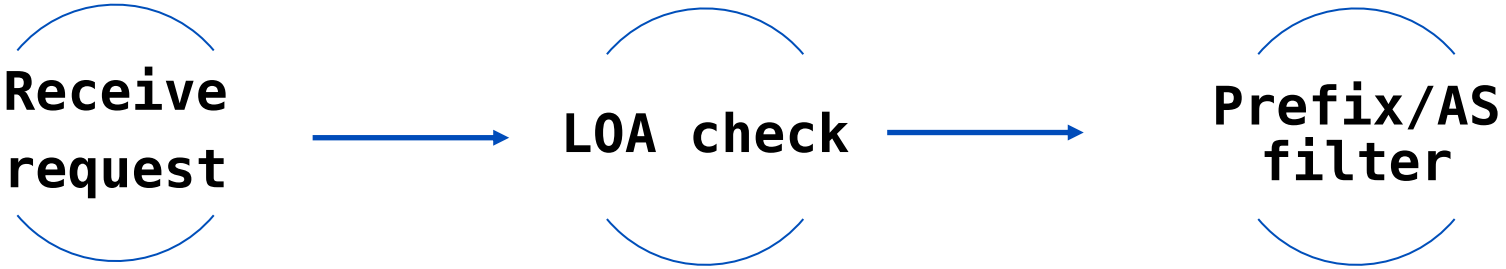

Tools and techniques

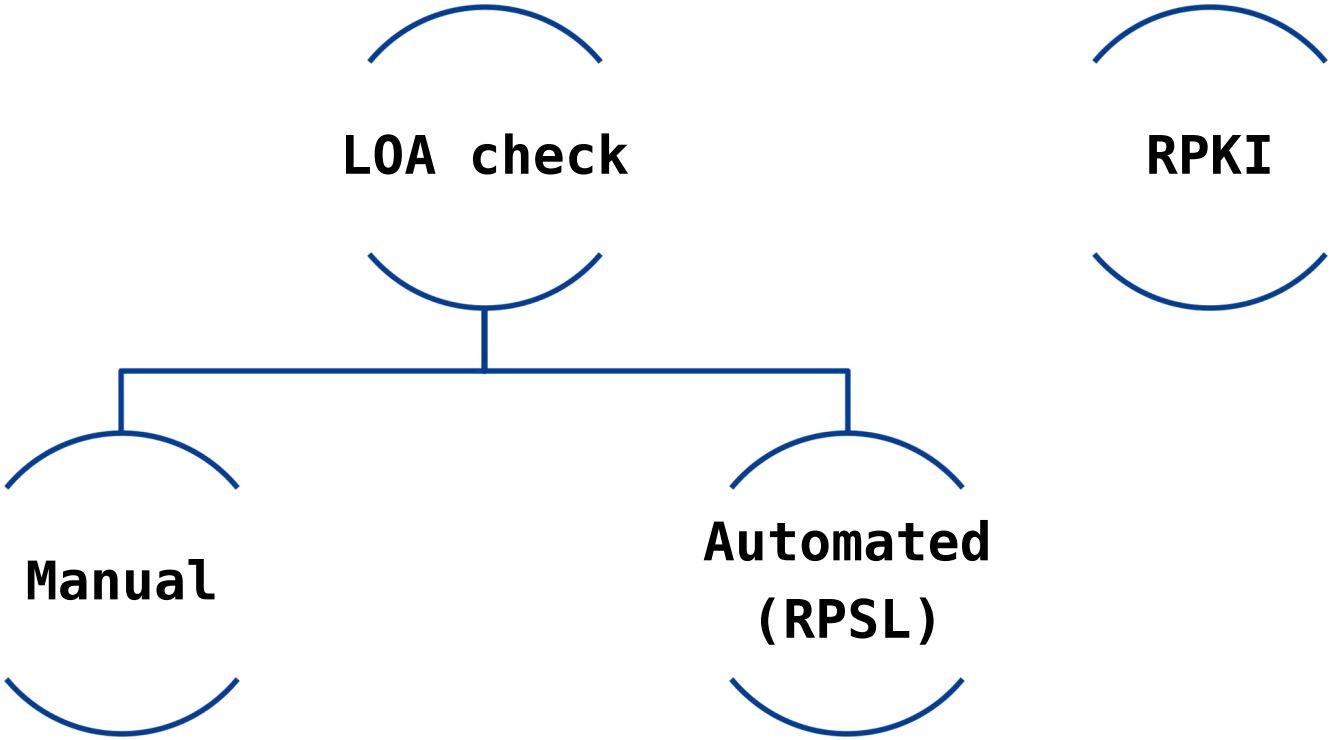

Learning curve… slow (BGP)

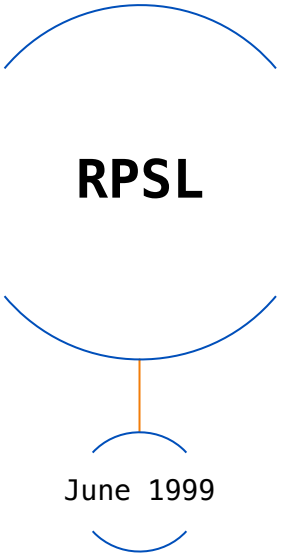

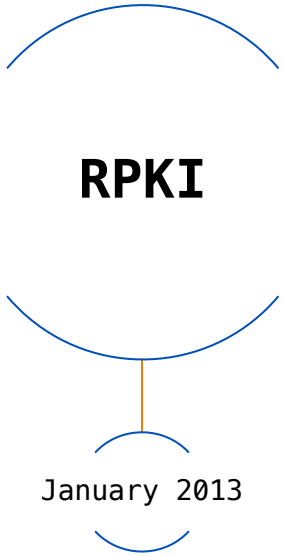

What is RPKI?

Goals of RPKI

- To authoritatively prove who owns an IP prefix and which AS(es) can originate
  - Attaching digital certificates to network resources (AS Number & IP Address)

Benefits of RPKI

- Prevents **route hijacking**
  - A prefix originated by an AS without authorization
  - *Reason:* ***malicious intent***

- Prevents **mis-origination**
  - A prefix that is mistakenly originated by an AS which does not own it
  - Also route leakage
  - *Reason:* ***configuration mistakes/fat-finger***

RPKI Building Blocks

1. Trust Anchors (RIRs)
2. Route Origination Authorizations (ROA)
3. RPKI Validator

Public Key Encryption Recap

- Public and private key mathematically related to each other
  - Cannot derive one from the other

- Encrypt with one and decrypt with the other
  - Encrypt with private, only public can decrypt
  - Encrypt with public, only private can decrypt

Public-key Infra

- Digital (X.509) certificates
  - associates a public key with an individual or organization
    - public key of the subject!

