RPKI

Resource Public Key Infrastructure

SANOG XXV | 16-24 January, 2015 | Kandy, Sri Lanka

Fakrul Alam fakrul@bdhub.com http://bd.linkedin.com/in/fakrulalam https://twitter.com/rapappu

Target Audience

- Knowledge of Internet Routing(specially BGP)
- Familiar with any IRR Database
- No need to know Cryptography
- Basic knowledge of PKI(Public Key Infrastructure)

Agenda

- BGP / RPKI
- Configuration
- Hands-on Lab (Juniper)

BGP

BGP (AS)

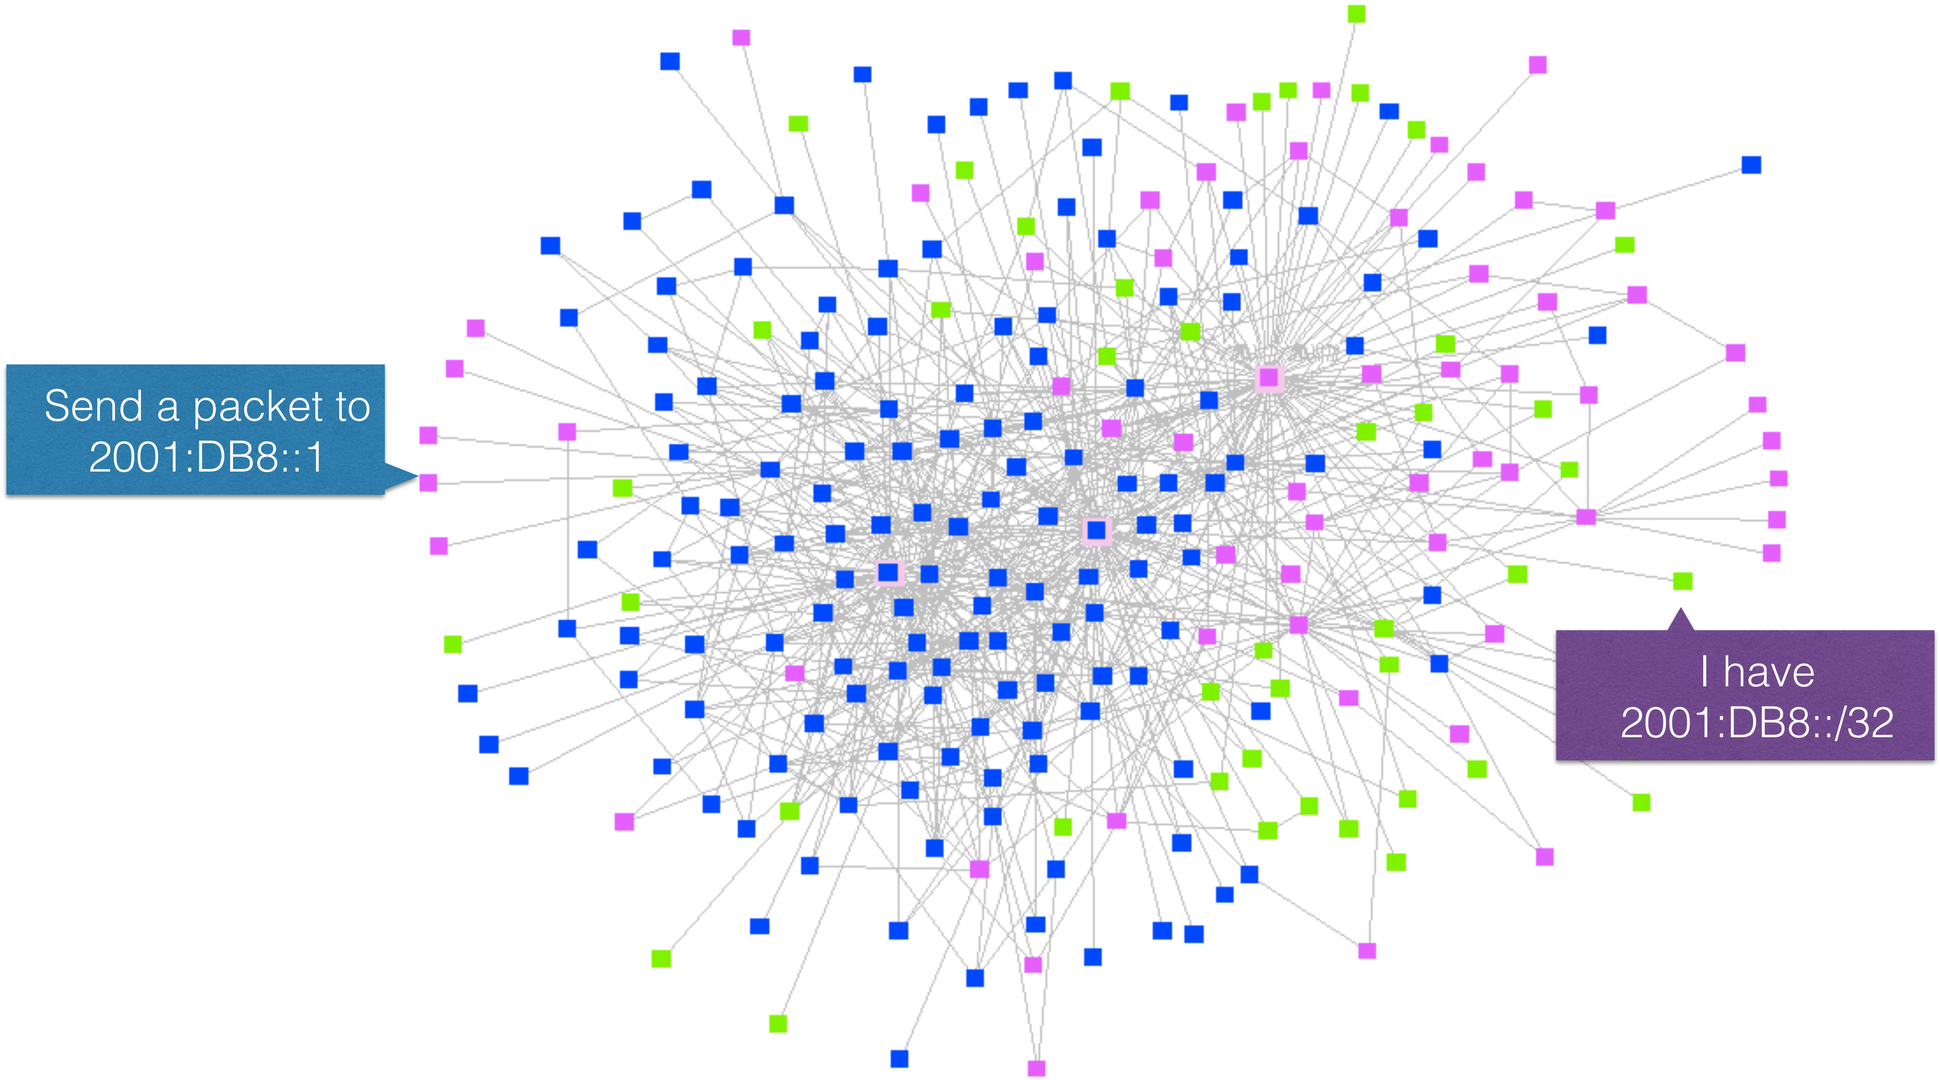

AS Path

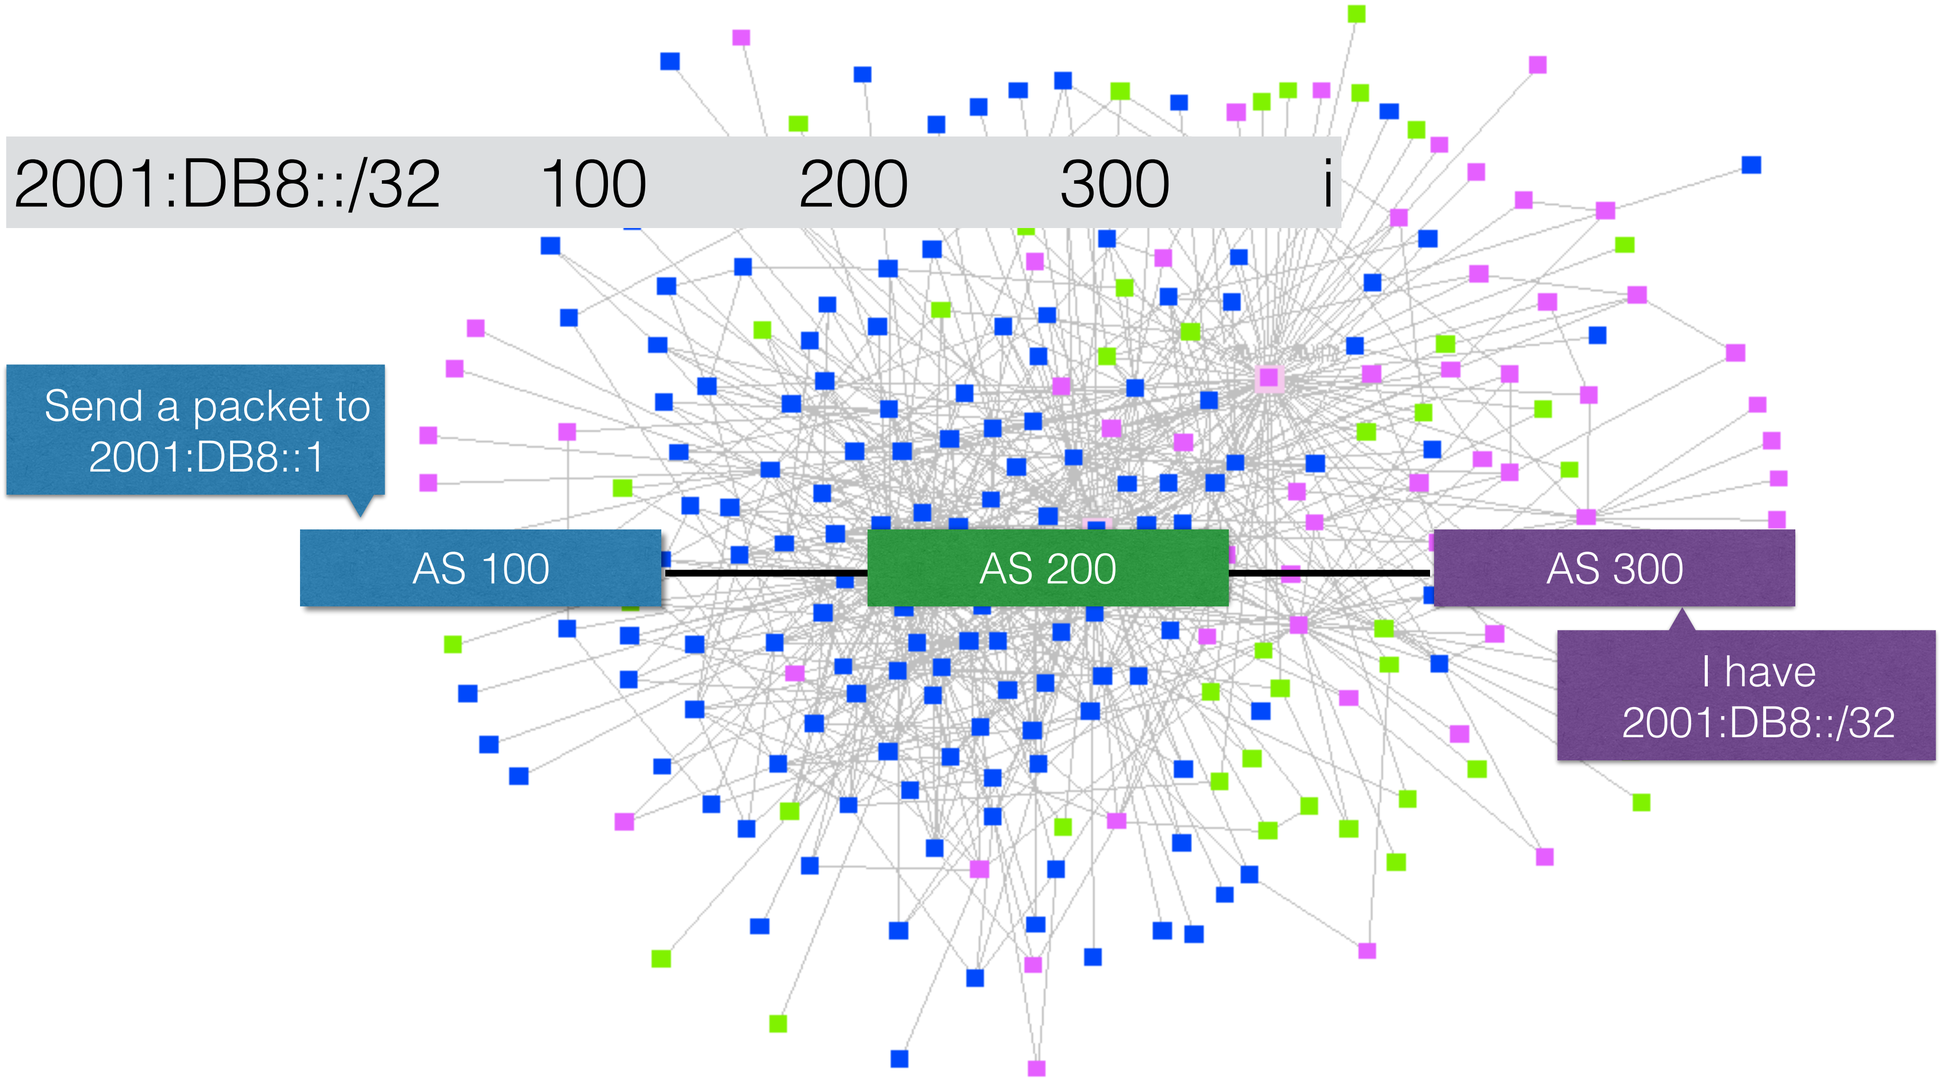
2001:DB8::/32 100 200 300 i
Send a packet to 2001:DB8::1
AS 100
AS 200
AS 300
I have 2001:DB8::/32

