

RPKI-Based Origin Validation, Routers, & Caches

AfNOG / Lusaka

2013.06.14

Randy Bush <randy@psg.com>
Cristel Pelsser <cristel@iij.ad.jp>
Rob Austein <sra@hactrn.net>
Michael Elkins <me@sigpipe.org>

Agenda

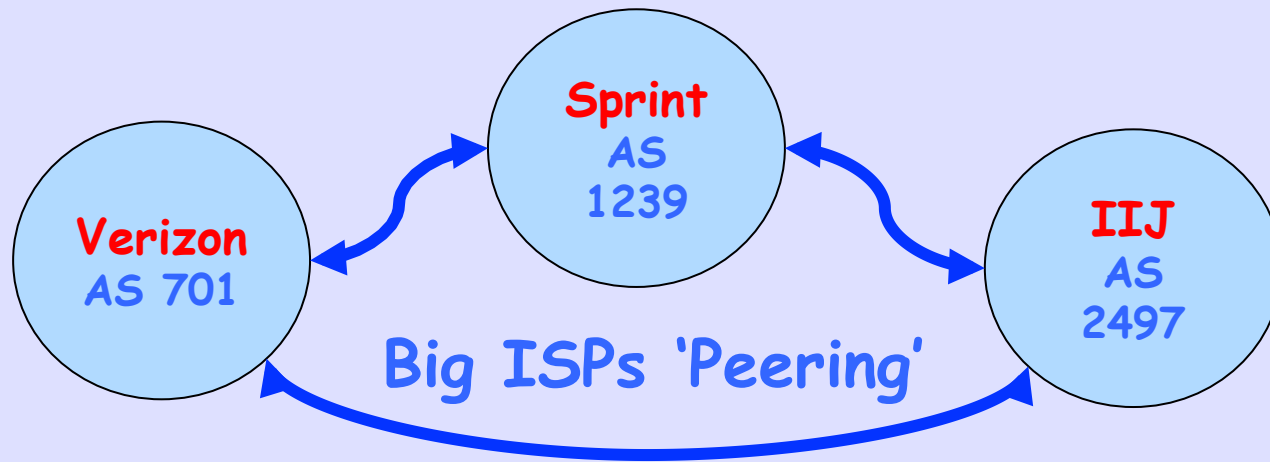
- **Some Technical Background**
- **Mis-Origination - YouTube Incident**
- **The RPKI - Needed Infrastructure**
- **RPKI-Based Origin Validation**
- **Use the GUI to make ROAs and look at the result on a router**
- **Build Your Own RPKI Cache**
- **Discussion**

This is Not New

- 1986 - Bellovin & Perlman identify the vulnerability in DNS and Routing
- 1999 - National Academies study called it out
- 2000 - S-BGP - X.509 PKI to support Secure BGP - Kent, Lynn, et al.
- 2003 - NANOG S-BGP Workshop
- 2006 - RPKI.NET(for ARIN) & APNIC start work on RPKI. RIPE starts in 2008.
- 2009 - RPKI.NET Open Testbed and running code in test routers


What is an AS?

An ISP or End Site



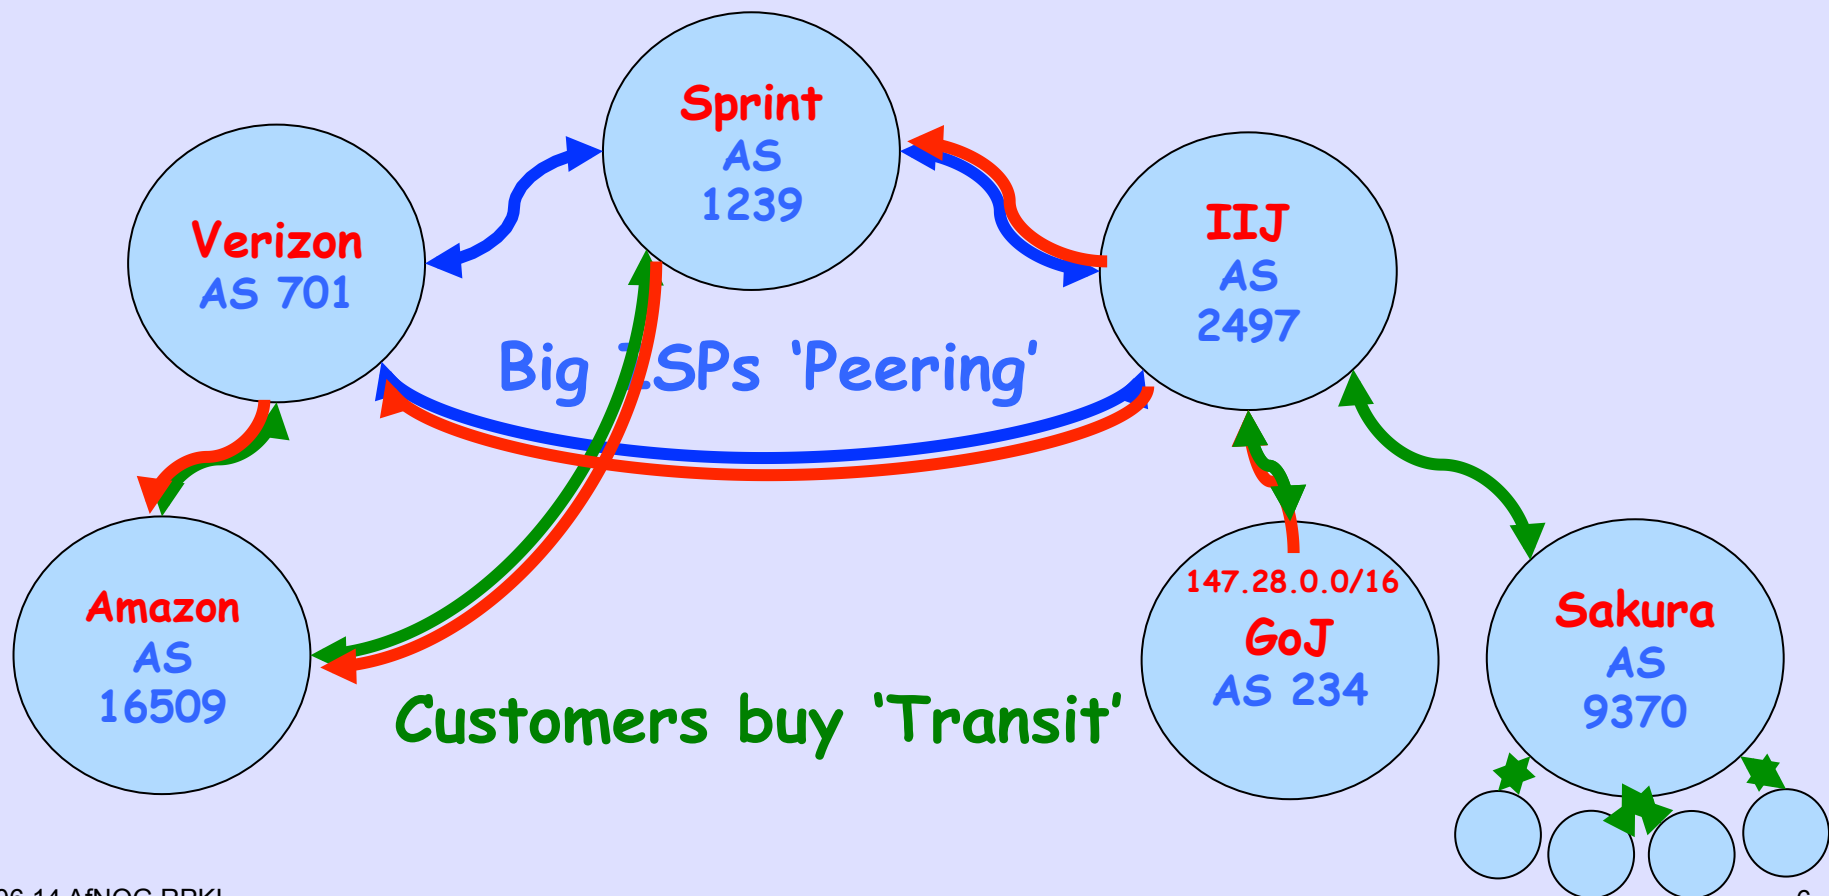
What is an AS?

An ISP or End Site



An IP Prefix is

Announced & Propagated



From Inside a Router

BGP routing table entry for **147.28.0.0/16**



Of Course it's Uglier 😊

```
r1.iad#sh ip bgp 147.28.0.0/16
```

```
BGP routing table entry for 147.28.0.0/16, version 21440610
```

```
Paths: (2 available, best #1, table default)
```

```
  Advertised to update-groups:
```

```
    1
```

```
Refresh Epoch 1
```

```
16509   1239   2497   234
```

```
  144.232.18.81 from 144.232.18.81 (144.228.241.254)
```

```
    Origin IGP, metric 841, localpref 100, valid, external, best
```

```
    Community: 3297:100 3927:380
```

```
    path 67E8FFCC RPKI State valid
```

```
Refresh Epoch 1
```

```
16509    701   2497   234
```

```
  129.250.10.157 (metric 11) from 198.180.150.253 (198.180.150.253)
```


```
    Origin IGP, metric 95, localpref 100, valid, internal
```

```
    Community: 2914:410 2914:1007 2914:2000 2914:3000 3927:380
```

```
    path 699A867C RPKI State valid
```


The YouTube Incident


The Plan



Poison Pakistan
Internal Routing

The YouTube Incident

What Happened



Poisoned the
Global Internet

We Call this *Mis-Origination*

a Prefix is Originated
by an AS Which Does
Not Own It

I Do Not Call it
Hijacking

Because that Assumes
Negative Intent

And These Accidents
Happen Every Day

Usually to Small Folk
Sometimes to Large

So,

What's the Plan?

Three Pieces

- **RPKI** - Resource Public Key Infrastructure, the Certificate Infrastructure to Support the other Pieces (deployed at all RIRs)
- **Origin Validation** - Using the RPKI to detect and prevent mis-originations of someone else's prefixes (in deployment)
- **AS-Path Validation AKA BGPsec** - Prevent Path Attacks on BGP (future work)

Why Origin Validation?

- Prevent YouTube accident & Far Worse
- Prevent 7007 accident, UU/Sprint 2 days!
- Prevents most accidental announcements
- Does not prevent malicious path attacks such as the Kapela/Pilosov DefCon attack
- That requires *Path Validation*, the third step, a few years away

We Need to be Able to
Authoritatively Prove
Who Owns an IP Prefix
And What AS(s) May
Announce It

Prefix Ownership
Follows the Allocation
Hierarchy
IANA, RIRs, ISPs, ...