Version	Version of X.509
Serial Number	Uniquely identifies the certificate
Signature Algorithm	Algorithms used by the CA to sign the cert
Issuer Name	Id of the CA (that issued the cert)
Validity Period	Cert validity
Subject Name	The cert owner
Subject Public Key	Owner's public key
Issuer ID	Extra info (Issuer of the cert)
Subject ID	Extra info (owner of the cert)
Extensions (CRL)	
CA Digital Signature	Digital Signature of the CA

RPKI Profile

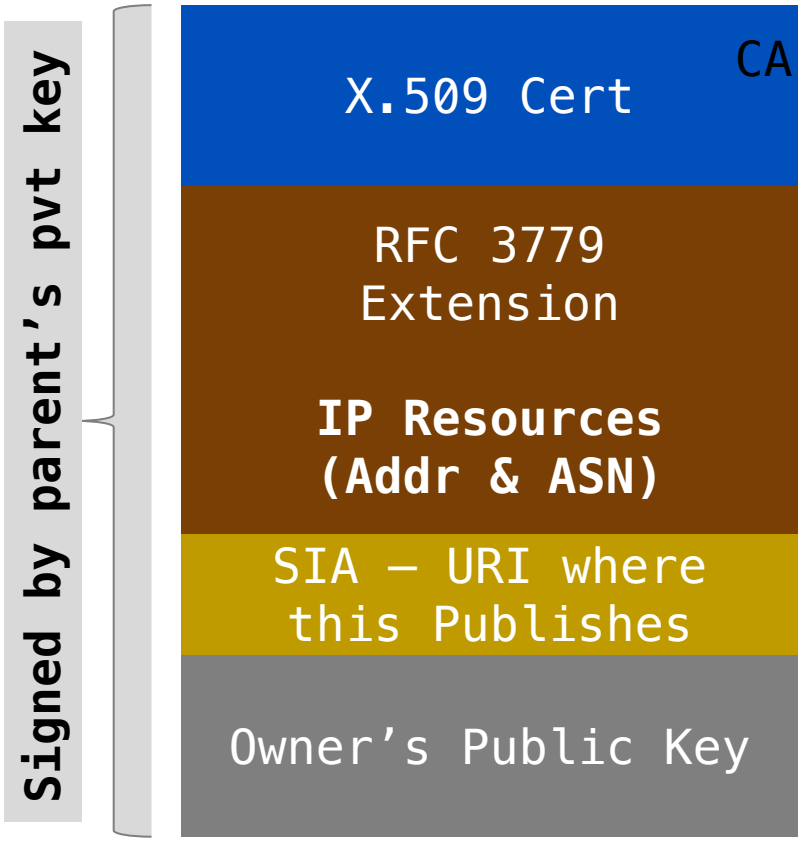

- Resource certificates are based on the X.509 v3 certificate format (RFC 5280)
- Extended by RFC 3779 – binds a list of resources **(IPv4/v6, ASN**) to the subject of the certificate
- SIA (subject information access) contains a URI that references the directory where it is published

Trust Anchor (TA)