AS Path

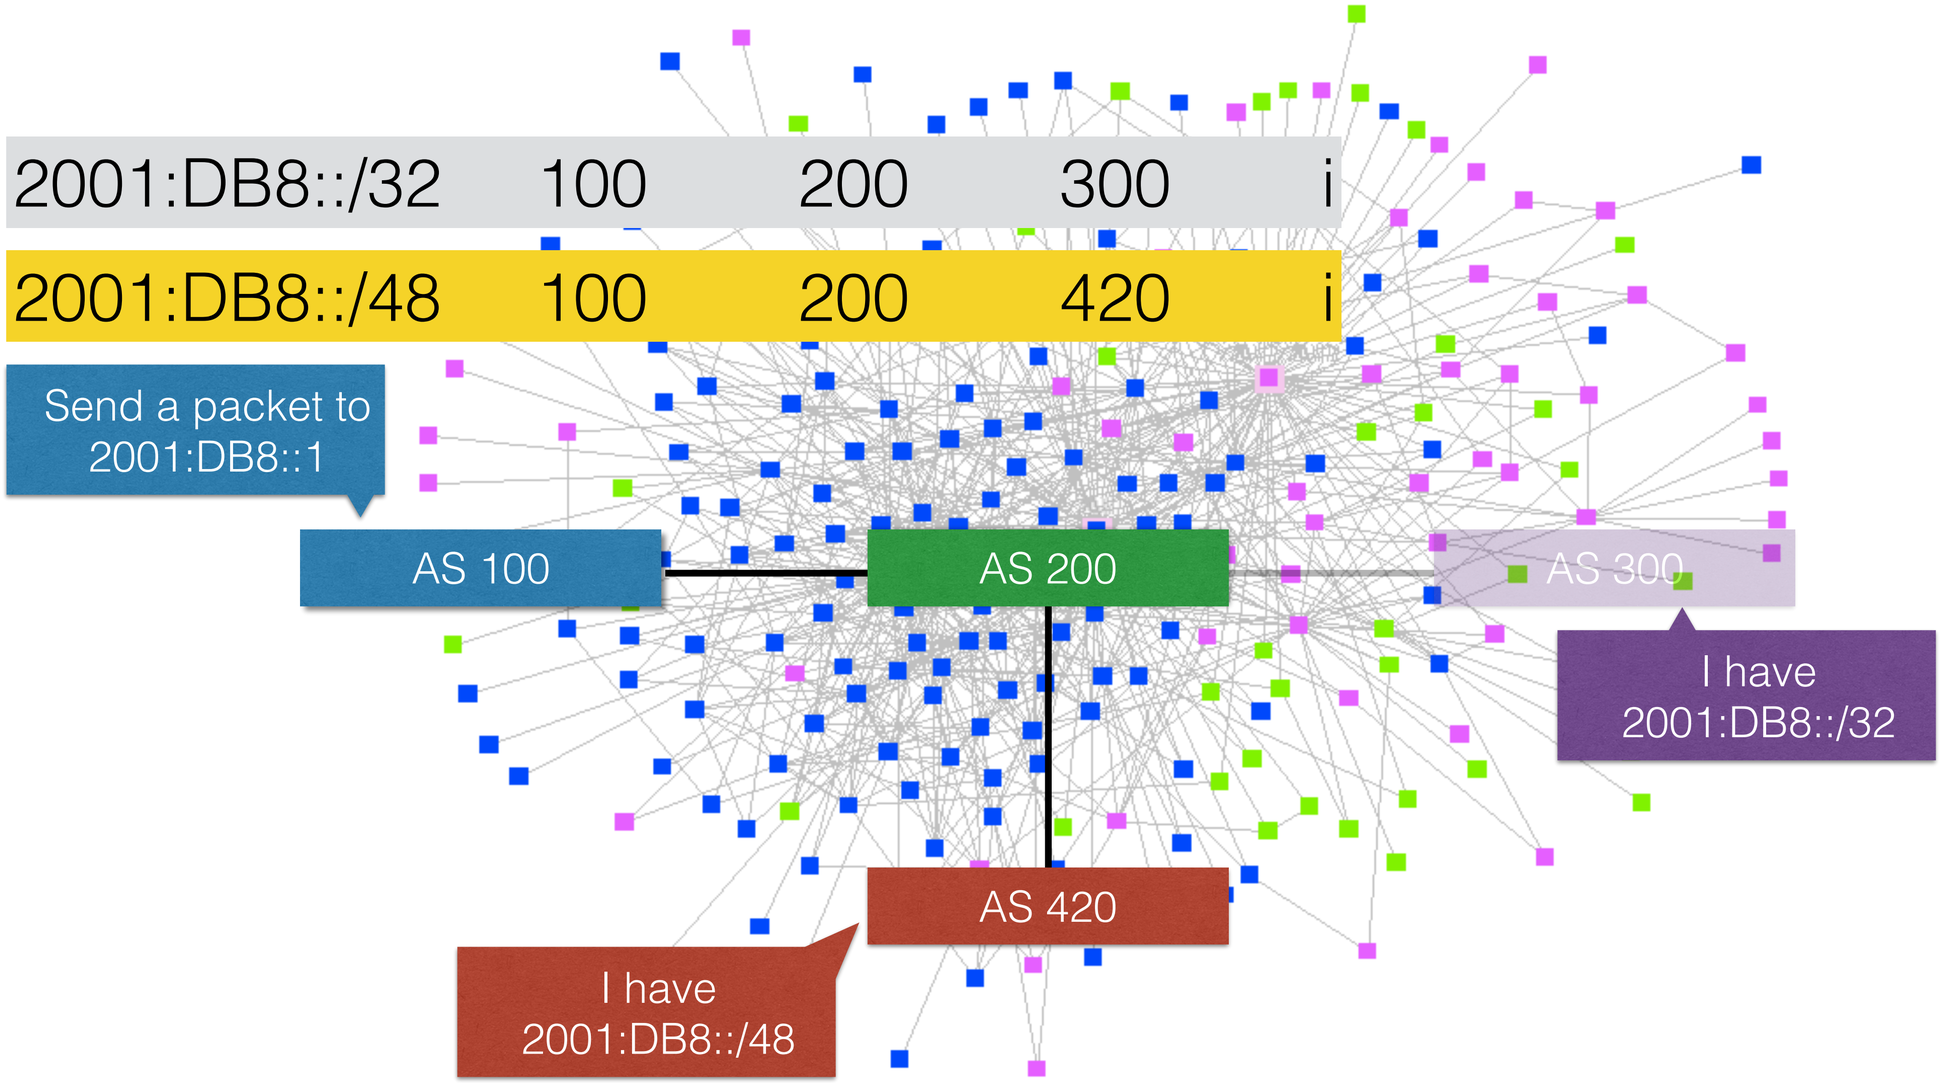
2001:DB8::/32 100 200 300 i
2001:DB8::/48 100 200 420 i
Send a packet to 2001:DB8::1
AS 100
AS 200
AS 300
I have 2001:DB8::/32
AS 420
I have 2001:DB8::/48

Historical Incident

- April 1997: The "AS 7007 incident" UU/Sprint for 2 days
- February 24, 2008: Pakistan's attempt to block YouTube access within their country takes down YouTube entirely.[6]
- November 11, 2008: The Brazilian ISP CTBC - Companhia de Telecomunicações do Brasil Central leaked their internal table into the global BGP table.
- April 8, 2010: China Telecom originated 37,000 prefixes not belonging to them in 15 minutes, causing massive outage of services globally.

source : http://en.wikipedia.org/wiki/IP_hijacking

Historical Incident

- For theory of positivity lets call all these as Mis-Origination
- Traffic Hijacking or Prefix Hijacking assumes Negative intent

Current Trend

- Filtering limited to the edges facing the customer
- Filters on peering and transit sessions are often too complex or take too many resources
- Check prefix before announcing it

Filter Where?

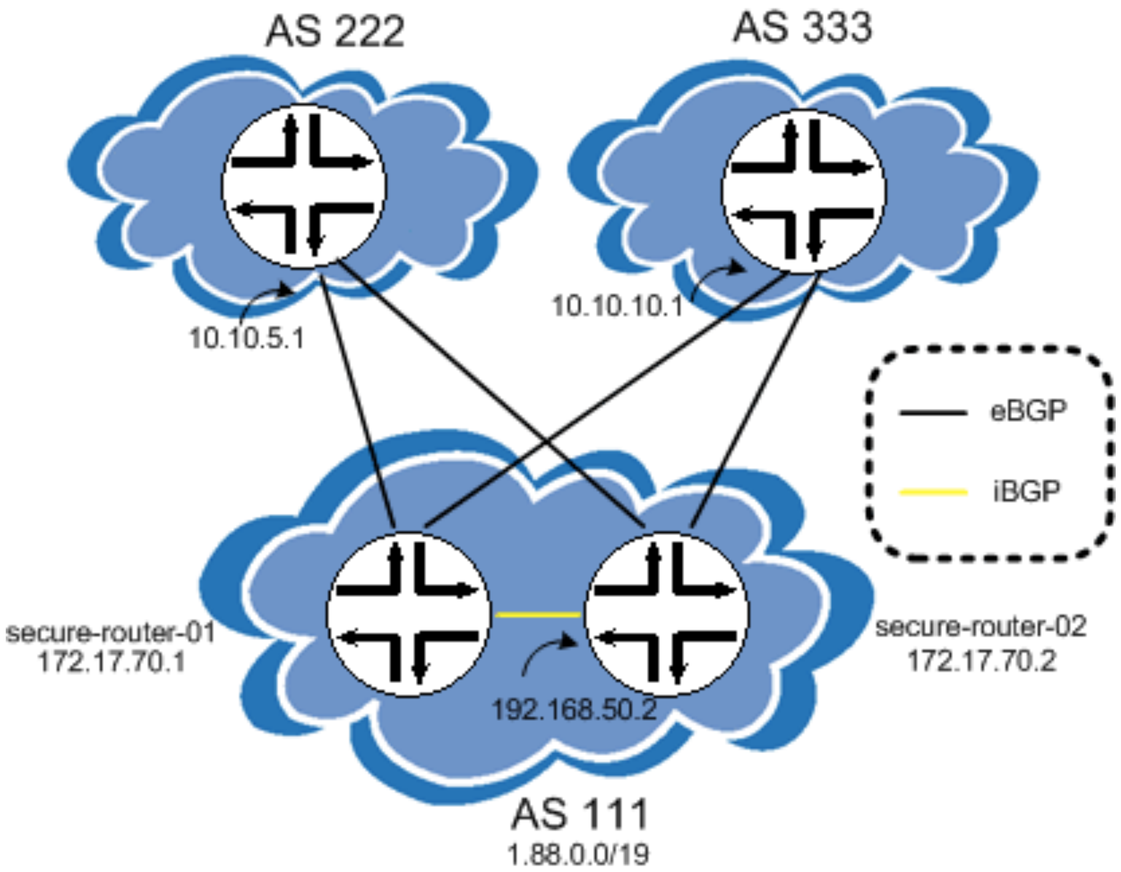

- Secure BGP Templates
  - http://www.cymru.com/gillsr/documents/junos-bgp-template.htm
  - https://www.team-cymru.org/ReadingRoom/Templates/secure-bgp-template.html

Internet Registry (IR)

- Maintains Internet Resources such as IP addresses and ASNs, and publish the registration information
  - Allocations for Local Internet Registries
  - Assignments for end-users
- APNIC is the Regional Internet Registry(RIR) in the Asia Pacific region
  - National Internet Registry(NIR) exists in several economies

The Eco-System

Internet Assigned Numbers Authority

Regional IR (RIR)

National IR (NIR)

Internet Service Provider

End User

Internet Routing Registry

- Maintains routing policy database
  - RADB is the most popular service, though some RIRs also provide similar services
  - Routing policy information is expressed in a series of objects
  - On RADB, a registered user can register any object
- route and route6 objects are used to indicate route origination
  - Prefix and origin AS