X.509-Based IP Resource PKI

RFCs Have Been
Long Published

Deployed by
All RIRs


**ROAs Registered
by > 1,000 Operators**

In Live Routers

Public-Key Concepts

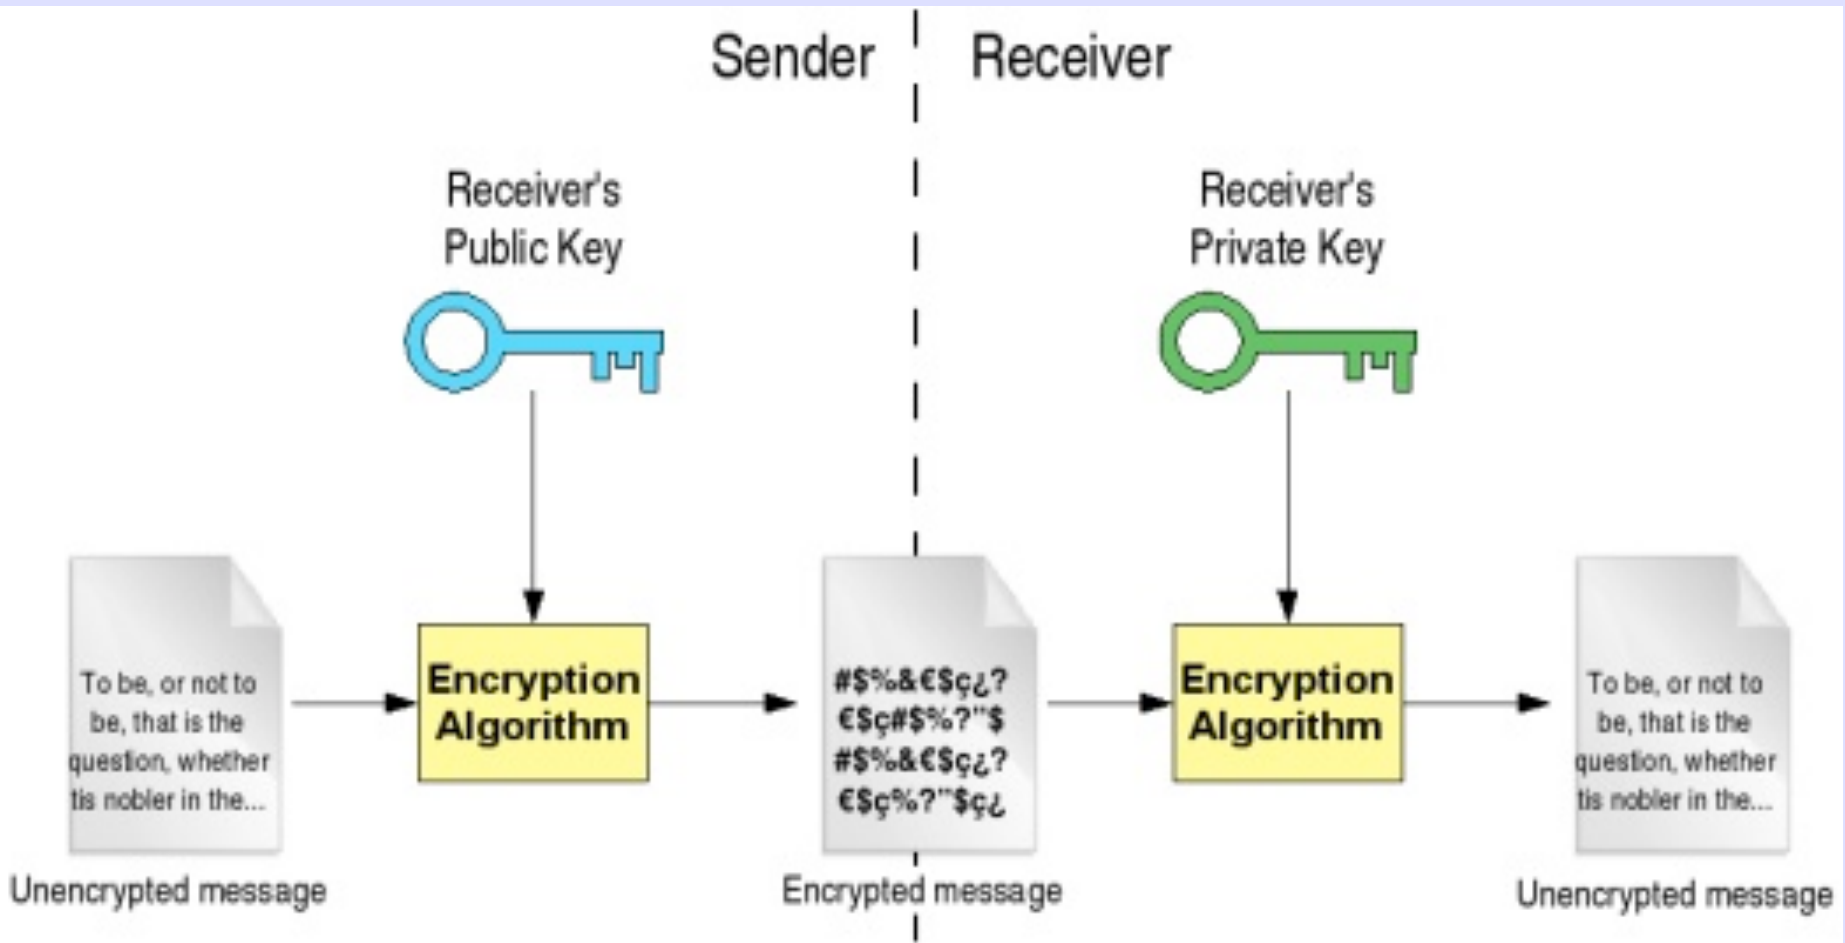
- **Private key:** This key must be known *only* by its owner.
- **Public key:** This key is known to everyone (it is *public*)
- **Relation between both keys:** What one key encrypts, the other one decrypts, and vice versa. That means that if you encrypt something with my public key (which you would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