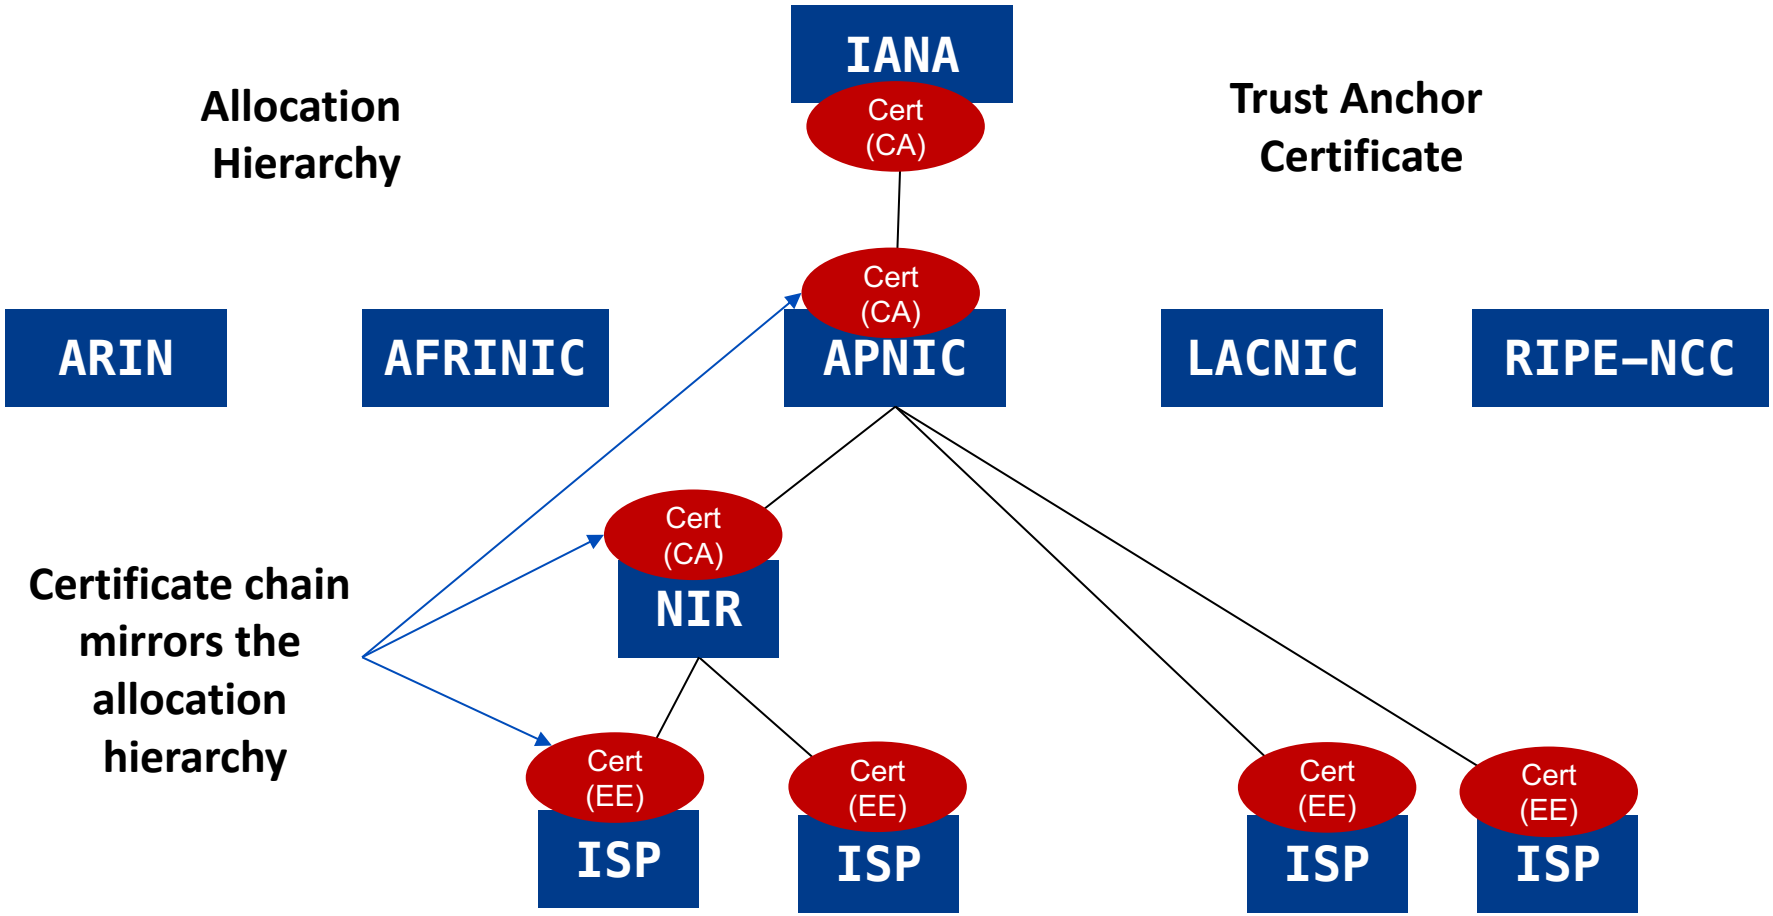

Source : http://isoc.org/wp/ietfjournal/?p=2438

Issuing Party

- Internet Registries (RIR, NIR, Large LIRs)
  - Certificate Authority
    - issues certificates for customers
  - Customers create their ROAs
    - Sign their resources with the certificates
- Publishes the ROA records

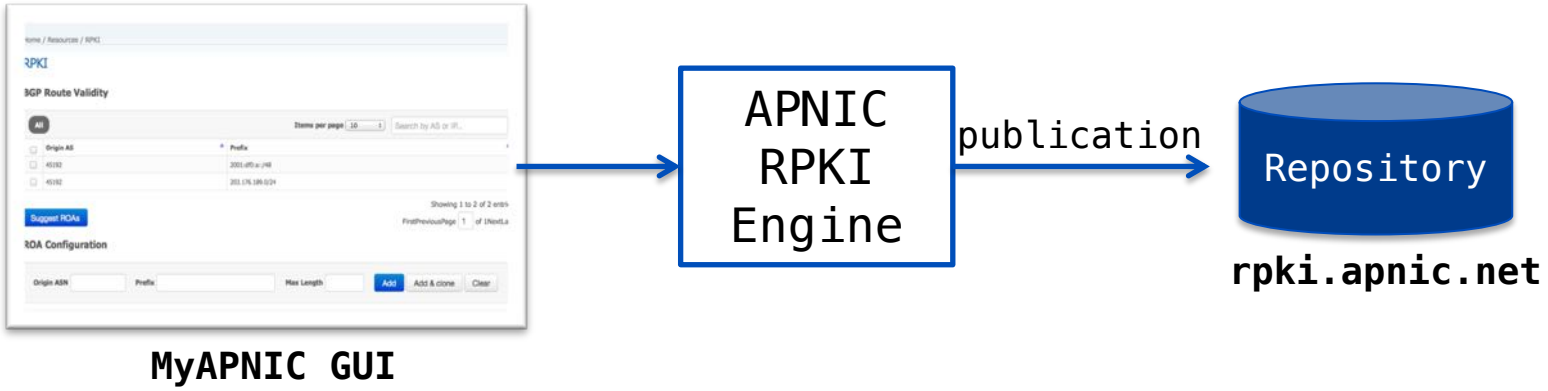

Single Trust anchor

- 27 Feb 2018: a single expanded trust anchor
  - https://blog.apnic.net/2018/02/27/updating-rpki-trust-anchor-configuration/

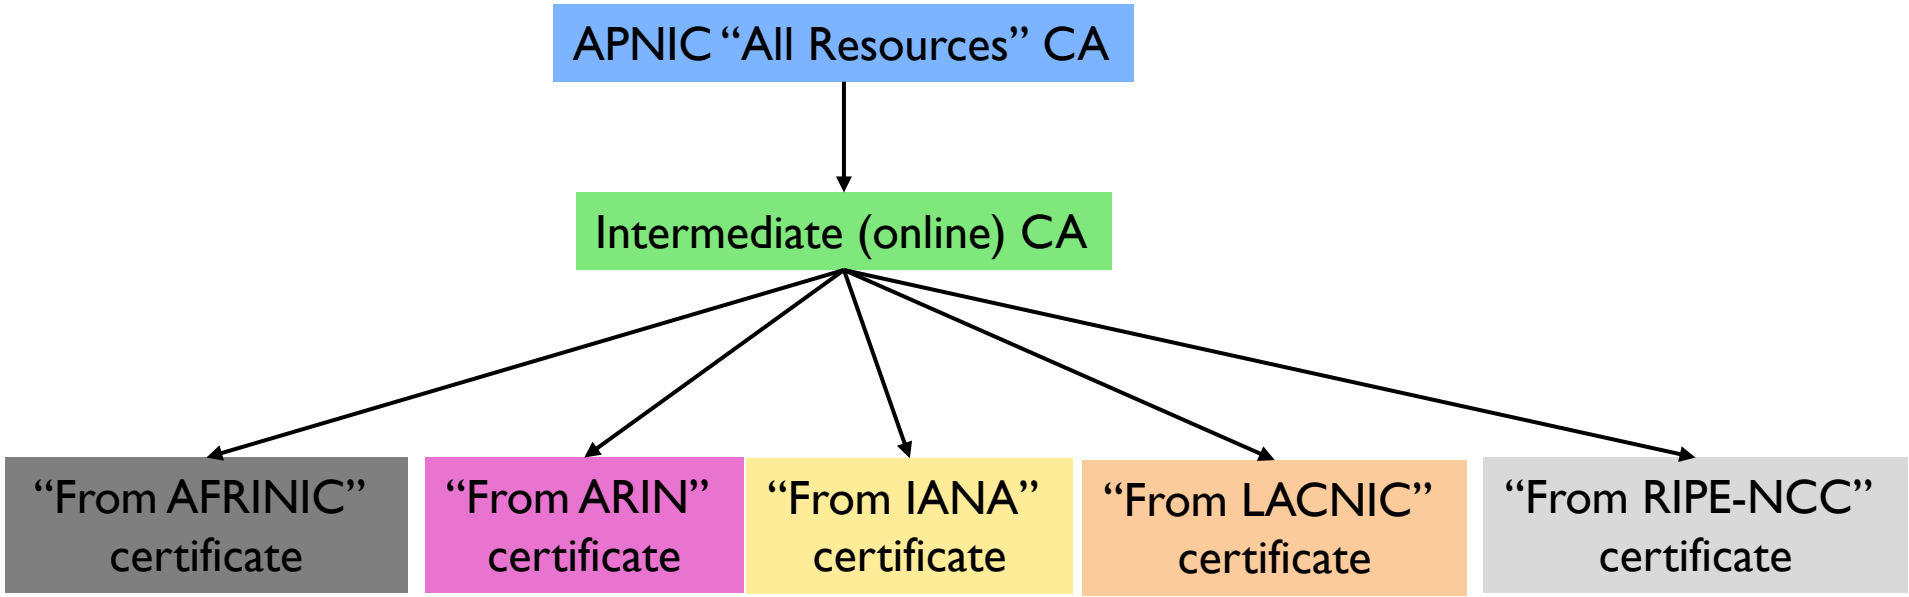

ROA- Route Origin Authorization

- A digitally signed object that contains a list of address prefixes and the nominated AS number
- It is an authority created by a prefix holder to authorize an ASN to originate one or more prefixes
  - Which can be verified cryptographically using RPKI