Still not enough
IRR is useful, but it’s not perfect

RPKI

Resource Pubic Key Infrastructure

IP Address & AS Numbers

Digital Certificate

RPKI Deployment

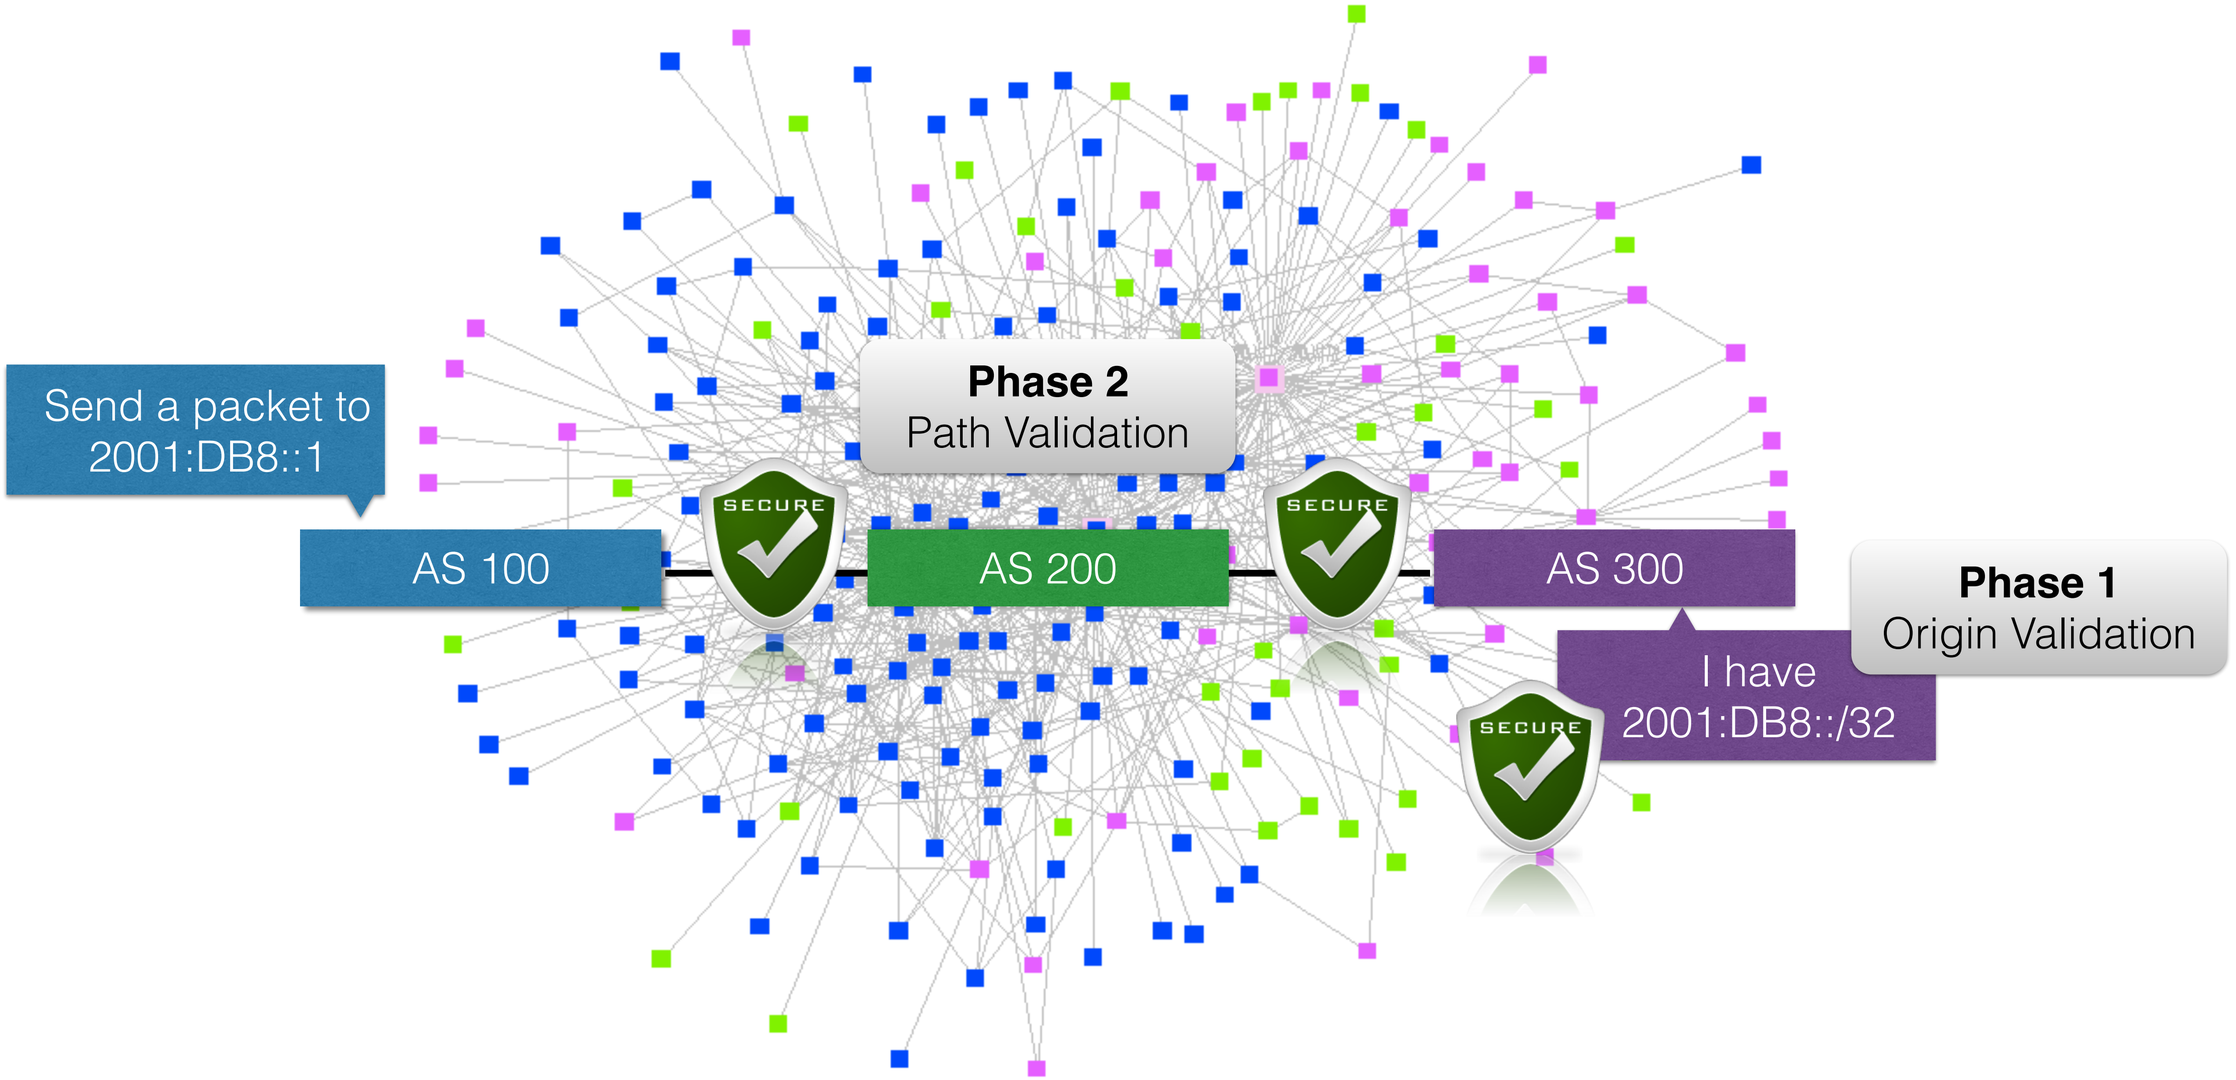

Goals of RPKI

- Able to authoritatively prove who owns an IP Prefix and what AS(s) may Announce It
  - Reducing routing leaks
  - Attaching digital certificates to network resources (AS Number & IP Address)
- Prefix Ownership Follows the Allocation Hierarchy IANA, RIRs, ISPs, …

RPKI Implementation

- Two RPKI implementation type
  - Delegated: Each participating node becomes a CA and runs their own RPKI repository, delegated by the parent CA.
  - Hosted: The RIR runs the CA functionality for interested participants.

