The RPKI and BGP Origin Validation

APRICOT / New Delhi

2012.02.27

Randy Bush <randy@psg.com>

Rob Austein <sra@isc.org>

Steve Bellovin <smb@cs.columbia.edu>

And a cast of thousands! Well, dozens :)

Why Origin Validation?

- Prevent YouTube accident
- 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' and locking the data plane to the control plane, the third step, a few years away

Prefix Has Origin AS

BGP routing table entry for 98.128.1.0/24

Paths: (32 available, best #21, table Default-IP-Routing-Table)

Three Pieces

- **RPKI** – Resource Public Key Infrastructure, the Certificate Infrastructure to Support the other Pieces (starting last year)
- **Origin Validation** – Using the RPKI to detect and prevent mis-originations of someone else's prefixes (early 2012)
- **AS-Path Validation** AKA **BGPsec** – Prevent Attacks on BGP (future work)

Resource Public Key Infrastructure (RPKI)

X.509 RPKI Being Developed & Deployed by IANA, RIRs, and Operators

Private/Public Keys

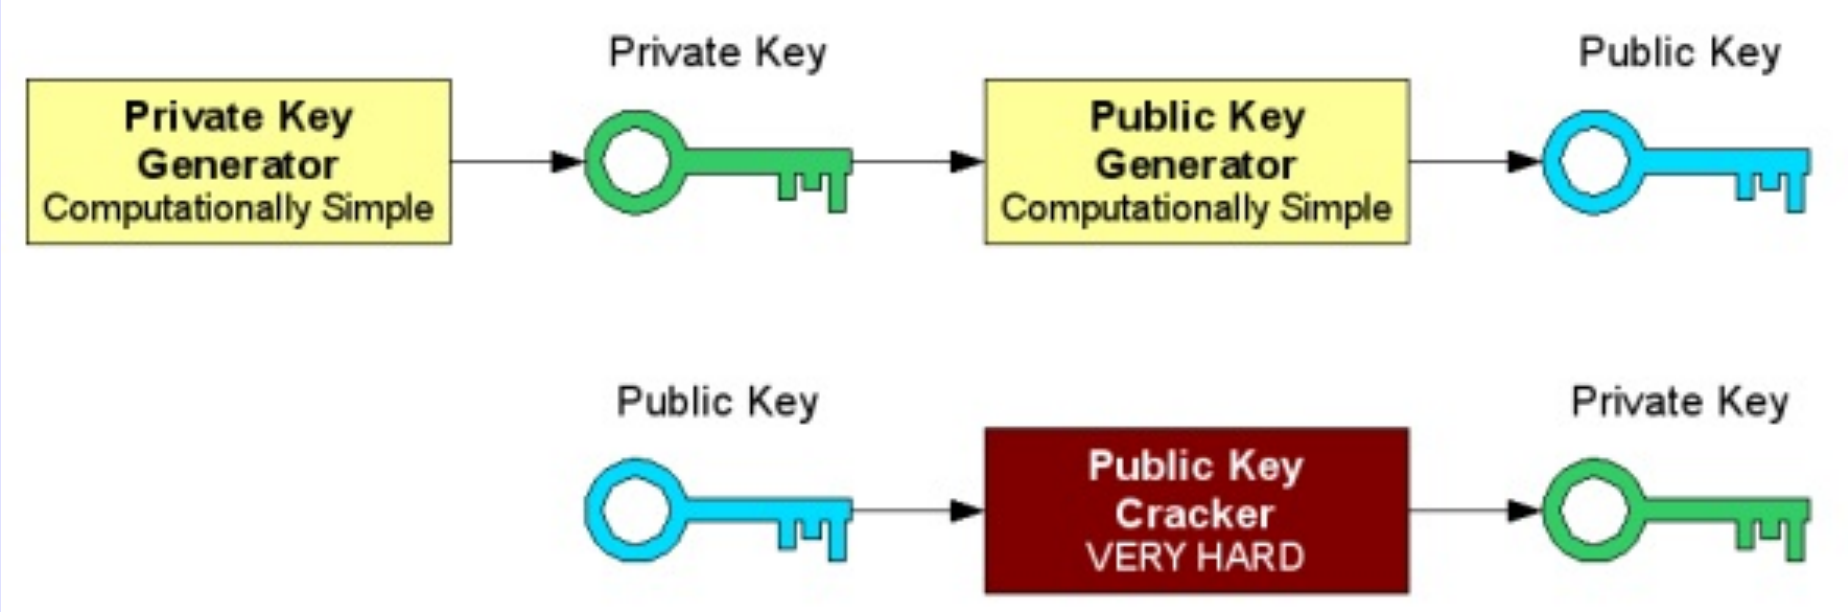

Stolen from - http://gdp.globus.org/gt4-tutorial/multiplehtml/ch09s03.html

En/DeCryption

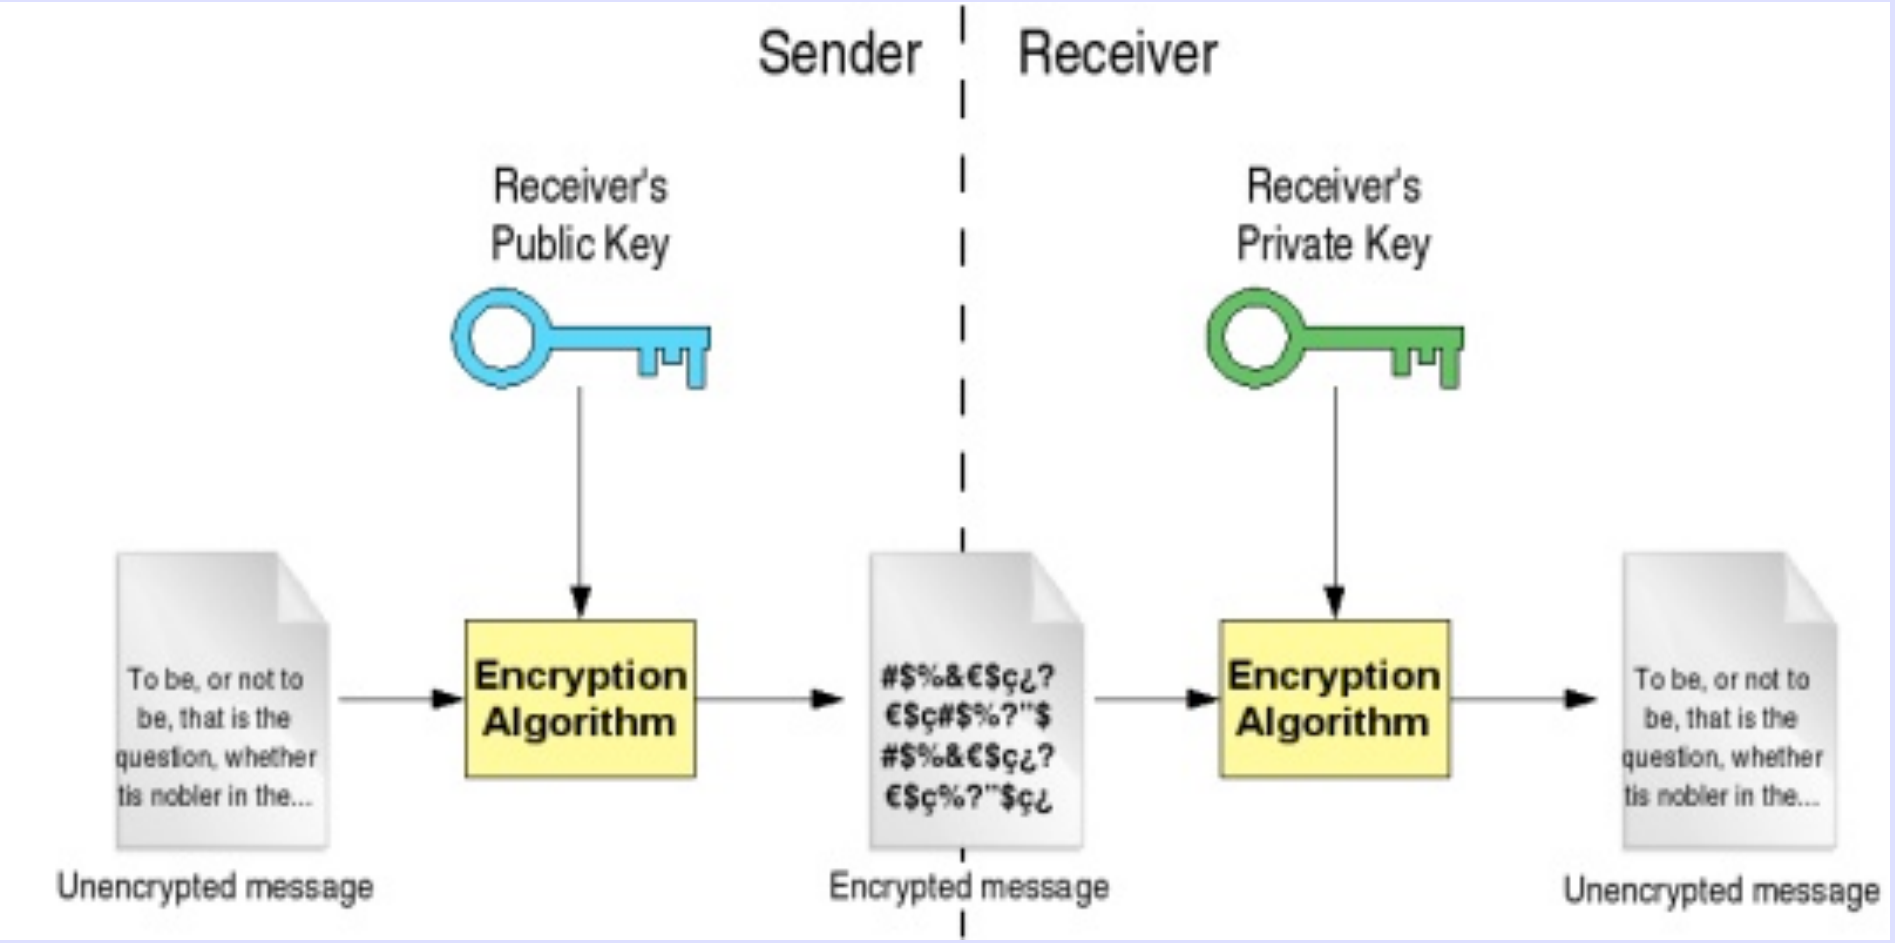

Digital Signature

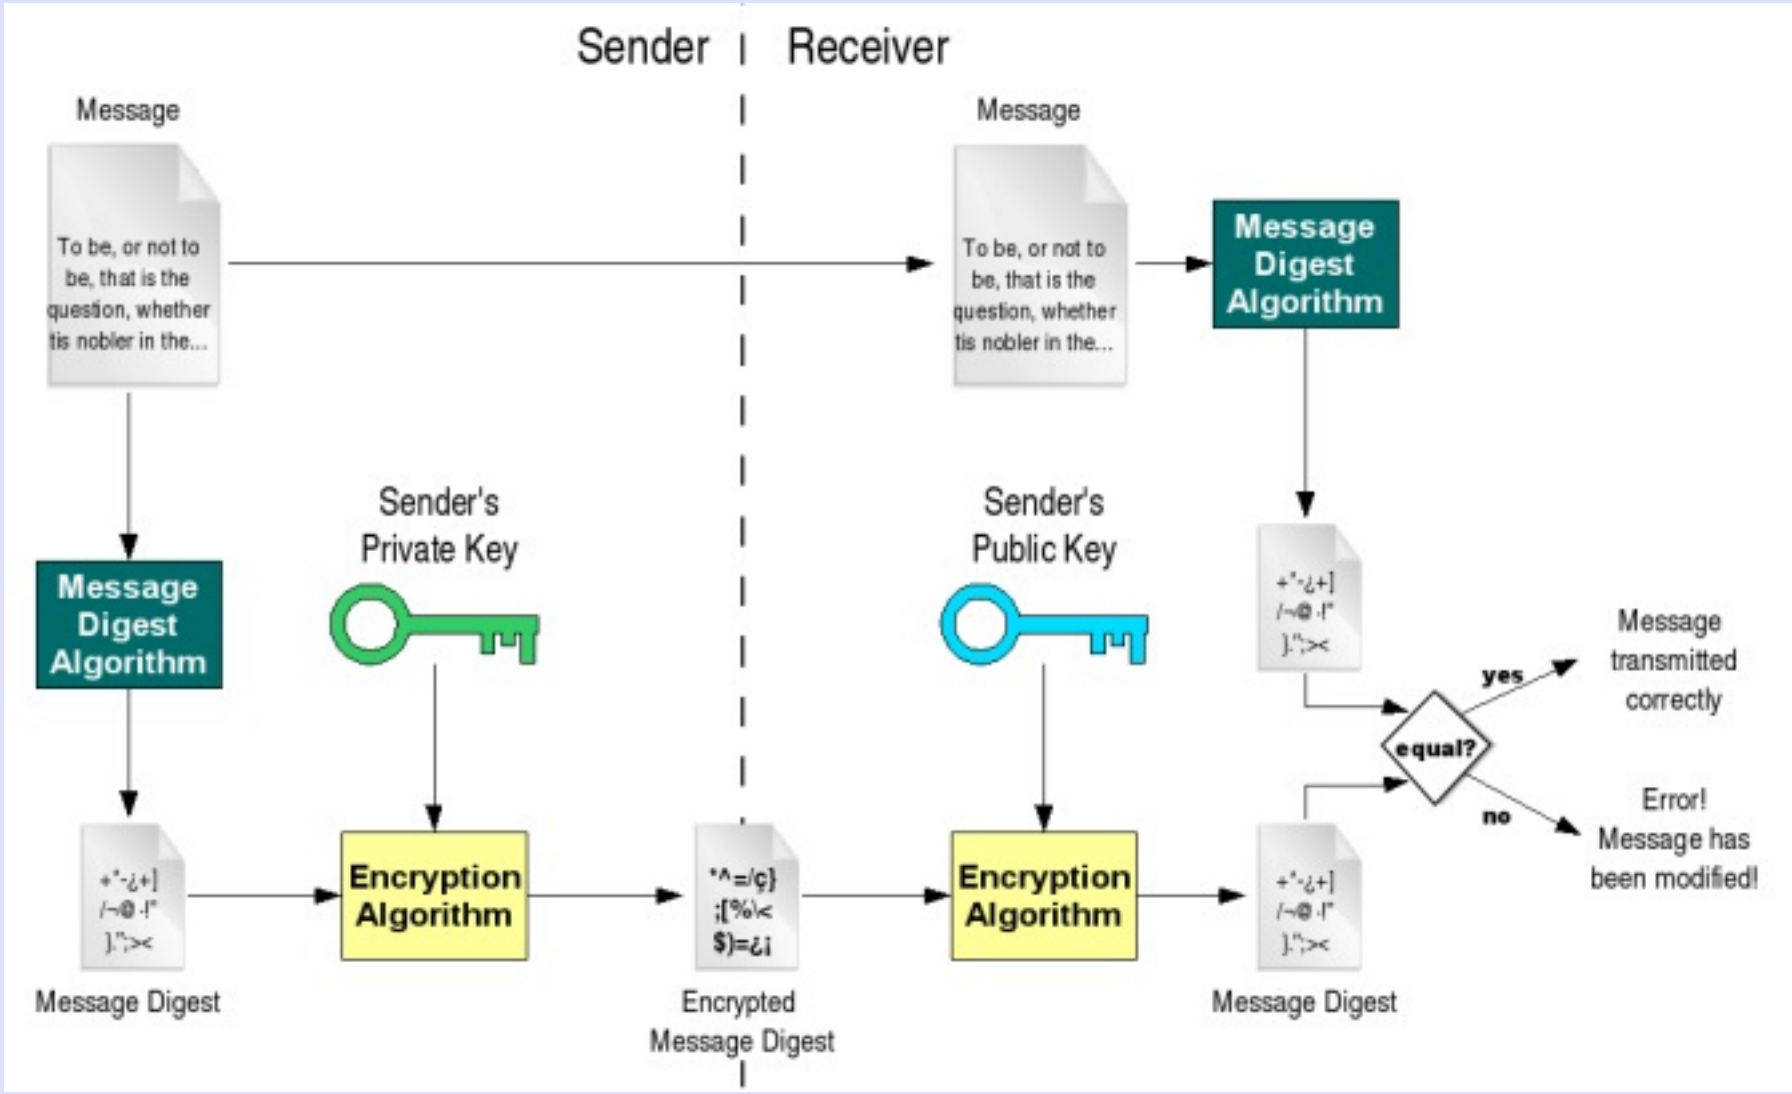

X.509 Certificate w/ 3779 Ext

Signed by Parent's Private Key

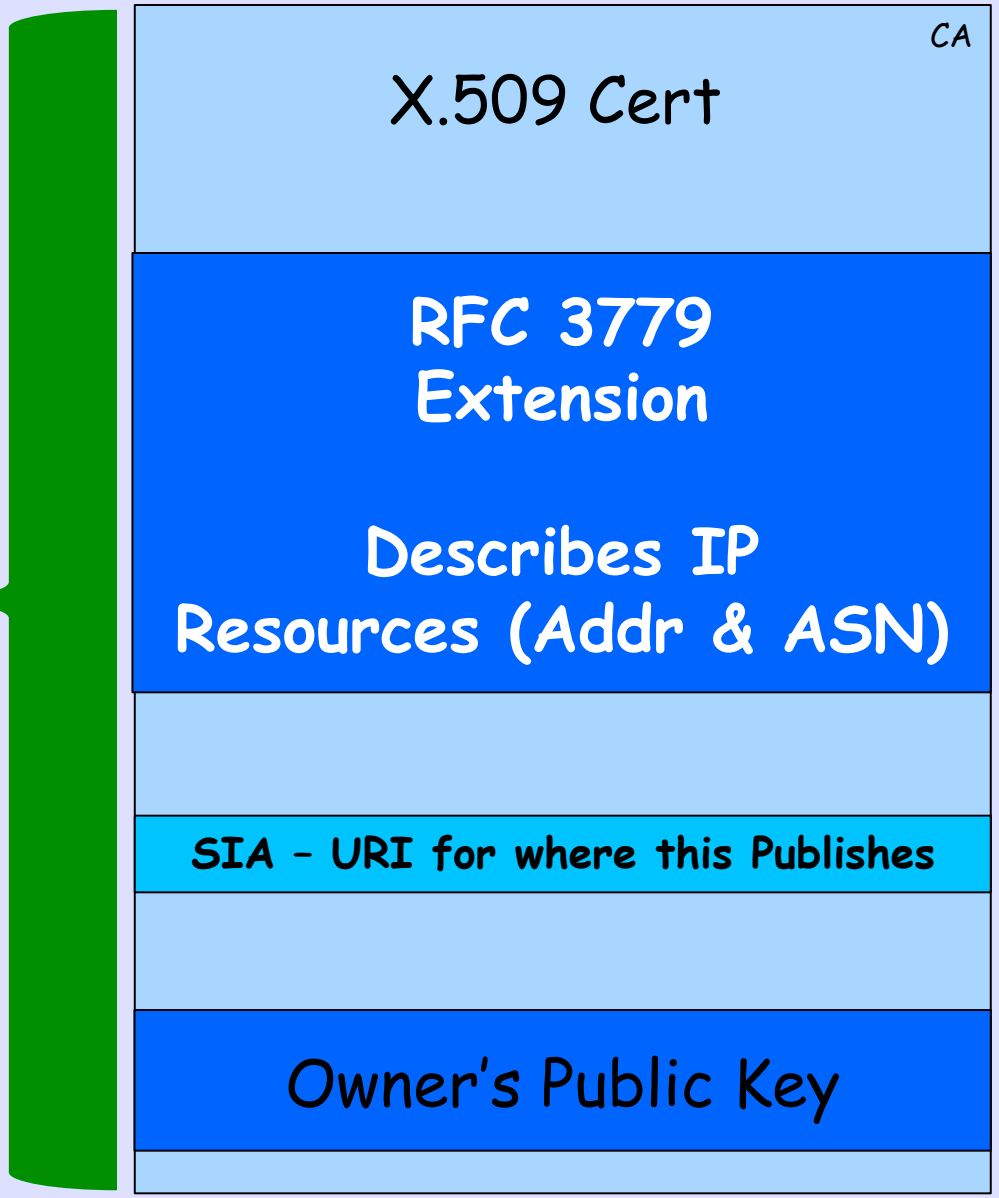

Certificate Hierarchy follows Allocation Hierarchy

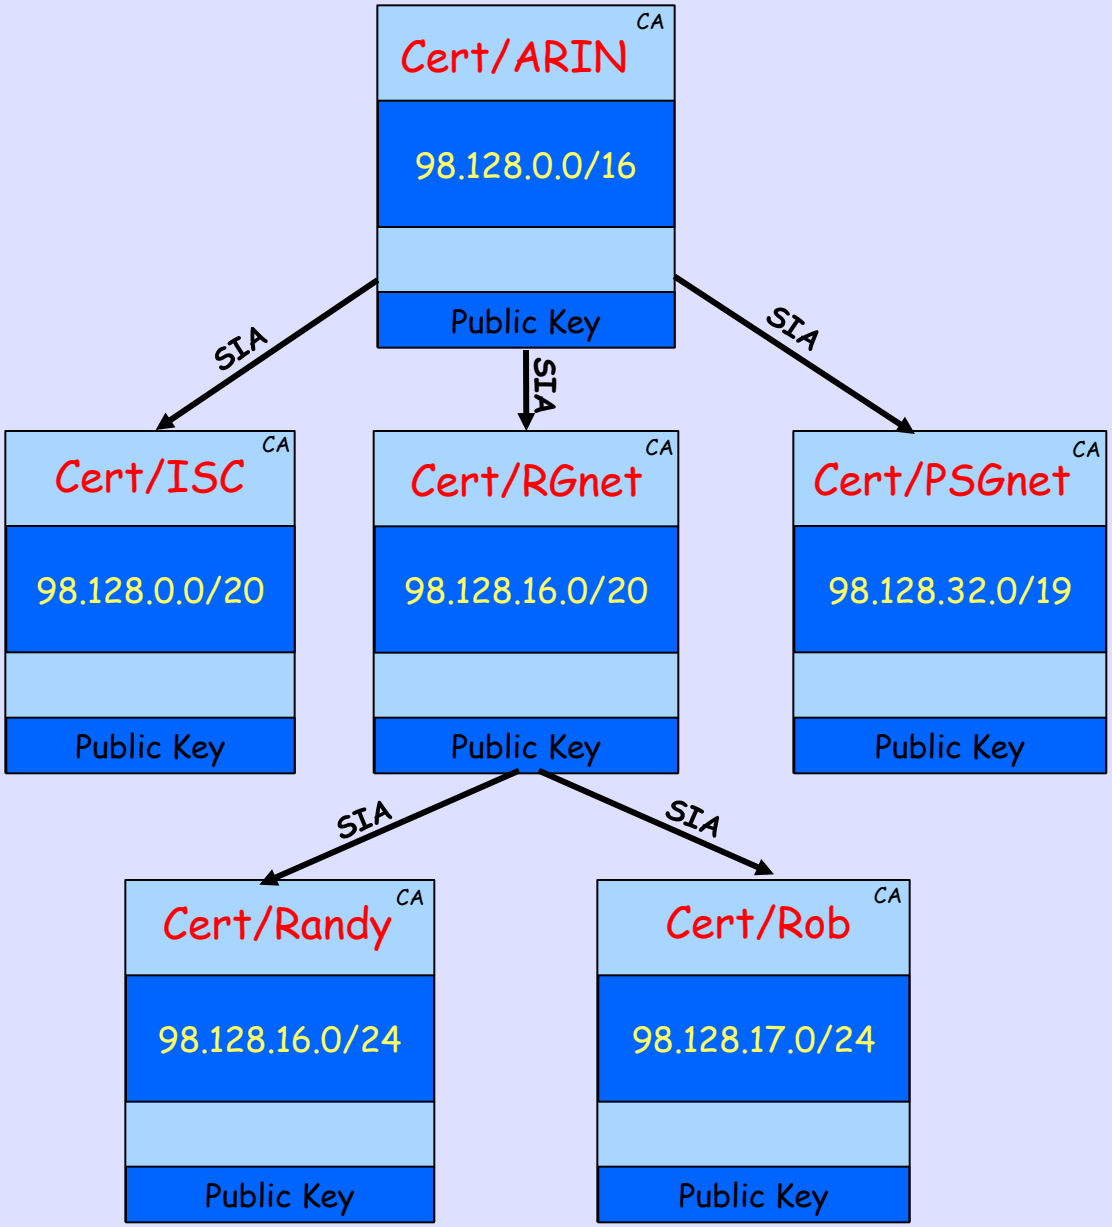

That's Who Owns It
but
Who May Route It?

Route Origin Authorization (ROA)

