

RPKI -> RTR Protocol

IETF- sidr / Maastricht

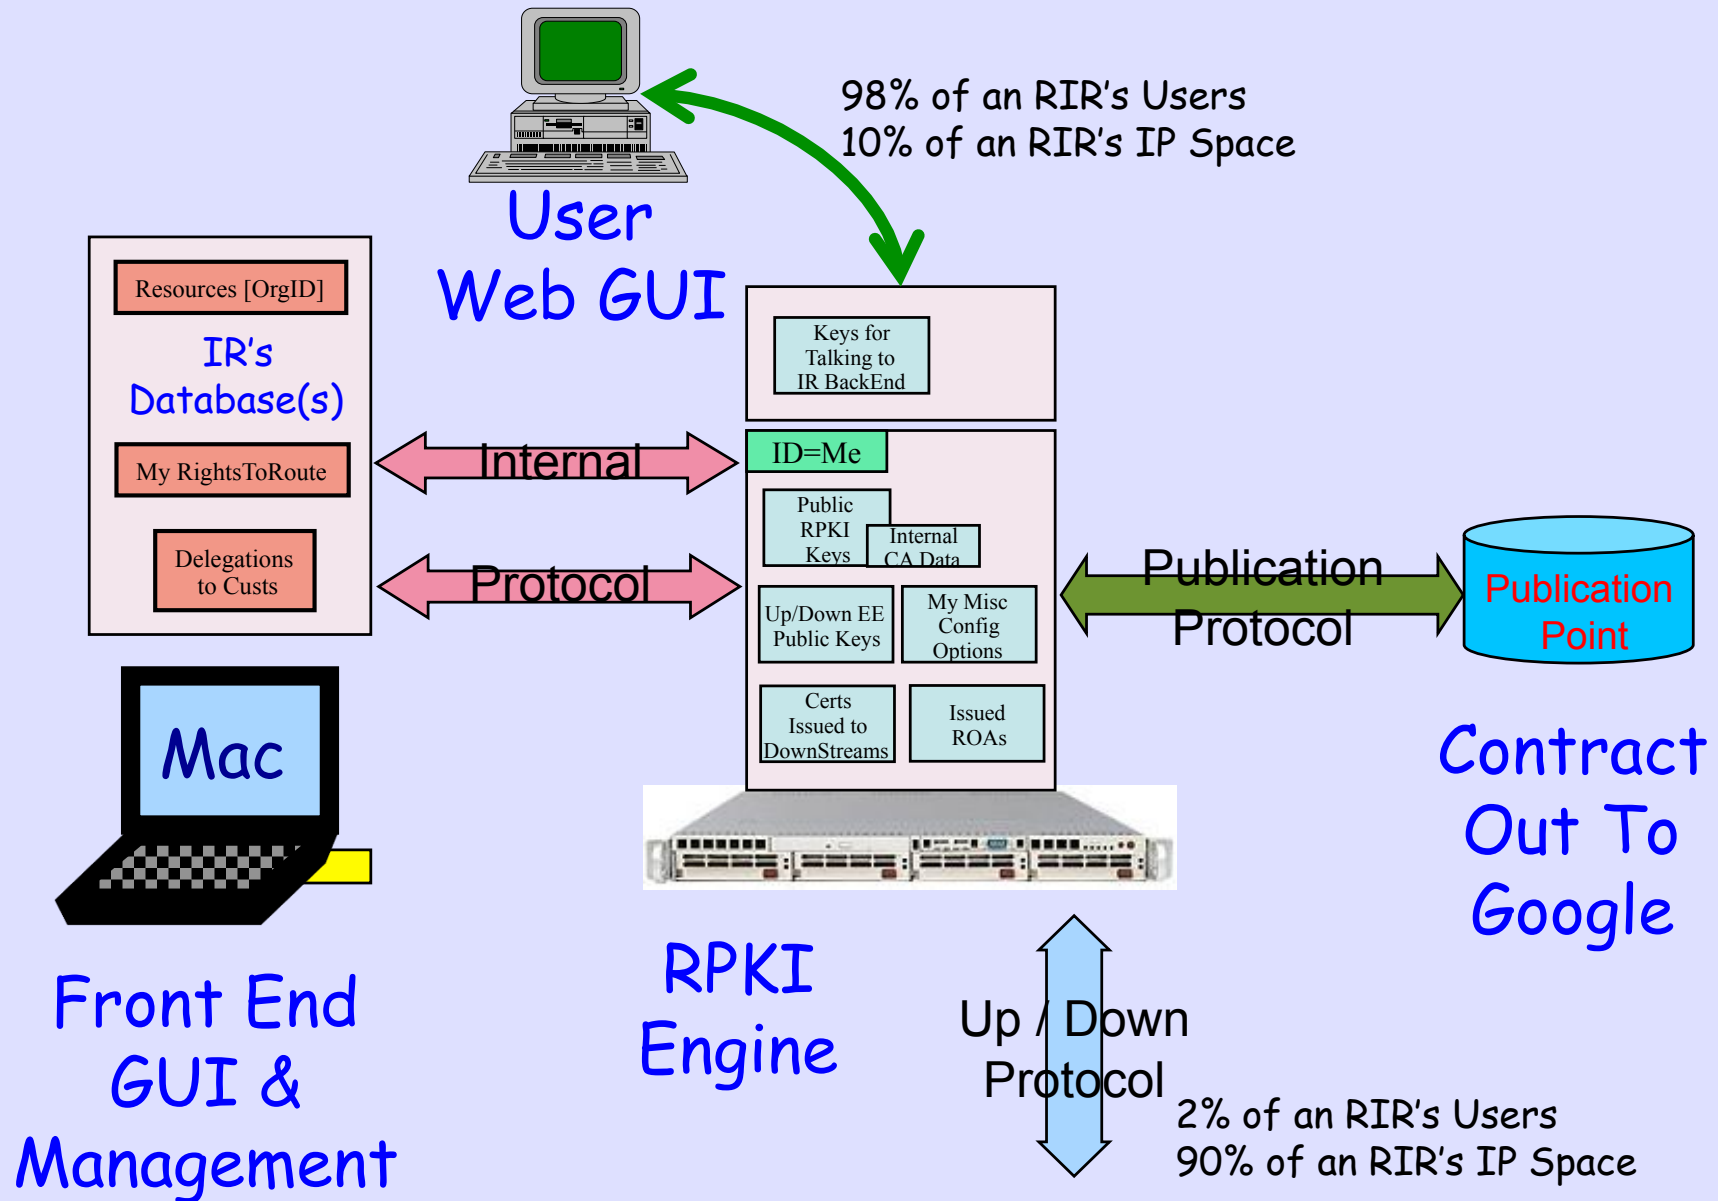
2010.07.28

Randy Bush <randy@psg.com>

Rob Austein <sra@isc.org>

draft-ymbk-rpki-rtr-protocol-06.txt

Using the RPKI



Received from One Parent

- [arin](#)

Accepted Resource	Not Before	Not After
ASN 3130	2010-05-11 03:43:36	2011-03-22 18:18:54
ASN 3927	2010-05-11 03:43:36	2011-03-22 18:18:54
ASN 3970	2010-05-11 03:43:36	2011-03-22 18:18:54
ASN 4128	2010-05-11 03:43:36	2011-03-22 18:18:54
67.21.36.0/24	2010-05-11 03:43:36	2011-03-22 18:18:54
69.166.11.0/24	2010-05-11 03:43:36	2011-03-22 18:18:54
98.128.0.0/16	2010-05-11 03:43:36	2011-03-22 18:18:54
147.28.0.0/16	2010-05-11 03:43:36	2011-03-22 18:18:54
192.83.230.0/24	2010-05-11 03:43:36	2011-03-22 18:18:54
192.169.0.0/23	2010-05-11 03:43:36	2011-03-22 18:18:54
198.133.206.0/24	2010-05-11 03:43:36	2011-03-22 18:18:54
198.180.150.0-198.180.153.255	2010-05-11 03:43:36	2011-03-22 18:18:54
207.34.0.0/24	2010-05-11 03:43:36	2011-03-22 18:18:54
216.21.0.0/24	2010-05-11 03:43:36	2011-03-22 18:18:54

What I Did With It

Children

- [Me](#)
- [loughborough](#)

Delegated resources:

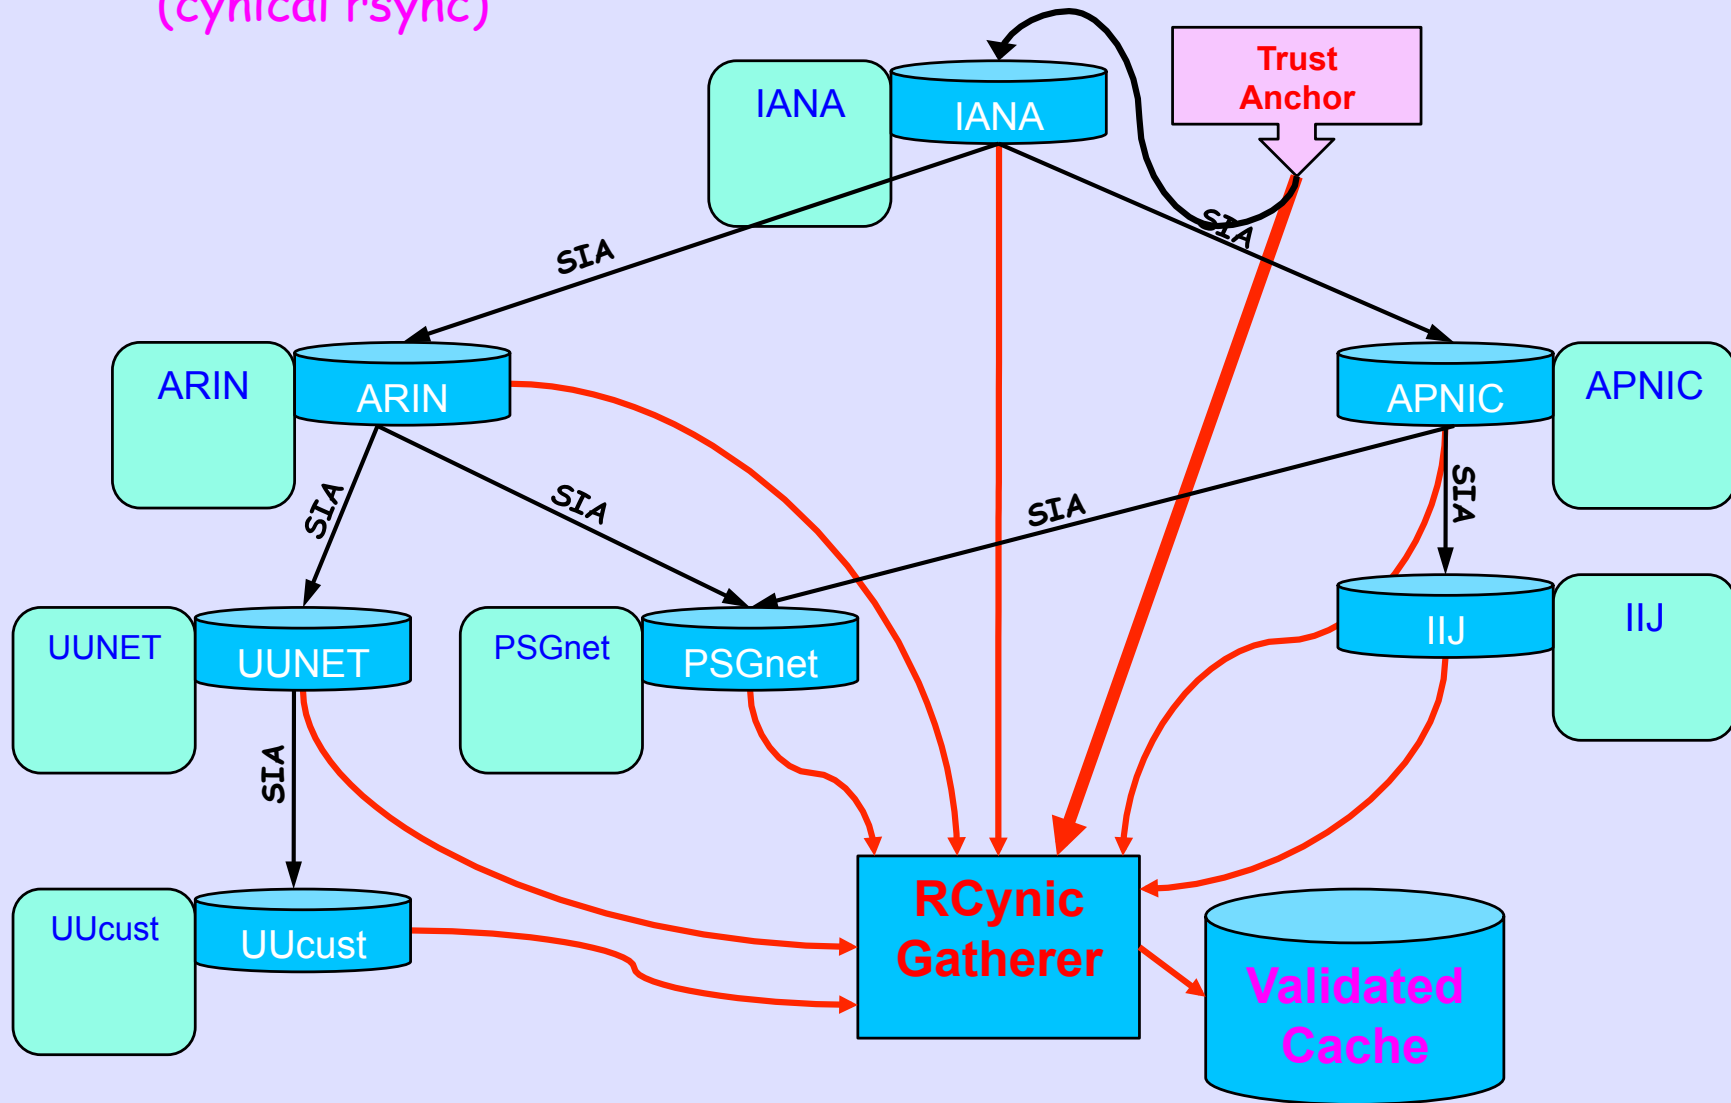
- [98.128.0.0/16](#)