Prefix	203.176.189.0
Max-length	/24
Origin ASN	AS17821

- Multiple ROAs can exist for the same prefix

Relying Party (RPKI Validator)

- The RPKI Validator
  - Gathers ROA from the distributed RPKI database
  - Validates each entry's signature
    - Validated cache

Relying Party (RPKI Validator)

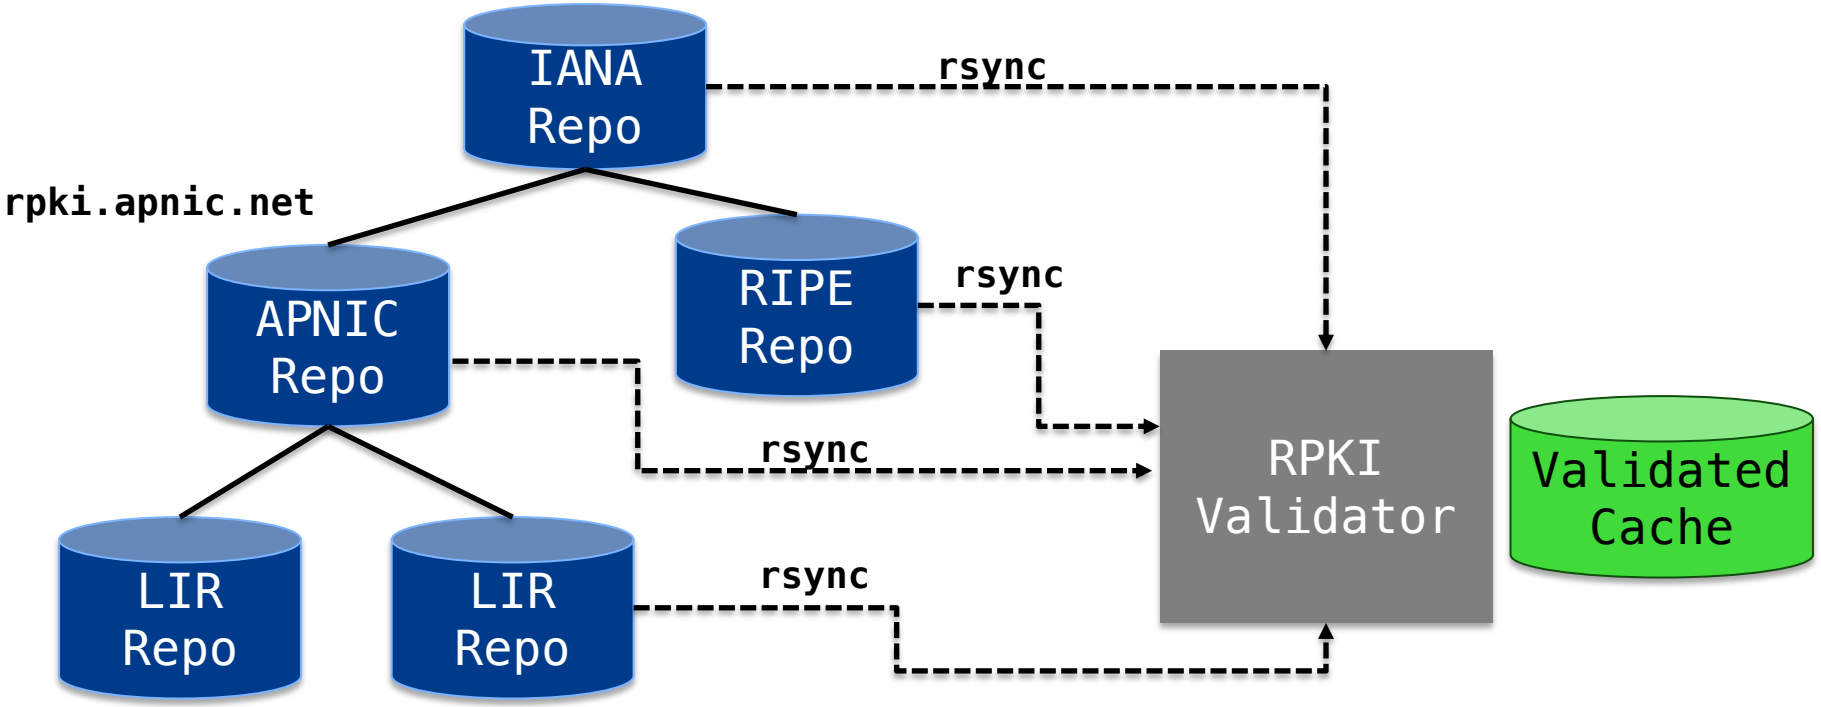

Origin Validation

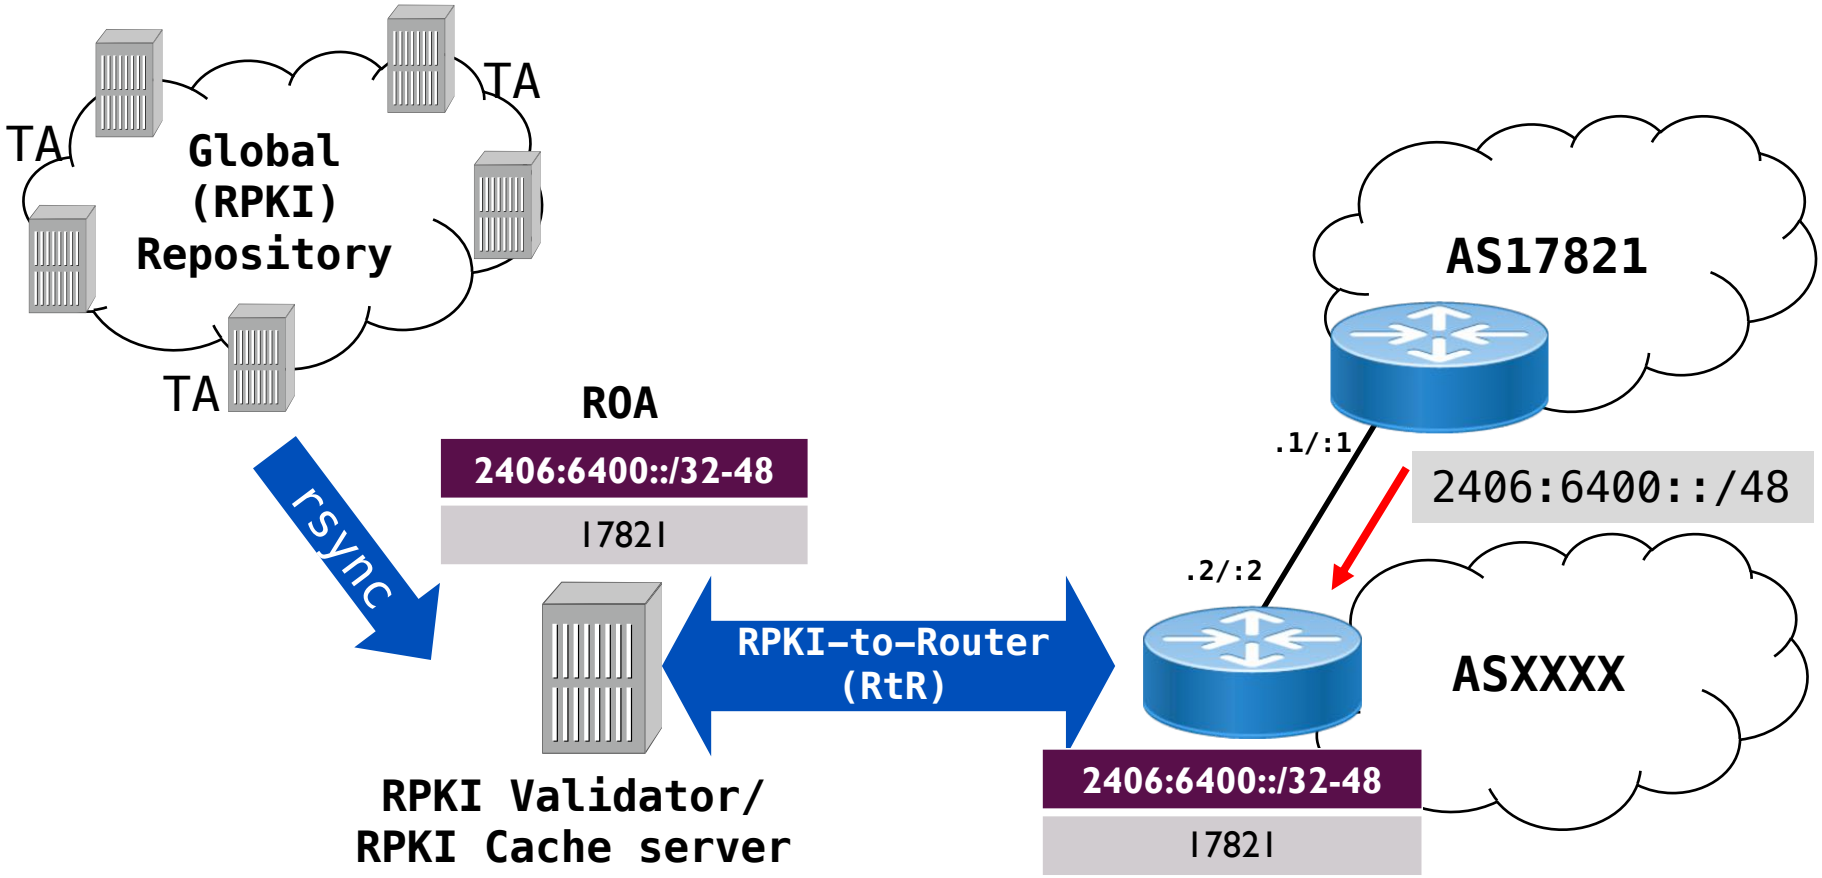

Origin Validation

- Router gets ROA information from the RPKI Cache
  - Crypto is stripped (by the validator)

- The BGP process will check each received BGP update against the ROA information and label
  - **Valid**
  - **Invalid**
  - **Not Found**

Validation States

- **Valid**
  - the prefix and AS pair are found in the database.
- **Invalid**
  - prefix is found, but origin AS is wrong, or
  - the prefix length is longer than the maximum length
- **Not Found / Unknown**
  - No valid ROA found
  - Neither valid nor invalid
    - Perhaps not created!

RPKI States

ROA =>	65420	10.0.0.0/16	/18
	Origin AS	Prefix	Max Length

VALID	AS65420	10.0.0.0/16
VALID	AS65420	10.0.128.0/17
INVALID	AS65421	10.0.0.0/16
INVALID	AS65420	10.0.10.0/24
UNKNOWN	AS65430	10.0.0.0/8

Policies based on validation

- Define your policy based on the validation state
  - Do nothing (observe)
  - Label BGP communities
  - **Modify preference values**
    - RFC7115
  - Drop Invalid announcements (paranoid!)
    - Invalid - but verify against other databases (IRR whois)

RPKI Caveats

- When RTR session goes down, the validation status changes to **NOT FOUND** for all routes after a while
  - ***Invalid => Not Found***
  - we need several RTR sessions (at least 2) and/or need careful filtering policies

- During router reload, which one is faster: receiving ROAs or receiving BGP updates?
  - If receiving BGP routes is faster than ROA, the router will propagate the invalid routes to its iBGP peers

RPKI Further Reading

- RFC 5280: X.509 PKI Certificates
- RFC 3779: Extensions for IP Addresses and ASNs
- RFC 6481-6493: Resource Public Key Infrastructure

Implementation

Create & publish your ROA

- MyAPNIC portal
  - Resources > RPKI
  - Here is a detailed guide:

https://www.apnic.net/wp-content/uploads/2017/12/ROUTE_MANAGEMENT_GUIDE.pdf

Create (publish) your ROA

- Available prefixes for which you can create ROA

Create (publish) your ROA

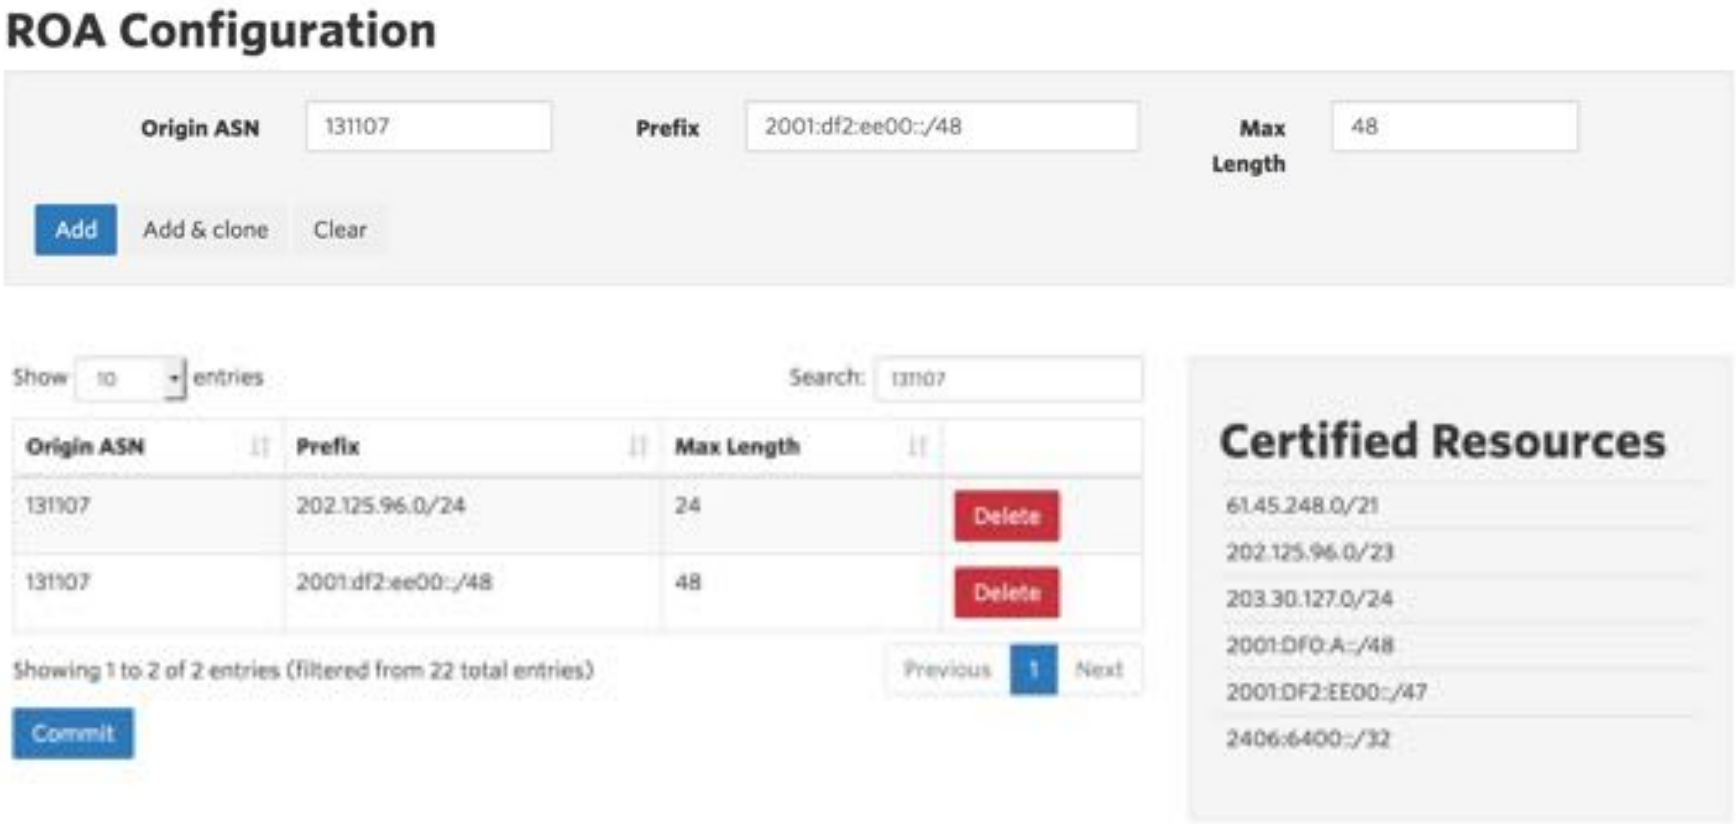

Check your ROA

```
# whois -h whois.bgpmon.net 2001:df2:ee00::/48

Prefix:                    2001:df2:ee00::/48
Prefix description:        APNICTRAINING-DC
Country code:              AU
Origin AS:                 131107
Origin AS Name:            APNICTRAINING LAB DC
RPKI status:               ROA validation successful
First seen:                2016-06-30
Last seen:                 2018-01-21
Seen by #peers:            97
```

```
# whois -h whois.bgpmon.net "--roa 131107 2001:df2:ee00::/48"

------------------------
ROA Details
------------------------
Origin ASN:       AS131107
Not valid Before: 2016-09-07 02:10:04
Not valid After:  2020-07-30 00:00:00  Expires in 2y190d9h34m23.2000000029802s
Trust Anchor:     rpki.apnic.net
Prefixes:         2001:df2:ee00::/48 (max length /48) 202.125.96.0/24 (max length /24)
```

Check your ROA

`https://bgp.he.net/`

Announced By		
Origin AS	**Announcement**	**Description**
AS131107	2001:df2:ee00::/48 	testing

Deploy RPKI Validator

- Two options:
  - RIPE RPKI Validator

```
https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources
```

  - Dragon Research Labs RPKI Toolkit

```
https://github.com/dragonresearch/rpki.net
```

RIPE - Validator

- Download RPKI Validator

```
# wget https://lirportal.ripe.net/certification/content/static/validator/rpki-validator-app-2.24-dist.tar.gz
```

- Installation

```
tar -zxvf rpki-validator-app-2.23-dist.tar.gz
cd rpki-validator-app-2.23
./rpki-validator.sh start
```

  - Need to download ARIN's TAL separately

```
wget https://www.arin.net/resources/rpki/arin-ripevalidator.tal
```