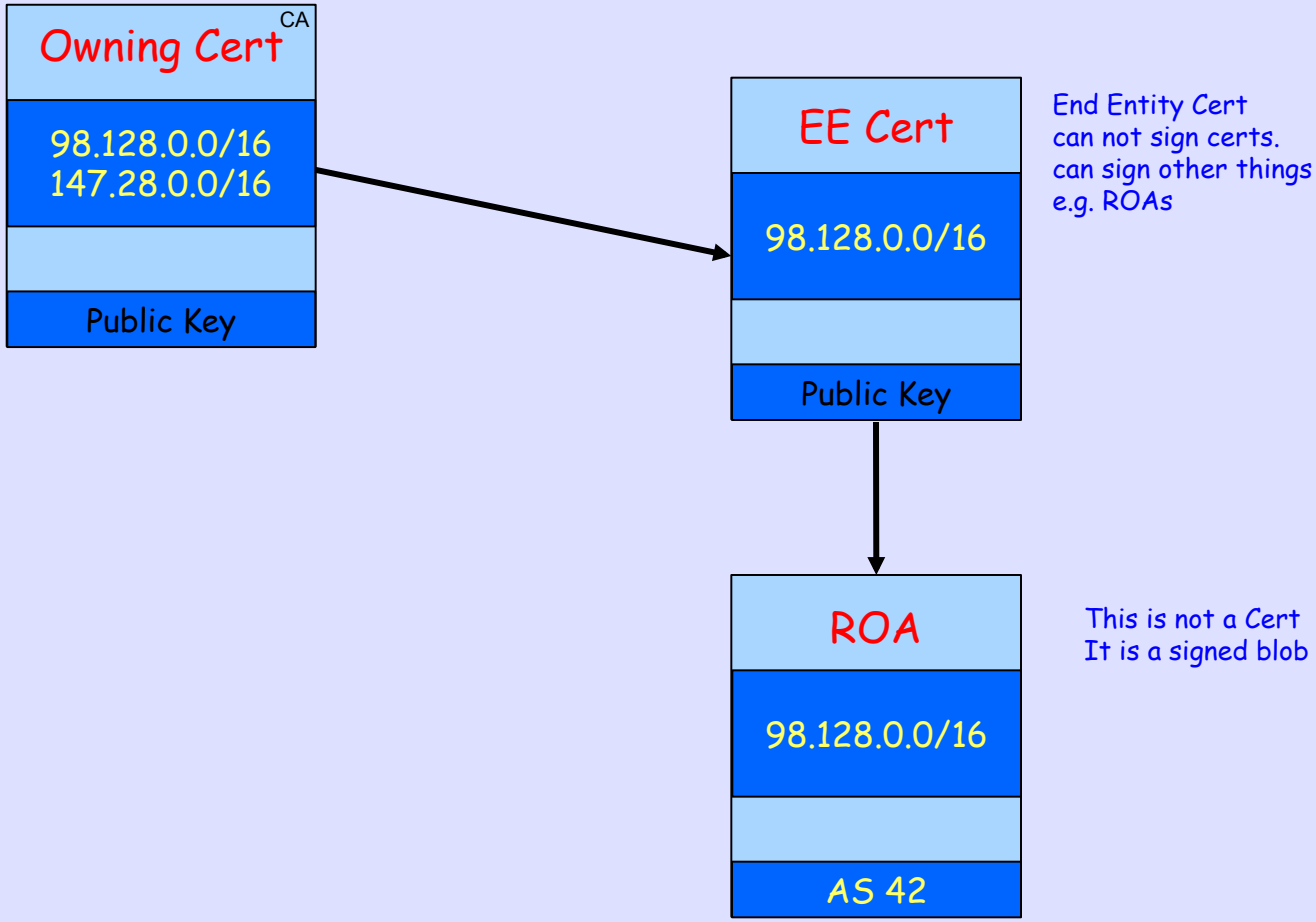

Multiple ROAs Make Before Break

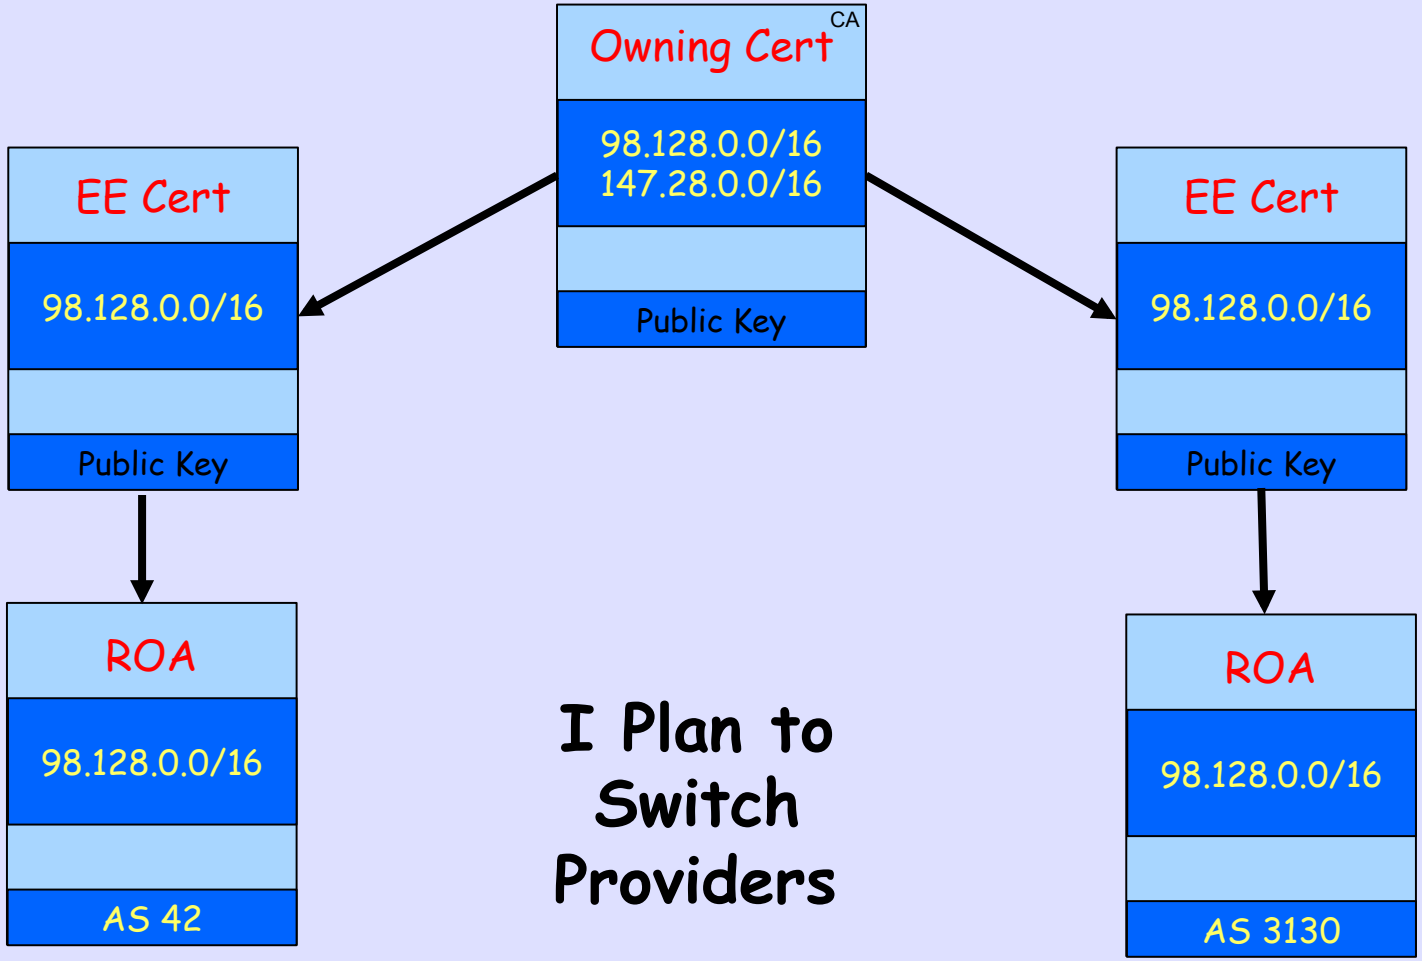

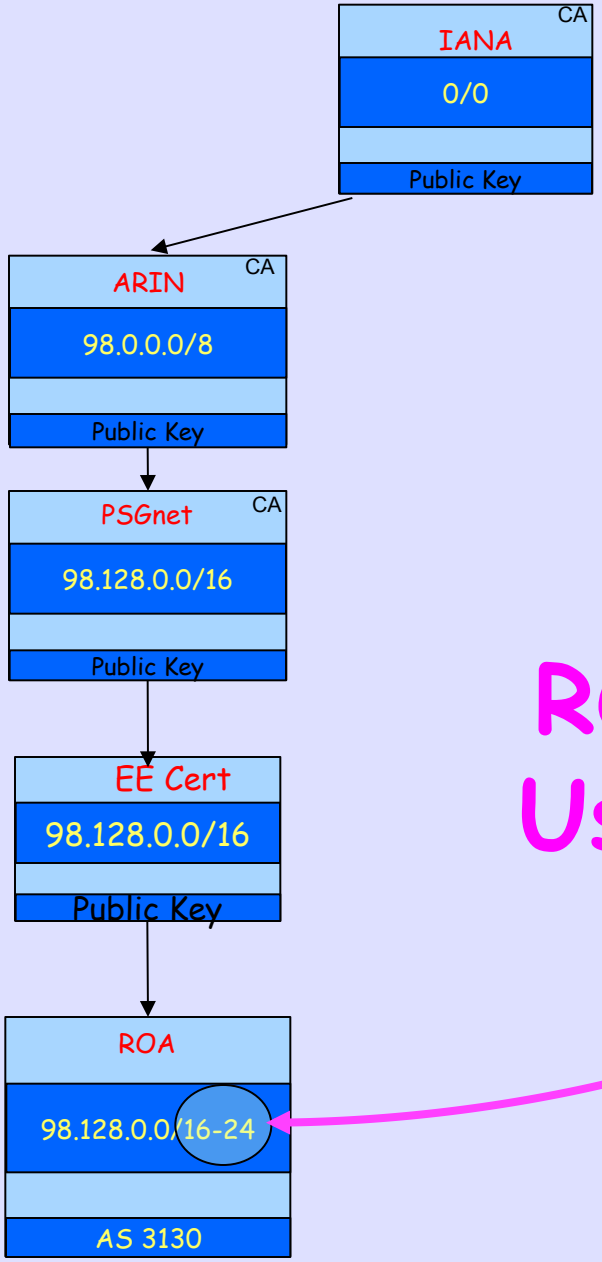

ROA Aggregation Using Max Length

RPKI-Based Origin Validation

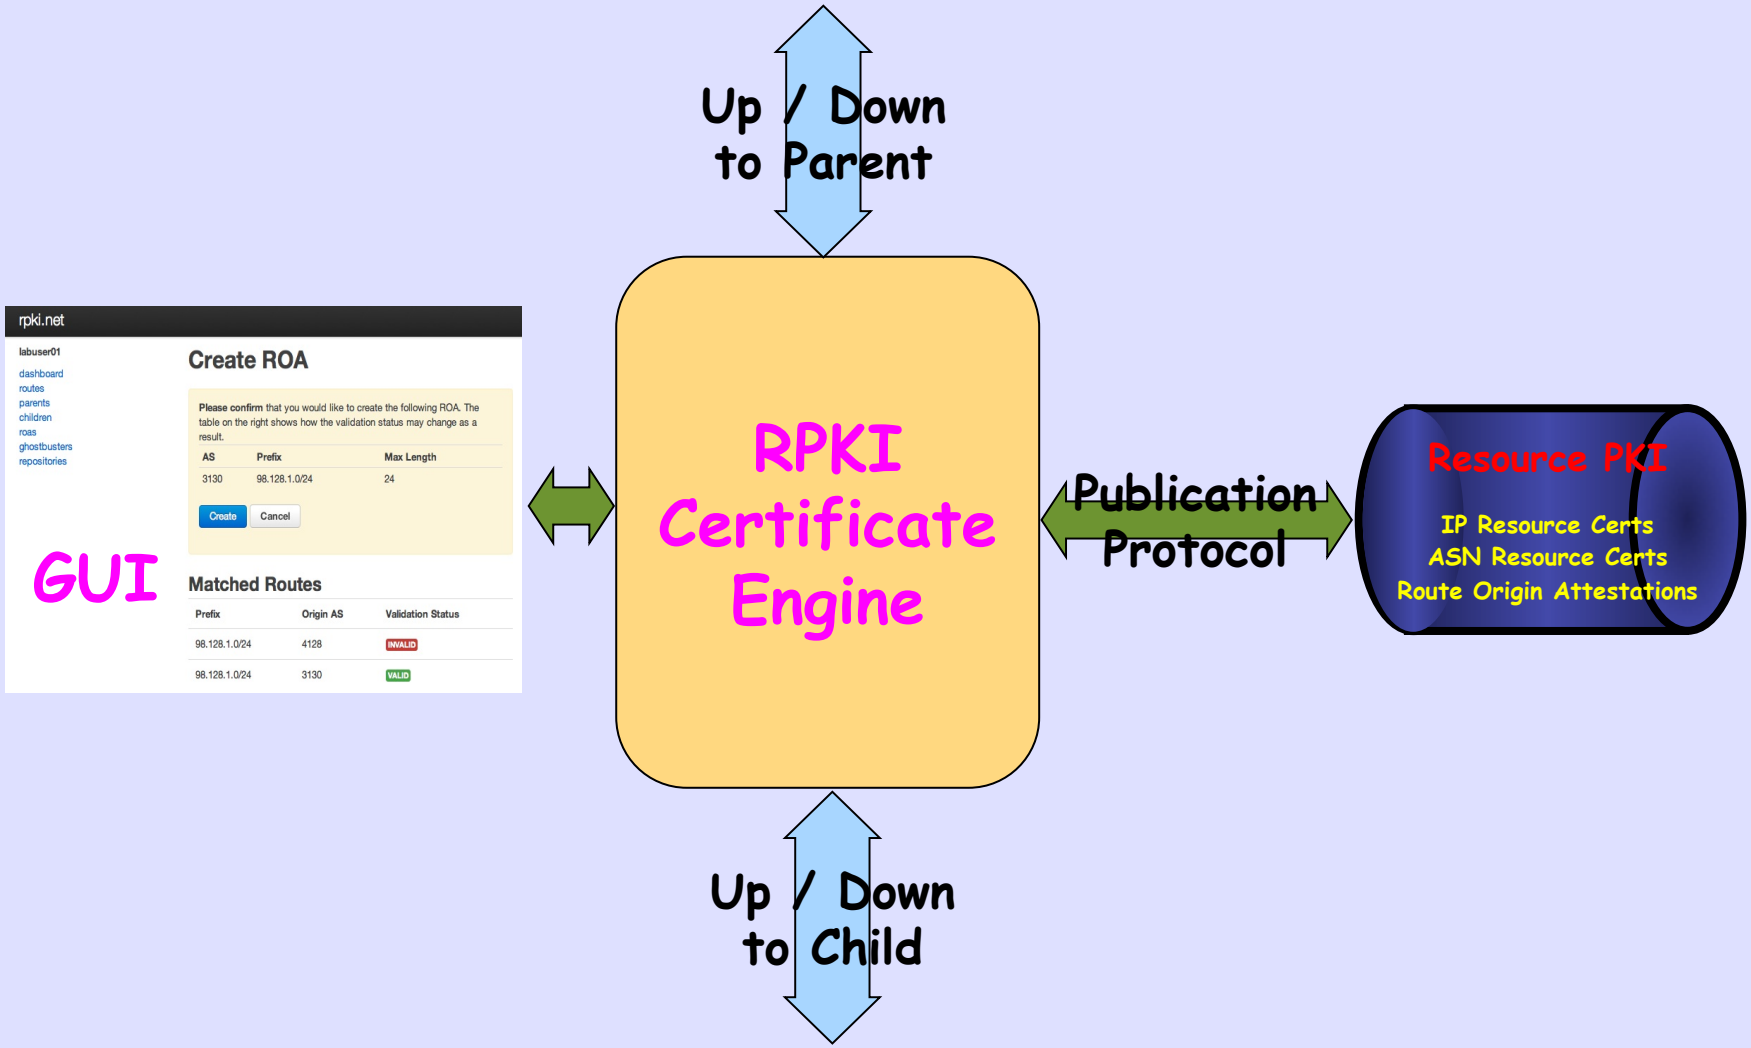
Up / Down
to Parent
rpki.net
labuser01
dashboard
routes
parents
children
roas
ghostbusters
repositories
Create ROA
Please confirm that you would like to create the following ROA. The table on the right shows how the validation status may change as a result.
AS
Prefix
Max Length
3130
98.128.1.0/24
24
Create
Cancel
GUI
Matched Routes
Prefix
Origin AS
Validation Status
98.128.1.0/24
4128
INVALID
98.128.1.0/24
3130
VALID
RPKI
Certificate
Engine
Publication
Protocol
Resource PKI
IP Resource Certs
ASN Resource Certs
Route Origin Attestations
Up / Down
to Child