Key Generation

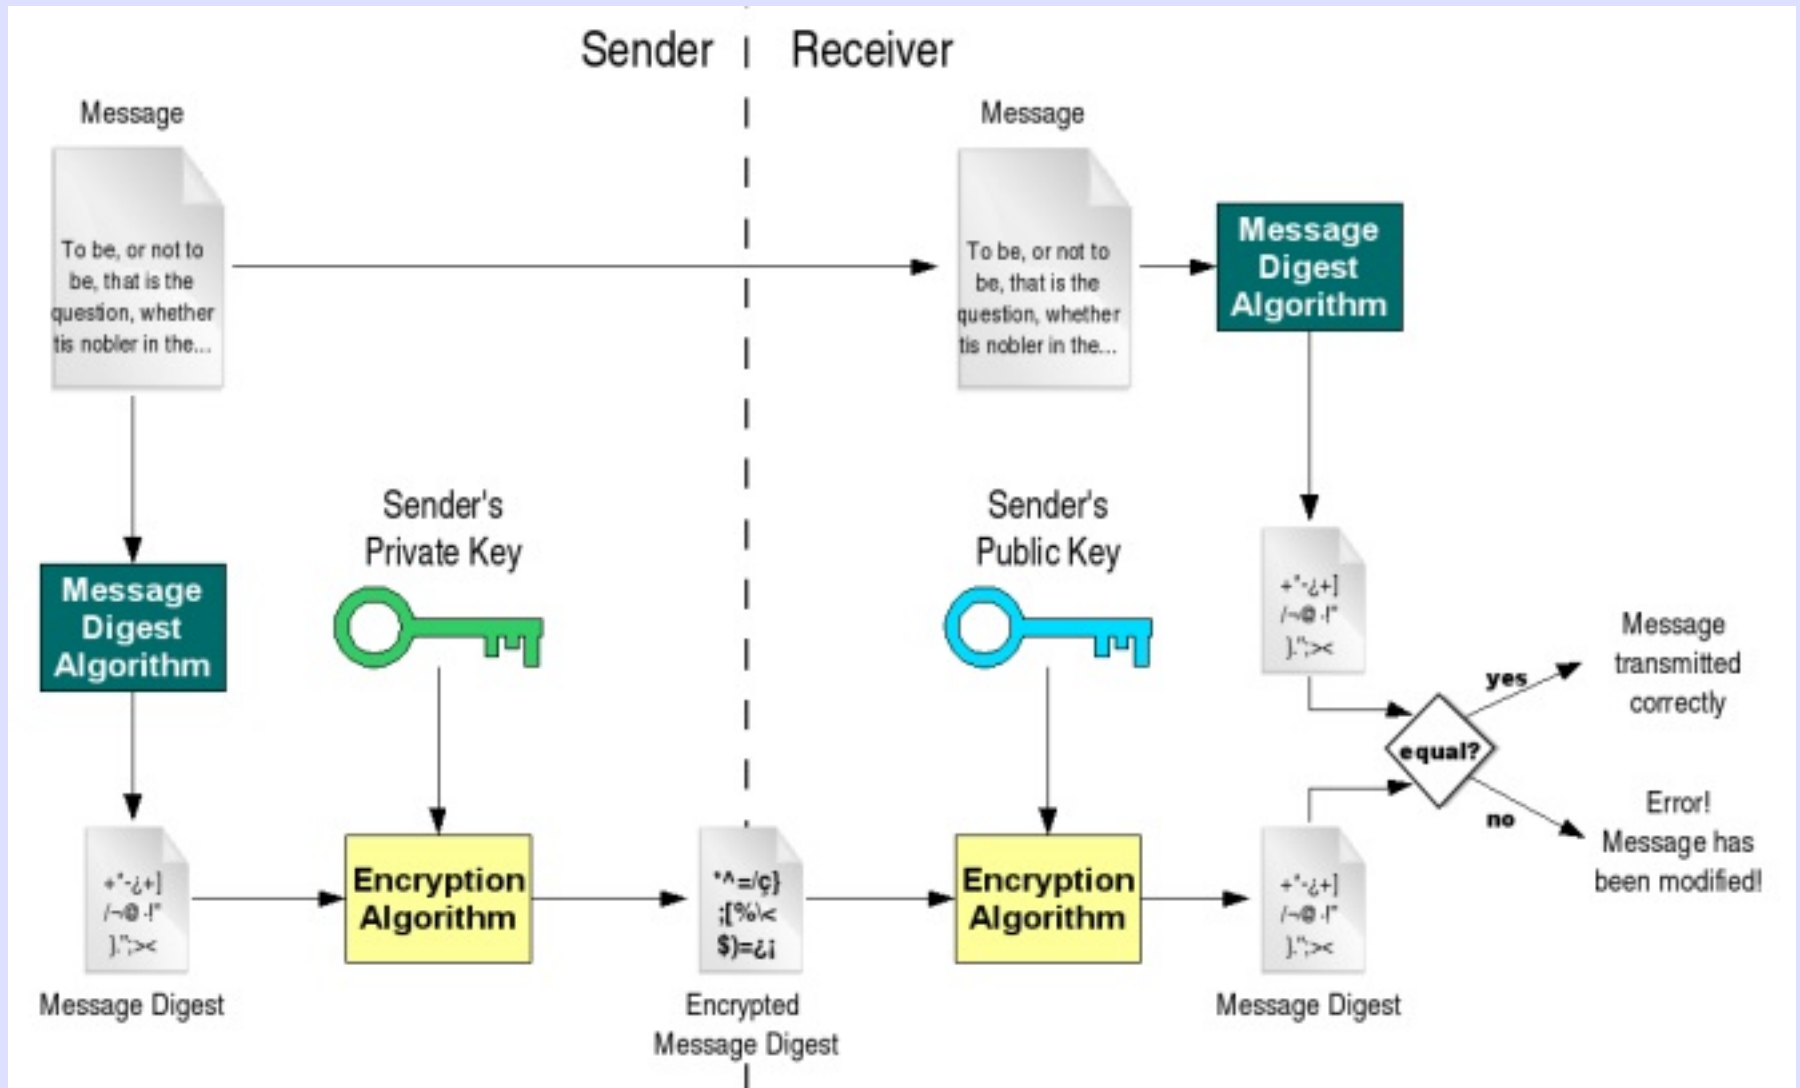


Stolen from - <http://gdp.globus.org/gt4-tutorial/multiplehtml/ch09s03.html>

En/DeCryption

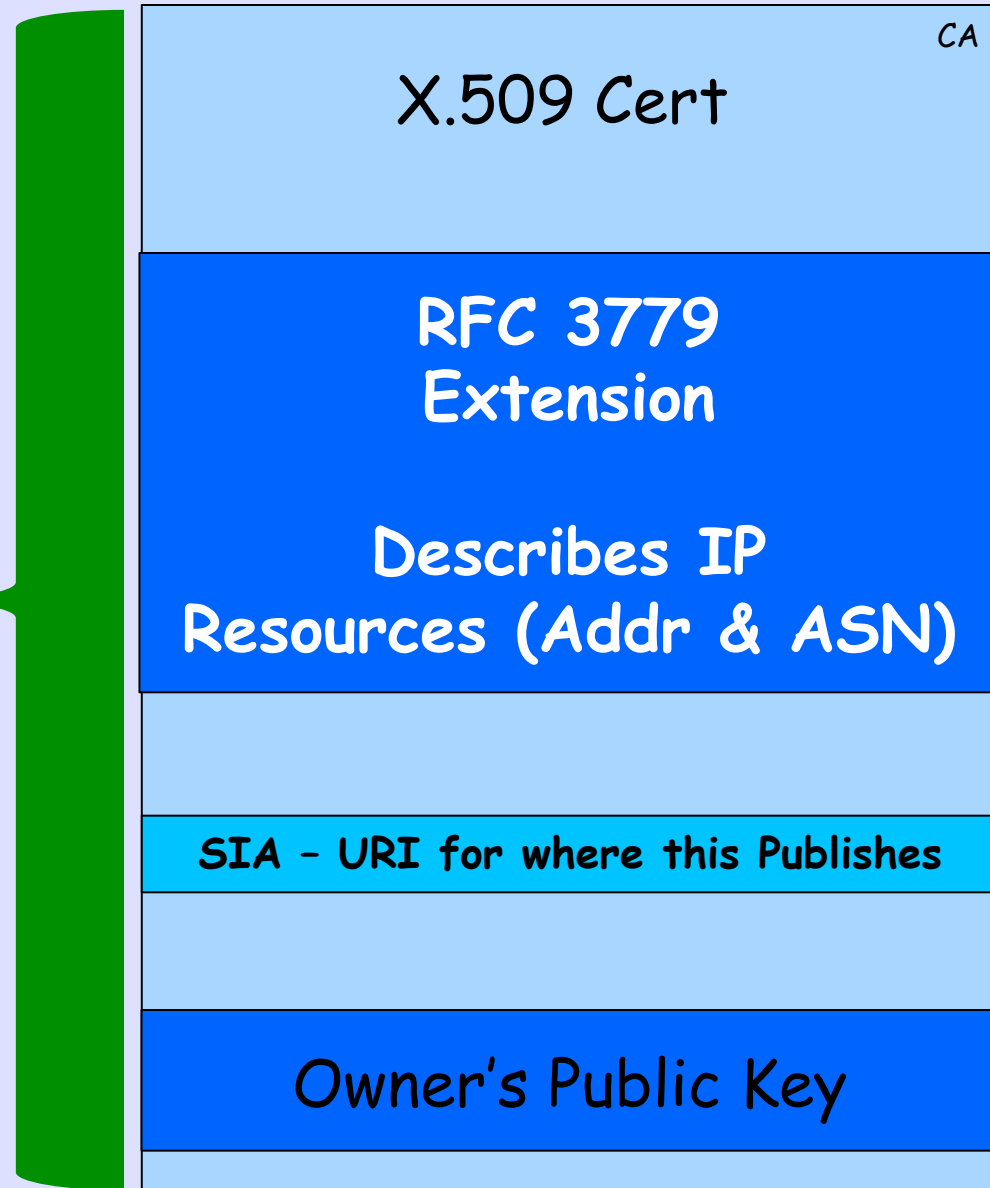


Digital Signature




X.509 Certificate w/ 3779 Ext

**Signed
by
Parent's
Private
Key**

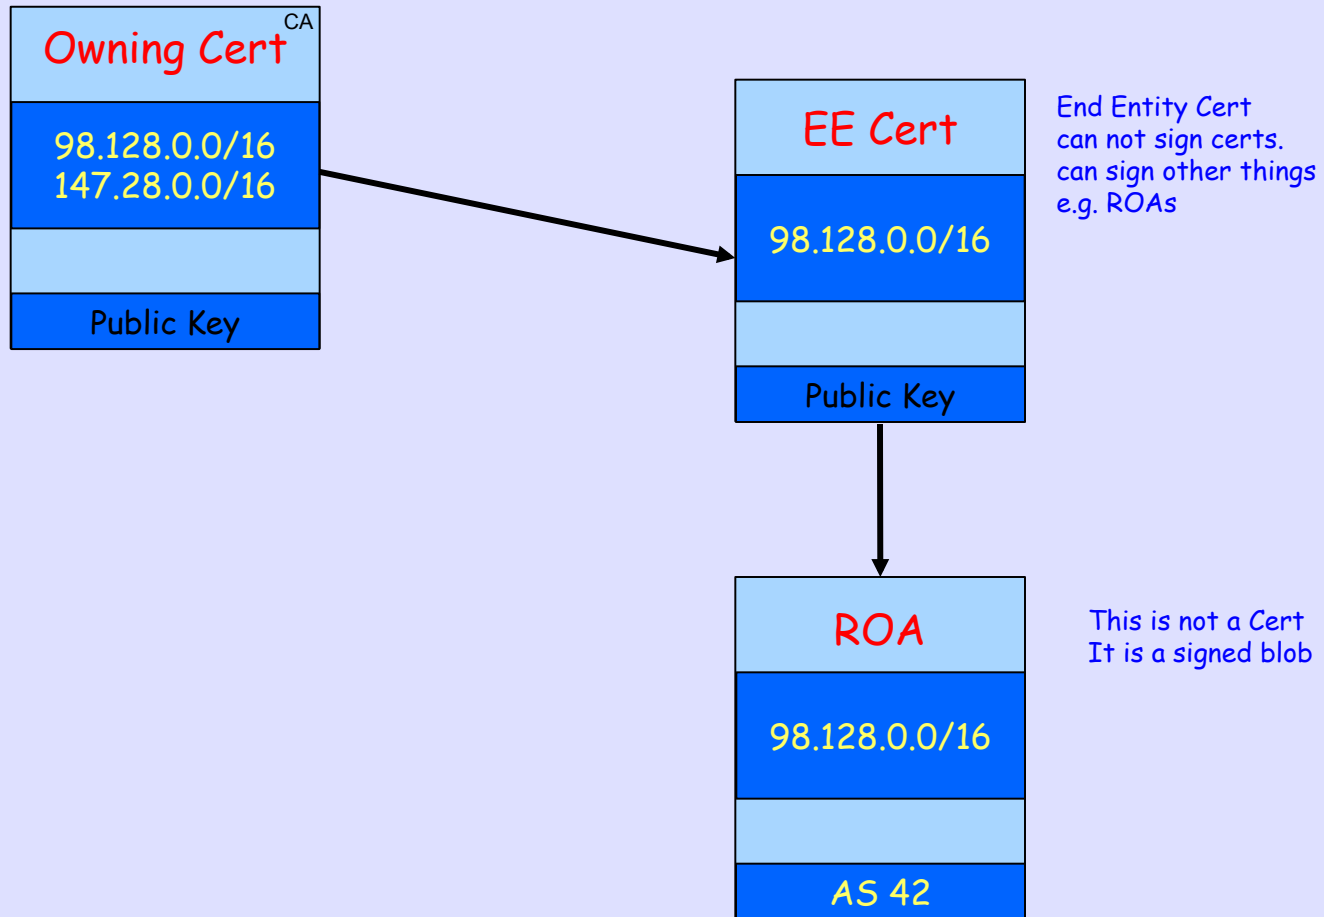



Certificate Hierarchy follows Allocation Hierarchy



That's Who Owns It
but
Who May Route It?

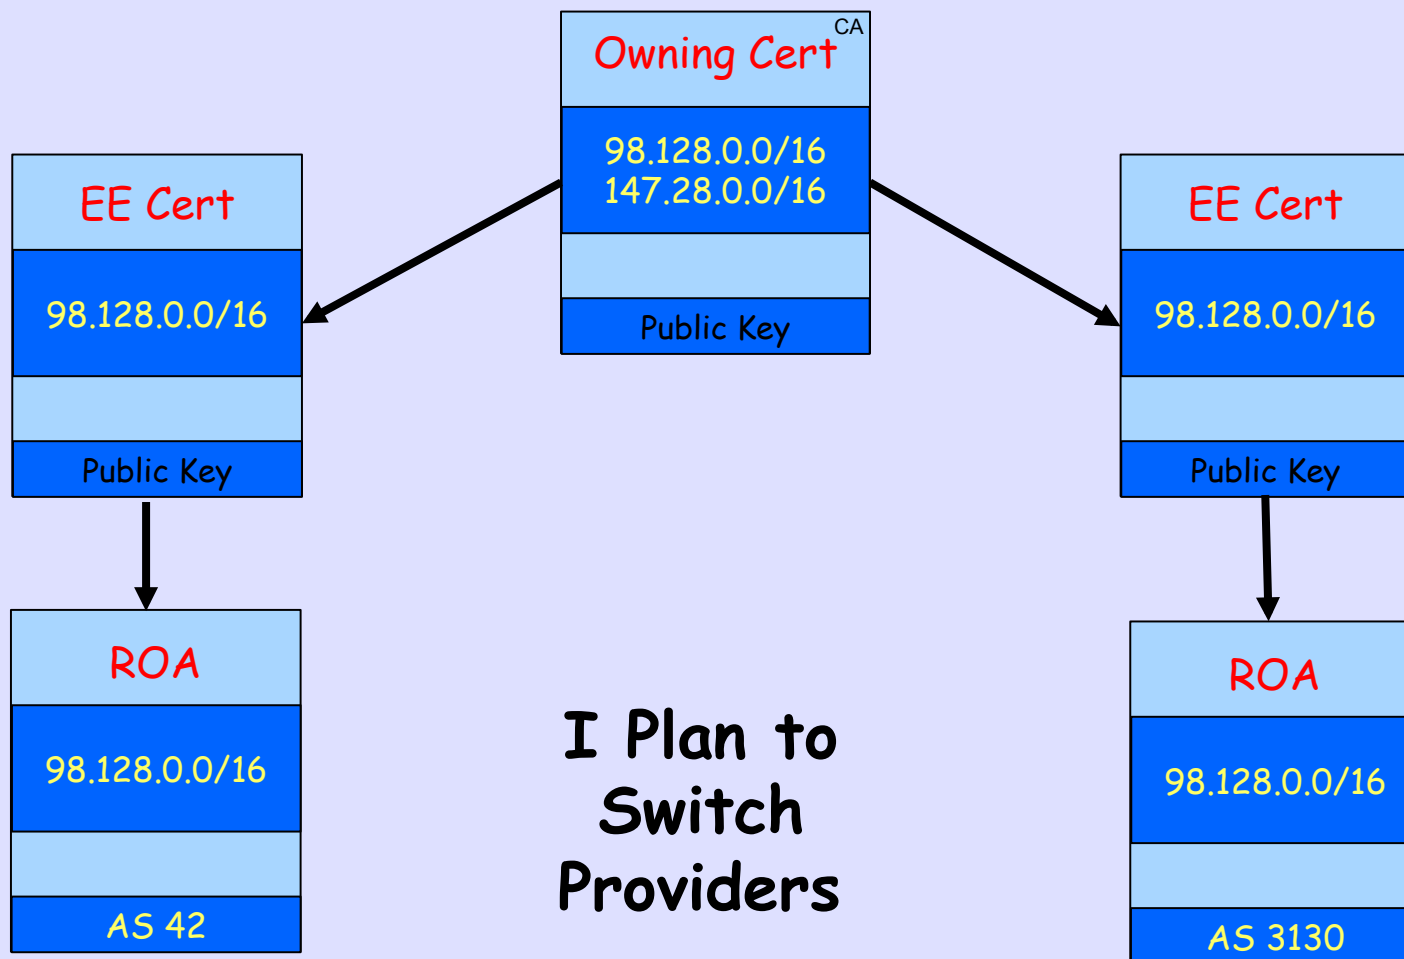
Route Origin Authorization (ROA)