My ROA [request]s

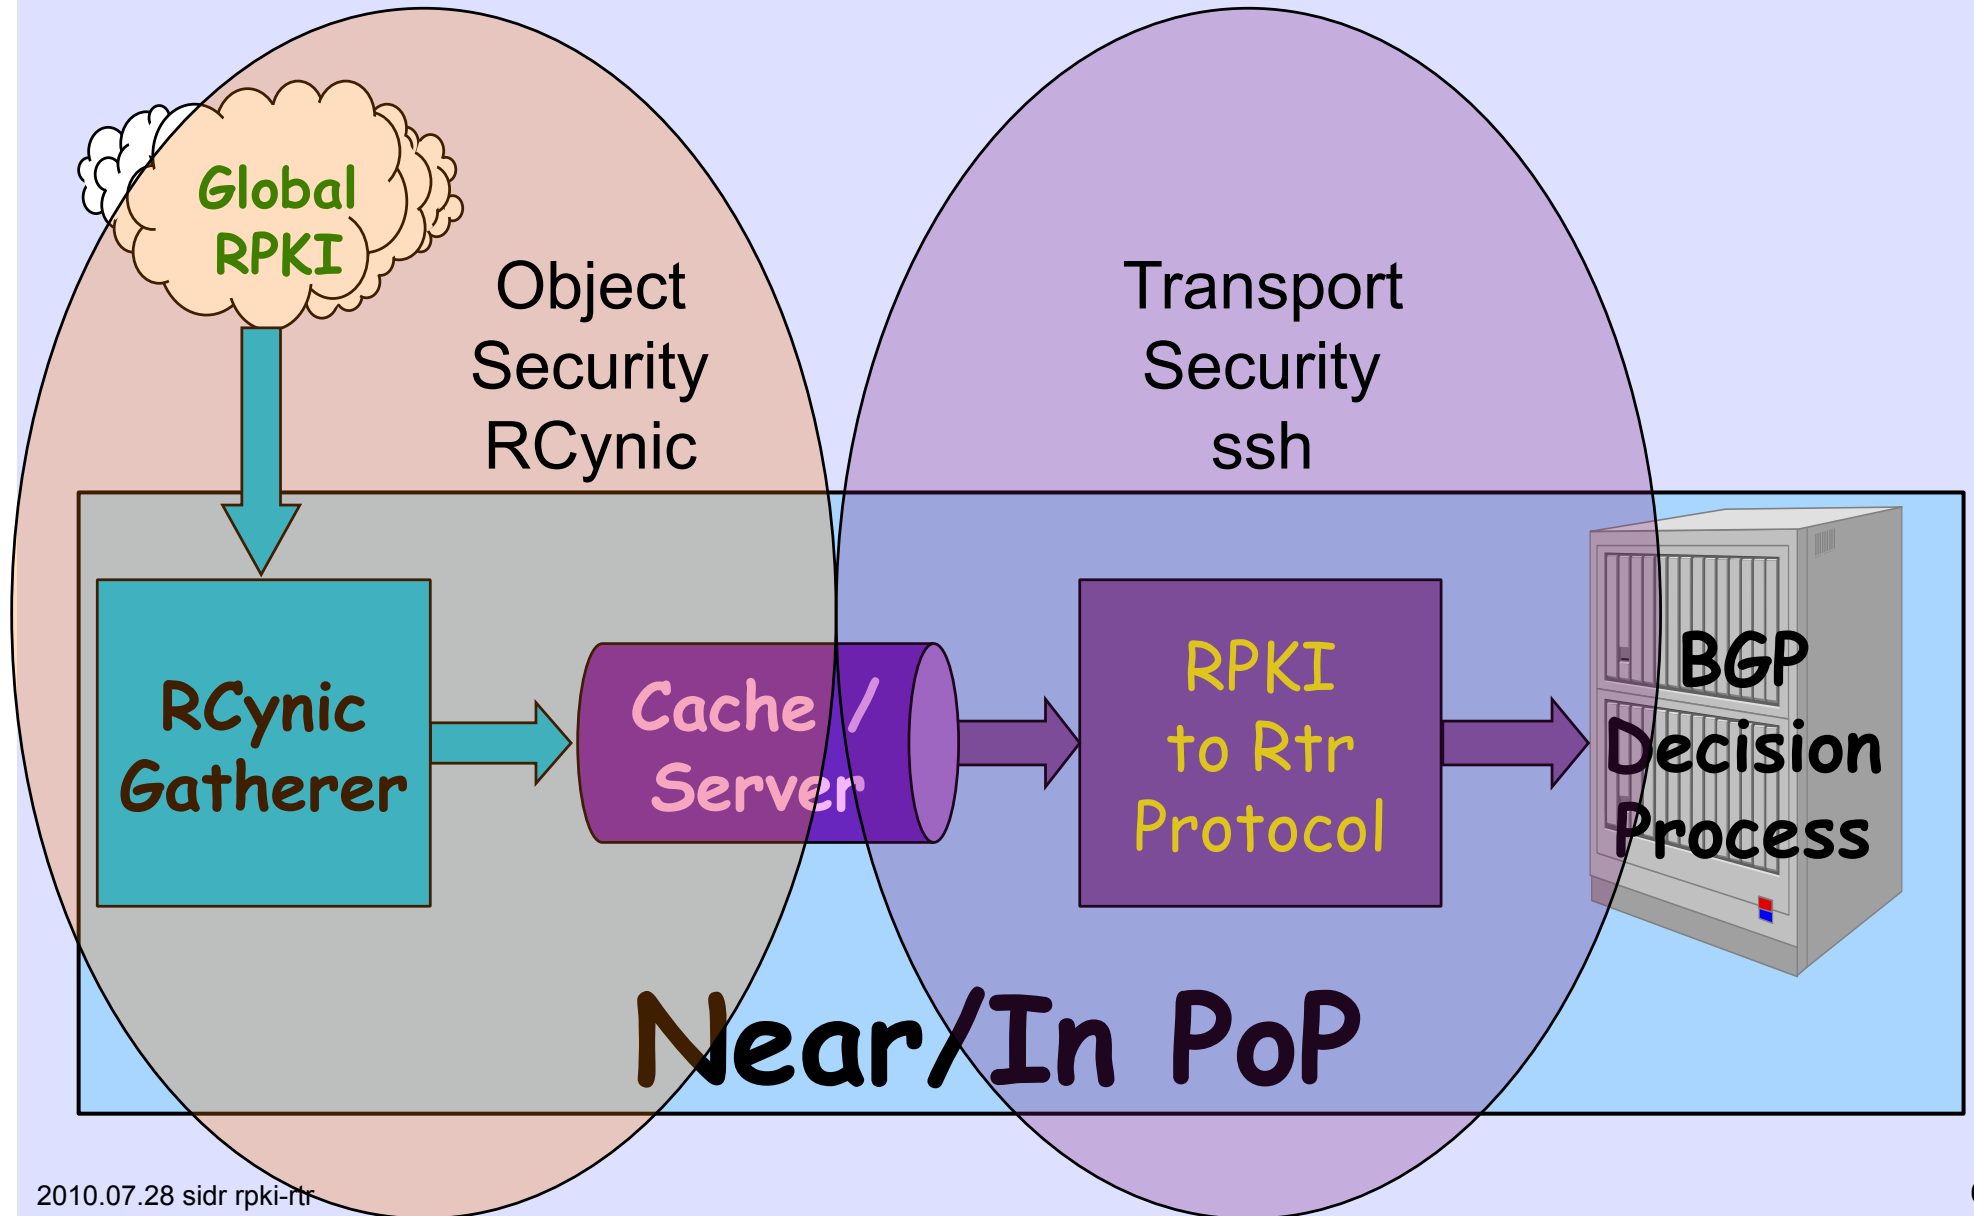
Prefix	ASN
<ul style="list-style-type: none">• 67.21.36.0/24• 69.166.11.0/24• 192.169.0.0/23• 207.34.0.0/24• 216.21.0.0/24• 216.21.14.0/24• 216.21.16.0/24	3970