Warning What ROA Will Do

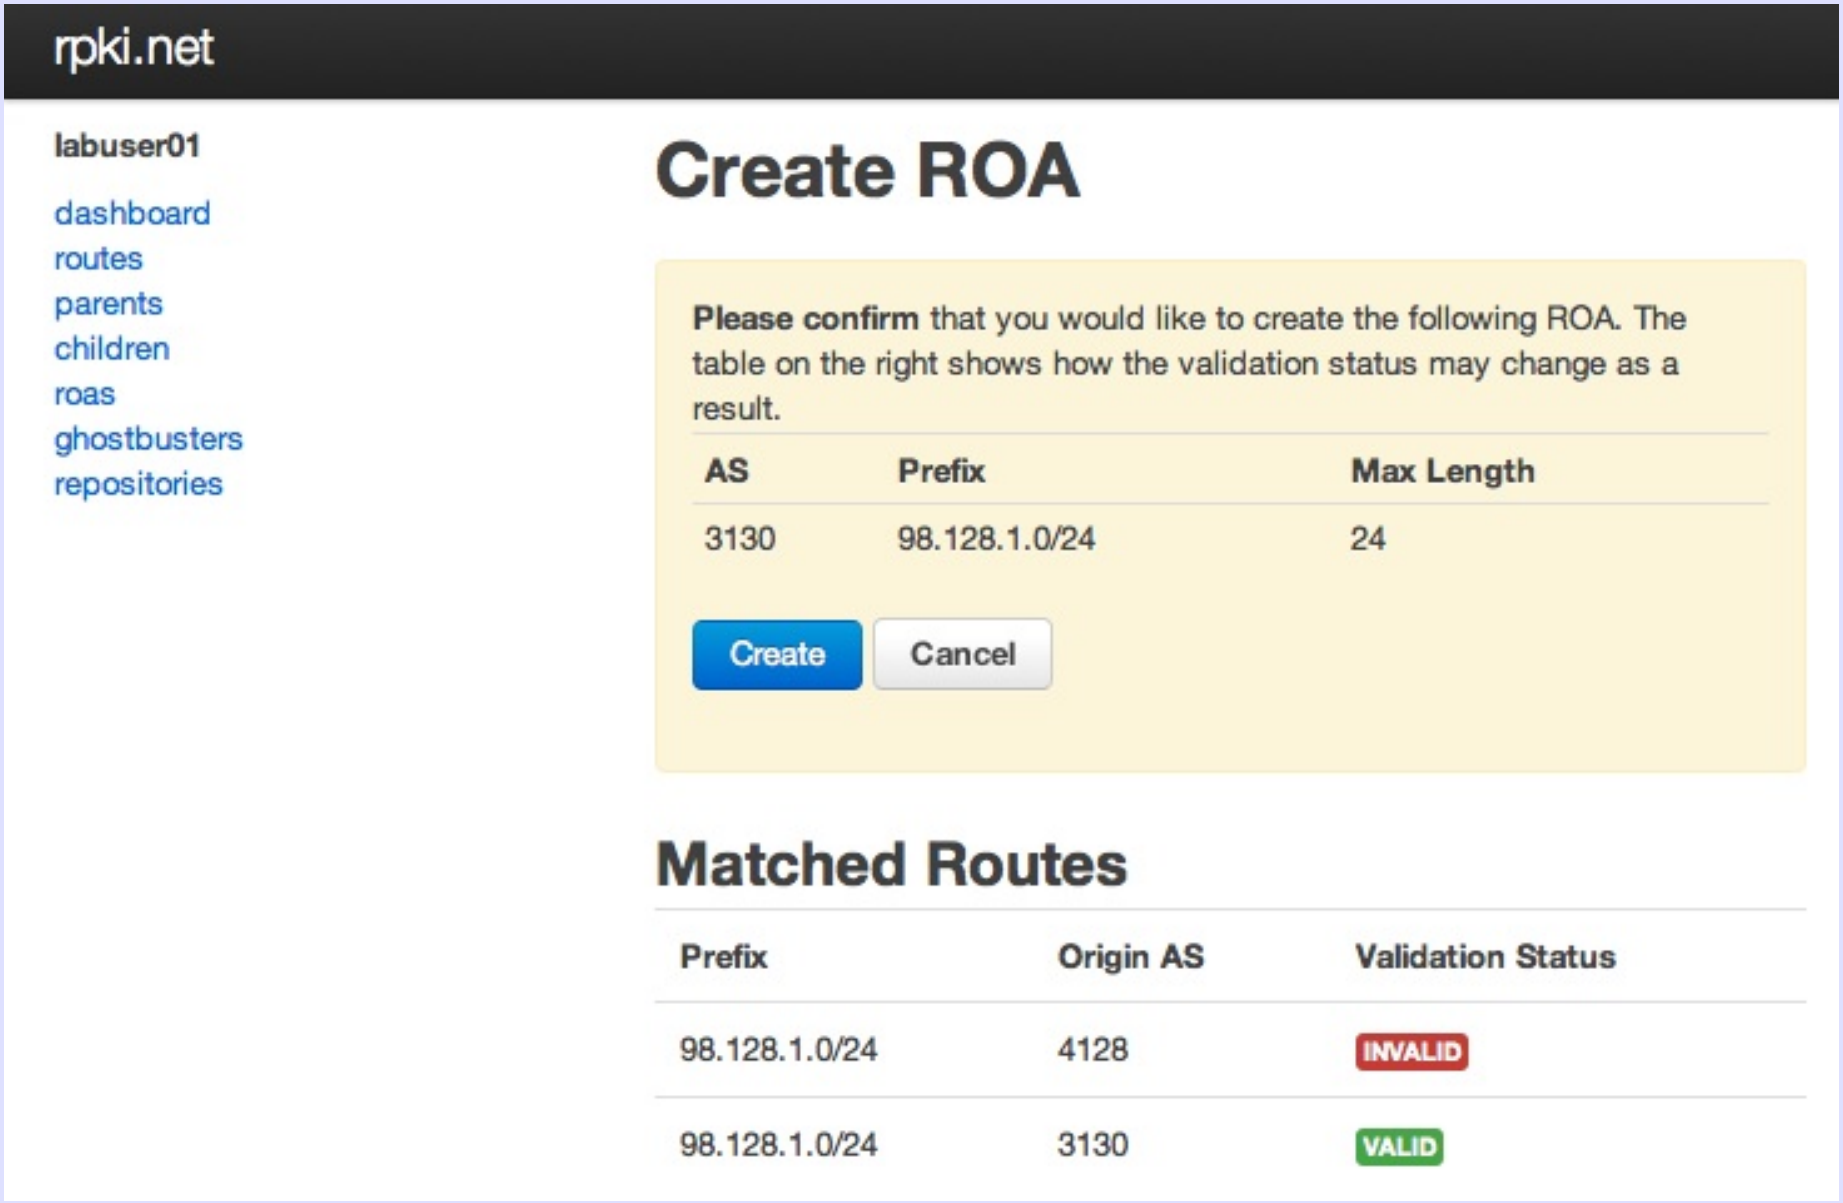

Issuing Parties

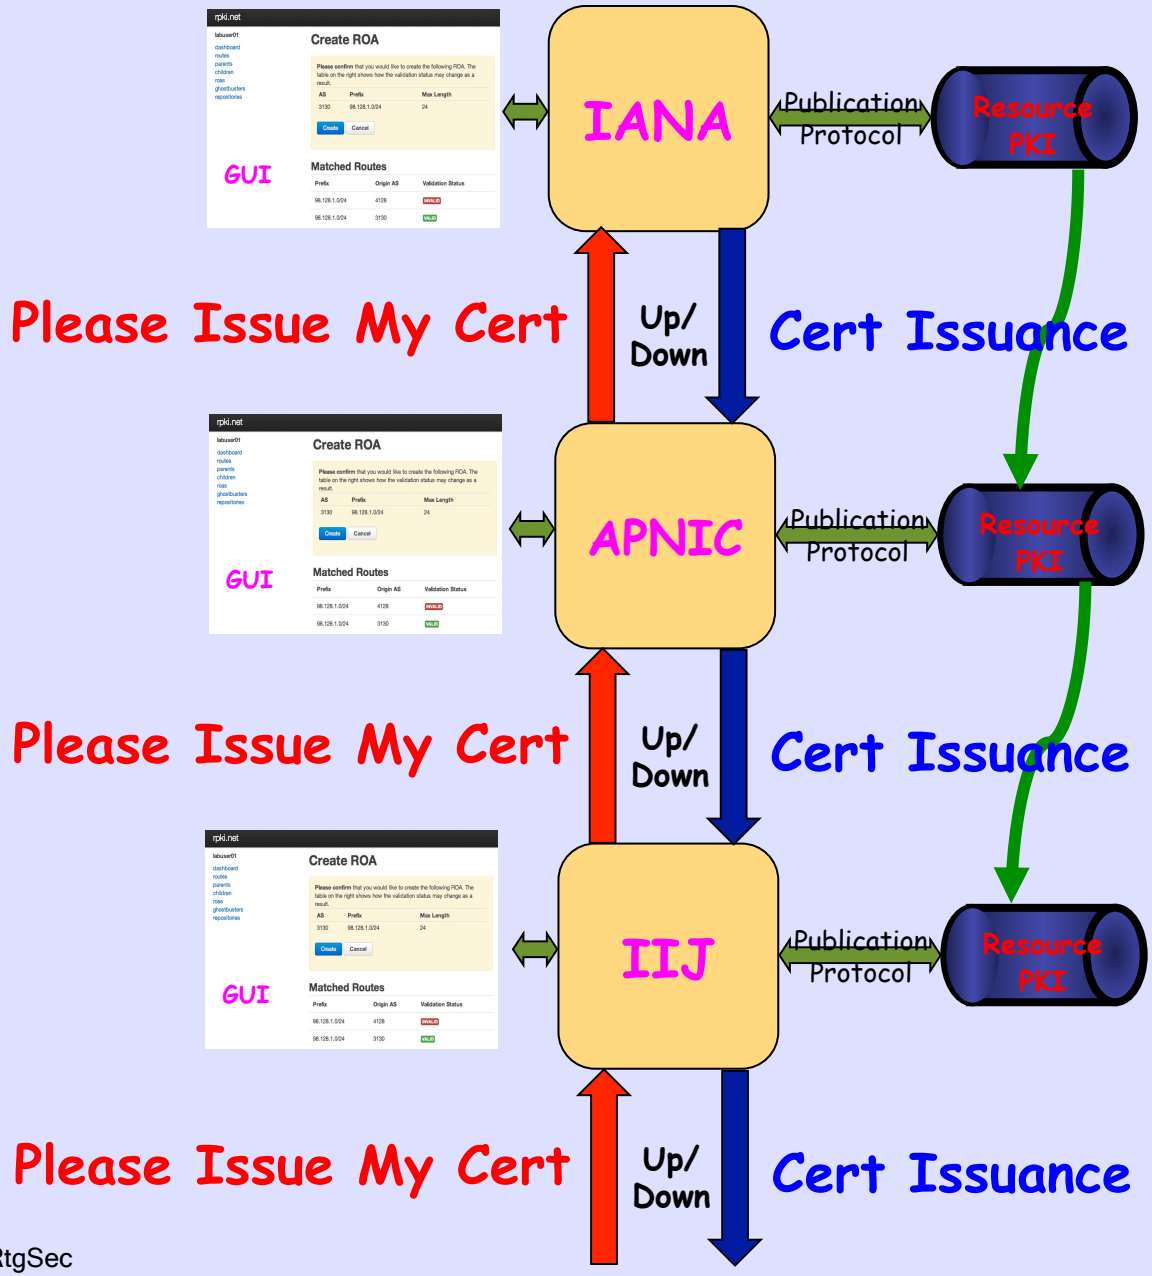

Issuing Parties

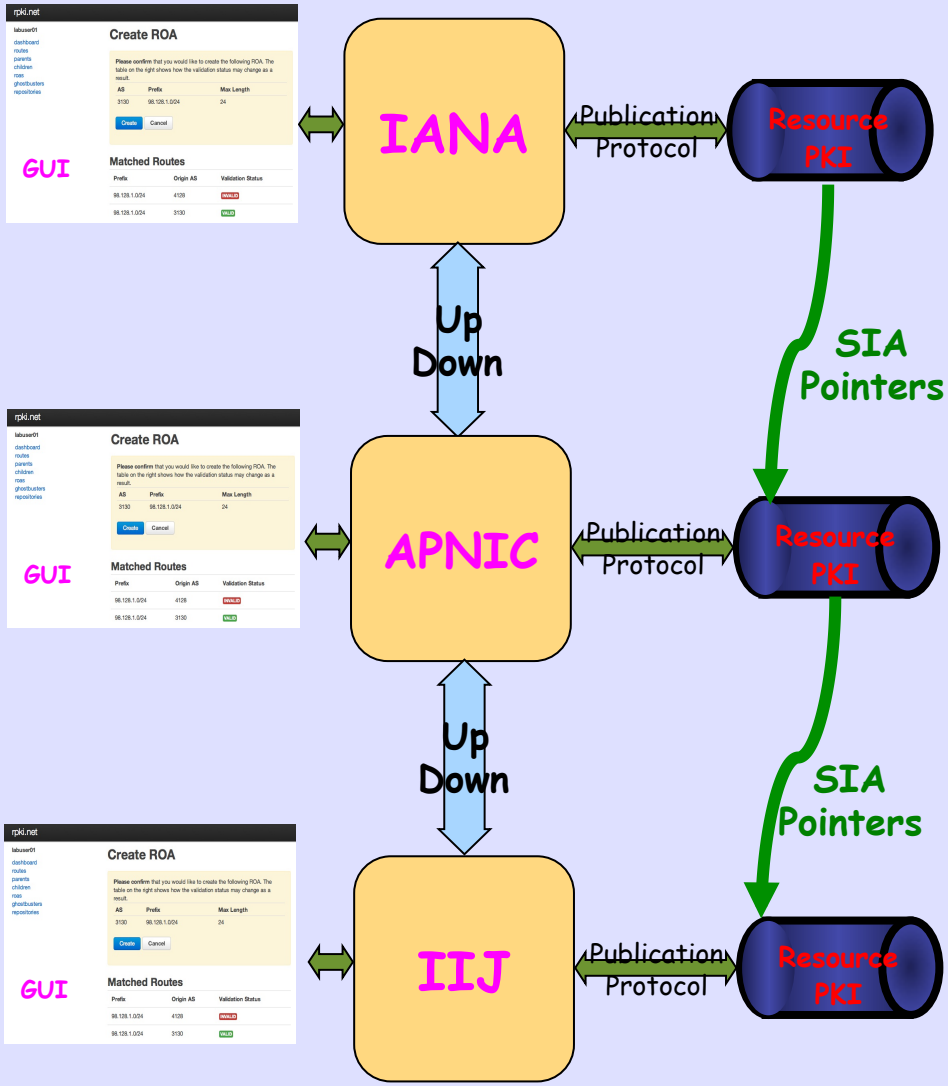

Issuing Parties Relying Parties

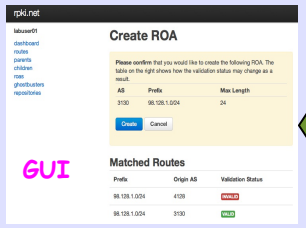

IANA

Publication Protocol

Trust Anchor

Resource PKI

Up Down

SIA Pointers

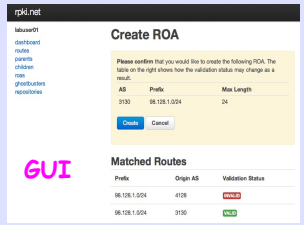

APNIC

Publication Protocol

Resource PKI

Up Down

SIA Pointers

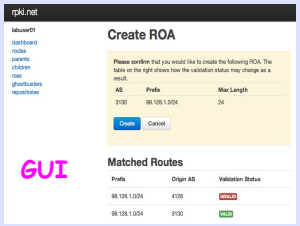

IIJ

Publication Protocol

Resource PKI

RCynic Gatherer

Validated Cache

Pseudo IRR

```
route:    147.28.0.0/16
descr:    147.28.0.0/16-16
origin:   AS3130
notify:   irr-hack@rpki.net
mnt-by:   MAINT-RPKI
changed:  irr-hack@rpki.net 20110606
source:   RPKI
```