ROA Aggregation Using Max Length

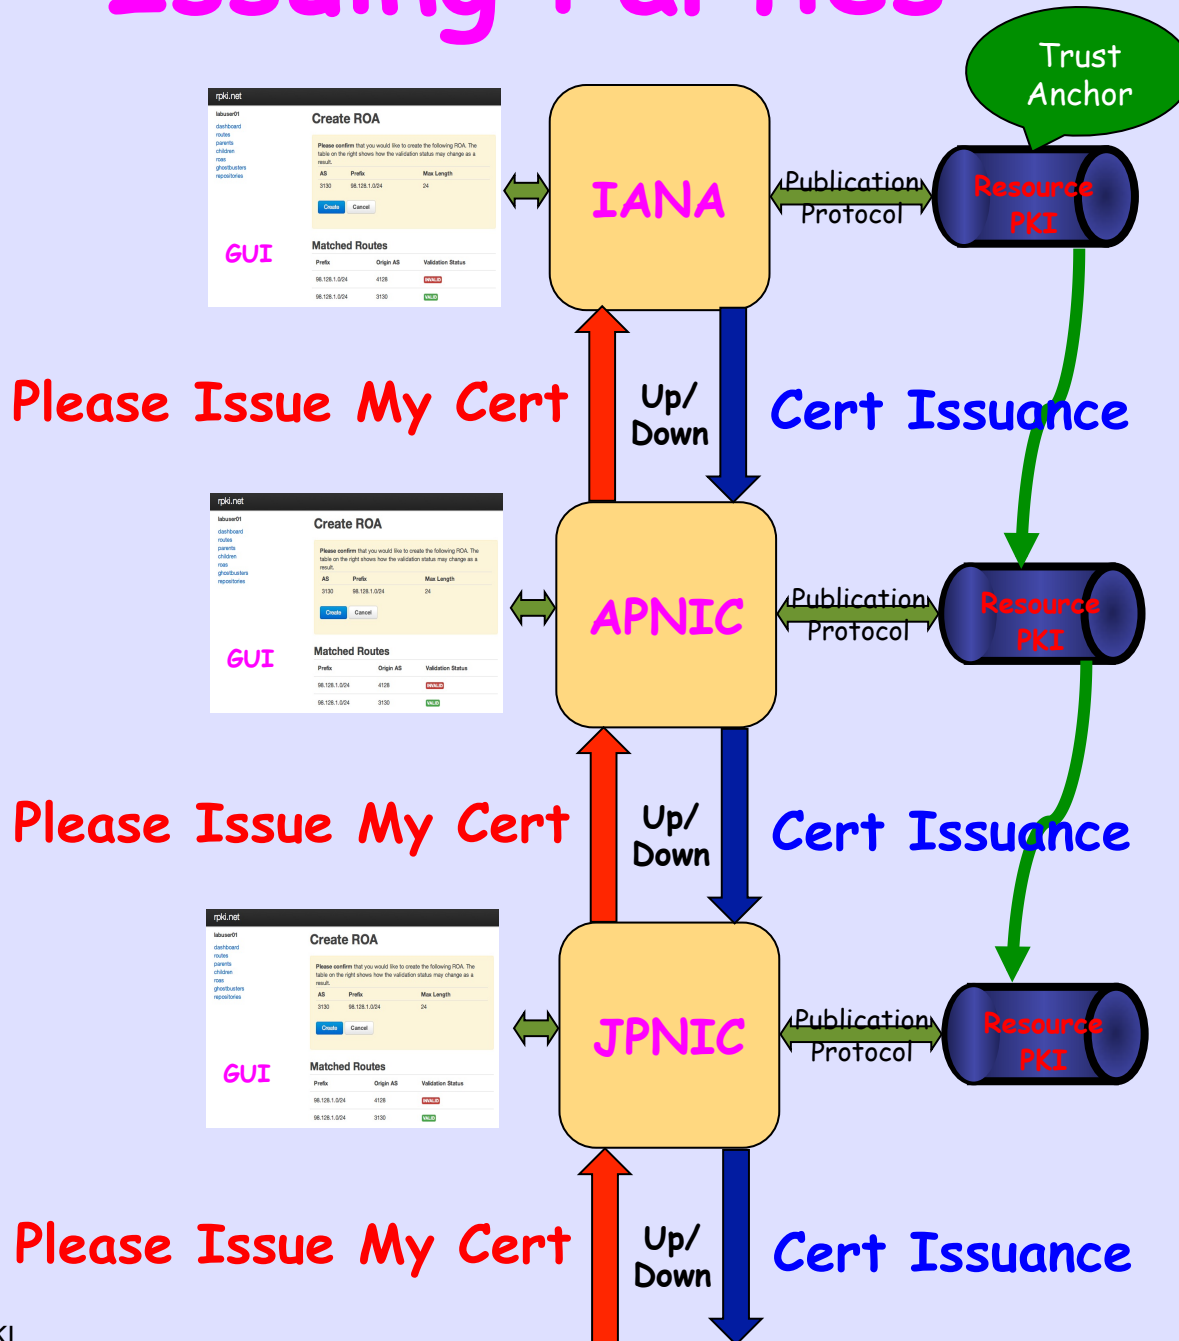
Multiple ROAs Make Before Break



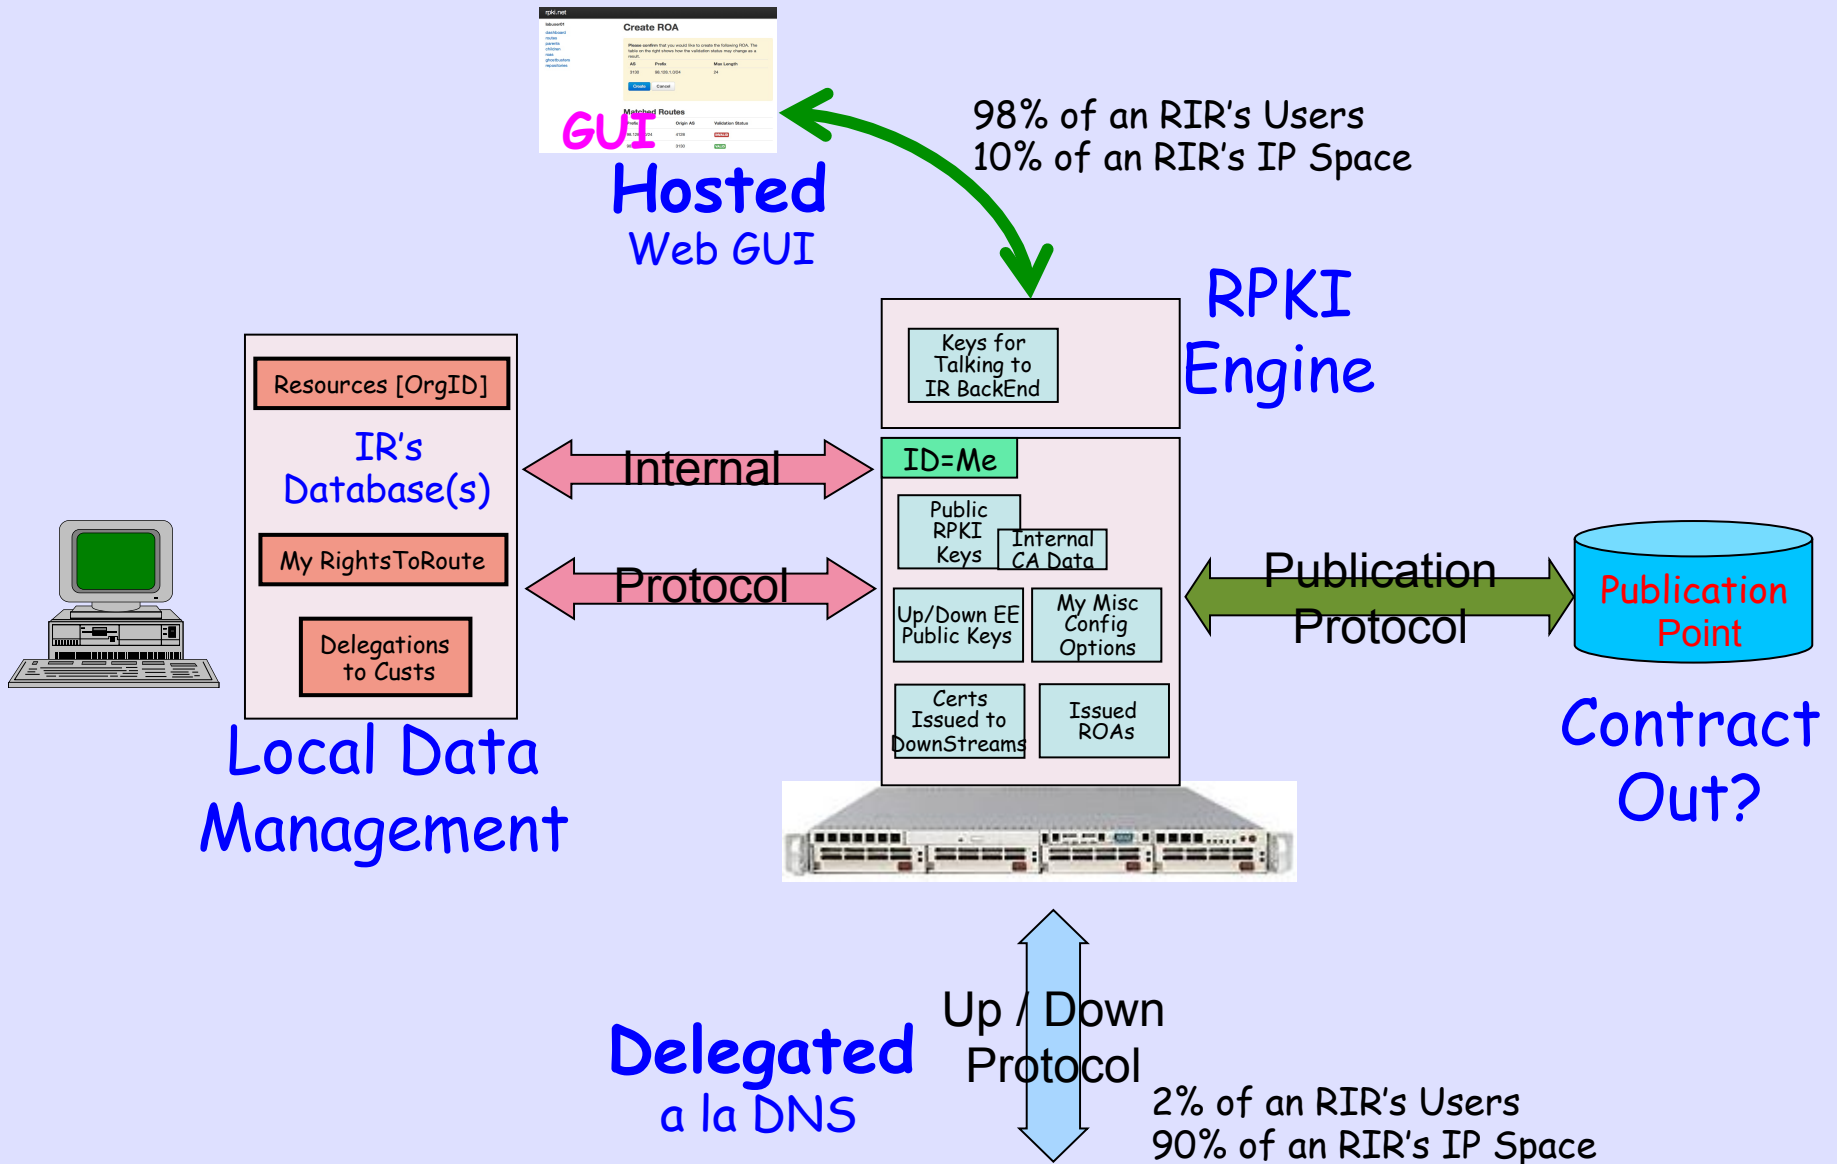
How RPKI is Generated



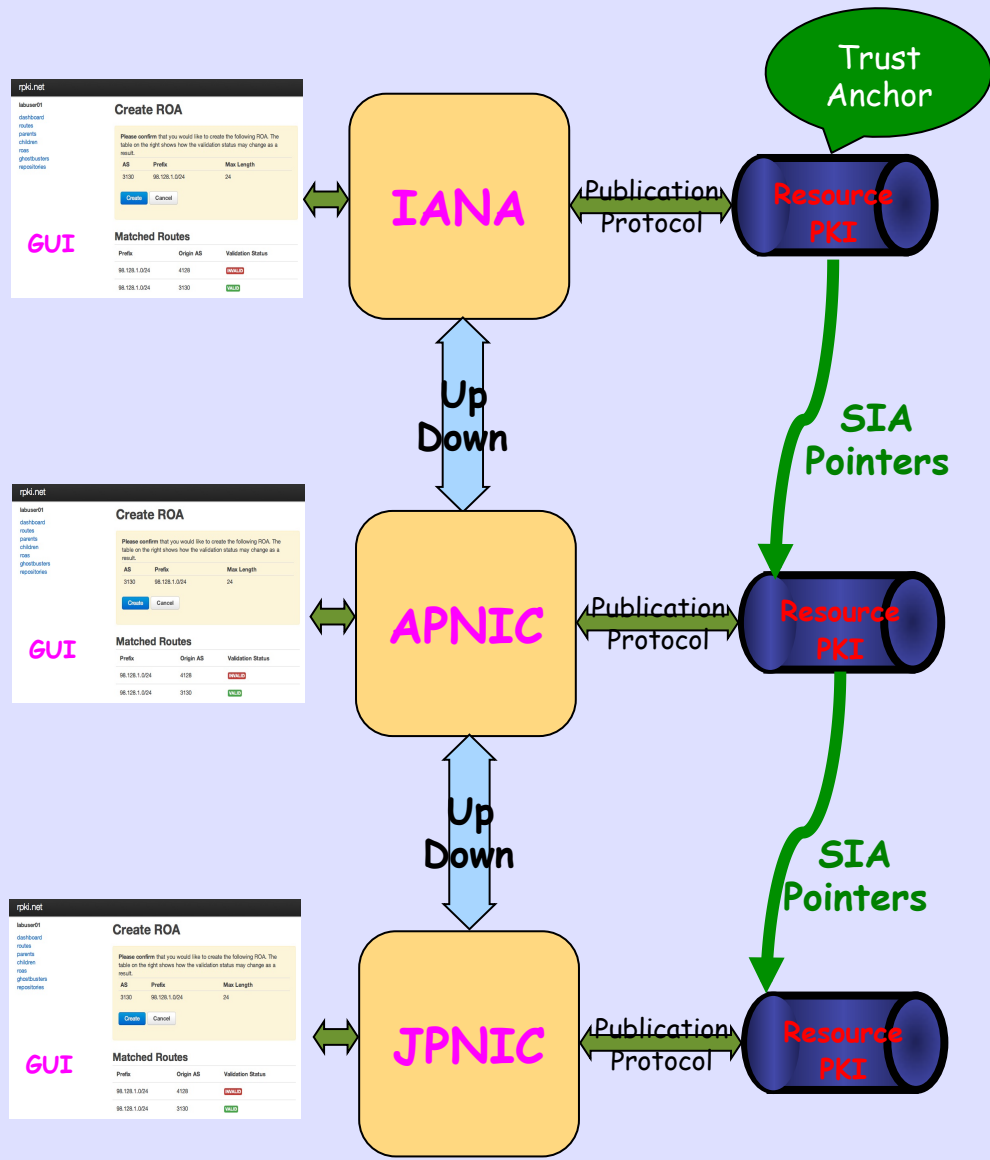
Issuing Parties



Hosted vs Delegated



Issuing Parties



Trust Anchor Locator

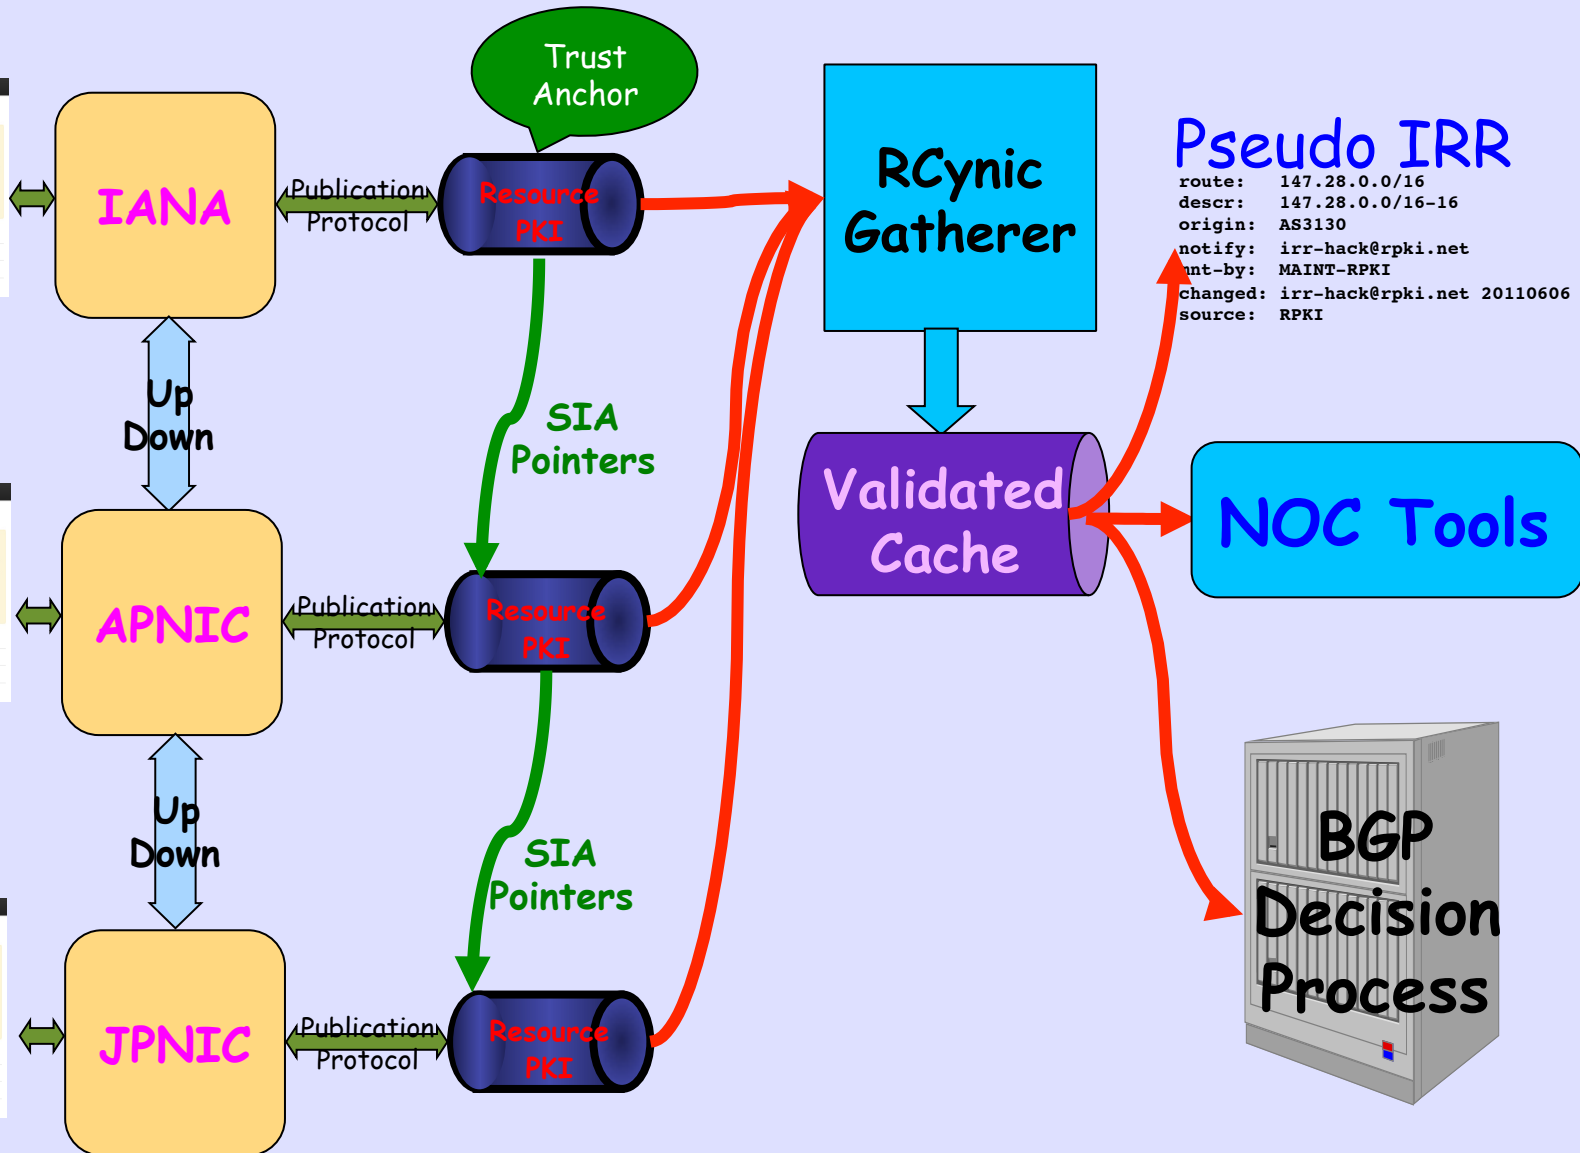



`rsync://ca0.rpki.net/ta1/root.cer`

```
MIIBIjANBgkqhkiG9w0BAQEFAA0CAQ8AMIIBCgKCAQEAzPSPpQxa0rxz9sbbvYGr
UlpgyBVxSI+t+k/WDKbr+VW7CjUoz6cc5KMFANKQWw3D6ER4kCwX4WJkD58AGGbw/
WeAe6m3aHc0RUVRkr45a4qSrYiG7Wq9RAXtwbh1XofB3zo+090ILXDaVP2U9bw+Q
yoJBJuAmZONt0bRgrktv8QhVtKvuYkH5ZIE7DkXJcJzBn6gv09dZsdwZm3xV3soX
HEKrz5pY6Sb2xoL1CyPqzG0fVfx10G5+dmcD/degPKxrEycAzjnHUzN1gus2jg26
dtkix7KG/Mn1h/k53j0FdQD+zqPwakgwqjvC0dSdHMRmsikj0EF9WrZIOjZUXV6q
6wIDAQAB
```

A URL and a Public Key that must be able to Decrypt the Certificate that is found at the URL so you know you can trust it


Issuing Parties

Relying Parties




Possible
Large ISP
Deployment

Caches
Feed
Caches



———— High Priority
———— Lower Priority

How Do ROAs Affect BGP Updates?



rpki.net

dashboard routes create update delete modify publications reports

Create ROA

Please confirm that you would like to create the following ROA. The data on the right shows how the validation status may change as a result.

AS	Prefix	Max Length
3133	96.128.1.0/24	24

Create Cancel

Matched Routes

Prefix	Origin AS	Validation Status
96.128.1.0/24	4128	Invalid
96.128.1.0/24	3130	Valid

rpki.net

dashboard routes create update delete modify publications reports

Create ROA

Please confirm that you would like to create the following ROA. The data on the right shows how the validation status may change as a result.

AS	Prefix	Max Length
3133	96.128.1.0/24	24

Create Cancel

Matched Routes

Prefix	Origin AS	Validation Status
96.128.1.0/24	4128	Invalid
96.128.1.0/24	3130	Valid

rpki.net

dashboard routes create update delete modify publications reports

Create ROA

Please confirm that you would like to create the following ROA. The data on the right shows how the validation status may change as a result.

AS	Prefix	Max Length
3133	96.128.1.0/24	24

Create Cancel

Matched Routes

Prefix	Origin AS	Validation Status
96.128.1.0/24	4128	Invalid
96.128.1.0/24	3130	Valid

ROAs Become Router VRPs

Want to Run on
Current Hardware




RPKI
Cache

RPKI-Rtr
Protocol

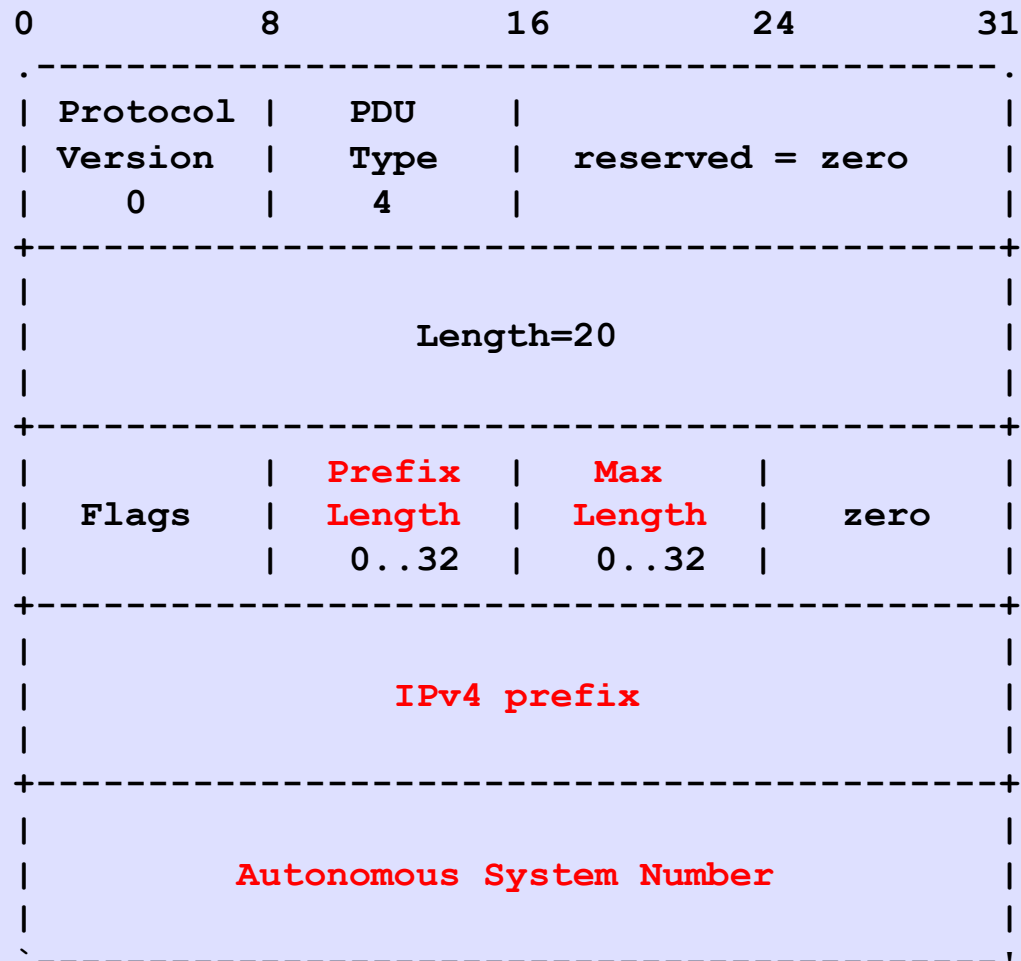


Crypto is