RCynic Cache Gatherer

(cynical rsync)




RPKI -> Router




Typical Exchange

```
Cache                                     Router
| <----- Reset Query -----> | R requests data
|                                     |
| ----- Cache Response -----> | C confirms request
| ----- IPvX Prefix -----> | C sends zero or more
| ----- IPvX Prefix -----> | IPv4 and IPv6 Prefix
| ----- IPvX Prefix -----> | Payload PDUs
| ----- End of Data -----> | C sends End of Data
|                                     | and sends new serial
~                                     ~
| ----- Notify -----> | (optional)
|                                     |
| <----- Serial Query -----> | R requests data
|                                     |
| ----- Cache Response -----> | C confirms request
| ----- IPvX Prefix -----> | C sends zero or more
| ----- IPvX Prefix -----> | IPv4 and IPv6 Prefix
| ----- IPvX Prefix -----> | Payload PDUs
| ----- End of Data -----> | C sends End of Data
|                                     | and sends new serial
~                                     ~
```


Reset Query



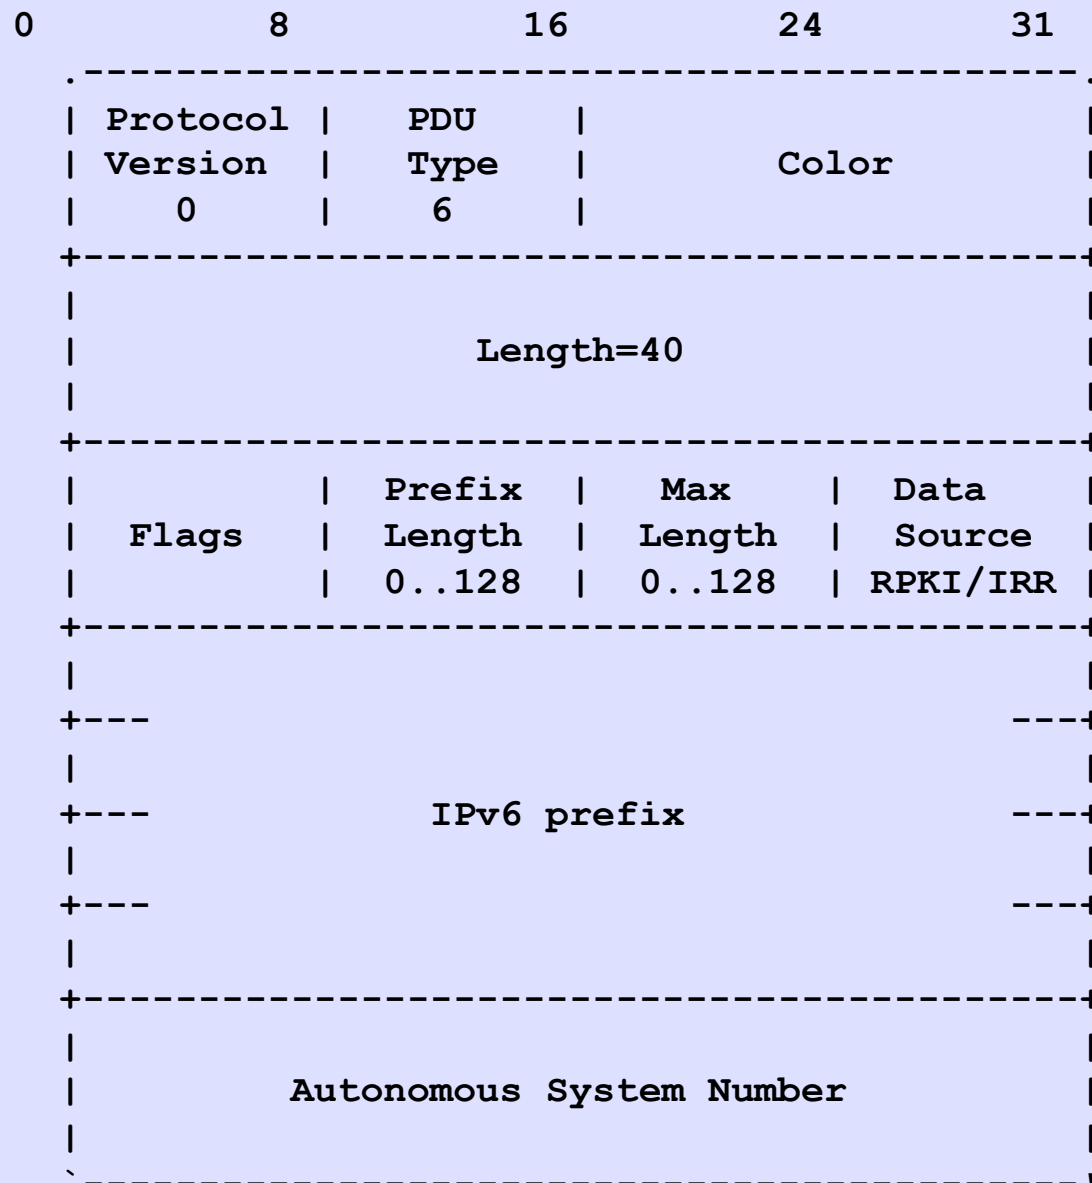
Cache Response




IPv4 Prefix



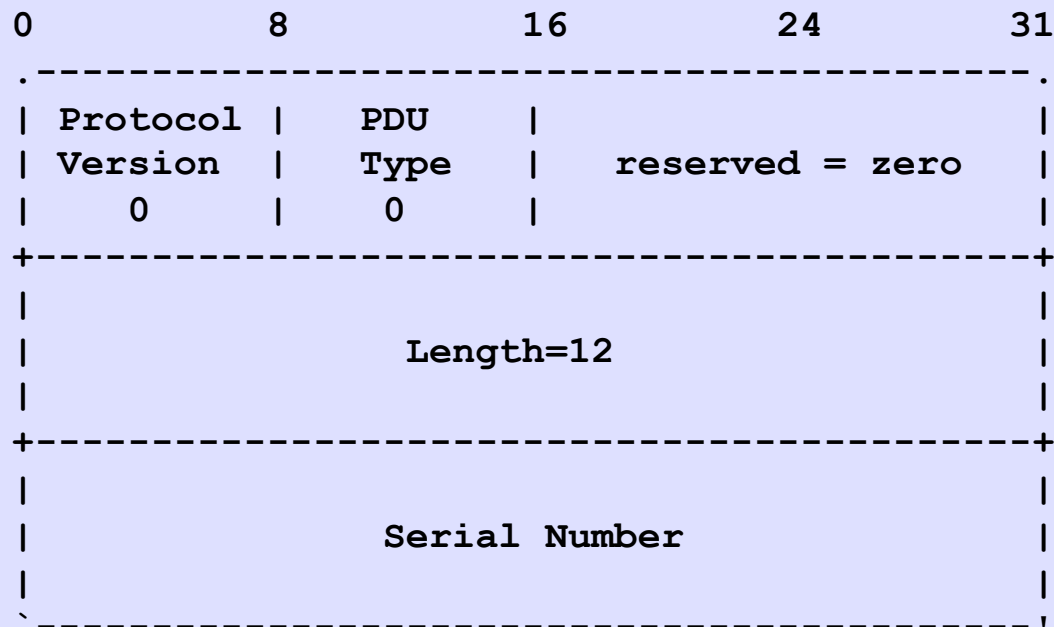
IPv6 Prefix




End of Data



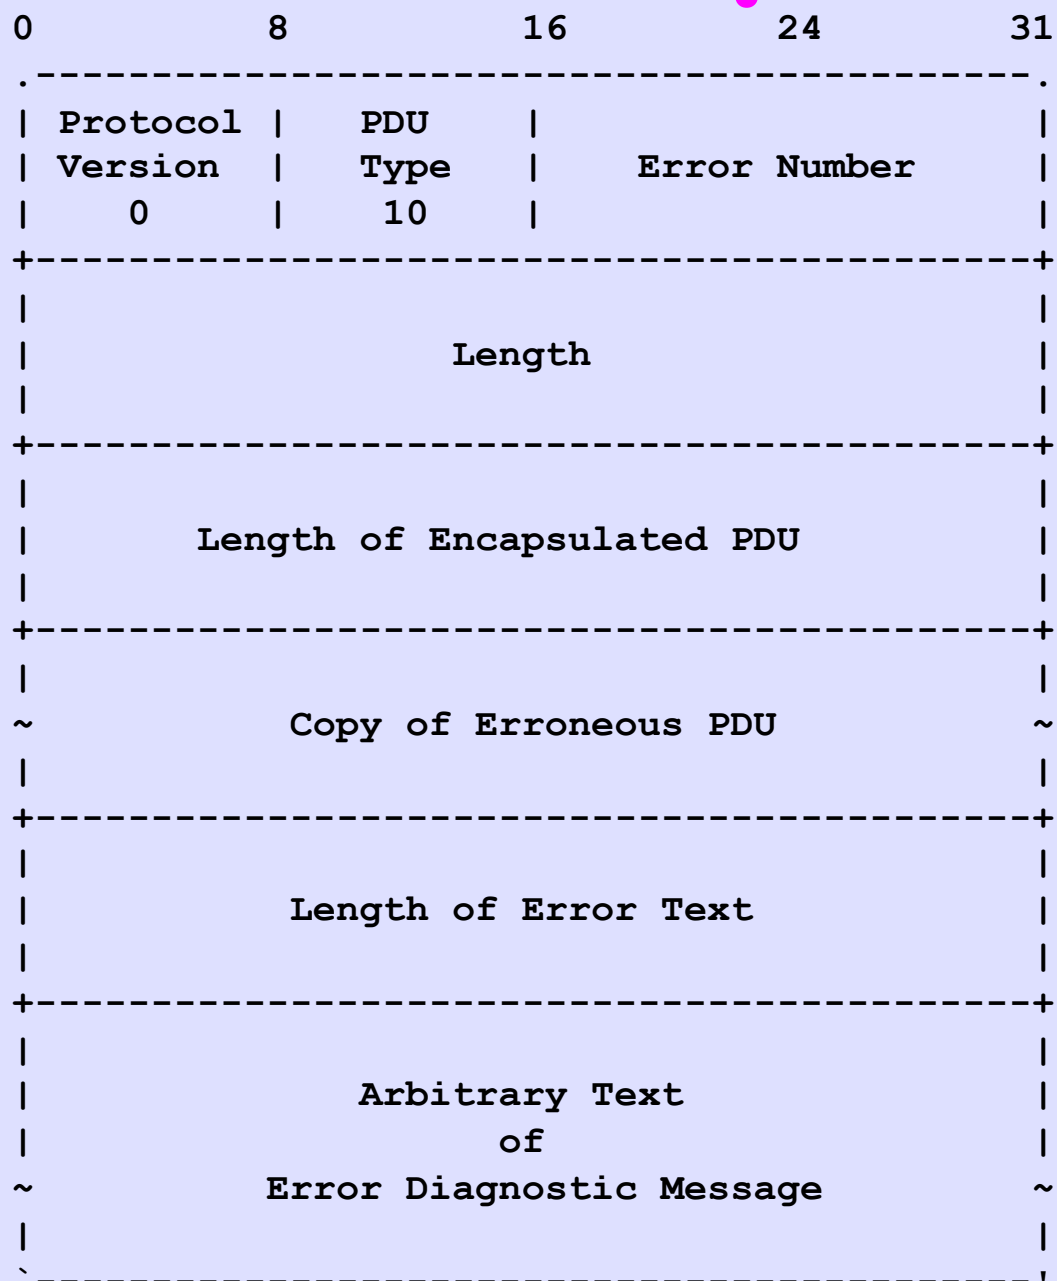
Notify (Think DNS)