NOC Tools

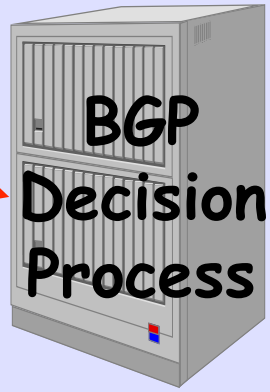

Extremely Large ISP Deployment

Caches Feed Caches

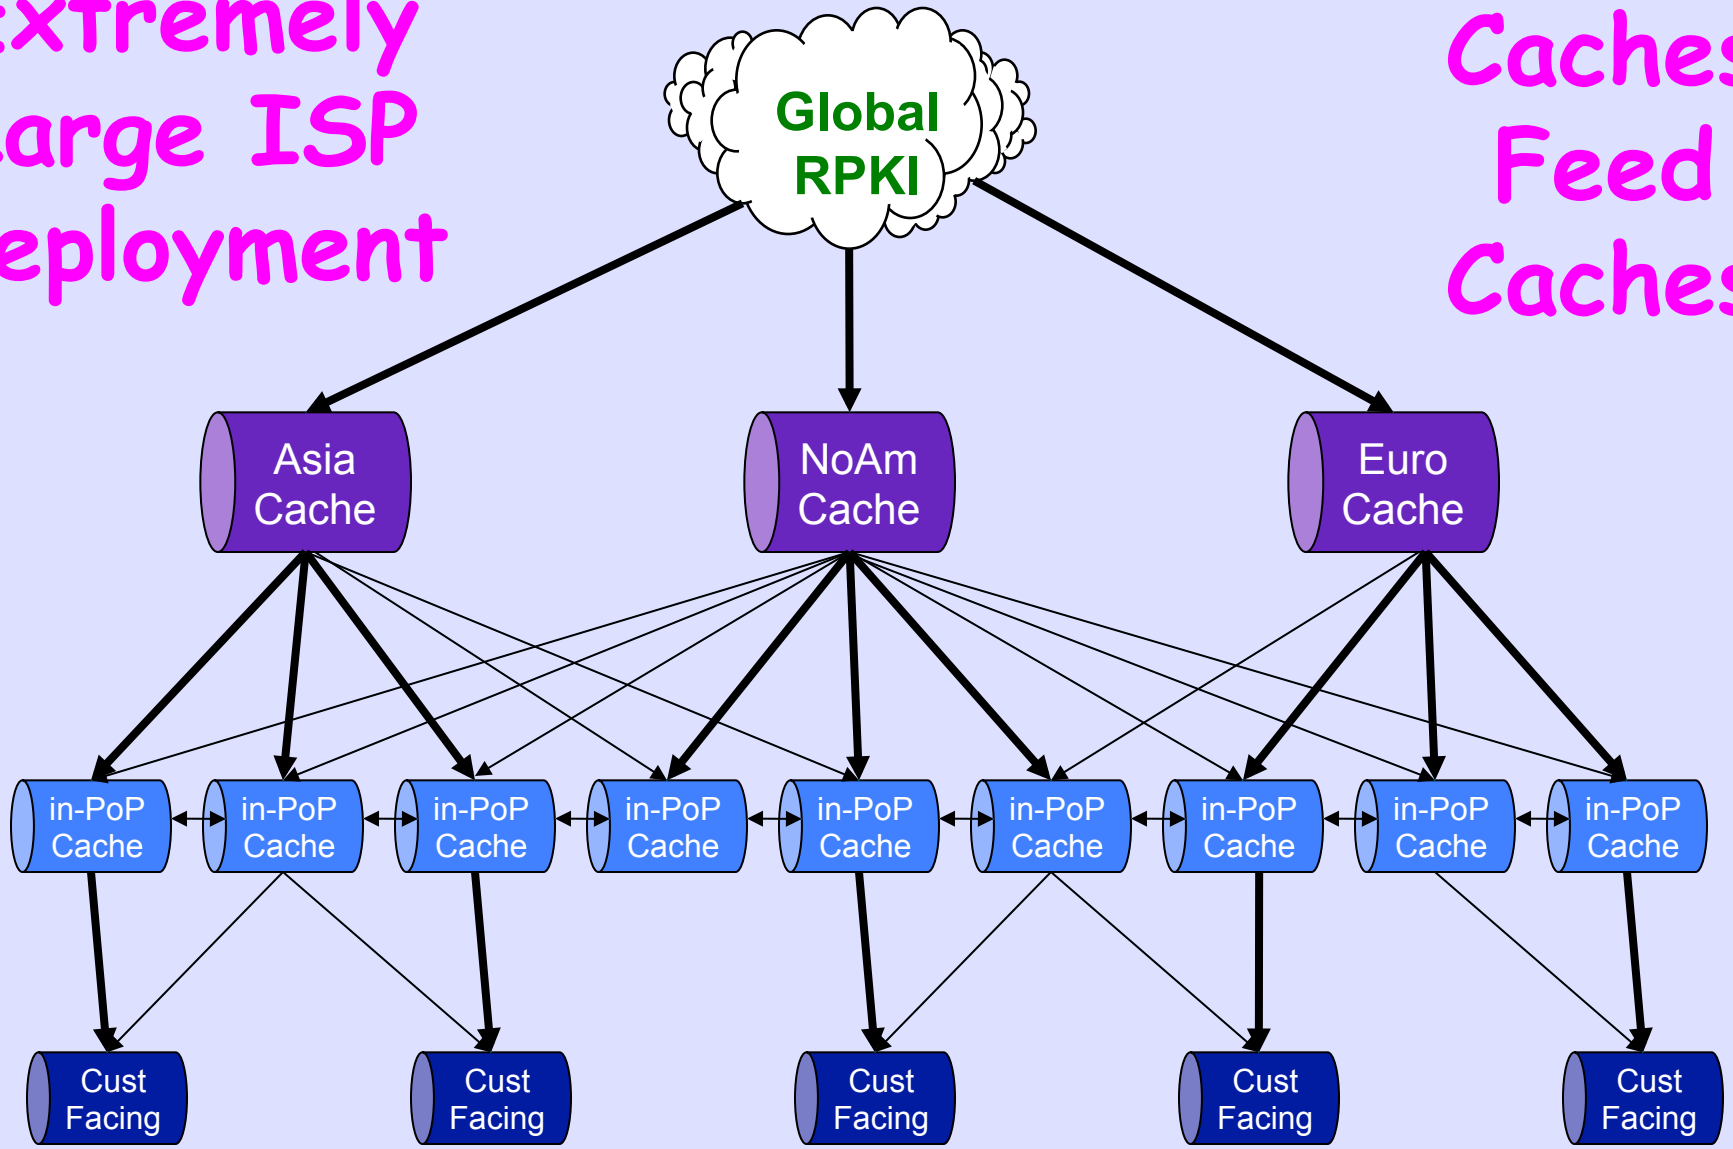

How Do ROAs Affect BGP Updates?

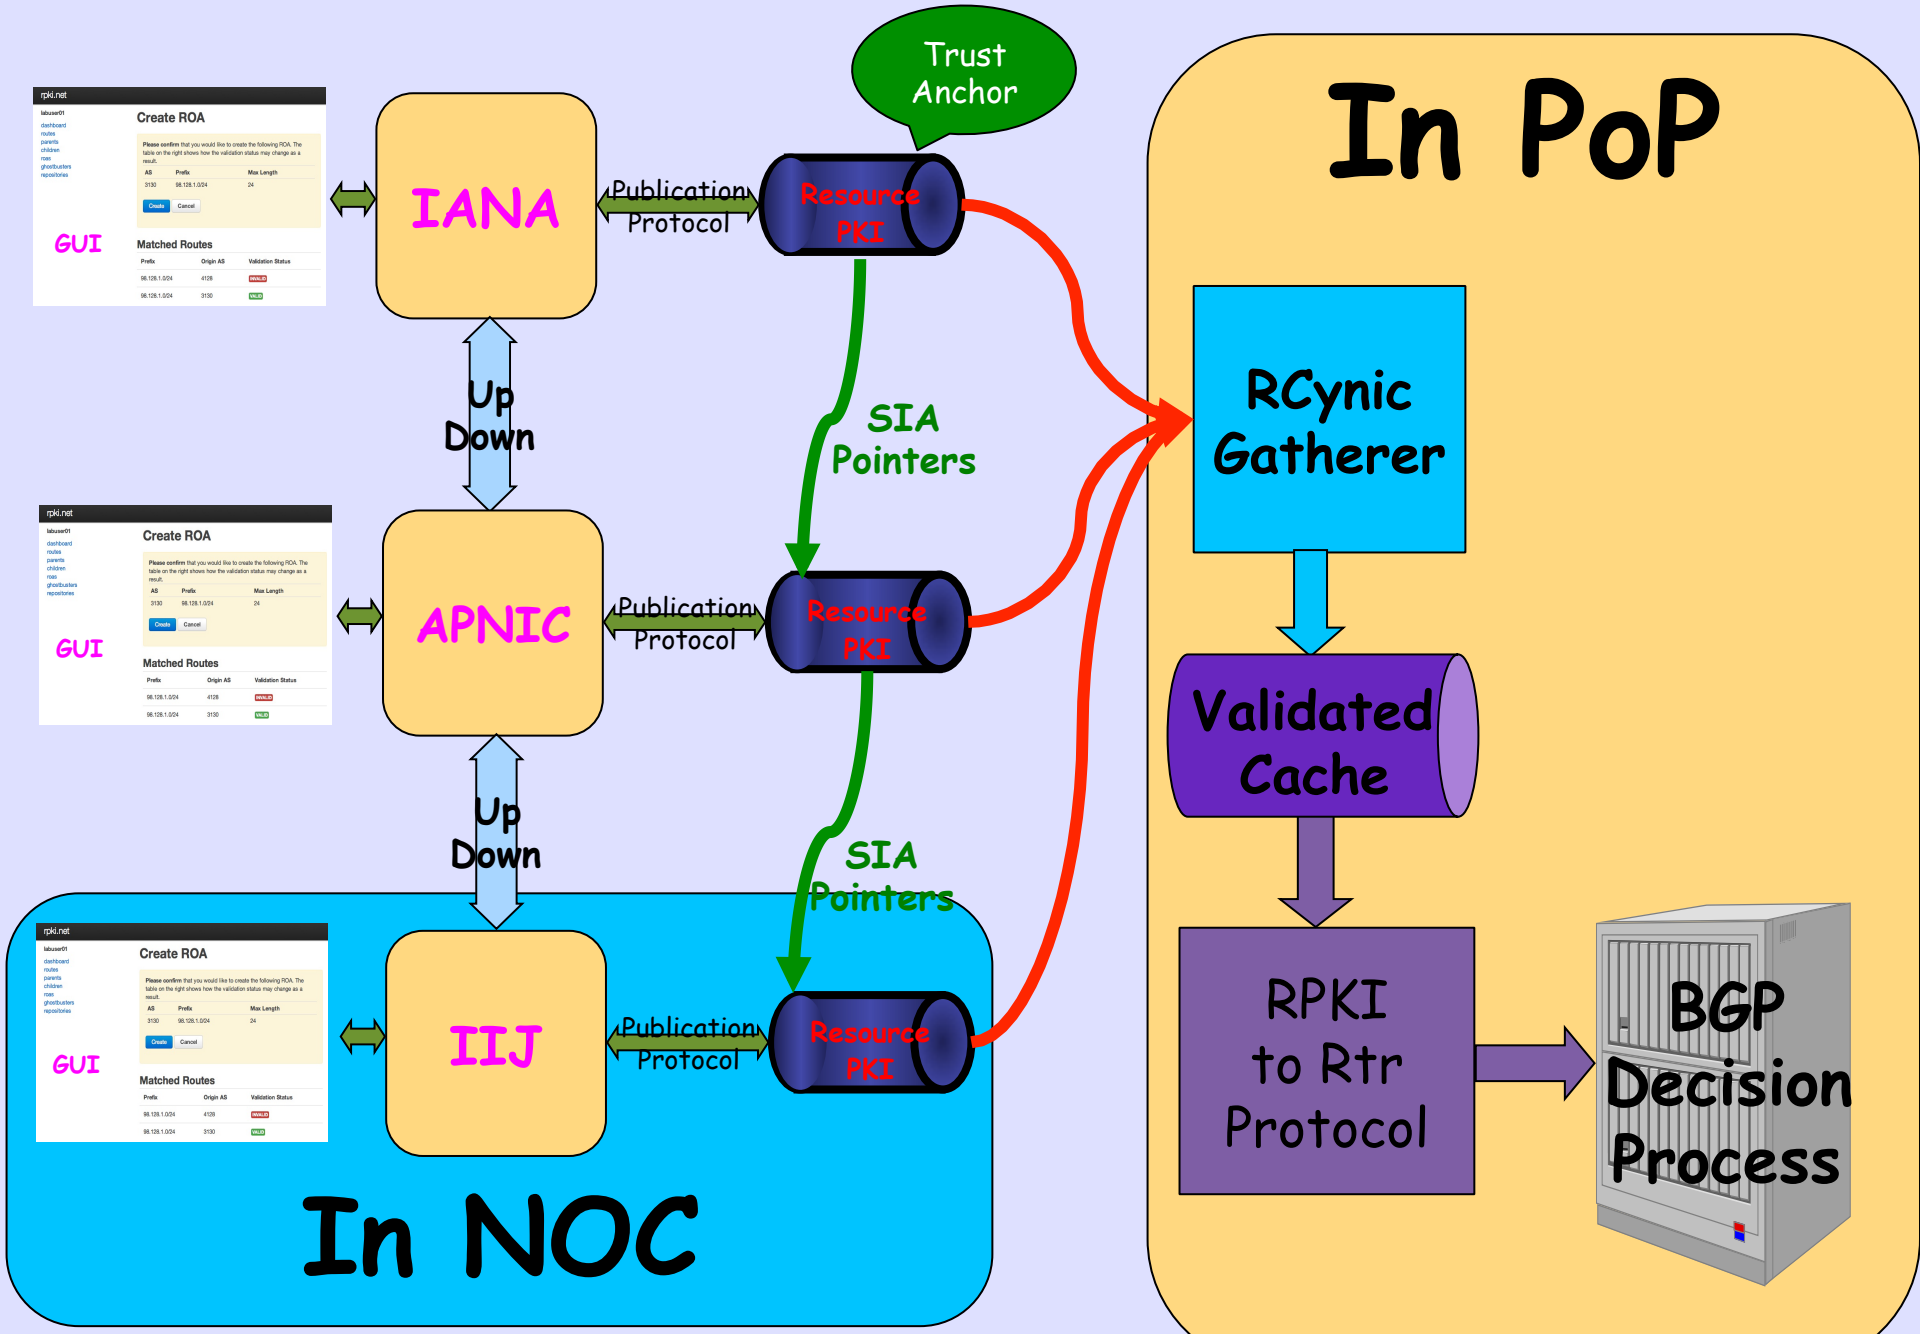
Trust Anchor
In PoP
IANA
APNIC
IIJ
GUI
Publication Protocol
Resource PKI
Up Down
SIA Pointers
RCynic Gatherer
Validated Cache
RPKI to Rtr Protocol
BGP Decision Process
In NOC