Stripped

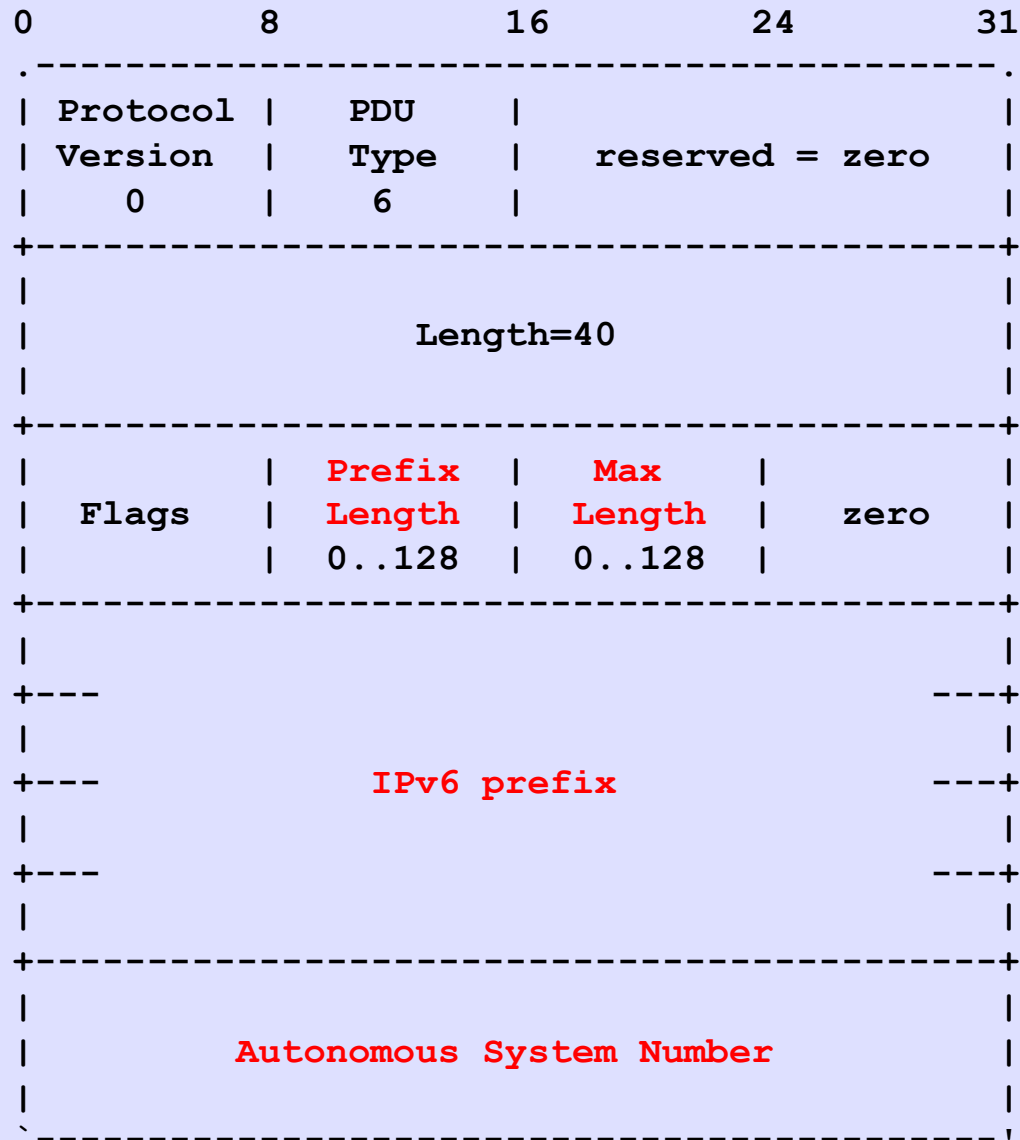


RPKI
VRPs

IPv4 Prefix




IPv6 Prefix



96 More Bits
No Magic

BGP Updates are
compared with
ROA Data loaded
from the RPKI

Marking BGP Updates



Result of Check

- **Valid** - A matching/covering VRP was found with a matching AS number
- **Invalid** - A covering VRP was found, but the AS number did not match, and there was no other matching one
- **NotFound** - No matching or covering VRP was found, same as today

Configure Router to Get ROAs

```
router bgp 651nn
```

```
...
```

```
bgp rpki server tcp 198.180.150.1 port 43779 refresh 3600
```

```
bgp rpki server tcp 147.28.0.35 port 93920 refresh 3600
```

```
...
```

Valid!

```
r0.sea#show bgp 192.158.248.0/24
```

```
BGP routing table entry for 192.158.248.0/24, version 3043542
```

```
Paths: (3 available, best #1, table default)
```

```
6939 27318
```

```
206.81.80.40 (metric 1) from 147.28.7.2 (147.28.7.2)
```

```
Origin IGP, metric 319, localpref 100, valid, internal,
```

```
best
```

```
Community: 3130:391
```

```
path 0F6D8B74 RPKI State valid
```

```
2914 4459 27318
```

```
199.238.113.9 from 199.238.113.9 (129.250.0.19)
```

```
Origin IGP, metric 43, localpref 100, valid, external
```

```
Community: 2914:410 2914:1005 2914:3000 3130:380
```

```
path 09AF35CC RPKI State valid
```

Invalid!

```
r0.sea#show bgp 198.180.150.0
```

```
BGP routing table entry for 198.180.150.0/24, version 2546236
```

```
Paths: (3 available, best #2, table default)
```

```
  Advertised to update-groups:
```

```
    2          5          6          8
```

```
Refresh Epoch 1
```

```
1239 3927
```

```
  144.232.9.61 (metric 11) from 147.28.7.2 (147.28.7.2)
```

```
    Origin IGP, metric 759, localpref 100, valid, internal
```

```
    Community: 3130:370
```

```
    path 1312CA90 RPKI State invalid
```

NotFound

```
r0.sea#show bgp 64.9.224.0
```

```
BGP routing table entry for 64.9.224.0/20, version 35201
```

```
Paths: (3 available, best #2, table default)
```

```
  Advertised to update-groups:
```

```
    2          5          6
```

```
Refresh Epoch 1
```

```
1239 3356 36492
```

```
  144.232.9.61 (metric 11) from 147.28.7.2 (147.28.7.2)
```

```
    Origin IGP, metric 4, localpref 100, valid, internal
```


```
    Community: 3130:370
```

```
    path 11861AA4 RPKI State not found
```

What are the BGP / VRRP¹ Matching Rules?