Serial Query



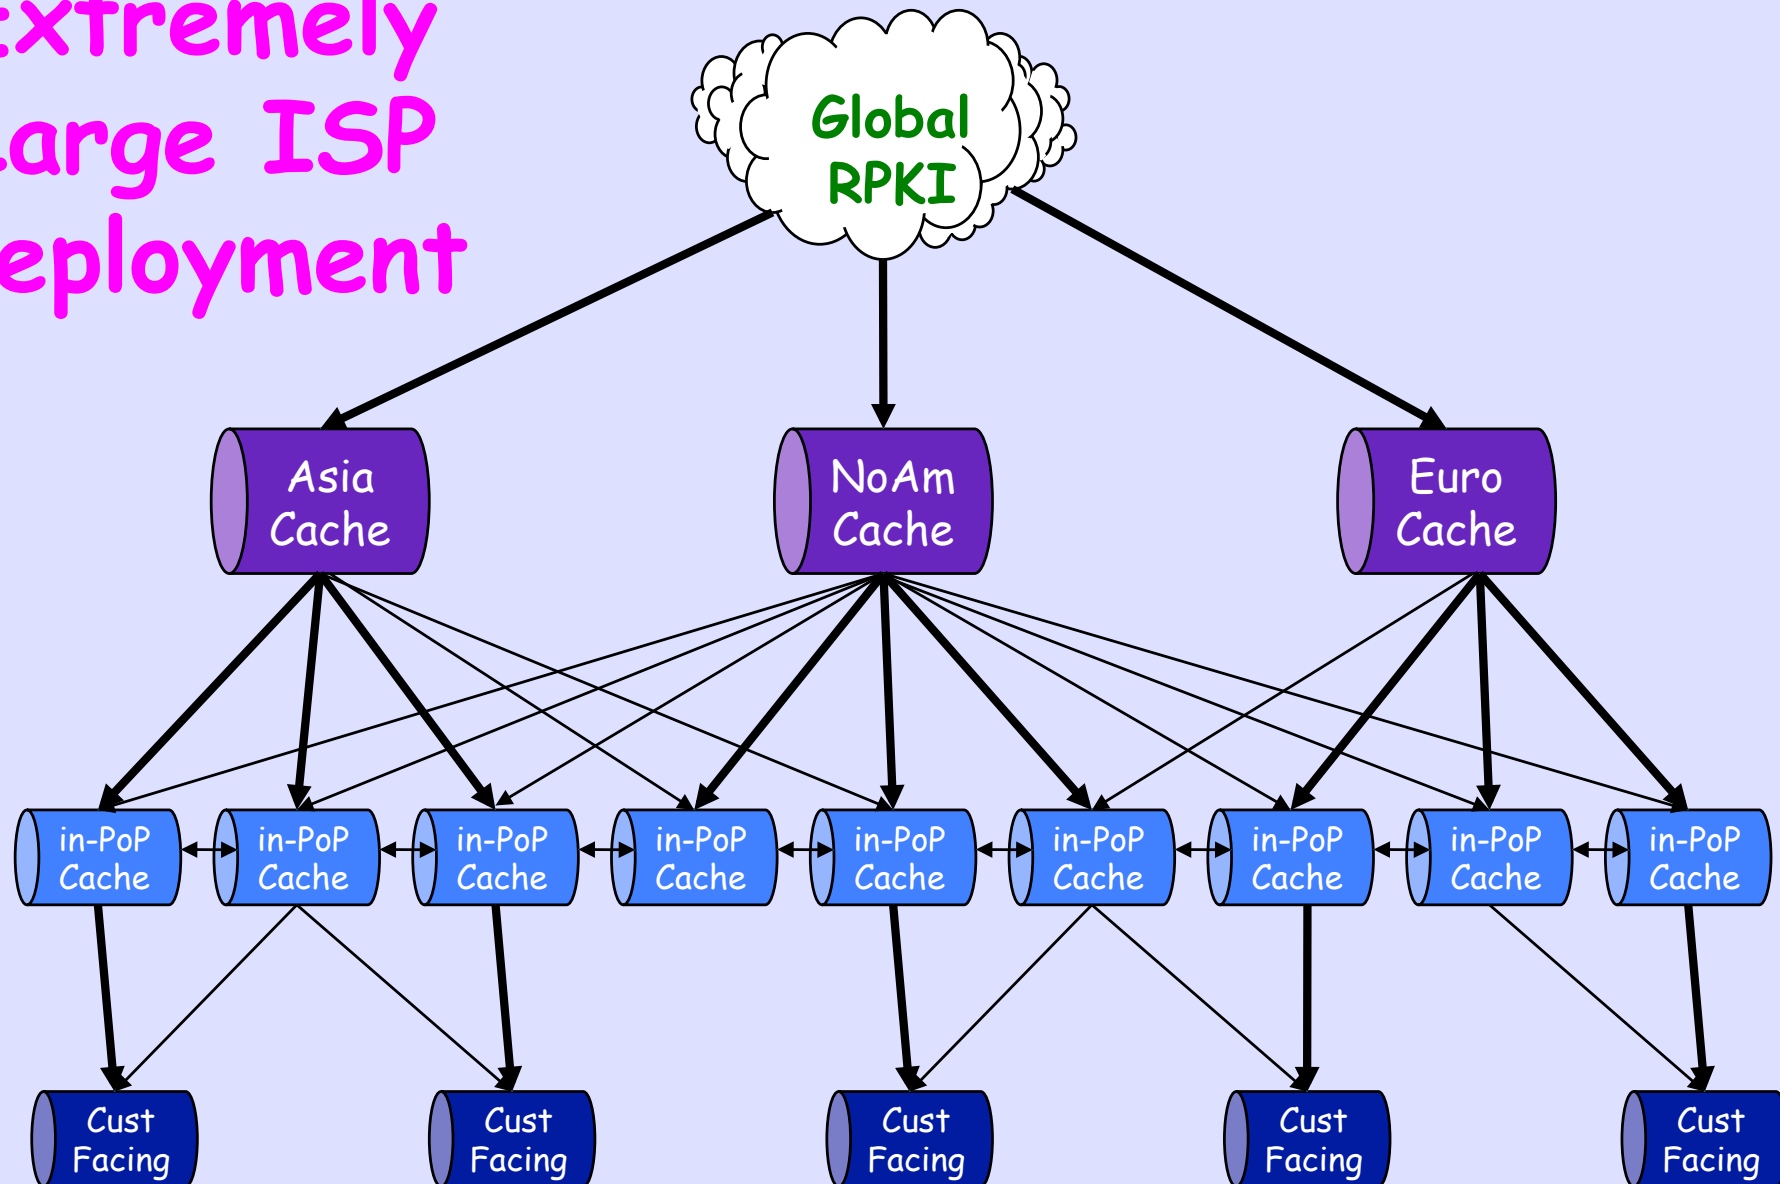
Error Response



Changing Caches

- Running on cache A happily
- A goes bad (A down, sends error, ...)
- Router decides to break off relationship with A
- Router keeps using old data from A
- Router tries other caches in priority order
- Router starts to load from B, in a separate buffer, but still runs on old data from A
- Router finishes loading data from B
- Router flushes all data from A and installs all data from B
- Router reevaluates installed prefixes against new data

Extremely Large ISP Deployment



— High Priority
— Lower Priority

Good Dog!

```
RP/0/1/CPU0:r0.dfw#show bgp 192.158.248.0/24
```

```
BGP routing table entry for 192.158.248.0/24
```

```
Versions:
```

Process	bRIB/RIB	SendTblVer
Speaker	132327	132327

```
Last Modified: Oct 2 01:06:47.630 for 13:33:12
```

```
Paths: (6 available, best #3)
```

```
Advertised to peers (in unique update groups):
```

```
204.69.200.26
```

```
Path #1: Received by speaker 0
```

```
2914 1299 6939 6939 27318
```

```
157.238.224.149 from 157.238.224.149 (129.250.0.85)
```

```
Origin IGP, metric 0, localpref 100, valid, external, \
```

```
origin validity state: valid
```

```
Community: 2914:420 2914:2000 2914:3000 4128:380
```

```
Path #2: Received by speaker 0
```

```
...
```

Bad Dog!

```
RP/0/1/CPU0:r0.dfw#sh bgp 64.9.224.0
```

```
BGP routing table entry for 64.9.224.0/20
```

```
Versions:
```