IPv4 Prefix

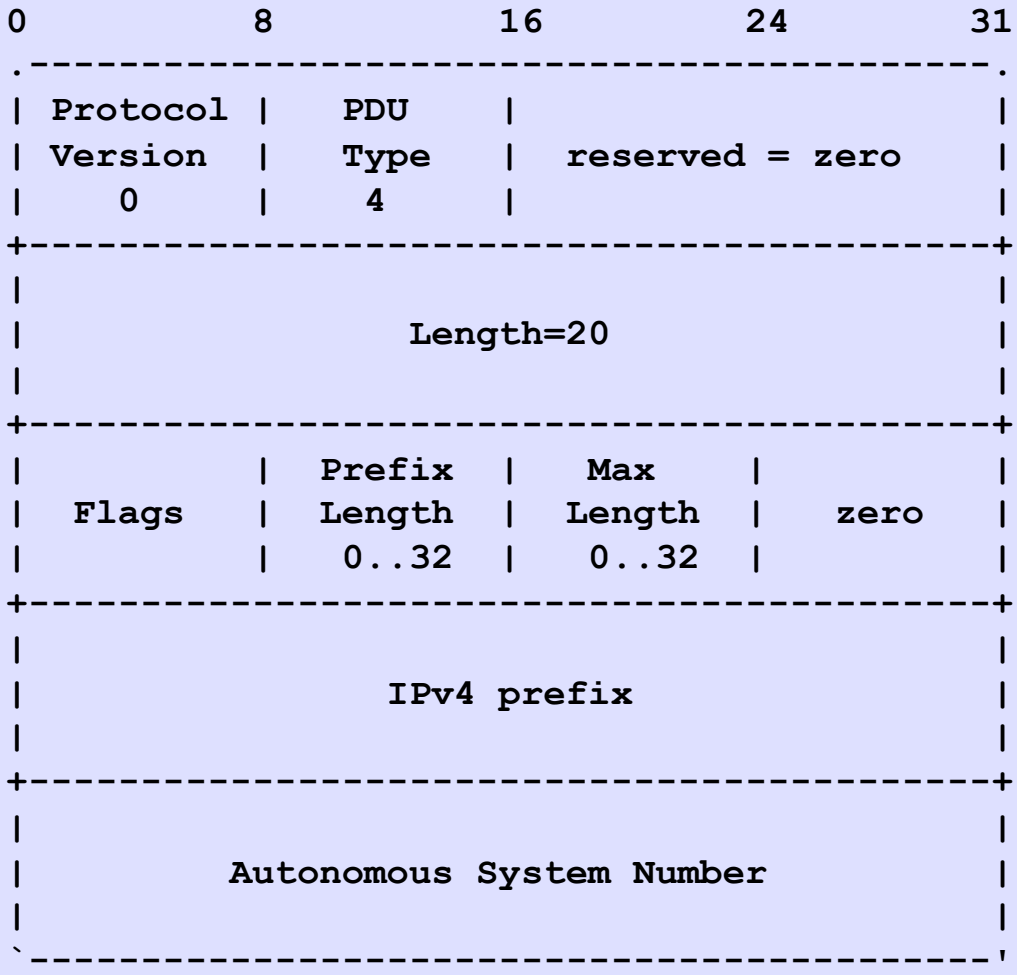
0
8
16
24
31
Protocol Version
0
PDU Type
4
reserved = zero
Length=20
Flags
Prefix Length
0..32
Max Length
0..32
zero
IPv4 prefix
Autonomous System Number

IPv6 Prefix

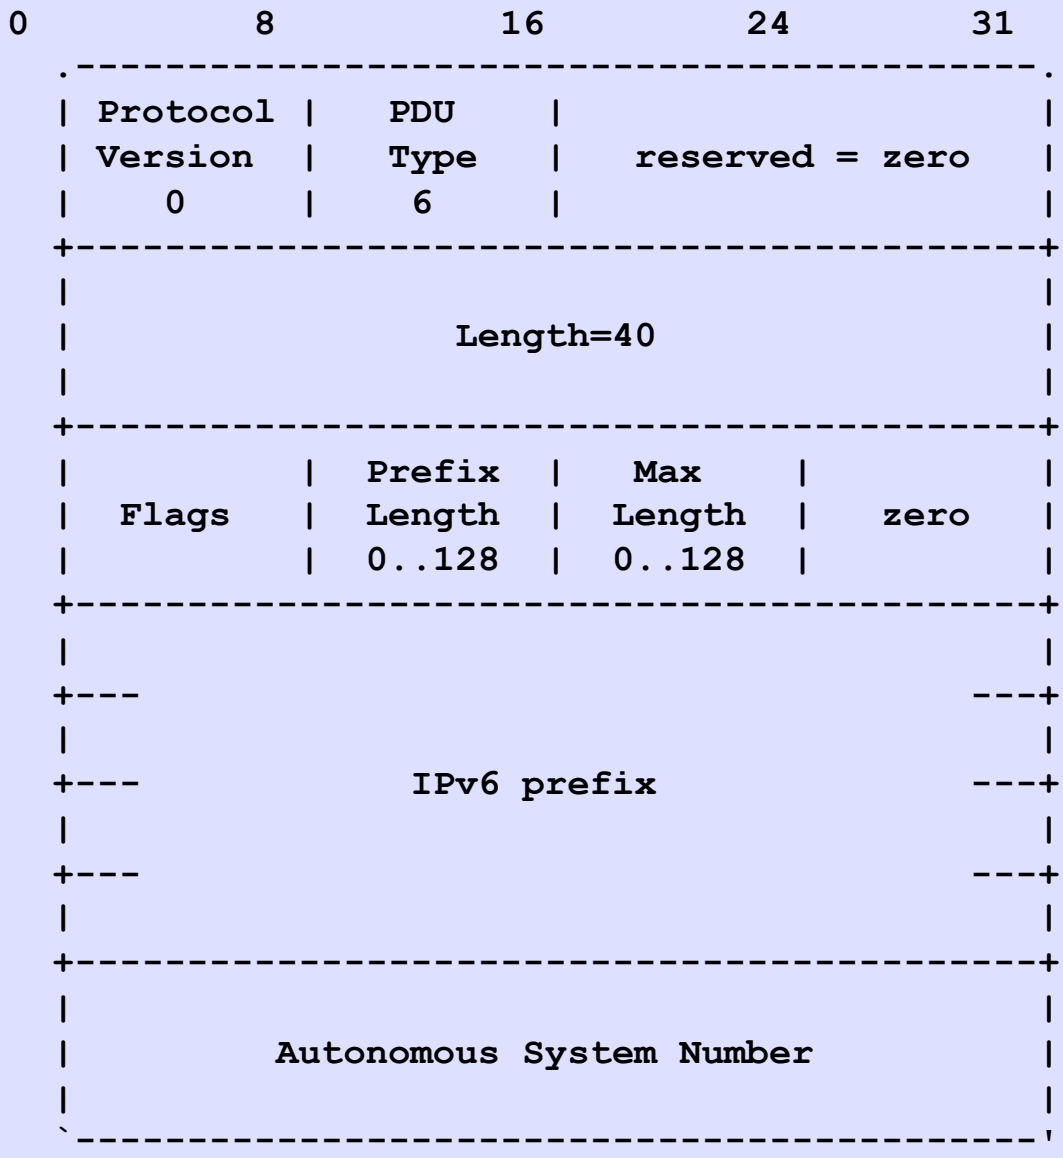
0
8
16
24
31
Protocol Version
0
PDU Type
6
reserved = zero
Length=40
Flags
Prefix Length
0..128
Max Length
0..128
zero
IPv6 prefix
Autonomous System Number

BGP Updates are compared with ROAs loaded from the RPKI

Marking BGP Updates

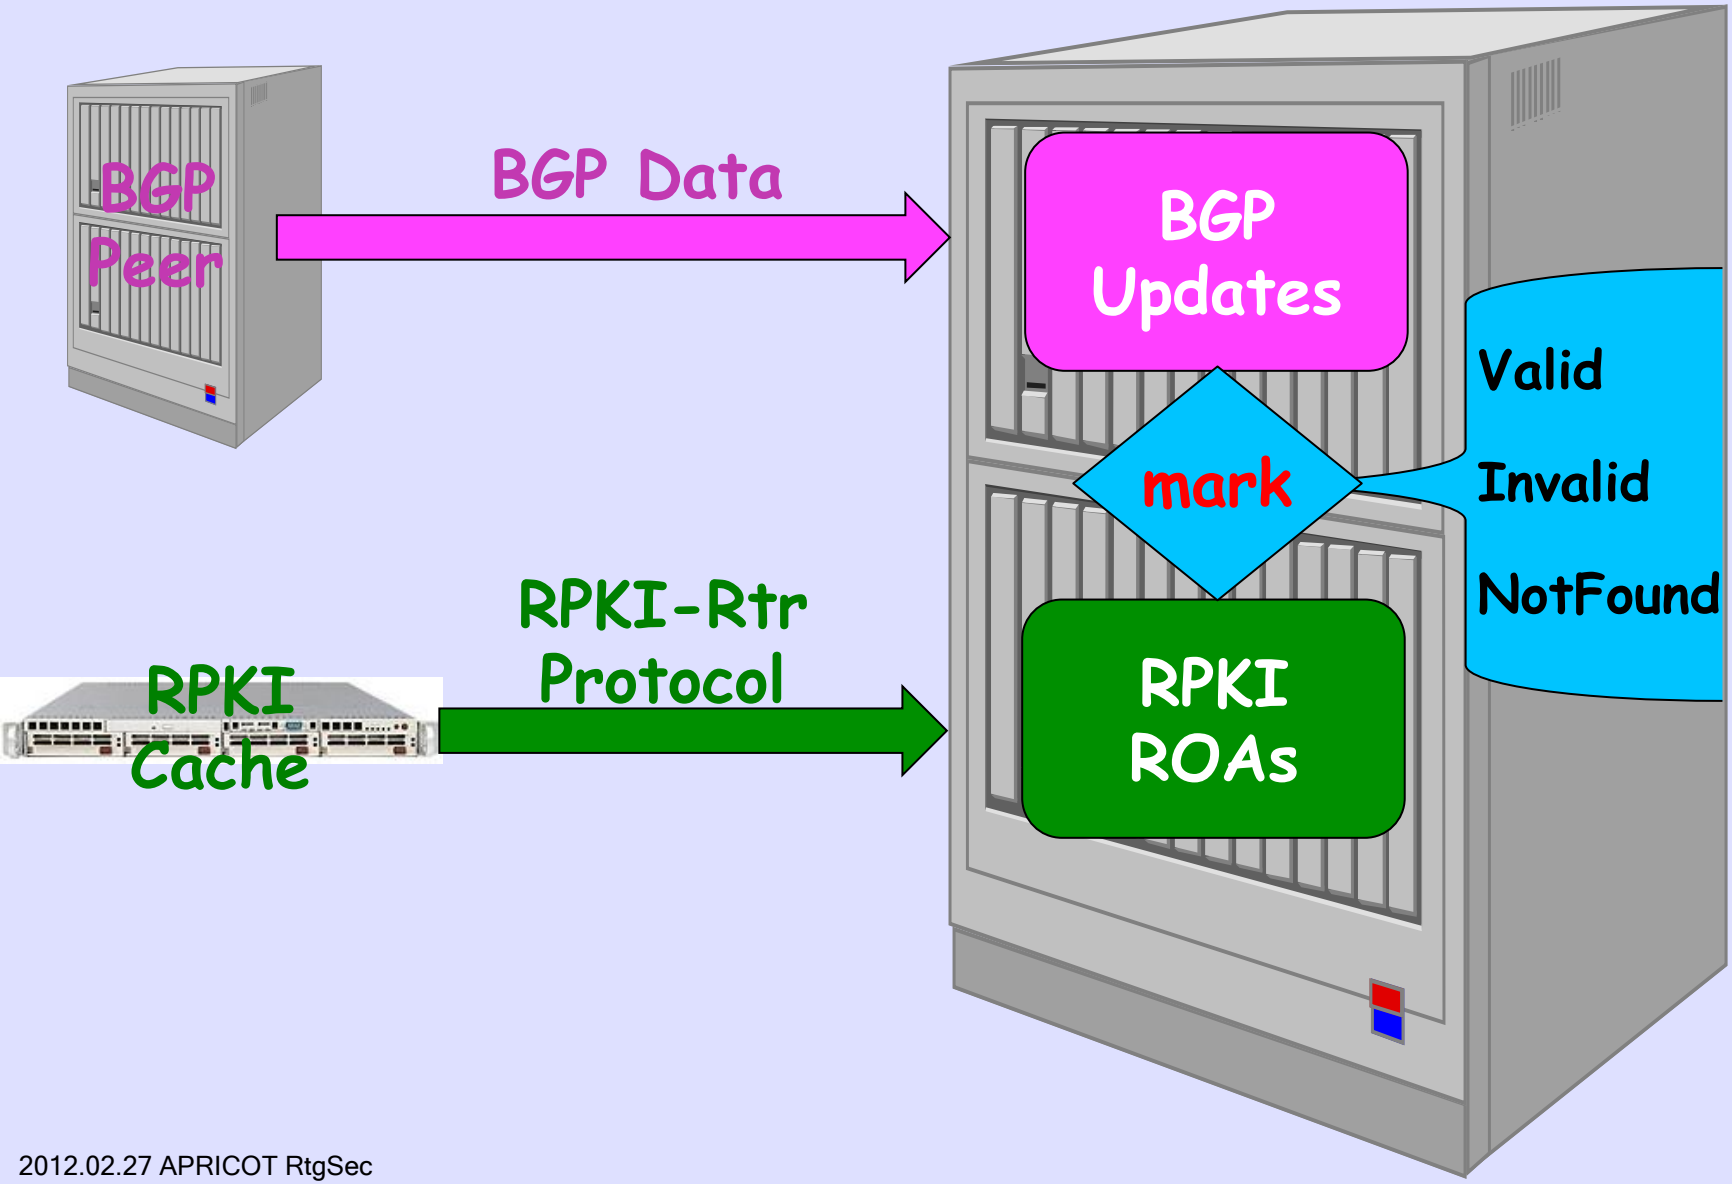

Result of Check

- **Valid** – A matching/covering ROA was found with a matching AS number
- **Invalid** – A matching or covering ROA was found, but AS number did not match, and there was no valid one
- **Not Found** – No matching or covering ROA was found, same as today