¹ Validated ROA Payload

A Prefix is Covered by a VRP when the VRP prefix length is less than or equal to the Route prefix length



Prefix is Matched by a VRP when the Prefix is Covered by that VRP , prefix length is less than or equal to the VRP max-len, and the Route Origin AS is equal to the VRP's AS

BGP	98.128.0.0/16 AS 42	
VRP	98.128.0.0/12-16 AS 42	Matched
VRP	98.128.0.0/16-24 AS 666	No. AS Mismatch
VRP	98.128.0.0/20-24 AS 42	No. VRP Longer

Matching and Validity

VRP₀

98.128.0.0/16-24 AS 6

VRP₁

98.128.0.0/16-20 AS 42

BGP	98.128.0.0/12	AS 42	NotFound, shorter than VRPs
BGP	98.128.0.0/16	AS 42	Valid, Matches VRP ₁
BGP	98.128.0.0/20	AS 42	Valid, Matches VRP ₁
BGP	98.128.0.0/24	AS 42	Invalid, longer than VRP with AS 42
BGP	98.128.0.0/24	AS 6	Valid, Matches VRP ₀

The Operator
Tests the Mark
and then
Applies Local Policy

Fairly Secure

```
route-map validity-0
```

```
  match rpki valid
```

```
  set local-preference 100
```

```
route-map validity-1
```

```
  match rpki not-found
```

```
  set local-preference 50
```

```
! invalid is dropped
```

Paranoid

```
route-map validity-0
```

```
  match rpki valid
```

```
  set local-preference 110
```

```
! everything else dropped
```

Security Geek

```
route-map validity-0
```

```
  match rpki invalid
```

```
  set local-preference 110
```

```
! everything else dropped
```

After AS-Path

```
route-map validity-0
```

```
  match rpki not-found
```

```
    set metric 100
```

```
route-map validity-1
```

```
  match rpki invalid
```

```
    set metric 150
```

```
route-map validity-2
```

```
  set metric 50
```

Set a Community

```
route-map validity-0
```

```
  match rpki valid
```

```
  set community 3130:400
```

```
route-map validity-1
```


```
  match rpki invalid
```

```
  set community 3130:200
```

```
route-map validity-2
```

```
  set community 3130:300
```

And it is All Monitored



But in the End, You Control Your Policy

"Announcements with Invalid origins SHOULD NOT be used, but MAY be used to meet special operational needs. In such circumstances, the announcement SHOULD have a lower preference than that given to Valid or NotFound."

-- draft-ietf-sidr-origin-ops

But if we do not
reject Invalid,
What is all this
work for?

RPKI at the Registries

- RIPE seriously deployed with a few thousand LIRs and thousands of ROAs
- APNIC is operational and moving forward, moving to RIPE's GUI
- ARIN is doing their best to make RPKI deployment very hard
- LACNIC is deployed and has $O(100)$ LIRs
- AFRINIC is deployed with $O(25)$ LIRs

Router Origin Validation

- **Cisco IOS - solid in 15.2**
- **Cisco IOS/XR - shipped in 4.3.2**
- **Juniper - shipped in 12.2**
- **AlcaLu - in development**

RPKI Implementations

- **RIPE/NCC - CA (partial closed) & RP (partial open)**
- **APNIC - CA only - Closed Source**
- **RTRlib/Berlin - RP only - Open Source**
- **BBN - RP Only - Open Source**
- **RPKI.NET - CA & RP - Open Source**

Today - RPKI.NET

- Open Source BSD License
- CA - Hosted and Delegated Models, GUI
- RP - RPKI-RTR, NOC Tools, IRR Gen
- FreeBSD, Ubuntu, Debian, ... Packaged
(docs still catching up)

Recent Hackathons

- **JaNOG January - LIRs bring up Relying Party. Few Attended, Most Succeeded**
- **JPNIC 20 Feb - RP only, 3 ISPs were successful**
- **APRICOT February - TW, & KR have CAs in internal test**
- **Beirut March - 18 LIRs as RPs**

 Search

logged in as randy | [Logout](#) | [Preferences](#) | [Help/Guide](#) | [About Trac](#)

Home	Documentation	Timeline	Roadmap	Browse Source	View Tickets	New Ticket	Search	Admin
----------------------	-------------------------------	--------------------------	-------------------------	-------------------------------	------------------------------	----------------------------	------------------------	-----------------------

wiki: [WikiStart](#) [Start Page](#) | [Index](#) | [History](#)
Last modified 2 months ago

rpki.net project site

This is the Trac site for the rpki.net project. The project provides a free, BSD License, open source, complete system for the Internet Registry or ISP. It includes separate components which may be combined to suit your needs:

- Certification Engine
- Relying Party Cache (sometimes called a 'validator')
- rpki-rtr protocol, to feed the data to routers doing RPKI-based origin validation
- GUI for use by users of the 'hosted' model (i.e. customers who do not run their own CA)
- Web Reporting Pages so you can see what your cache has found
- Creation of pseudo-IRR data for those who wish to feed RPSL toolchains

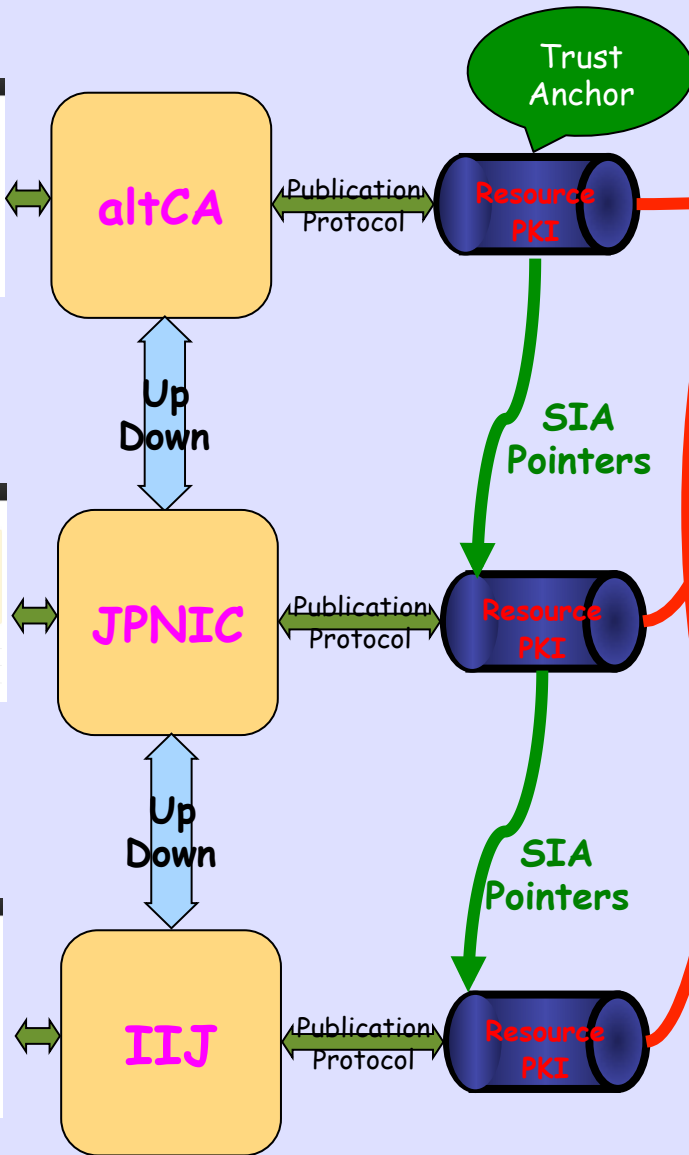



Documentation

Primary [documentation for the code](#) is here, in the Trac wiki. PDF and flat text forms derived from are available in the source code repository.

Bug Reports

Issuing Parties

Relying Parties



Today

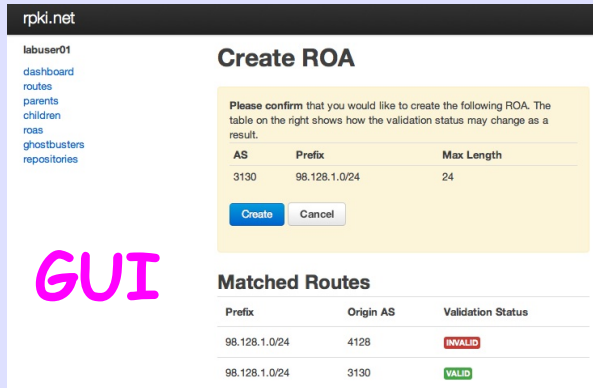
- Register Our Prefixes in CA
- Issue ROAs Using CA's Web Portal
- Configure Routers to get ROAs from Caches
- Build RP Caches to Fetch from CAs