RPKI Origin Validation

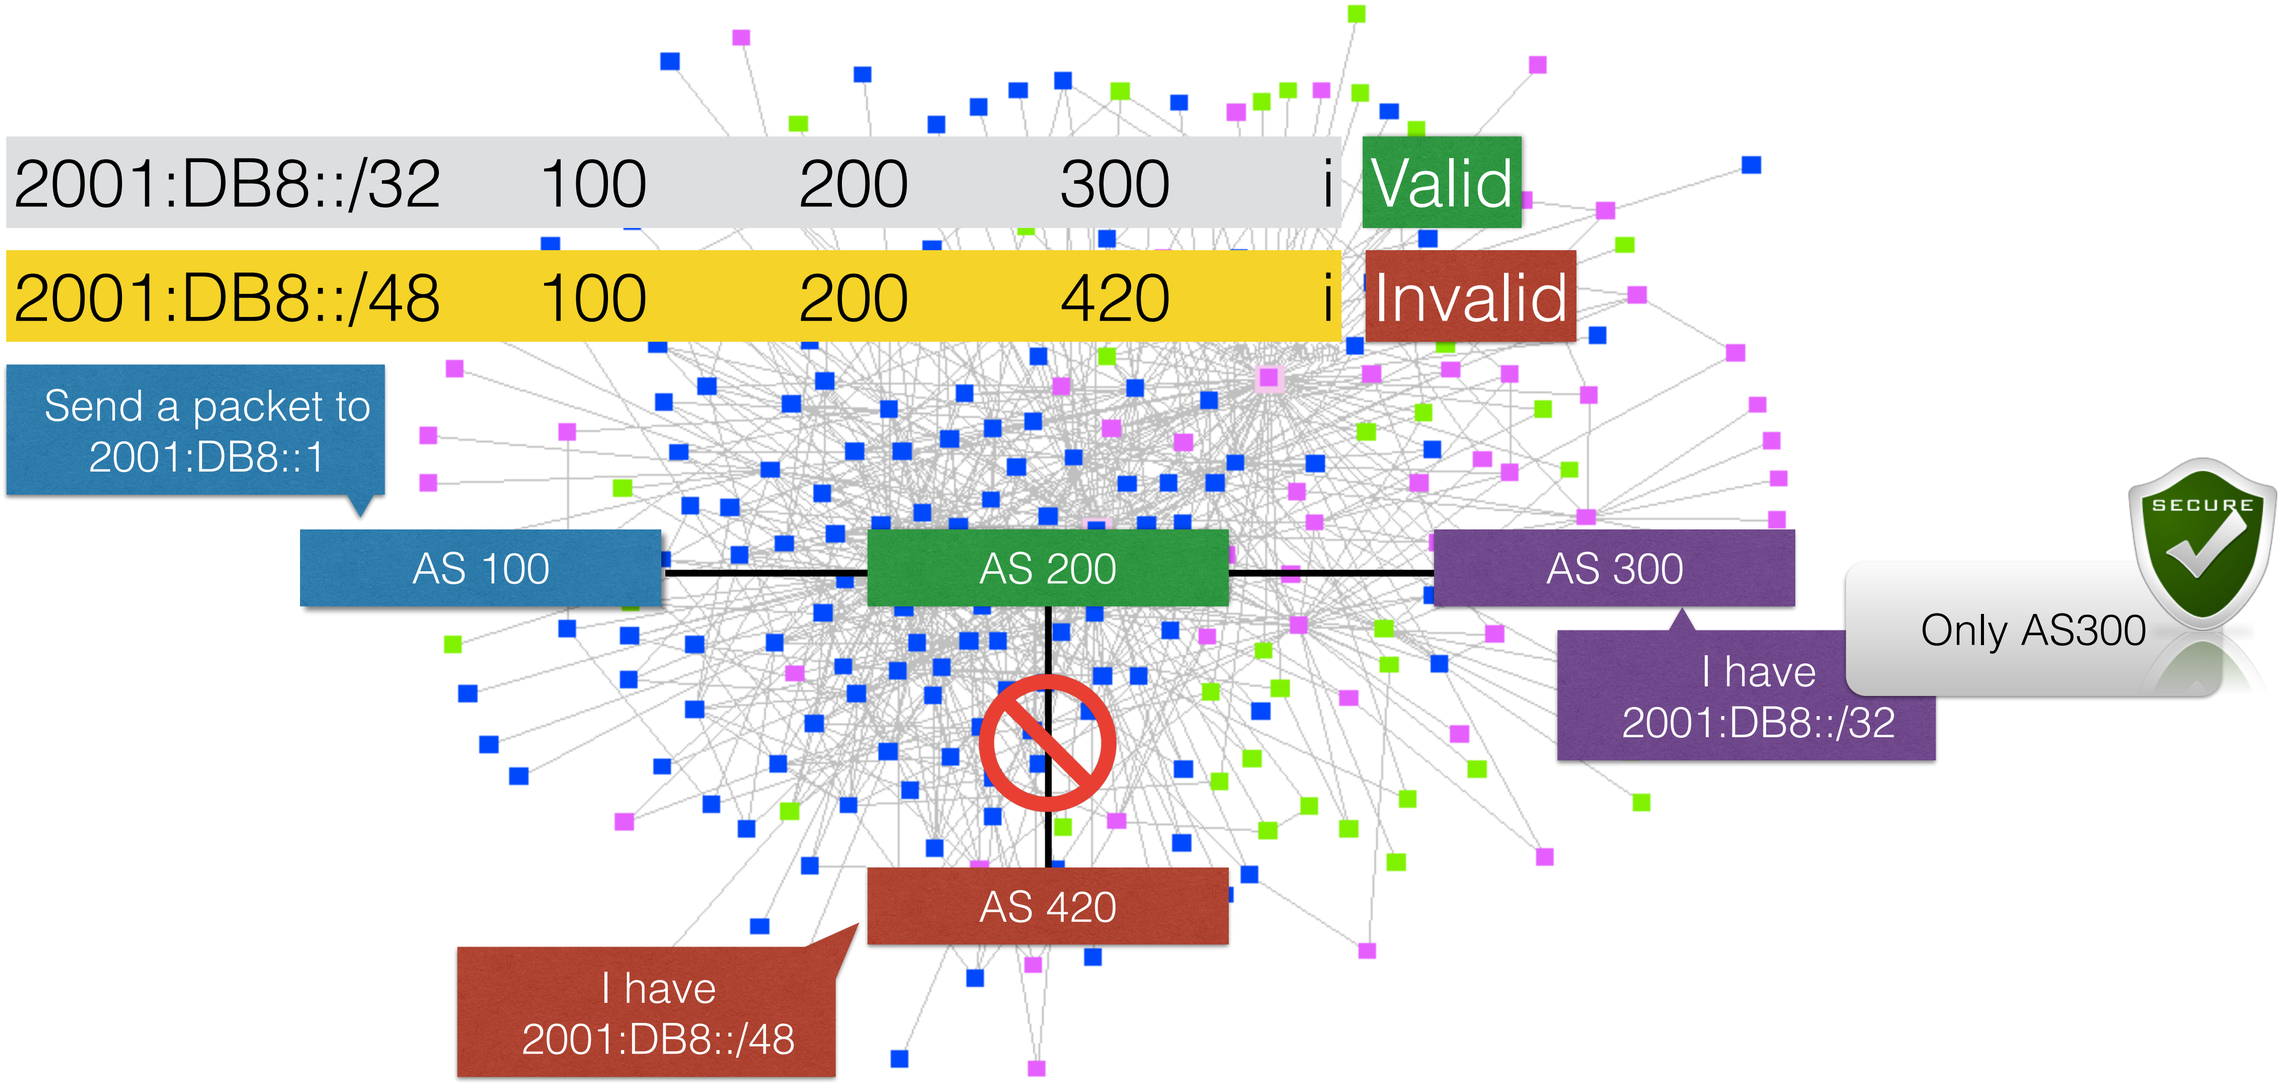

RPKI Building Blocks

- Trust Anchors (RIR's)
- Route Origination Authorizations (ROA)
- Validators

Let’s discuss these building blocks in details

PKI & Trust Anchors

Public Key Concept

- **Private key**: This key must be known only by its owner.
- **Public key**: This key is known to everyone (it is public)
- **Relation between both keys**: What one key encrypts, the other one decrypts, and vice versa. That means that if you encrypt something with my public key (which you would know, because it's public :-), I would need my private key to decrypt the message.
- Same alike http with SSL aka https

X.509 Certificates 3779 EXT

Signed by Parent's Private Key

Certificates are X.509 certificates that conform to the PKIX profile [PKIX]. They also contain an extension field that lists a collection of IP resources (IPv4 addresses, IPv6 addresses and AS Numbers) [RFC3779]