Configure Router to Get ROAs

```
router bgp 3130
 …
 bgp rpki server tcp 198.180.150.1 port 42420 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 / ROA Matching Rules?

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

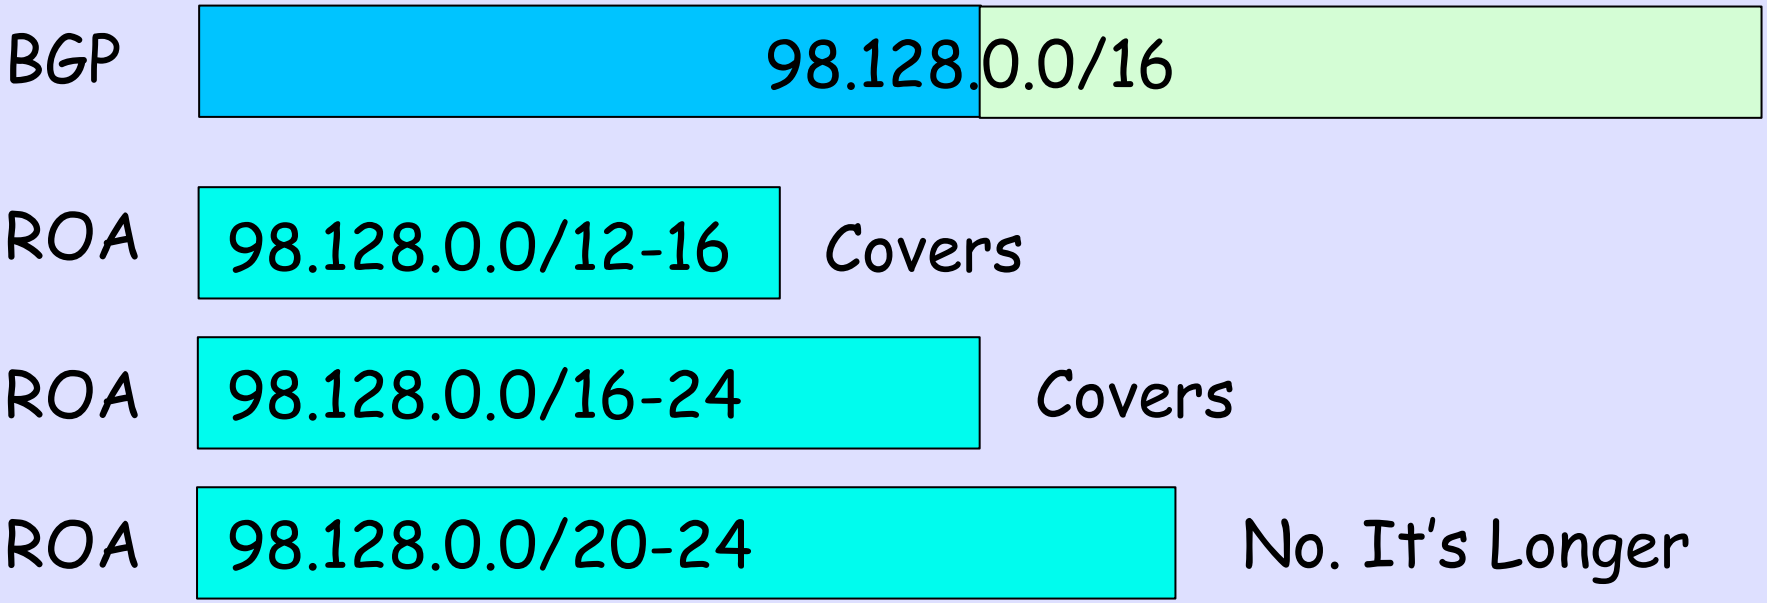

Prefix is Matched by a ROA when the Prefix is Covered by that ROA, prefix length is less than or equal to the ROA max-len, and the Route Origin AS is equal to the ROA's AS

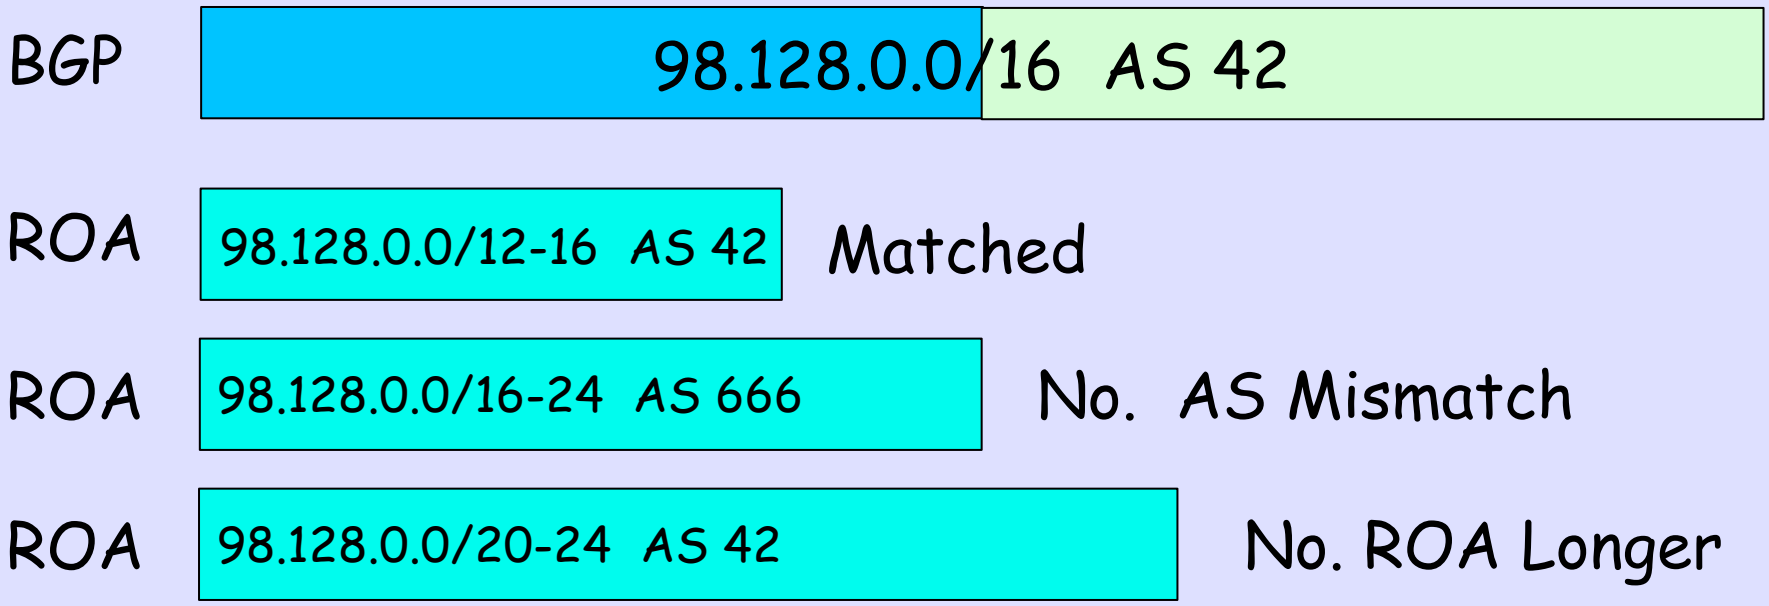

Matching and Validity

ROA_0	98.128.0.0/16-24 AS 6
ROA_1	98.128.0.0/16-20 AS 42

BGP	98.128.0.0/12 AS 42	NotFound, shorter than ROAs
BGP	98.128.0.0/16 AS 42	Valid, Matches ROA_1
BGP	98.128.0.0/20 AS 42	Valid, Matches ROA_1
BGP	98.128.0.0/24 AS 42	Invalid, longer than ROAs
BGP	98.128.0.0/24 AS 6	Valid, Matches ROA_0

The Operator Tests and then Sets 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
```