Get Copy of This Preso

<https://psg.com/130614.pdf>

So You Can
Copy and Paste

Lab Overview



GUI


django





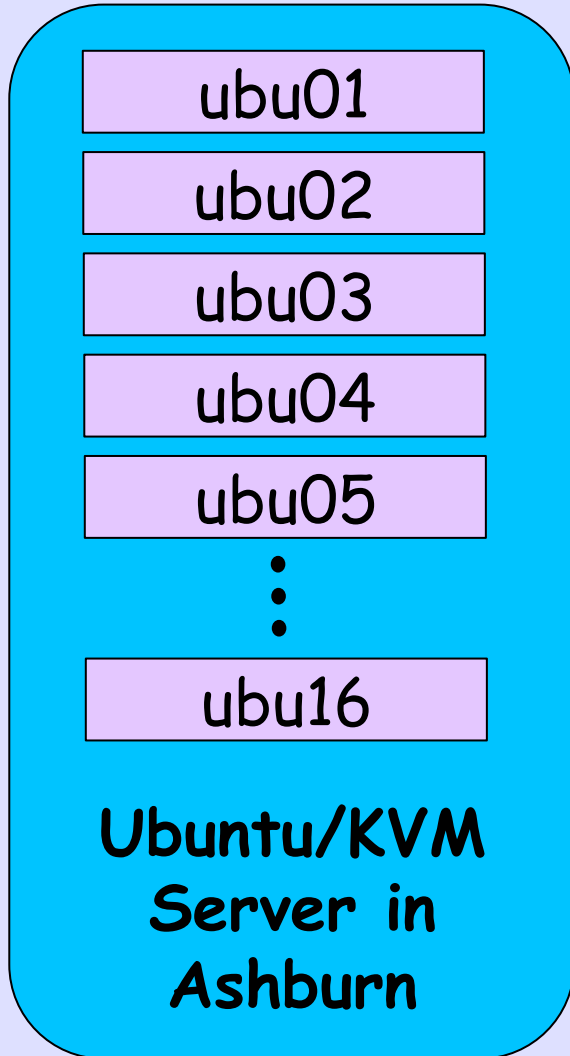
Publication Protocol



RPKI to Rtr Protocol



Lab Environment



DynaMIPS on MacMini

10.0.0.0/8

RPKI-Rtr Protocol

AS65000

AS65001

194.126.10.110
Global Internet

RPKI Cache

98.128.0.0/16
98.128.0.0/24
98.128.1.0/24
...
98.128.31.0/24




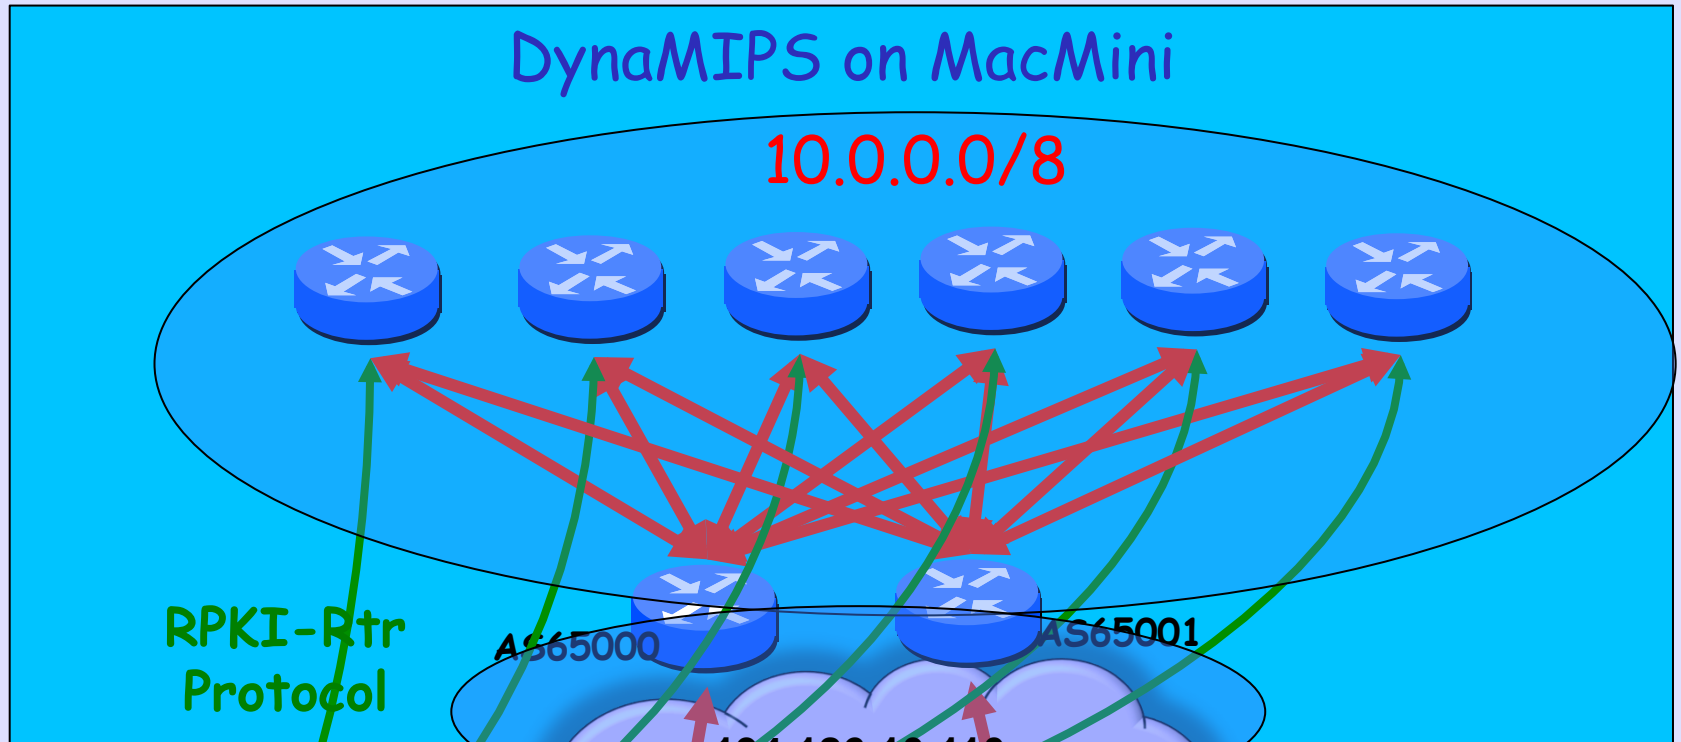
AS3130

Seattle

AS4128

Dallas

98.128.0.0/16
98.128.0.0/24
98.128.1.0/24
...
98.128.31.0/24



IP Address Allocation

98.128.0.0/16 ARIN Experimental Allocation

98.128.0.0/24 Instructors Play

98.128.1.0/24 labuser01

98.128.2.0/24 labuser02

...

98.128.32.0/24 labuser32

GUI Accounts

<https://ca0.vmini.rpki.net/>

<u>UserID</u>	<u>Password</u>
labuser01	fnord
labuser02	fnord
labuser03	fnord
...	
labuser16	fnord

https://ca0.vmini.rpki.net/

The screenshot shows a web interface for rpki.net. At the top, there is a navigation bar with the following items: "rpki.net", "Home", "Logged in as labuser00", and "Log Out". Below the navigation bar, there is a login form with the following elements:

- A label "Username" above a text input field containing "labuserNN".
- A label "Password" above a text input field containing "fnord".
- A "Login" button below the password field.

The Dashboard

labuser01

[dashboard](#)
[routes](#)

[export identity](#)


Resources

Resource	Valid Until	Parent
98.128.1.0/24	April 29, 2013, 10:18 p.m.	rgnet

 refresh


ROA Requests

Prefix	Max Length	AS
--------	------------	----

 Create

Children


Handle

 Import

Unallocated Resources


The following resources have not been allocated to a child, nor appear in a ROA.

IPv4

Prefix	Action
98.128.1.0/24	 ROA


Ghostbuster Requests

Full Name	Organization	Email	Telephone
-----------	--------------	-------	-----------

 Create

Parents

Handle
rgnet 

 Import

Create a ROA

labuser01

Create ROA Request

[dashboard](#)
[routes](#)

Prefix

Max Prefix Length

AS



Preview

Cancel

Routes matching your prefix:

Prefix	AS
--------	----

What Will Happen?

Confirm ROA Request

Please confirm that you would like to create the following ROA. The accompanying table indicates how the validation status may change as a result.

AS	Prefix	Max Length
4128	98.128.1.0/24	24

Matched Routes

Prefix	Origin AS	Validation Status
98.128.1.0/24	4128	valid
98.128.1.0/24	3130	invalid

Create

Cancel



Router Accounts

```
% telnet vmini.rpki.net 20NN
```

```
user: isplab
```

```
password: lab-PW
```

```
enable: lab-PW
```

Ask Me to Assign a Router to You

Configure RPKI Server

```
router bgp 651nn
```

```
  bgp rpki server tcp 172.16.10.192 port  
  43779 refresh 180
```

! inject your prefix into BGP

```
network 98.128.nn.0 mask 255.255.255.0
```

Check Server

```
r0.sea#show ip bgp rpki servers
```

```
BGP SOVC neighbor is 198.180.150.1/43779 connected to port 43779
```

```
Flags 0, Refresh time is 600, Serial number is 1304239609
```

```
InQ has 0 messages, OutQ has 0 messages, formatted msg 345
```

```
Session IO flags 3, Session flags 4008
```

```
Neighbor Statistics:
```

```
Nets Processed 624
```

```
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

```
Connection is ECN Disabled
```

```
Minimum incoming TTL 0, Outgoing TTL 255
```

```
Local host: 199.238.113.10, Local port: 57932
```

```
Foreign host: 198.180.150.1, Foreign port: 43779
```

```
Connection tableid (VRF): 0
```

Look at Table

```
router1#show ip bgp rpk table
```

```
76 BGP sovc network entries using 6688 bytes of memory
```

```
78 BGP sovc record entries using 1560 bytes of memory
```

Network	Maxlen	Origin-AS	Source	Neighbor
98.128.0.0/24	24	3130	0	198.180.150.1/4242
98.128.0.0/16	16	3130	0	198.180.150.1/4242
98.128.6.0/24	24	4128	0	198.180.150.1/4242
98.128.9.0/24	24	3130	0	198.180.150.1/4242
98.128.30.0/24	24	1234	0	198.180.150.1/4242
128.224.1.0/24	24	3130	0	198.180.150.1/4242
129.6.0.0/17	17	49	0	198.180.150.1/4242
129.6.112.0/24	24	10866	0	198.180.150.1/4242
129.6.128.0/17	17	49	0	198.180.150.1/4242
147.28.0.0/16	16	3130	0	198.180.150.1/4242

Look at BGP Table

```
r0.sea#sh ip bgp
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
* i	I198.180.150.0	144.232.9.61	100	0	1239	3927 i
*>	I	199.238.113.9	0	2914	3927	i
*	I	129.250.11.41	0	2914	3927	i
*>	V198.180.152.0	199.238.113.9	0	2914	4128	i
*	V	129.250.11.41	0	2914	4128	i
*>	N198.180.155.0	199.238.113.9	0	2914	22773	i
*	N	129.250.11.41	0	2914	22773	i
*>	N198.180.160.0	199.238.113.9	0	2914	23308	13408 5752 i
*	N	129.250.11.41	0	2914	23308	13408 5752 i

Look at a Prefix

```
R3#show ip bgp 98.128.0.0/24
```

```
BGP routing table entry for 98.128.0.0/24, version 360
```

```
Paths: (2 available, best #1, table default)
```

```
65000 3130
```

```
10.0.0.1 from 10.0.0.1 (193.0.24.64)
```


```
Origin IGP, localpref 100, valid, external, best  
path 680D859C RPKI State valid
```

```
65001 4128
```


```
10.0.1.1 from 10.0.1.1 (193.0.24.65)
```

```
Origin IGP, localpref 100, valid, external  
path 680D914C RPKI State invalid
```


Fat-Finger Detected



ROA Controls Validity



Try Your Own /24



Mis-Origination Caught

```
R3#sh ip bgp 98.128.0.0/24
```

```
BGP routing table entry for 98.128.0.0/24, version 94
```

```
Paths: (2 available, best #2, table default)
```

```
  Advertised to update-groups:
```

```
    1
```

```
  Refresh Epoch 1
```

```
 65000 3130
```

```
    10.0.0.1 from 10.0.0.1 (65.38.193.12)
```

```
      Origin IGP, localpref 100, valid, external
```

```
      path 6802D4DC RPKI State invalid
```

```
  Refresh Epoch 1
```

```
 65001 4128
```

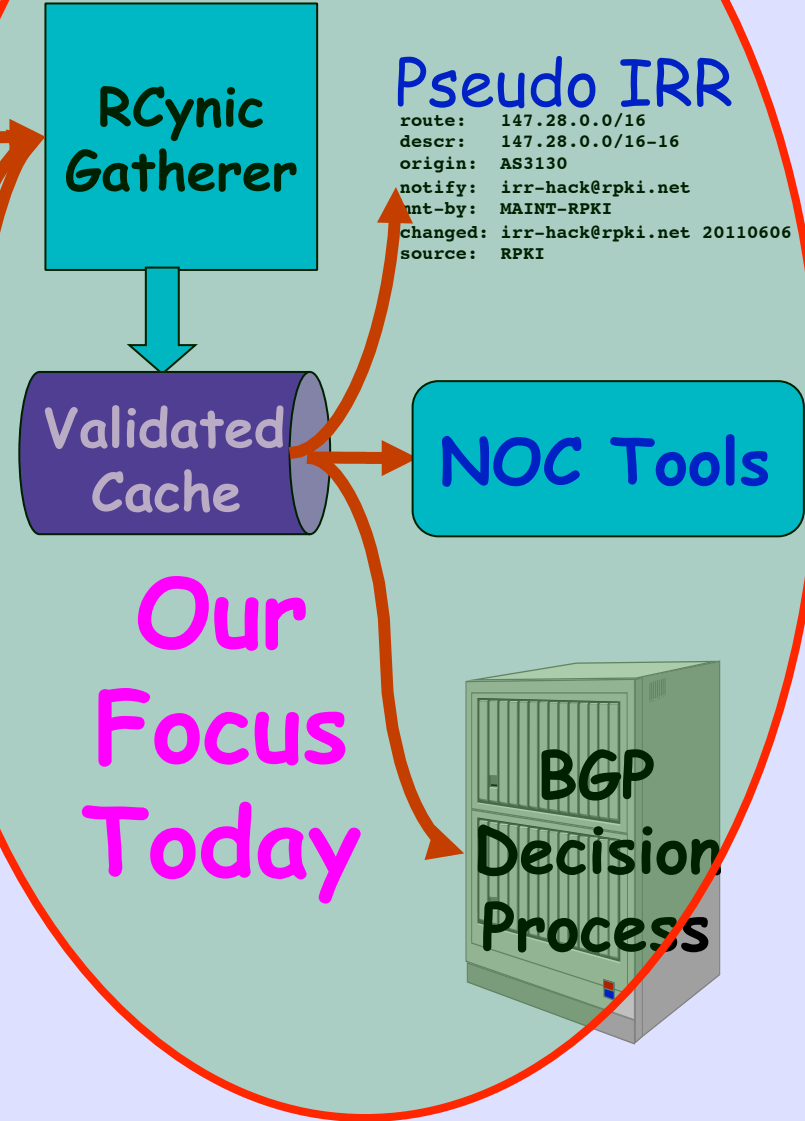
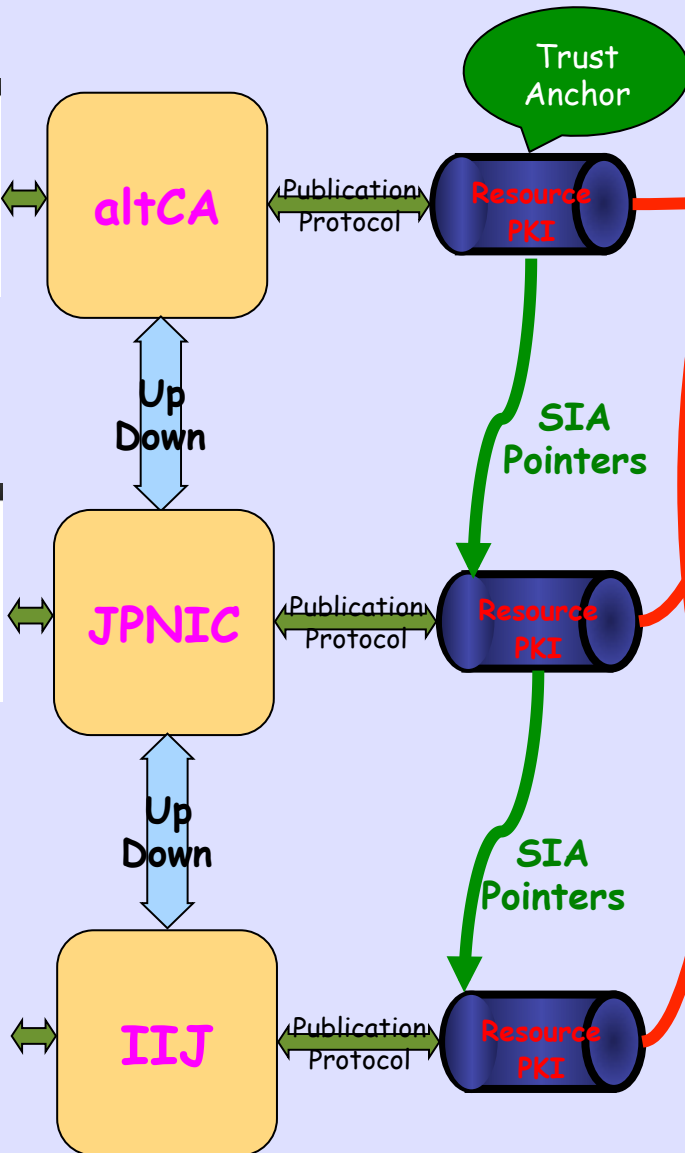



```
    10.0.1.1 from 10.0.1.1 (65.38.193.13)
```

```
      Origin IGP, localpref 100, valid, external, best
```