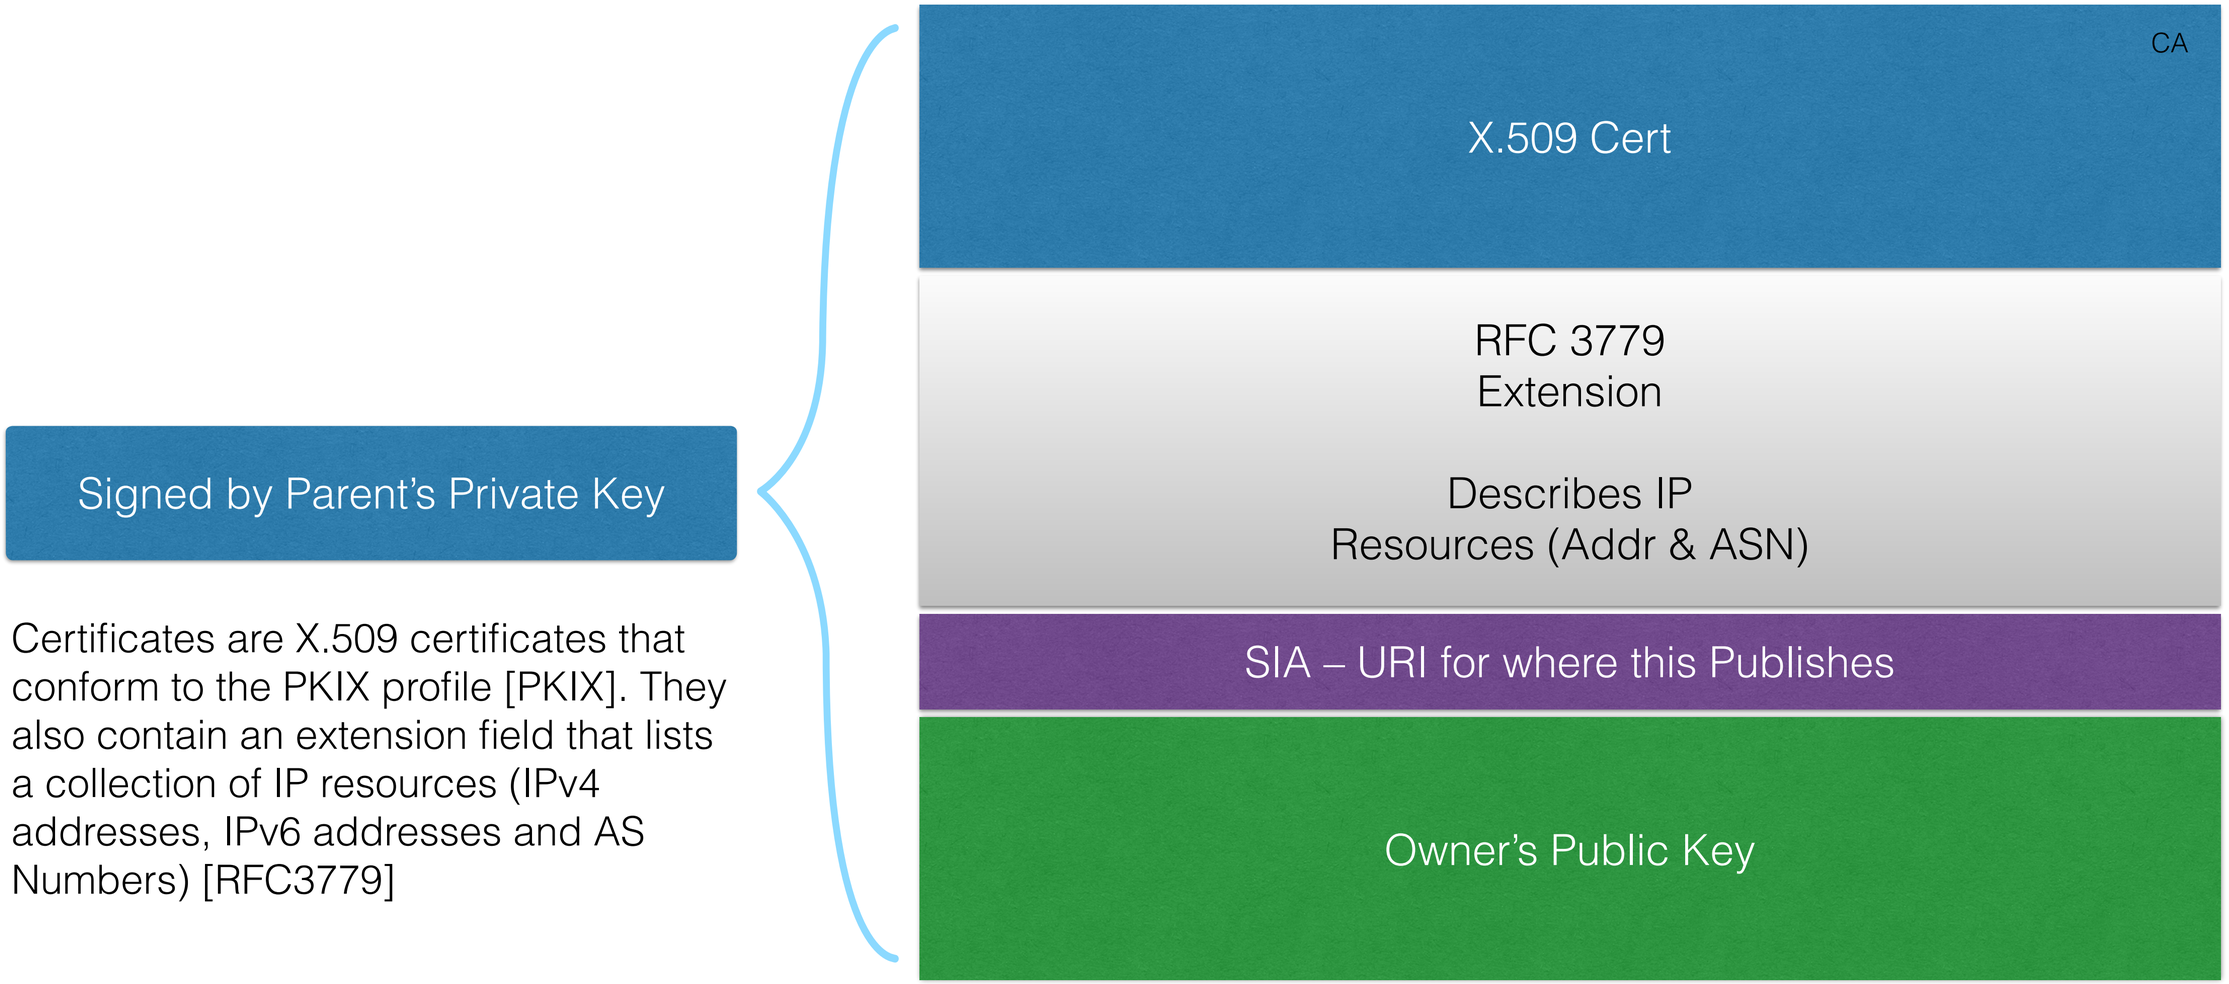

Trust Anchor

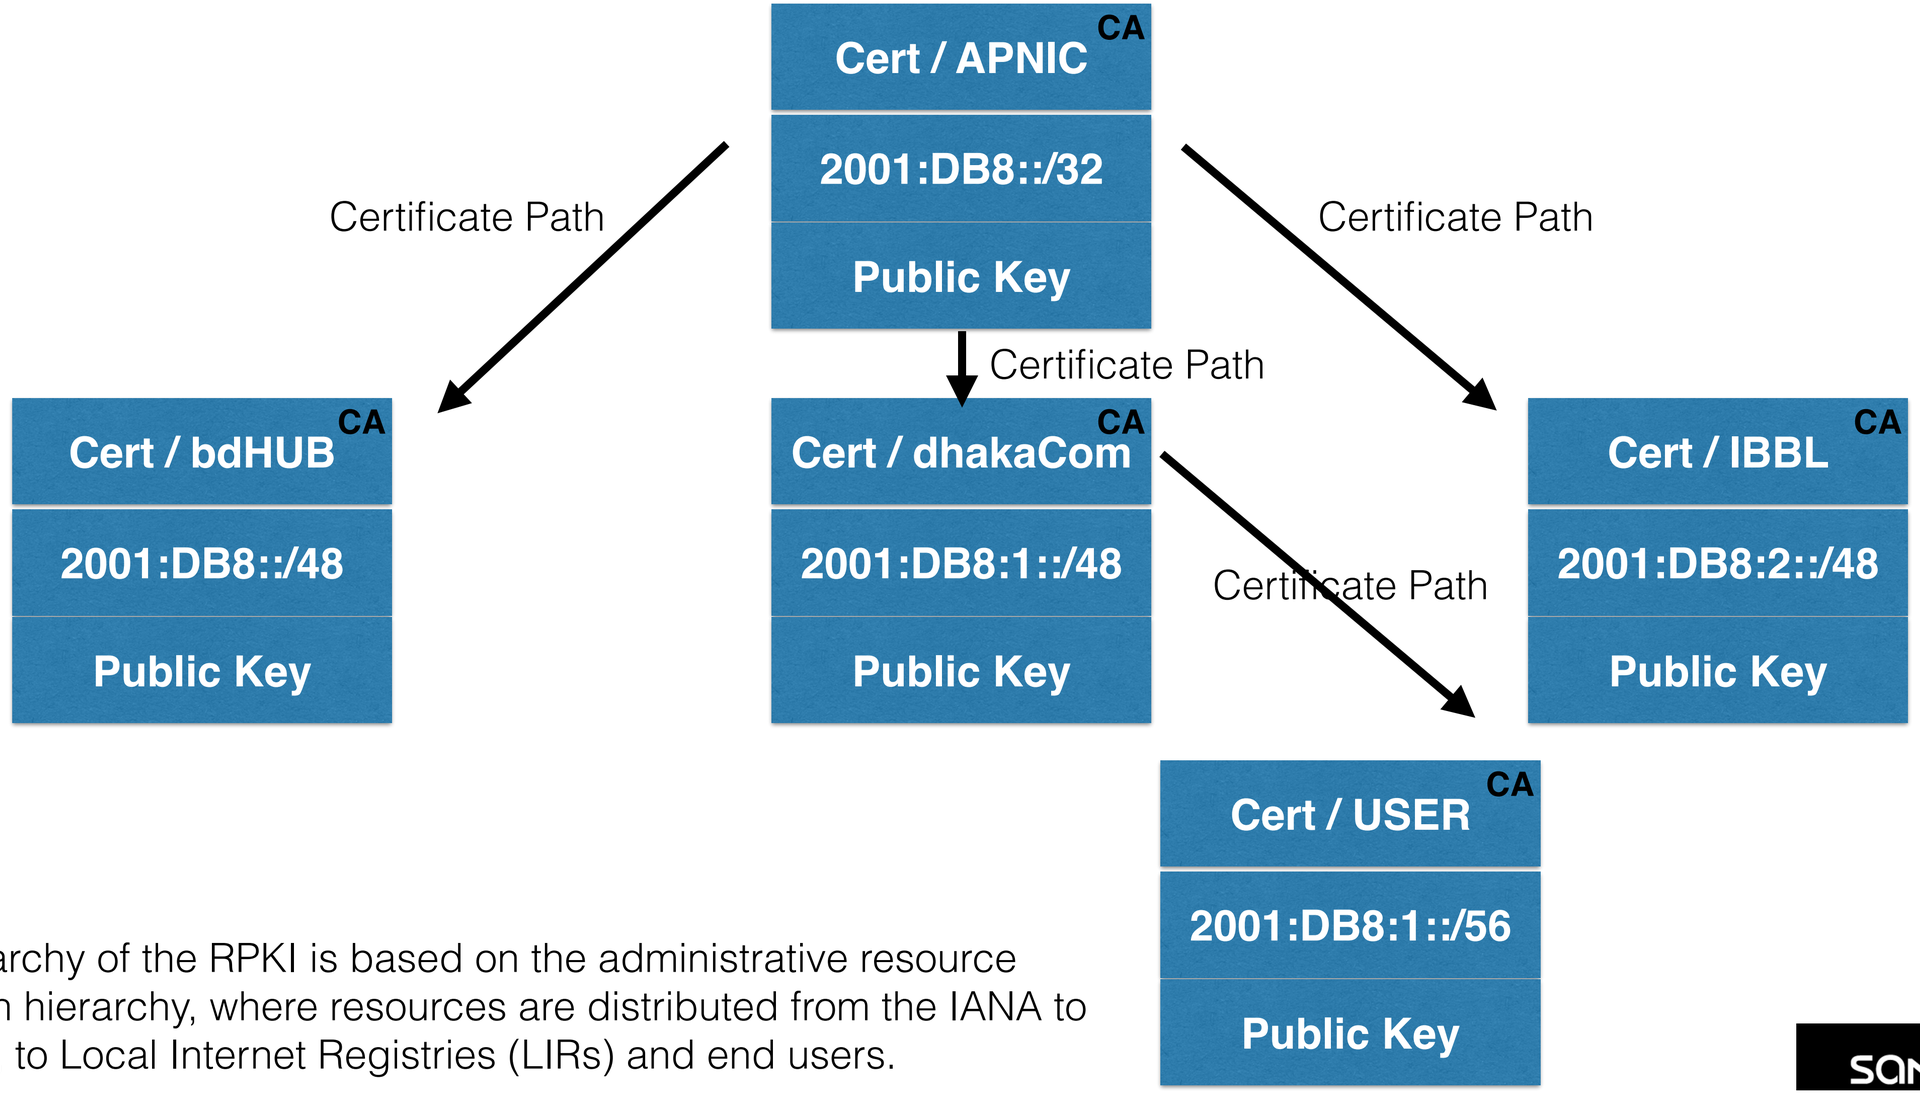

The hierarchy of the RPKI is based on the administrative resource allocation hierarchy, where resources are distributed from the IANA to the RIRs, to Local Internet Registries (LIRs) and end users.

Trust Anchor Locator (TALs)

- In cryptographic systems with hierarchical structure, a Trust anchor is an authoritative entity for which trust is assumed and not derived.
- In X.509 architecture, a root certificate would be the trust anchor from which whole chain of trust is derived. The trust anchor must be in possession of the trusting party beforehand to make any further certificate path validation possible.
- RPKI uses Internet Assigned Numbers Authority(IANA) as the trust anchor, and Regional Internet Registries(RIR) as immediately subordinate nodes to that anchor.

PKI in IRR

- The RIRs hold a self-signed root certificate for all the resources that they have in the registry
  - They are the trust anchor for the system
- That root certificate is used to sign a certificate that lists your resources
- You can issue child certificates for those resources to your customers
  - When making assignments or sub allocations

ROA
Route Origin Authorizations

Route Origination Authorizations (ROA)

- Next to the prefix and the ASN which is allowed to announce it, the ROA contains:
  - A minimum prefix length
  - A maximum prefix length
  - An expiry date
  - Origin ASN
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

Validators

Origin Validation

- Router gets ROA information from the RPKI Cache
  - RPKI verification is done by the RPKI Cache
- The BGP process will check each announcement with the ROA information and label the prefix

Result of Check

- **Valid** – Indicates that the prefix and AS pair are found in the database.
- **Invalid** – Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database.
- **Not Found / Unknown**– Indicates that the prefix is not among the prefixes or prefix ranges in the database.

Valid > Unknown > Invalid

ROA Example

Prefix: 10.0.0.0/16
ASN: 65420

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

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

Local Policy

- You can define your policy based on the outcomes
  - Do nothing
  - Just logging
  - Label BGP communities
  - Modify preference values
  - Rejecting the announcement

RPKI Support in Routers

- The RPKI-RTR Protocol is an IETF Internet Draft
- Production Cisco Support:
  - ASR1000, 7600, ASR903 and ASR901 in releases 15.2(1)S or XE 3.5
  - Cisco Early Field Trial (EFT):
  - ASR9000, CRS1, CRS3 and c12K (IOS-XR 4.3.2)
- Juniper has support since version 12.2
- Quagga has support through BGP-SRX

RPKI Caveats

- When RTR session goes down, the RPKI status will be not found for all the bgp route after a while
  - Invalid => not found
  - we need several RTR sessions or care your filtering policy
- In case of the router reload, which one is faster, receiving ROAs or receiving BGP routes?
  - If receiving BGP is match faster than ROA, the router propagate the invalid route to others
  - We need to put our Cache validator within our IGP scope

Who do we trust?

- Can we trust the *IR for hosting our Private Keys?

Two digital certificates have been mistakenly issued in Microsoft's name that could be used by virus writers to fool people into running harmful programs, the software giant warned Thursday.

According to Microsoft, someone posing as a Microsoft employee tricked VeriSign, which hands out so-called digital signatures, into issuing the two certificates in the software giant's name on Jan. 30 and Jan. 31.

Such certificates are critical for businesses and consumers who download patches, updates and other pieces of software from the Internet, because they verify that the software is being supplied from a particular company, such as Microsoft.

RPKI Further Reading

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

RPKI Configuration

Topology for Origin Validation

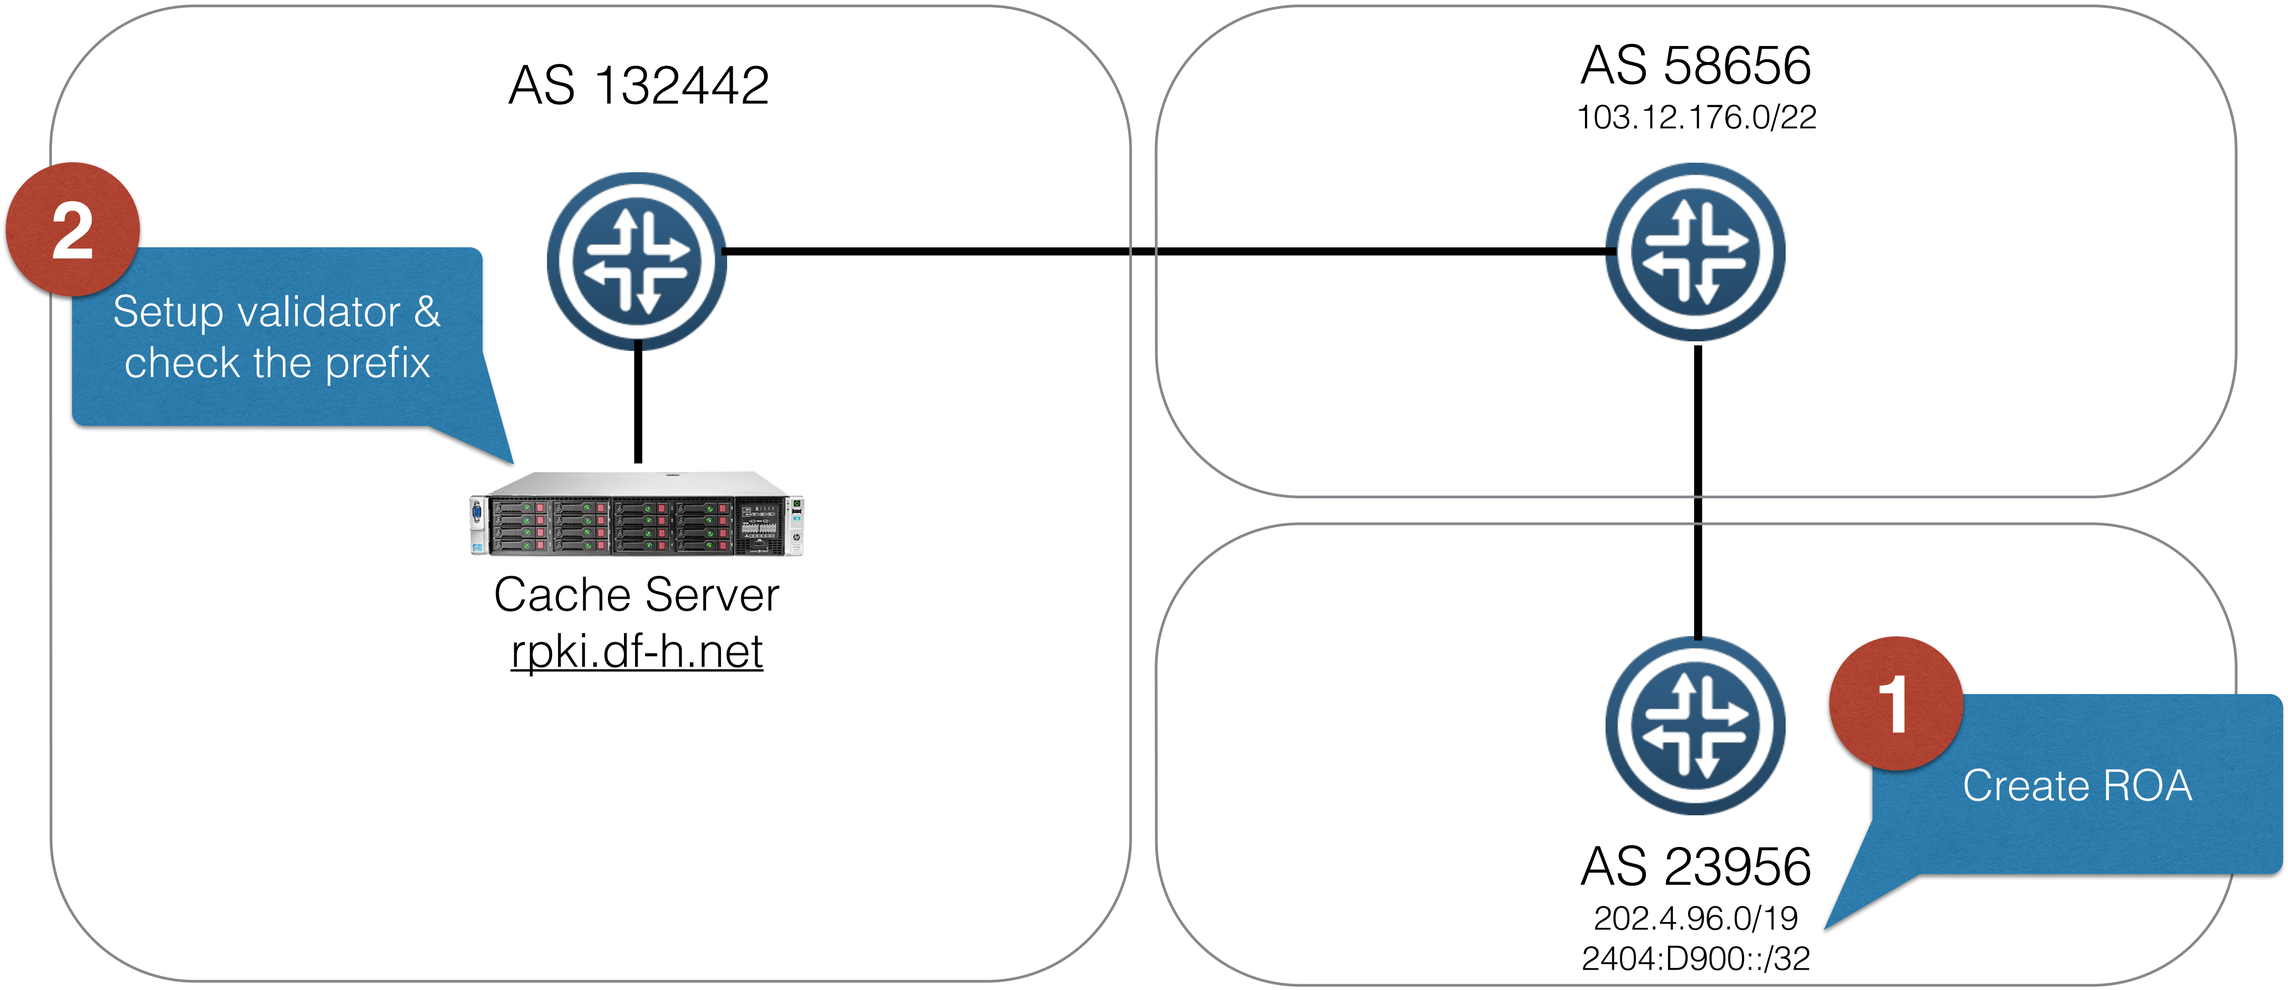

Phase I - Publishing ROA

- Login to your MyAPNIC portal
- Required valid certificate
- Go to Resources > Certification Tab