Process	bRIB/RIB	SendTblVer
Speaker	0	0

```
Last Modified: Oct 2 17:38:27.630 for 4d22h
```

```
Paths: (6 available, no best path)
```

```
Not advertised to any peer
```

```
Path #1: Received by speaker 0
```

```
2914 3356 36492
```

```
157.238.224.149 from 157.238.224.149 (129.250.0.85)
```

```
Origin IGP, metric 2, localpref 100, valid, external,  
origin validity state: invalid
```

```
Community: 2914:420 2914:2000 2914:3000 4128:380
```

Strange Dog!

```
RP/0/1/CPU0:r0.dfw#sh bgp 147.28.0.0
```

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

```
Versions:
```

Process	bRIB/RIB	SendTblVer
Speaker	337691	337691

```
Last Modified: Oct 2 17:40:16.630 for 4d22h
```

```
Paths: (6 available, best #1)
```

```
Advertised to peers (in unique update groups):
```

```
204.69.200.26
```

```
Path #1: Received by speaker 0
```

```
2914 3130
```

```
157.238.224.149 from 157.238.224.149 (129.250.0.85)
```

```
Origin IGP, metric 68, localpref 100, valid, external, \
```

```
origin validity state: not found
```

```
Community: 2914:410 2914:2000 2914:3000 4128:380
```

Open Source (BSD Lisc)

Running Code

<https://subvert-rpki.hactrn.net/>

Test Code in Routers

Talk to Ed Kern

Work Supported By

- **US Government**

THIS PROJECT IS SPONSORED BY THE DEPARTMENT OF HOMELAND SECURITY UNDER AN INTERAGENCY AGREEMENT WITH THE AIR FORCE RESEARCH LABORATORY (AFRL).

- **Internet Initiative Japan**

- **Cisco, Juniper, Google, NTT, Equinix**