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
```

Allocation in Reality

/16 Assignment from RIR

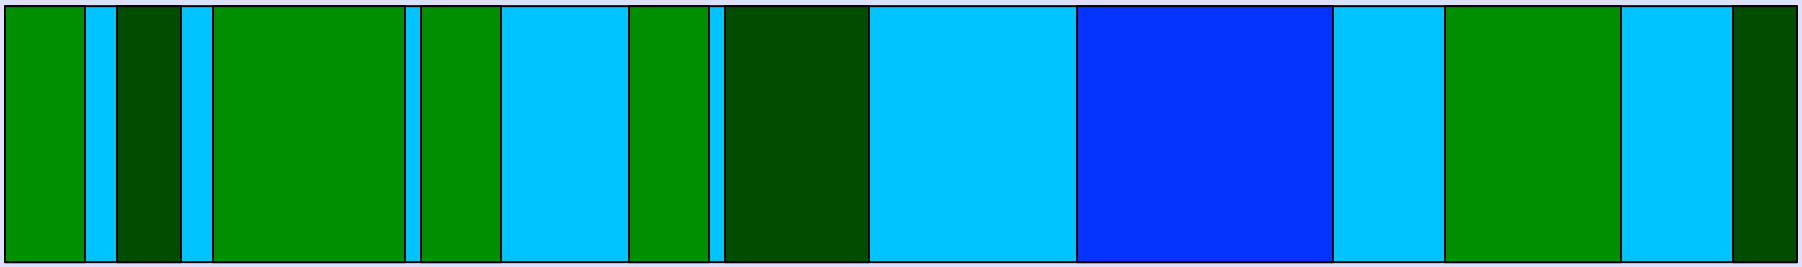

My Infrastructure

BGP Cust

Static (non BGP) Cust

Unused

ROA Use

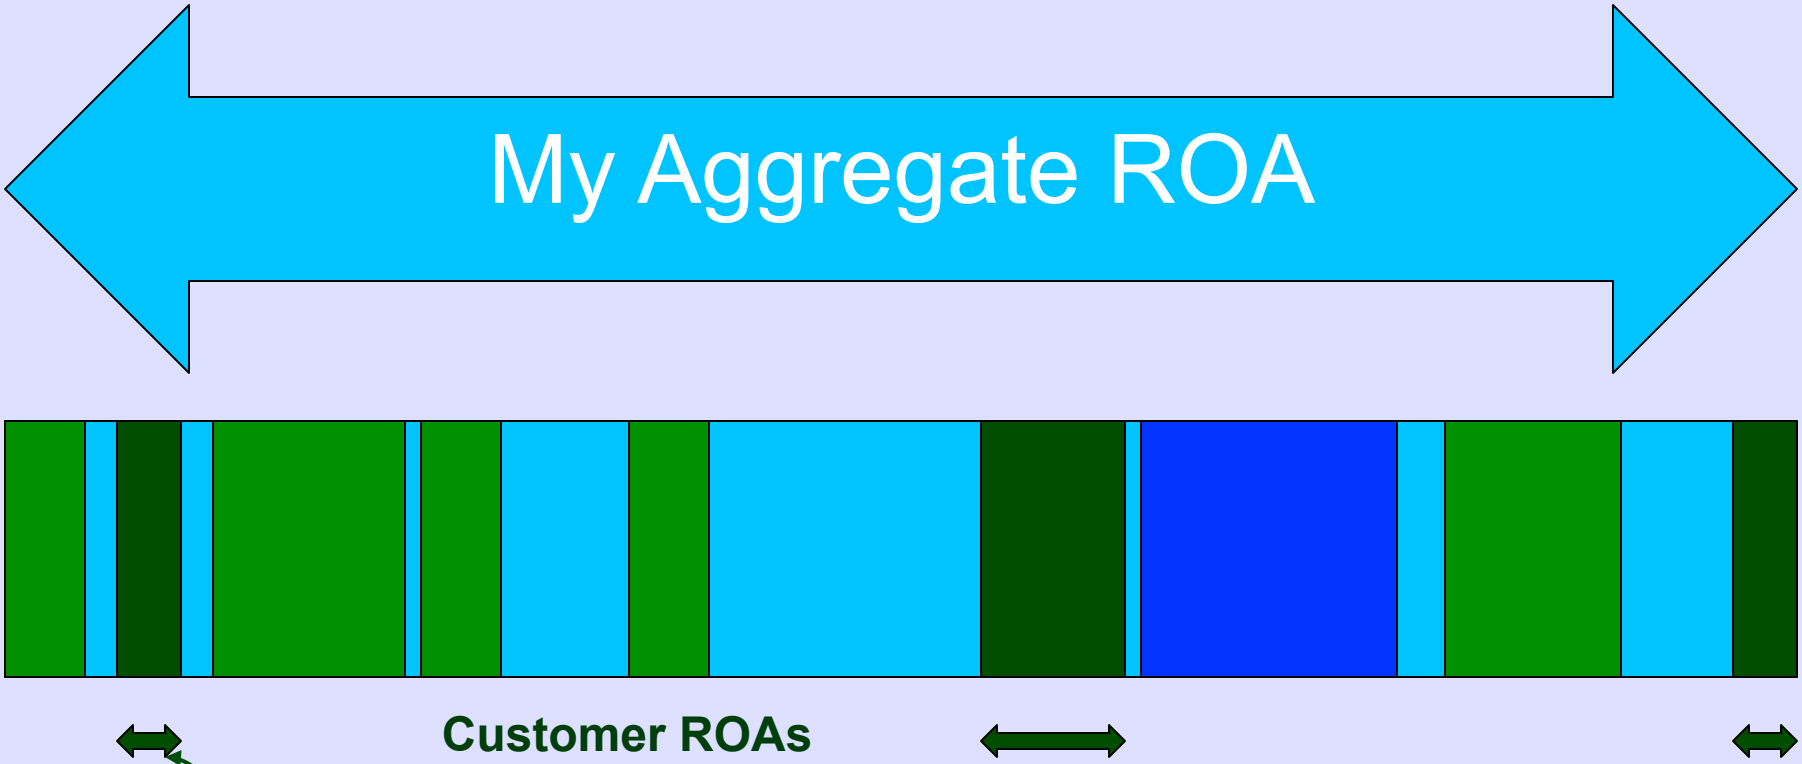
My Aggregate ROA
Customer ROAs
I Generate for 'Lazy' Customer
My Infrastructure
BGP Cust
Static (non BGP) Cust
Unused

Covering a Customer

I Issue a ROA for the Covering Prefix

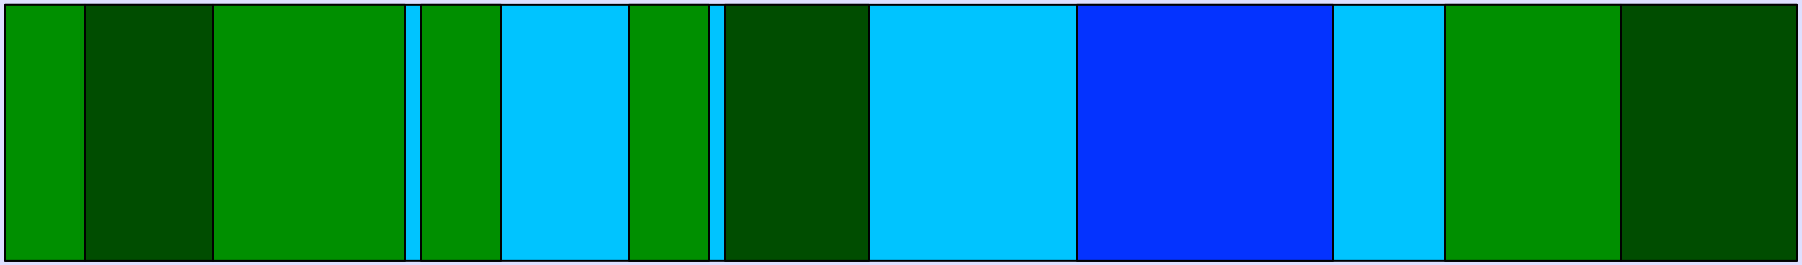

I need to do this to protect
Static Customers and my Infrastructure

My Infrastructure

BGP Cust

Static (non BGP) Cust

Unused

Covering a Customer

But if I Issue a ROA for the Covering Prefix

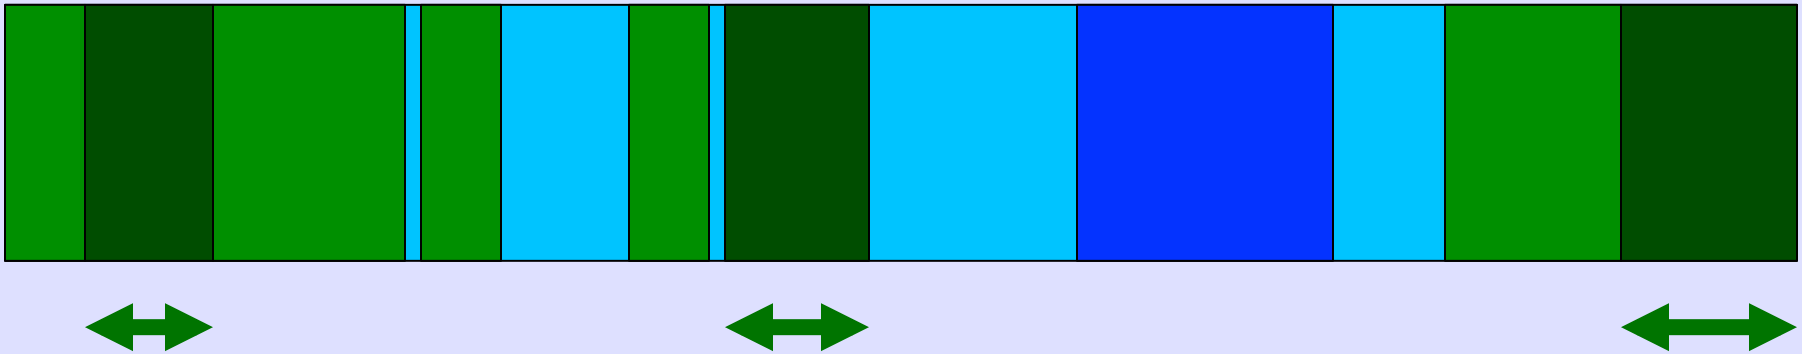

Before My Customers issue ROAs for These

My Infrastructure

BGP Cust

Static (non BGP) Cust

Unused

Covering a Customer

If I Issue a ROA for the Covering Prefix

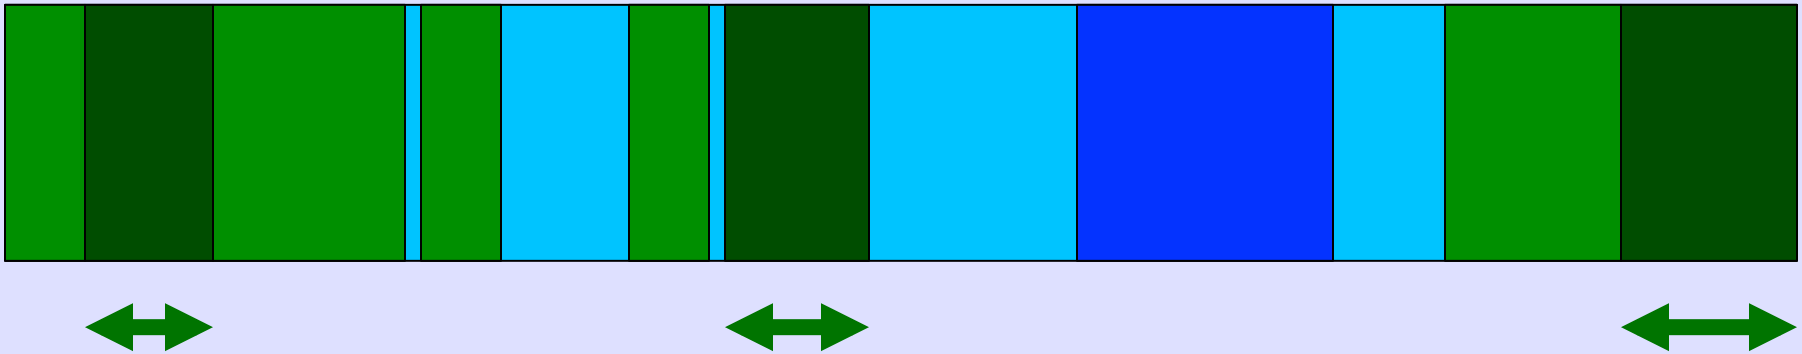

Before My Customers issue ROAs for These

Their Routing Becomes Invalid!

My Infrastructure

BGP Cust

Static (non BGP) Cust

Unused

Up-Chain Expiration

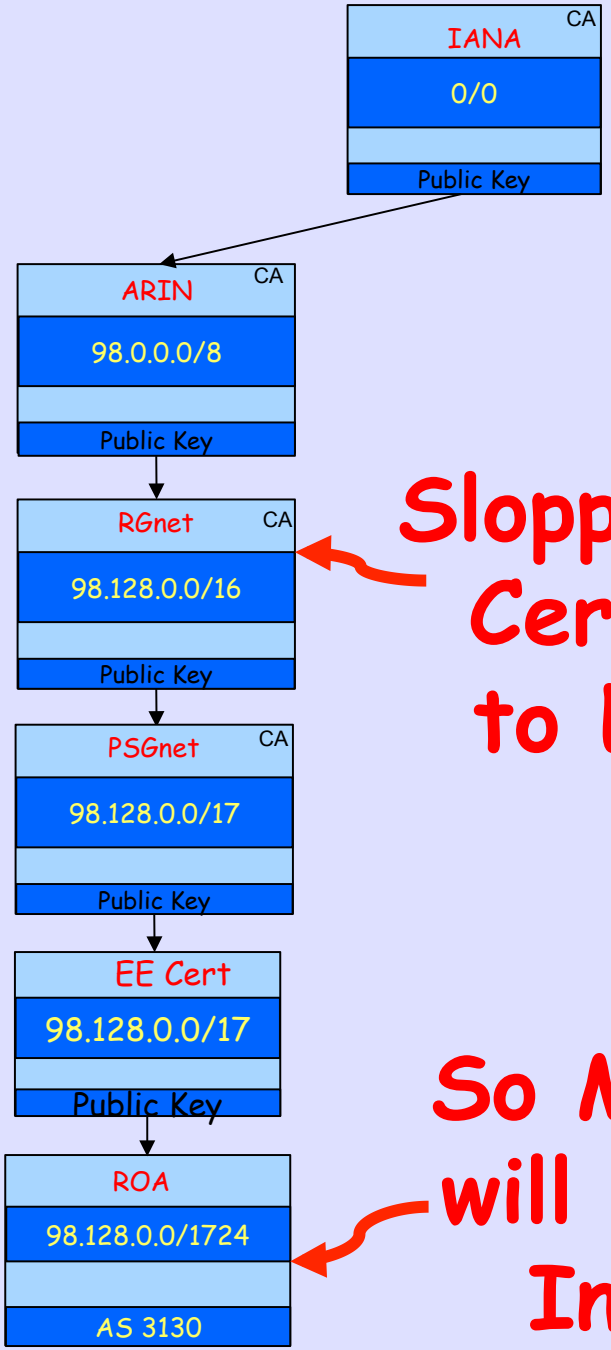

These are not Identity Certs

So Who You Gonna Call?

Sloppy Admin Cert Soon to Expire!

So My ROA will become Invalid!

ROA Invalid but I Can Route

- The ROA will become Invalid
- My announcement will just become NotFound, not Invalid
- Unless my upstream has a ROA for the covering prefix, which is likely

So Who You Gonna Call?

Ghostbusters!

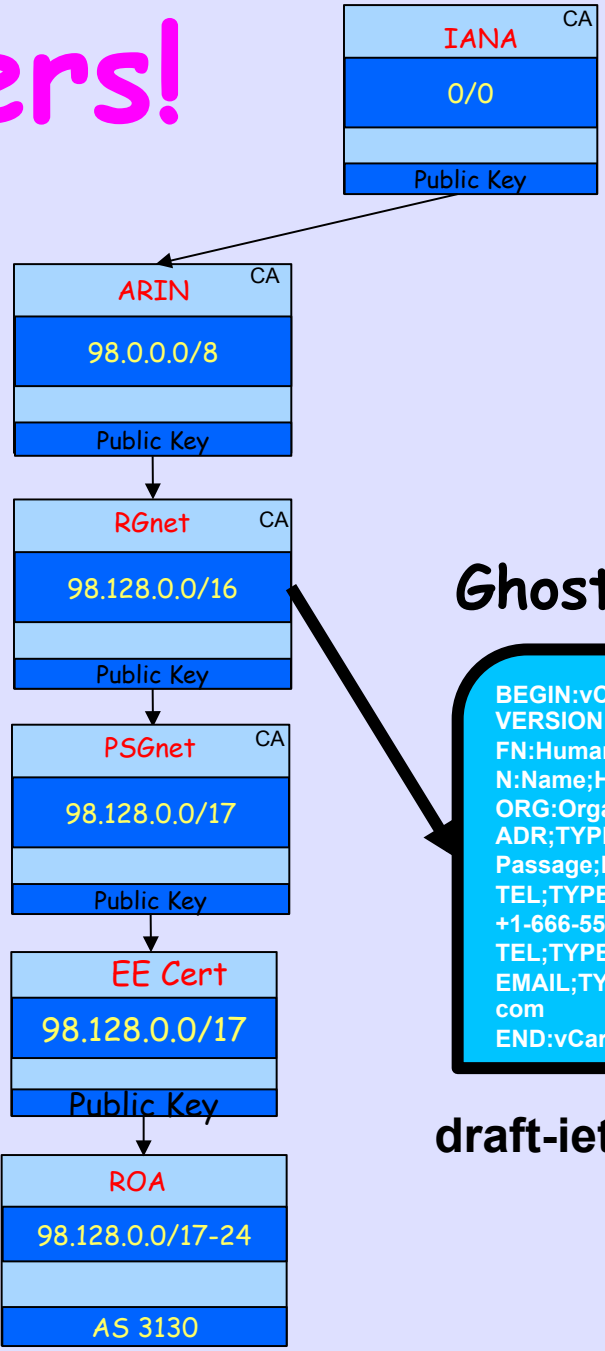

Ghostbusters Record

```
BEGIN:vCard
VERSION:3.0
FN:Human's Name
N:Name;Human's;Ms.;Dr.;OCD;ADD
ORG:Organizational Entity
ADR;TYPE=WORK:;;42 Twisty
Passage;Deep Cavern; WA; 98666;U.S.A.
TEL;TYPE=VOICE,MSG,WORK:
+1-666-555-1212
TEL;TYPE=FAX,WORK:+1-666-555-1213
EMAIL;TYPE=INTERNET:human@example.
com
END:vCard
```

draft-ietf-sidr-ghostbusters

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."

-- draft-ietf-sidr-origin-ops

But if I do not reject Invalid, what is all this for?

Open Source (BSD Lisc) Running Code

https://rpki.net/

Test Code in Routers

Talk to C & J

Vendor Code

- Cisco IOS and XR test code have Origin Validation now, shipping some code now
- Juniper has test code now, ship 2Q2012
- Work continues daily in test routers
- Compute load much less than ACLs from IRR data, **10µsec per update!**

BGPsec AS-Path Validation

Future Work

Origin Validation is Weak

- RPKI-Based Origin Validation only stops accidental misconfiguration, which is very useful. But ...
- A malicious router may announce as any AS, i.e. forge the ROAed origin AS.
- This would pass ROA Validation as in draft-ietf-sidr-pfx-validate.

Full Path Validation

- Rigorous per-prefix AS path validation is the goal
- Protect against origin forgery and AS-Path monkey in the middle attacks
- Not merely showing that a received AS path is not impossible

Forward Path Signing

AS hop N signing (among other things) that it is sending the announcement to AS hop N+1 by AS number, is believed to be fundamental to protecting against monkey in the middle attacks

Forward Path Signing