Phase I - Publishing ROA

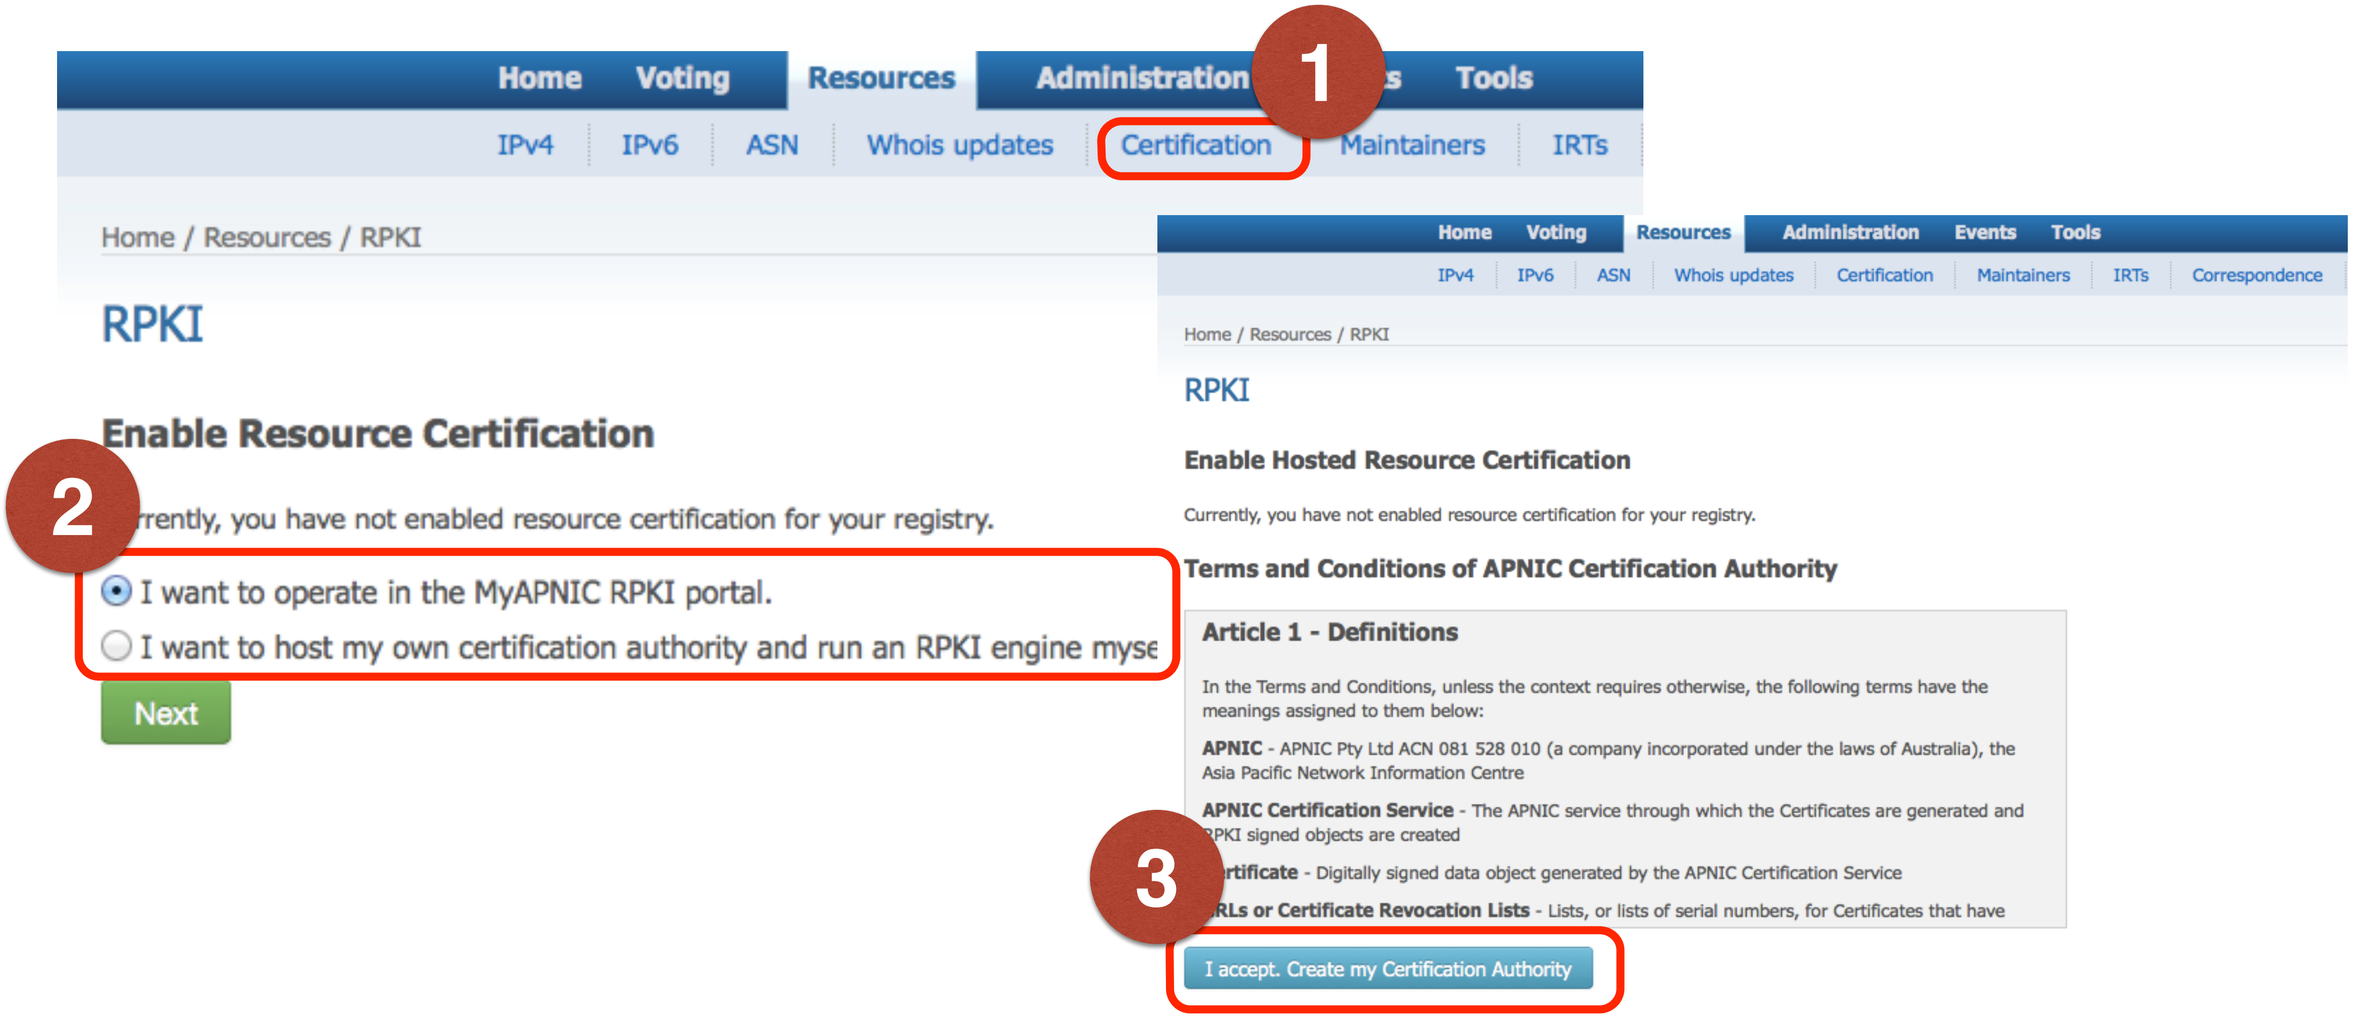

Phase I - Publishing ROA

RPKI

BGP Route Validity

All

Items per page 10

Search by AS or IP...

	Origin AS	Prefix
☐	23956	118.179.192.0/19
☐	23956	202.4.96.0/19
☐	23956	2405:7600::/32

- Show available prefix for which you can create ROA

Phase I - Publishing ROA - IPv4

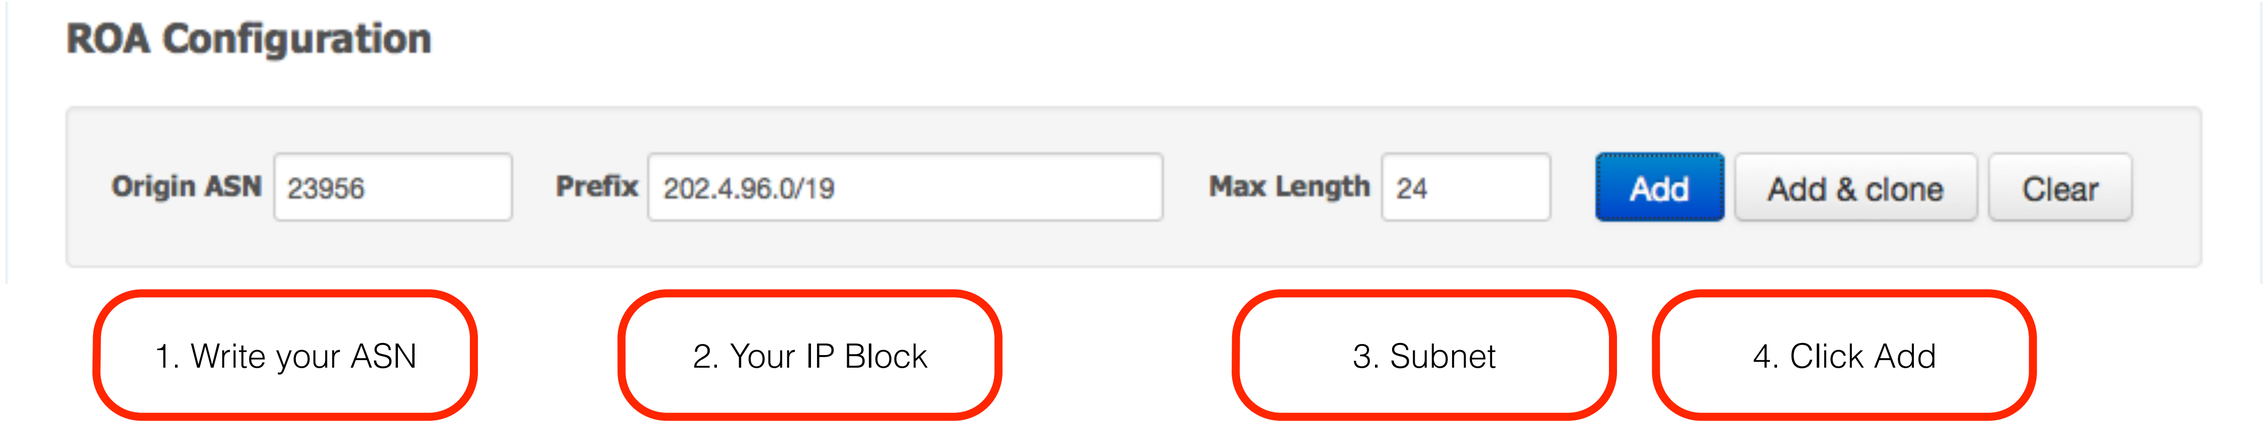

- Create ROA for smaller block.