    - Move it to **"<base-folder>/conf/tal"** and restart

RIPE - Validator

`http://rpki-validator.apnictraining.net:8080/`

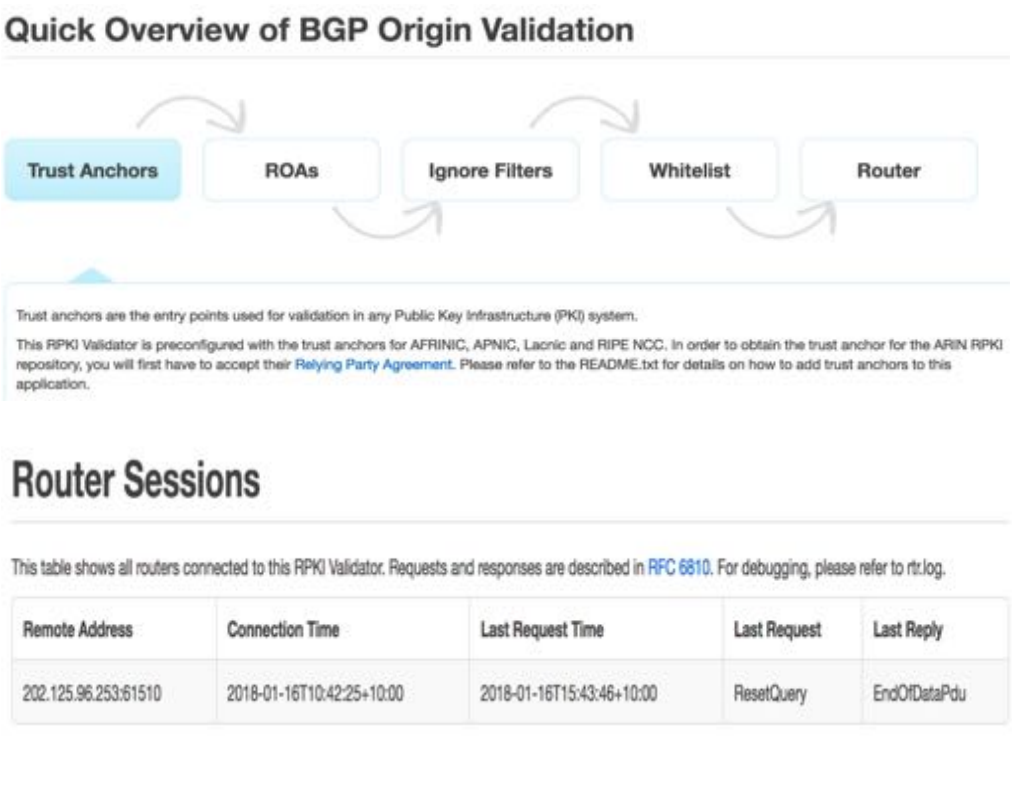

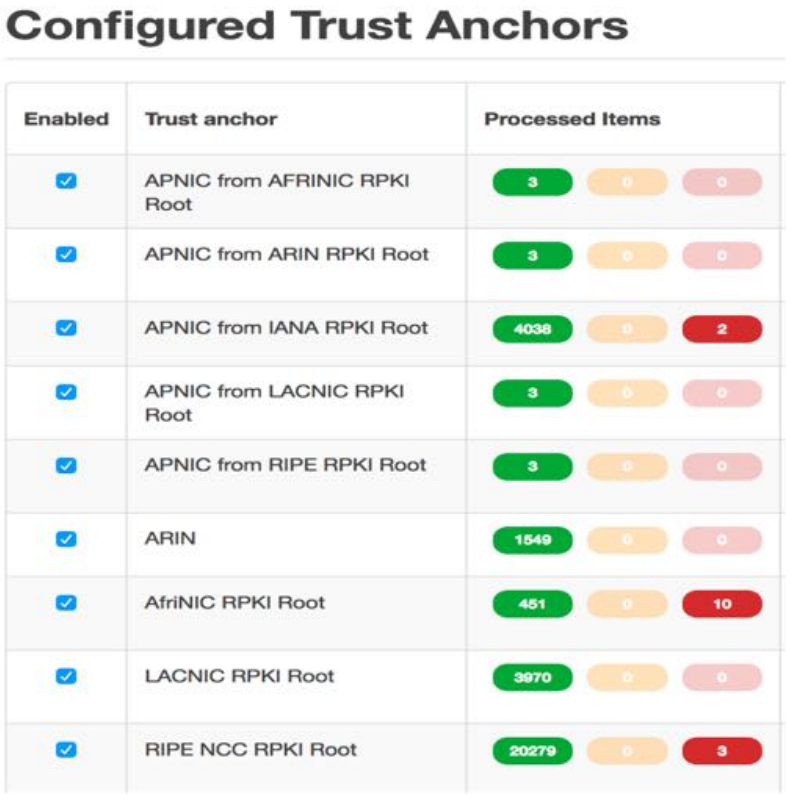

Dragon Research - Validator

- Installation on Ubuntu 16.04 Xenial

```
https://github.com/dragonresearch/rpki.net/blob/master/doc/quickstart/xenial-rp.md
```

- Installation

```
# wget -q -O /etc/apt/trusted.gpg.d/rpki.gpg https://download.rpki.net/APTng/apt-gpg-key.gpg
```

```
# wget -q -O /etc/apt/sources.list.d/rpki.list https://download.rpki.net/APTng/rpki.xenial.list

-q: quite (wget output)
-O: output to <file>
```

```
# apt update
```

```
# apt install rpki-rp
```

Dragon Research - Validator

http://rpki-dragonresearch.apnictraining.net/rcynic/

rcynic summary 2017-01-03T01:07:37Z

Overview Repositories Problems All Details

Grand totals for all repositories

	Tainted by stale CRL	Object accepted	Manifest interval overruns c
None .cer	28	5981	
None .crl		5948	
None .gbr		3	
None .mft		5948	1
None .roa		5923	
Total	**28**	**23803**	**1**

Current total object counts (distinct URIs)

Repository	.cer	.crl	.gbr	.mft	.roa
ca.rg.net					
ca0.rpki.net					
localcert.ripe.net					
repository.lacnic.net					
rpki-pilot.lab.dtag.de					
rpki-repository.nic.ad.jp					
rpki.afrinic.net					
rpki.apnic.net					
rpki.ripe.net					
Total	**0**	**0**	**0**	**0**	**0**

Overview for repository rpki.apnic.net

	Tainted by stale CRL	Object accepted	Manifest interval over
None .cer		752	
None .crl		748	
None .mft		748	
None .roa		492	
Total		**2740**	

Configuration - IOS

- Establishing session with the validator

```
router bgp 131107
 bgp rpki server tcp <validator-IP> port 323 refresh 120
```

- Policies based on validation:

```
route-map ROUTE-VALIDATION permit 10
 match rpki valid
 set local-preference 110
!
route-map ROUTE-VALIDATION permit 20
 match rpki not-found
 set local-preference 100
!
route-map ROUTE-VALIDATION permit 10
 match rpki invalid
 set local-preference 90
!
```

Configuration - IOS

- Apply the route-map to inbound updates

```
router bgp 131107
!---output omitted-----!
 address-family ipv4
  neighbor X.X.X.169 activate
  neighbor X.X.X.169 route-map ROUTE-VALIDATION in
 exit-address-family
 !
 address-family ipv6
  neighbor X6:X6:X6:X6::151 activate
  neighbor X6:X6:X6:X6::151 route-map ROUTE-VALIDATION in
 exit-address-family
!
```

Configuration - JunOS

- Establishing session with the validator

```
routing-options {
   autonomous-system 131107;
   validation {
       group rpki-validator {
           session <validator-IP> {
               refresh-time 120;
               port 8282;
               local-address X.X.X.253;
           }
       }
   }
}
```

Router Configuration - JunOS

- Define policies based on the validation states

```
policy-options {
   policy-statement ROUTE-VALIDATION {
       term valid {
           from {
               protocol bgp;
               validation-database valid;
           }
           then {
               local-preference 110;
               validation-state valid;
               accept;
           }
       }
       term invalid {
           from {
               protocol bgp;
               validation-database invalid;
           }
           then {
               local-preference 90;
               validation-state invalid;
               accept;
           }
       }
```

```
term unknown {
            from {
                protocol bgp;
                validation-database unknown;
            }
            then {
                local-preference 100;
                validation-state unknown;
                accept;
            }
        }
    }
}
```