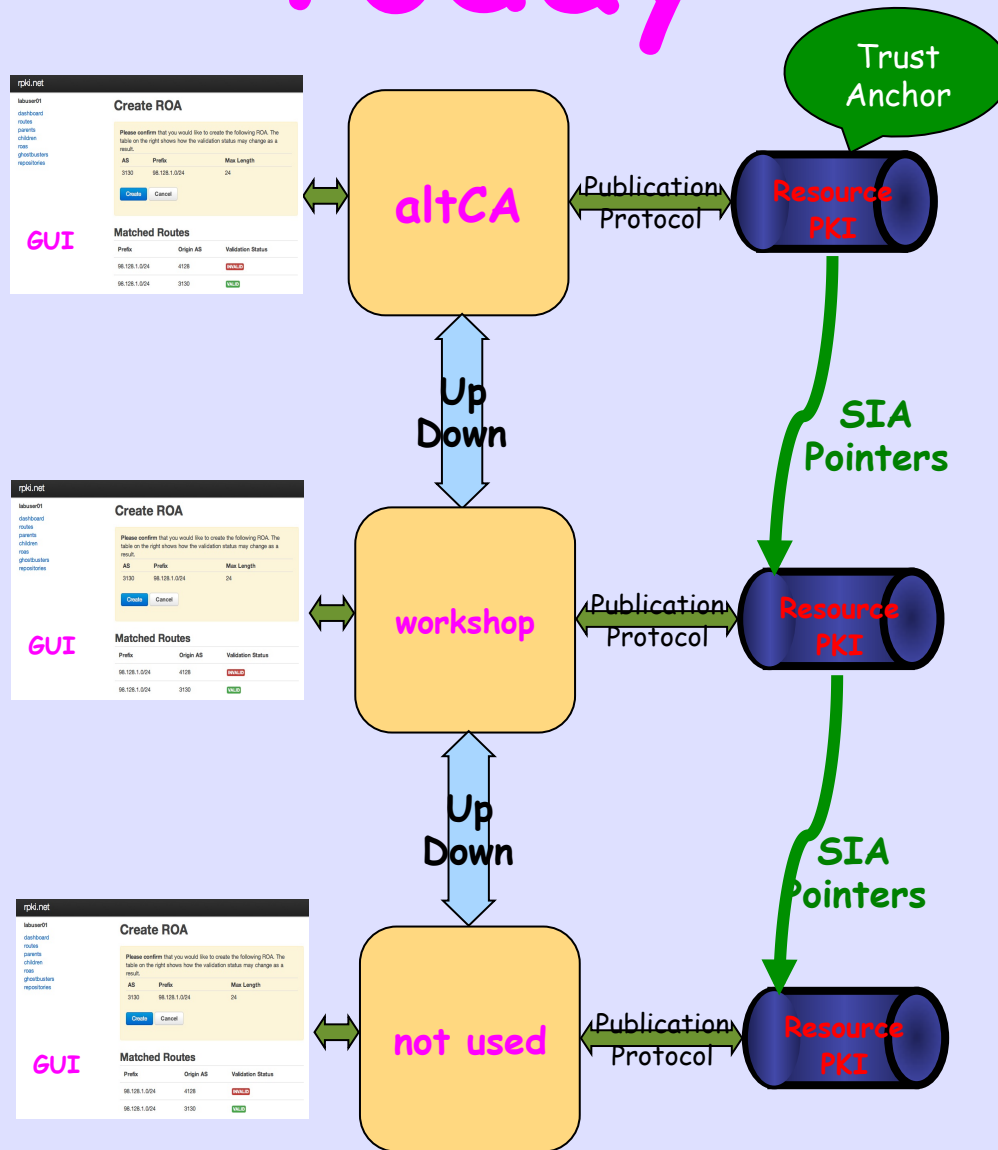
```
      path 6802D7C8 RPKI State valid
```

Issuing Parties

Relying Parties



Today



altCA TAL

So You Can Copy & Paste

`rsync://ca0.rpki.net/tal/root.cer`

```
MIIBIjANBgkqhkiG9w0BAQEFAA0CAQ8AMIIBCgKCAQEAzPSPpQxa0rxz9sbbvYGr
UlpgyBVxSI+t+k/WDKbr+VW7CjUoz6cc5KMFANKQWw3D6ER4kCwX4WJkD58AGGbw/
WeAe6m3aHc0RUVRkr45a4qSrYiG7Wq9RAXtwbhLXofB3zo+090I1XDaVP2U9bw+Q
yoJBJuAmZONt0bRgrktv8QhVtKvuYkH5ZIE7DkXJcJzBn6gv09dZsdwZm3xV3soX
HEKrz5pY6Sb2xoL1CyPqzG0fVFXl0G5+dmcD/degPKxrEycAzjnHUzN1gus2jg26
dtkix7KG/Mn1h/k53j0FdQD+zqPwakgwqjvC0dSdHMRmsikj0EF9WrZIOjZUXV6q
6wIDAQAB
```


altCA TAL



Also At
<https://wiki.rg.net/wiki/RPKI>

Routers

- **Use Your Own! (in production images from C&J)**
- **16 DynaMIPS 7200s**
- **7200s in Live Internet**

Be Careful !

- Some Caches Have a LOT of ROAs
- Do Not Configure DynaMIPS to a Server With RIR TALs Because RIPE Data Has Thousands of ROAs
- dfw0, 198.180.152.11 Has Full BGP Table if you want to crash DynaMIPS

7200s in Live Internet

```
$ ssh demo@r0.sea.rg.net
```

```
Password: demo
```

or

```
$ ssh demo@r1.iad.rg.net
```

```
Password: demo
```

Sorry, no enable 😊

BGP Configuration

```
router bgp 651nn
```

```
  bgp rpki server tcp 172.16.10.192 \
```

```
    port 43779 refresh 600
```

! inject your prefix into BGP

```
  network 98.128.nn.0 255 255 255.0
```

That's All!!

Take a Look!

```
r0.sea#show ip bgp 198.180.150.0
```

```
1239 3927
```

```
144.232.9.61 (metric 11) from 147.28.7.2 (147.28.7.2)
```

```
Origin incomplete, metric 750, localpref 110, valid,  
internal, best
```

```
path 14CB1C54 RPKI State valid
```

```
2914 3927
```

```
199.238.113.9 from 199.238.113.9 (129.250.0.167)
```

```
Origin IGP, metric 90, localpref 100, valid, external
```

```
Community: 2914:410 2914:1001 2914:2000 2914:3000 3130:380
```

```
path 196F0E70 RPKI State valid
```

Cisco Adventure

```
r0.sea#show ip bgp rpki ?
```

```
servers  Display RPKI cache server information
```

```
table    Display RPKI table entries
```

Install RP Cache SW

On Your Laptop

(FreeBSD, Ubuntu, ...)

or

Remote (Virtual) Ubuntu Server

Borrowing a VM

- Servers in Ashburn, Virginia US
- Ubuntu 12.04.2 LTS Server
- ssh labuserNN@ubuNN.rpki.net
- Password is Vattood1
- You are in sudoers
- Do not worry, machines are snapshotted

Installing RP Package

Copy and Paste from

<https://trac.rpki.net/wiki/doc/RPKI/Installation/UbuntuPackages>

```
$ wget -q -O - http://download.rpki.net/APT/apt-gpg-key.asc | sudo apt-key add -
```

```
$ sudo wget -q -O /etc/apt/sources.list.d/rpki.list http://download.rpki.net/APT/rpki.ubuntu.list
```

```
$ sudo apt-get update
```

```
$ sudo apt-get install rpki-rp
```

Is RPKI-Rtr Running?

```
$ rtr-origin --client tcp localhost 43779
```

```
2013-03-09 23:10:19 rtr-origin/client[17089]: [Startup]
```

```
2013-03-09 23:10:19 rtr-origin/client[17089]: [Starting raw TCP  
connection to localhost:43779]
```

```
2013-03-09 23:10:19 rtr-origin/client[17089]: [reset_query]
```

```
2013-03-09 23:10:19 rtr-origin/client[17089]: [cache_response,  
nonce 37394]
```

```
2013-03-09 23:10:19 rtr-origin/client[17089]: + 12322
```

```
78.192.0.0/10-10
```

```
00:04:00:00:00:00:00:14:01:0A:0A:00:4E:C0:00:00:00:00:30:22
```

```
2013-03-09 23:10:19 rtr-origin/client[17089]: + 3320
```

```
79.192.0.0/10-10
```

```
00:04:00:00:00:00:00:14:01:0A:0A:00:4F:C0:00:00:00:00:0C:F8
```

Select Trust Anchors

- We are going to use emulated routers on DynaMIPS, and it is memory limited
- So we can not feed it all the ROAs in the Global RPKI
- Edit `/etc/rcynic.conf`
- Leave only these TALs

```
trust-anchor-locator.1 = /etc/rpki/trust-anchors/altCA.tal
```

```
trust-anchor-locator.2 = /etc/rpki/trust-anchors/rpki.net-testbed.tal
```

Speed it Up

- We do not want to wait an hour for ROAs to propagate
- Make rcynic run frequently
- `$ sudo crontab -e -u rcynic`

```
0/5 * * * * exec /usr/bin/rcynic-cron
```

Configure Router to Use Your Cache

```
router bgp 651nn
```

```
bgp rpki server tcp 147.28.0.11n \  
port 43779 refresh 180
```

! inject your prefix into BGP

```
network 98.128.nn.0 255 255 255.0
```

Check Server

```
r0.sea#show ip bgp rpki servers
```

```
BGP SOVC neighbor is 198.180.150.1/43779 connected to port 43779
```

```
Flags 0, Refresh time is 600, Serial number is 1304239609
```

```
InQ has 0 messages, OutQ has 0 messages, formatted msg 345
```

```
Session IO flags 3, Session flags 4008
```

```
Neighbor Statistics:
```

```
Nets Processed 624
```

```
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

```
Connection is ECN Disabled
```

```
Minimum incoming TTL 0, Outgoing TTL 255
```

```
Local host: 199.238.113.10, Local port: 57932
```

```
Foreign host: 198.180.150.1, Foreign port: 43779
```

```
Connection tableid (VRF): 0
```

Look at ROA/VRF Table

```
router1#show ip bgp rpki table
```

```
76 BGP sovc network entries using 6688 bytes of memory
```

```
78 BGP sovc record entries using 1560 bytes of memory
```

Network	Maxlen	Origin-AS	Neighbor
98.128.0.0/24	24	3130	198.180.150.1/42420
98.128.0.0/16	16	3130	198.180.150.1/42420
98.128.6.0/24	24	4128	198.180.150.1/42420
98.128.9.0/24	24	3130	198.180.150.1/42420
98.128.30.0/24	24	1234	198.180.150.1/42420
128.224.1.0/24	24	3130	198.180.150.1/42420
129.6.0.0/17	17	49	198.180.150.1/42420
129.6.112.0/24	24	10866	198.180.150.1/42420
129.6.128.0/17	17	49	198.180.150.1/42420
147.28.0.0/16	16	3130	198.180.150.1/42420

Look at BGP Table

```
r0.sea#sh ip bgp
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
* i	I198.180.150.0	144.232.9.61	100	0	1239 3927	i
*>	I	199.238.113.9	0	2914	3927	i
*	I	129.250.11.41	0	2914	3927	i
*>	V198.180.152.0	199.238.113.9	0	2914	4128	i
*	V	129.250.11.41	0	2914	4128	i
*>	N198.180.155.0	199.238.113.9	0	2914	22773	i
*	N	129.250.11.41	0	2914	22773	i
*>	N198.180.160.0	199.238.113.9	0	2914	23308 13408	5752 i
*	N	129.250.11.41	0	2914	23308 13408	5752 i

Look at a Prefix

```
R3#show ip bgp 98.128.0.0/24
```

```
BGP routing table entry for 98.128.0.0/24, version 360
```

```
Paths: (2 available, best #1, table default)
```

```
65000 3130
```

```
10.0.0.1 from 10.0.0.1 (193.0.24.64)
```

```
Origin IGP, localpref 100, valid, external, best  
path 680D859C RPKI State valid
```

```
65001 4128
```

```
10.0.1.1 from 10.0.1.1 (193.0.24.65)
```

```
Origin IGP, localpref 100, valid, external  
path 680D914C RPKI State invalid
```

Fun Things to Do

- Look at your neighbors' prefixes
- Use the GUI to change the ROA for your prefix, and wait for it to propagate
- Attack your neighbor's prefix by configuring your bgp to announce it
`router bgp 651nn`
`network 98.128.<neighbor>.0 255 255 255.0`

Notice it Did Not Work

```
router bgp 651nn  
  network 98.128.<neighbor>.0 255 255 255.0
```

```
show ip bgp 98.128.<neighbor>.0
```

- It is marked Invalid on your own router! It caught you injecting a bad prefix. To cheat you need to

```
router bgp 651nn  
  no bgp bestpath prefix-validate local
```

- Now ask others to look for your announcement. They should see it as Invalid

Now You Know
How to Prevent
YouTube Incident-2
And Stay Out of
The Newspapers

Please Do
Try This
At Home