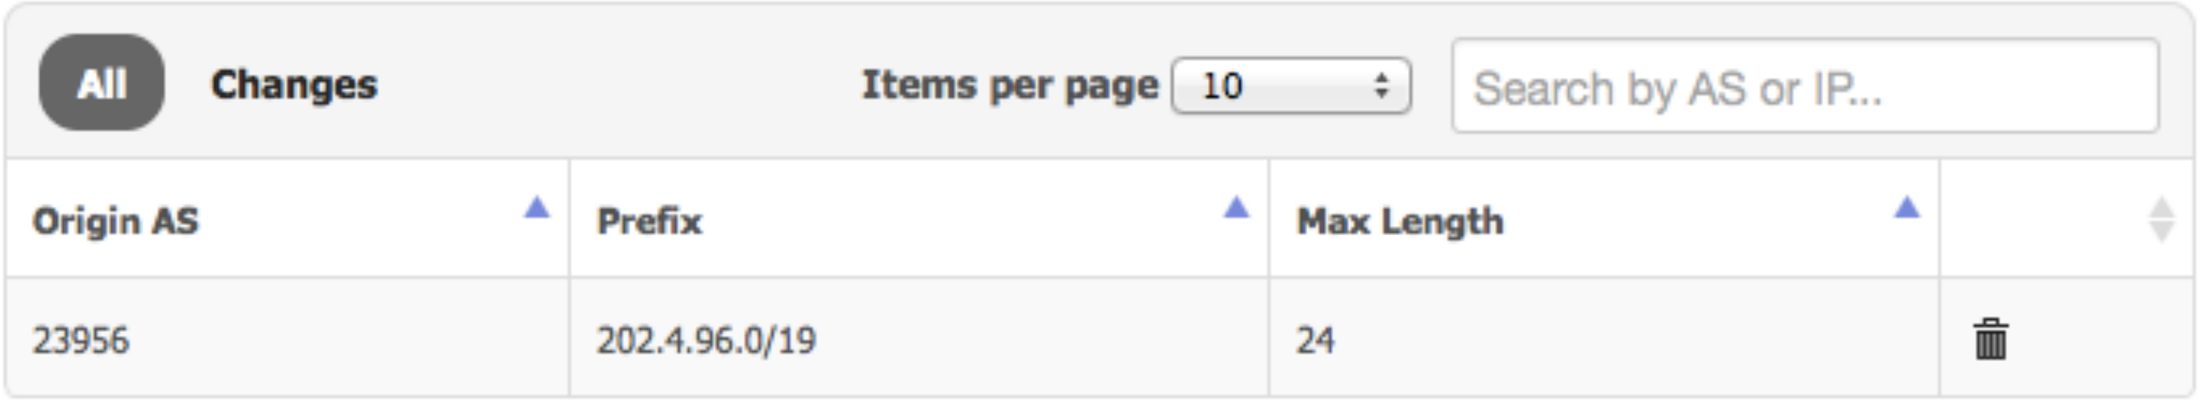

Phase I - Publishing ROA - IPv6

ROA Configuration

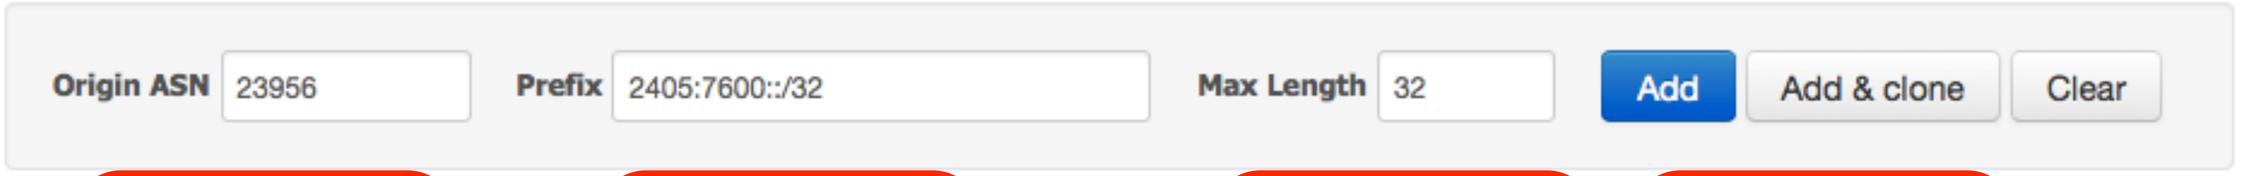

1. Write your ASN
2. Your IP Block
3. Subnet
4. Click Add

- ROA for your IPv6 prefix

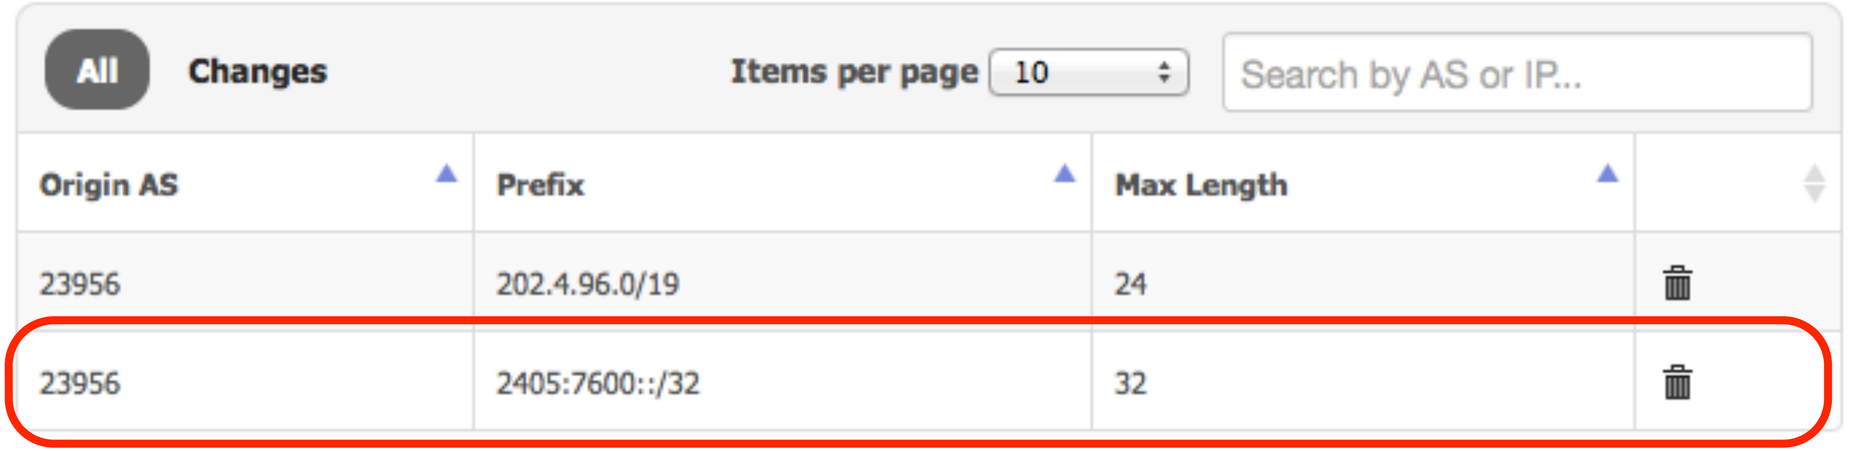

Phase I - Check your ROA

```
# whois -h whois.bgpmon.net 202.4.96.0/24


Prefix:              202.4.96.0/24
Prefix description:  APT (Dhakacom)
Country code:        BD
Origin AS:           23956
Origin AS Name:      DHAKACOM-BD-AS dhakaCom Limited,BD
RPKI status:         ROA validation successful
First seen:          2013-12-23
Last seen:           2014-07-20
Seen by #peers:      203
```

Phase I - Check your ROA

```
# whois -h whois.bgpmon.net " --roa 23956 202.4.96.0/24"

0 - Valid

------------------------

ROA Details

------------------------

Origin ASN:       AS23956

Not valid Before: 2014-07-20 15:20:10

Not valid After:  2014-12-30 00:00:00  Expires in 161d12h52m42s

Trust Anchor:     rpki.apnic.net

Prefixes:         202.4.96.0/19 (max length /24)

                  2405:7600::/32 (max length /32)
```

Phase II - RPKI Validator

- Download RPKI Validator

http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources

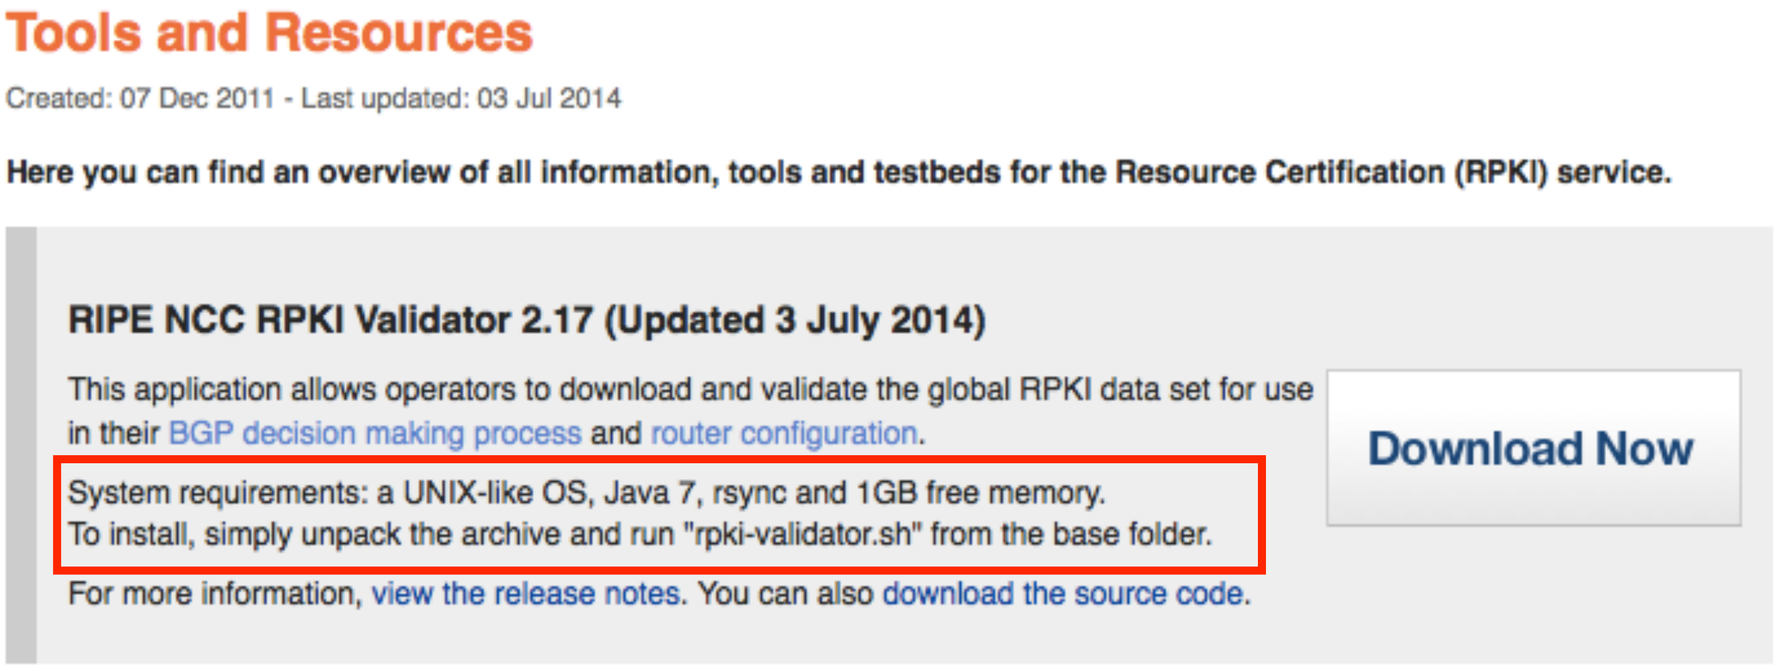

Phase II - RPKI Validator

```
# tar -zxvf rpki-validator-app-2.17-dist.tar.gz

# cd rpki-validator-app-2.17

# ./rpki-validator.sh start
```