Router Configuration - JunOS

- Apply the policy to inbound updates

```
protocols {
   bgp {
       group external-peers {
           #output-ommitted
           neighbor X.X.X.1 {
               import ROUTE-VALIDATION;
               family inet {
                   unicast;
               }
           }
       }
       group external-peers-v6 {
           #output-ommitted
           neighbor X6:X6:X6:X6::1 {
               import ROUTE-VALIDATION;
               family inet6 {
                   unicast;
               }
           }
       }
   }
```

RPKI Verification - IOS

- IOS has only

```
#sh bgp ipv6 unicast rpki ?
  servers Display RPKI cache server information
  table   Display RPKI table entries


#sh bgp ipv4 unicast rpki ?
  servers Display RPKI cache server information
  table   Display RPKI table entries
```

RPKI Verification - IOS

- Check the RTR session

```
#sh bgp ipv4 unicast rpki servers

BGP SOVC neighbor is X.X.X.47/323 connected to port 323
Flags 64, Refresh time is 120, Serial number is 1516477445, Session ID is 8871
InQ has 0 messages, OutQ has 0 messages, formatted msg 7826
Session IO flags 3, Session flags 4008
 Neighbor Statistics:
 Prefixes 45661
 Connection attempts: 1
 Connection failures: 0
 Errors sent: 0
 Errors received: 0

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 255
Local host: X.X.X.225, Local port: 29831
Foreign host: X.X.X.47, Foreign port: 323
```

RPKI Verification - IOS

- Check the RPKI cache

```
#sh bgp ipv4 unicast rpki table
37868 BGP sovc network entries using 6058880 bytes of memory
39655 BGP sovc record entries using 1268960 bytes of memory

Network            Maxlen Origin-AS Source Neighbor
1.9.0.0/16         24     4788      0      202.125.96.47/323
1.9.12.0/24        24     65037     0      202.125.96.47/323
1.9.21.0/24        24     24514     0      202.125.96.47/323
1.9.23.0/24        24     65120     0      202.125.96.47/323
```

```
#sh bgp ipv6 unicast rpki table
5309 BGP sovc network entries using 976856 bytes of memory
6006 BGP sovc record entries using 192192 bytes of memory

Network             Maxlen Origin-AS Source Neighbor
2001:200::/32       32     2500      0      202.125.96.47/323
2001:200:136::/48   48     9367      0      202.125.96.47/323
2001:200:900::/40   40     7660      0      202.125.96.47/323
2001:200:8000::/35  35     4690      0      202.125.96.47/323
```

Check routes - IOS

```
#sh bgp ipv4 unicast 202.144.128.0/19
BGP routing table entry for 202.144.128.0/19, version 3814371
Paths: (1 available, best #1, table default)
 Advertised to update-groups:
    2
 Refresh Epoch 15
 4826 17660
   49.255.232.169 from 49.255.232.169 (114.31.194.12)
     Origin IGP, metric 0, localpref 110, valid, external, best
     Community: 4826:5101 4826:6570 4826:51011 24115:17660
     path 7F50C7CD98C8 RPKI State valid
     rx pathid: 0, tx pathid: 0x0
```

```
#sh bgp ipv6 unicast 2402:7800::/32
BGP routing table entry for 2402:7800::/32, version 1157916
Paths: (1 available, best #1, table default)
 Advertised to update-groups:
    2
 Refresh Epoch 15
 4826
   2402:7800:10:2::151 from 2402:7800:10:2::151 (114.31.194.12)
     Origin IGP, metric 0, localpref 100, valid, external, best
     Community: 4826:1000 4826:2050 4826:2110 4826:2540 4826:2900 4826:5203
     path 7F50B266CBD8 RPKI State not found
     rx pathid: 0, tx pathid: 0x0
```

RPKI Verification - JunOS

- Check the RTR session

```
>show validation session
Session                                 State Flaps   Uptime #IPv4/IPv6 records
202.125.96.46                            Up       75 09:20:59 40894/6747


>show validation session 202.125.96.46
Session                                 State Flaps   Uptime #IPv4/IPv6 records
202.125.96.46                            Up       75 09:21:18 40894/6747
```

RPKI Verification - JunOS

- Check the RPKI cache

```
>show validation database
RV database for instance master

Prefix                Origin-AS    Session                  State      Mismatch
1.9.0.0/16-24              4788 202.125.96.46               valid
1.9.12.0/24-24            65037 202.125.96.46               valid
1.9.21.0/24-24            24514 202.125.96.46               valid
1.9.23.0/24-24            65120 202.125.96.46               valid

----------
2001:200::/32-32           2500 202.125.96.46               valid
2001:200:136::/48-48       9367 202.125.96.46               valid
2001:200:900::/40-40       7660 202.125.96.46               valid
2001:200:8000::/35-35      4690 202.125.96.46               valid
2001:200:c000::/35-35     23634 202.125.96.46               valid
2001:200:e000::/35-35      7660 202.125.96.46               valid
```

- *Would have been nice if they had per AF!*

RPKI Verification - JunOS

- Can filter per origin ASN

```
>show validation database origin-autonomous-system 45192
RV database for instance master


Prefix                Origin-AS     Session                   State      Mismatch
202.125.97.0/24-24       45192      202.125.96.46              valid
203.176.189.0/24-24      45192      202.125.96.46              valid
2001:df2:ee01::/48-48    45192      202.125.96.46              valid


 IPv4 records: 2
 IPv6 records: 1
```

- *IOS should have something similar!*

Check routes - JunOS

```
>show route protocol bgp 202.144.128.0

inet.0: 693024 destinations, 693024 routes (693022 active, 0 holddown, 2
hidden)
+ = Active Route, - = Last Active, * = Both

202.144.128.0/20 *[BGP/170] 1w4d 21:03:04, MED 0, localpref 110, from
202.125.96.254
                      AS path: 4826 17660 I, validation-state: valid
                    >to 202.125.96.225 via ge-1/1/0.0
```

```
>show route protocol bgp 2001:201::/32

inet6.0: 93909 destinations, 93910 routes (93909 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

2001:201::/32     *[BGP/170] 21:18:14, MED 0, localpref 100, from
2001:df2:ee00::1
                      AS path: 65332 I, validation-state: unknown
                    >to fe80::dab1:90ff:fedc:fd07 via ge-1/1/0.0
```

Configuration - Reference Link

- **Cisco**
  - https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/irg-origin-as.pdf

- **Juniper**
  - https://www.juniper.net/documentation/en_US/junos/topics/topic-map/bgp-origin-as-validation.html

- **RIPE:**
  - https://www.ripe.net/manage-ips-and-asns/resource-management/certification/router-configuration

https://www.apnic.net/community/security/resource-certification/#routing


Questions