Phase II - RPKI Validator

http://ip_address:8080

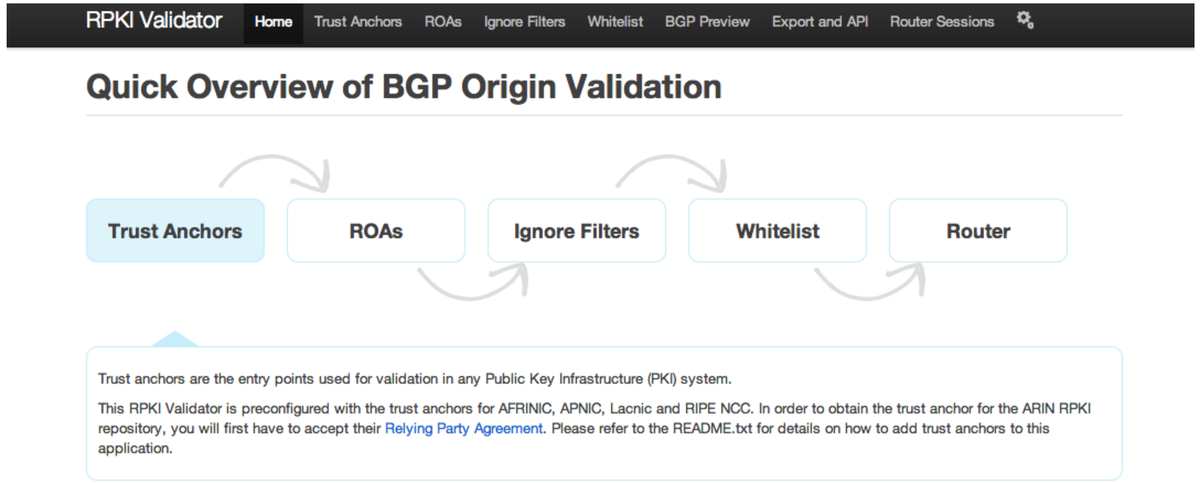

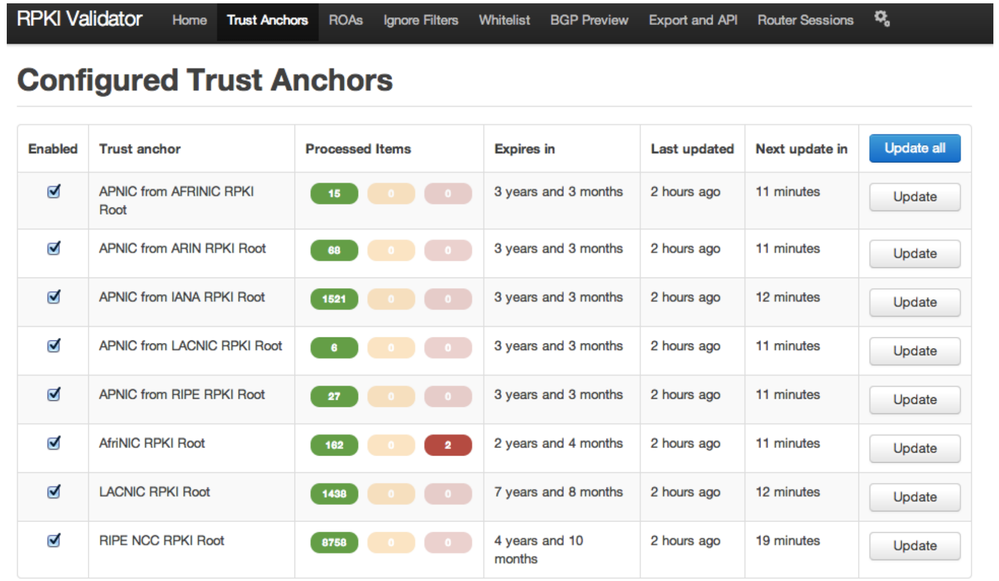

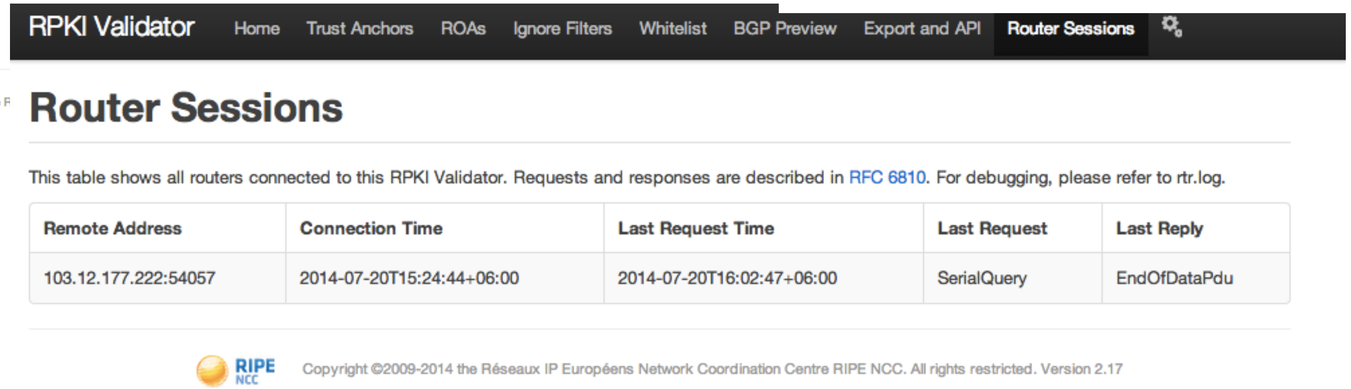

Phase III - Router Configuration (Juniper)

1. Establish session with RPKI Validator

```
routing-options {
    validation {
        group RPKI {
            session 103.21.75.10 {
                refresh-time 120;
                hold-time 180;
                port 8282;
                local-address 103.12.75.1;
            }
        }
    }
}
```

Phase III - Router Configuration (Juniper)

2. Configure policy to tag valid ROA

```
policy-options {
    policy-statement route-validation {
        term valid {
            from {
                protocol bgp;
                validation-database valid;
            }
            then {
                validation-state valid;
                accept;
            }
        }
    }
}
```

Phase III - Router Configuration (Juniper)

3. Push policy to the BGP neighbor

```
protocols {
    bgp {
        log-updown;
        import route-validation;
        group EBGP {
            type external;
        |
        | other cofigurations
        |
        }
    }
}
```

Check your prefix

```
fakrul@rpki-test> show route protocol bgp 202.4.96.0/24

inet.0: 506658 destinations, 506659 routes (506656 active, 0 holddown, 2
hidden)

+ = Active Route, - = Last Active, * = Both

202.4.96.0/24          *[BGP/170] 01:42:11, localpref 100
                          AS path: 58656 23956 I, validation-state: valid
                        > to 103.12.177.221 via ge-1/0/9.0
```

Command

#show validation session

```
fakrul@rpki-test> show validation session

Session                                 State    Flaps     Uptime #IPv4/IPv6 records

103.21.75.10                            Up           0 1d 09:33:54 9728/1431
```

Command

#show validation statistics

```
fakrul@rpki-test> Total RV records: 13529

Total Replication RV records: 13529

  Prefix entries: 13050

  Origin-AS entries: 13529

Memory utilization: 2626782 bytes

Policy origin-validation requests: 0

  Valid: 0

  Invalid: 0

  Unknown: 0

BGP import policy reevaluation notifications: 37818

  inet.0, 37818

  inet6.0, 0
```

Command

#show validation database

```
fakrul@rpki-test> show validation database

RV database for instance master

Prefix                  Origin-AS Session                                   State   Mismatch

2.0.0.0/12-16                3215 202.4.96.100                              valid

2.0.0.0/16-16                3215 202.4.96.100                              valid

2.1.0.0/16-16                3215 202.4.96.100                              valid

2.2.0.0/16-16                3215 202.4.96.100                              valid

2.3.0.0/16-16                3215 202.4.96.100                              valid

2.4.0.0/16-16                3215 202.4.96.100                              valid

2.5.0.0/16-16                3215 202.4.96.100                              valid

2.6.0.0/16-16                3215 202.4.96.100                              valid
```

Command

#show route protocol bgp validation-state valid

```
fakrul@rpki-test> show route protocol bgp validation-state valid


inet.0: 506561 destinations, 506562 routes (506559 active, 0 holddown, 2 hidden)

+ = Active Route, - = Last Active, * = Both


2.0.0.0/16          *[BGP/170] 1d 10:26:39, localpref 100

                       AS path: 58656 6453 5511 3215 I, validation-state: valid

                     > to 103.12.177.221 via ge-1/0/9.0

2.1.0.0/16          *[BGP/170] 1d 10:26:39, localpref 100

                       AS path: 58656 6453 5511 3215 I, validation-state: valid
```

!Caution!

- Make sure that your router IOS is bug free for RPKI; other wise….

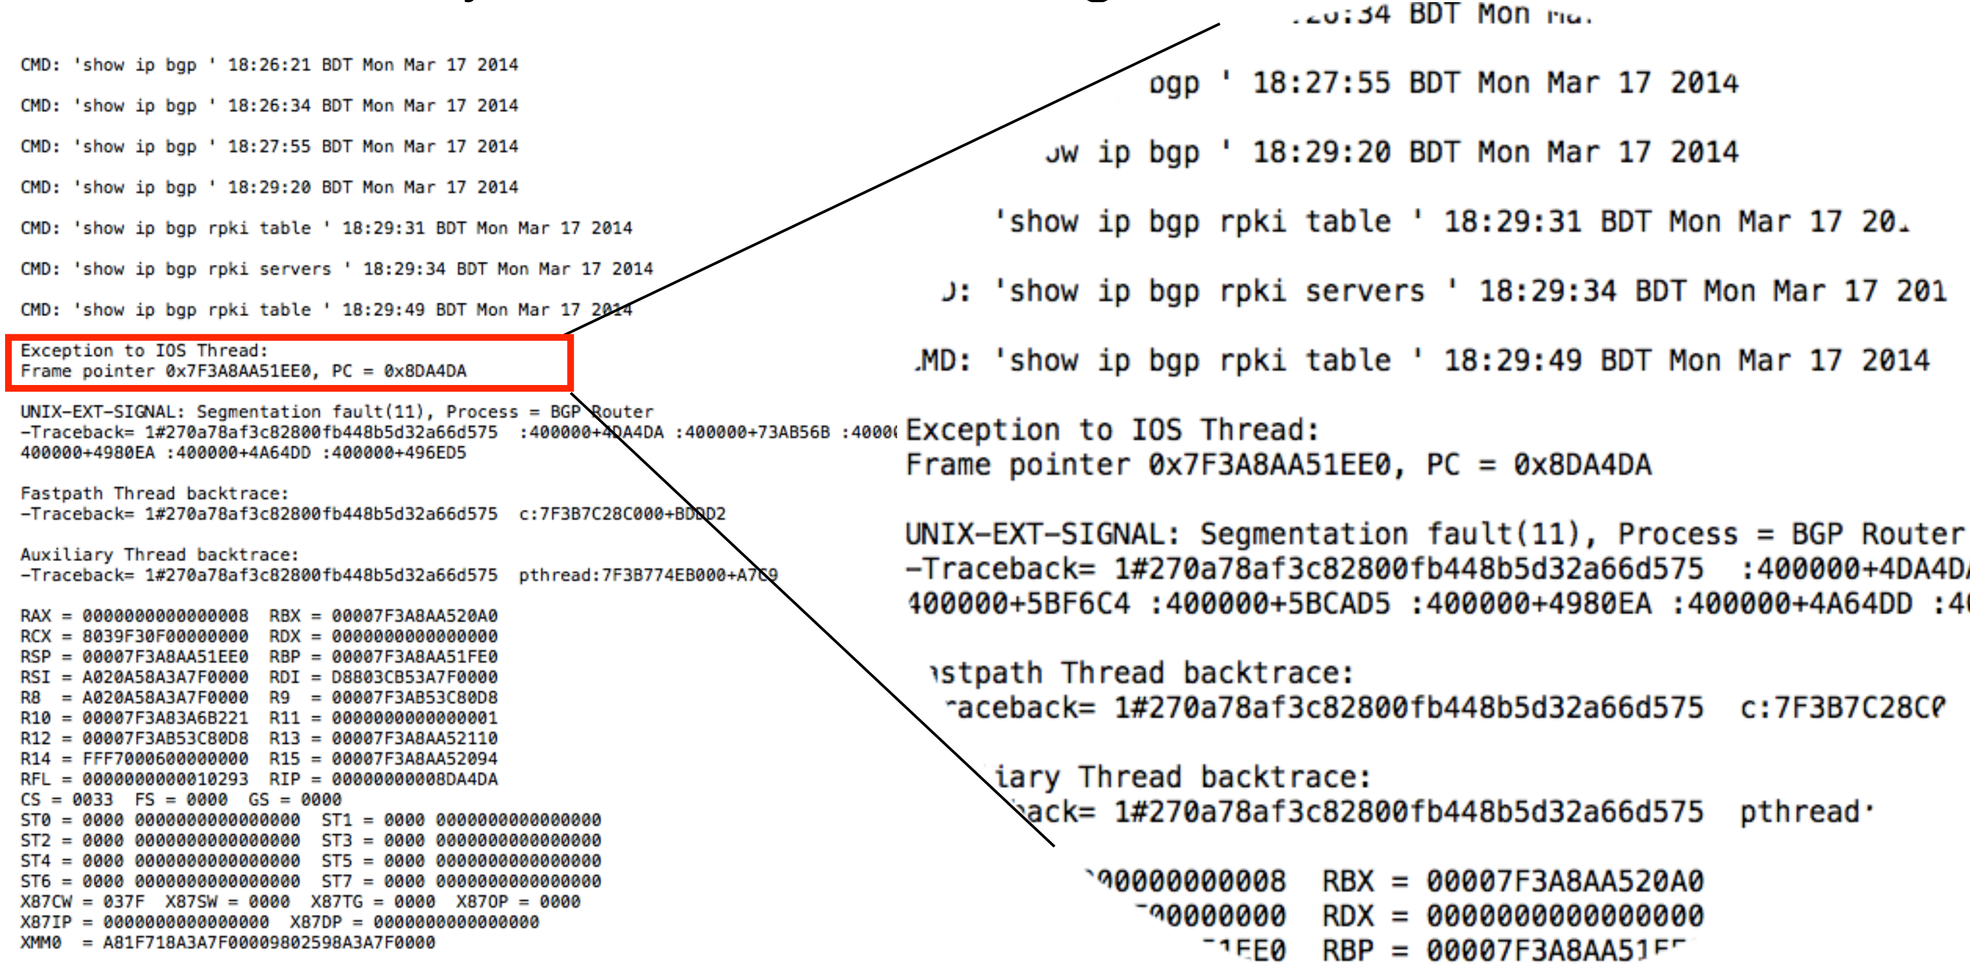

Check your prefix

Cisco (hosted by the RIPE NCC)

Public Cisco router: rpki-rtr.ripe.net

Telnet username: ripe / No password

Juniper (hosted by Kaia Global Networks)

Public Juniper routers: 193.34.50.25, 193.34.50.26

Telnet username: rpki / Password: testbed

source : http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources

Configuration - Reference Link

Cisco

http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html#wp3677719851

Juniper

http://www.juniper.net/techpubs/en_US/junos12.2/topics/topic-map/bgp-origin-as-validation.html


READY TO
ROA
http://www.apnic.net/roa
APNIC

RPKI Demo