The RPKI, IPv4, ...
The News at Eleven

AfriNIC / Rabat

2008.06.04

Randy Bush <randy@psg.com>

Internet Initiative Japan

- Originally, an initiative to get Japan on the Internet
- Asian and some US backbone
- Commercial customer base
- Internet, not telephant, MPLS, ...
- First commercial IPv6 deployment
- WIDE, Kame, ...

We're Old Fashioned

- Internet, not ATM-2 == MPLS, etc
- VoIP etc over IP, it is possible!
- IPSec is a big seller, the P in VPN
- High touch, a lot of services
- Quality, quality, and quality
- And we're profitable!

Agenda

- RPKI (some details) and why I care
- BGP Security
- IPv4 free pool run-out
- Policy, Fairness, and Best Use
- Routing Table Growth
- What I want
- What's next?

I have been working on this RPKI X.509 Certification of Resource Stuff

X.509 Cert w/ 3779

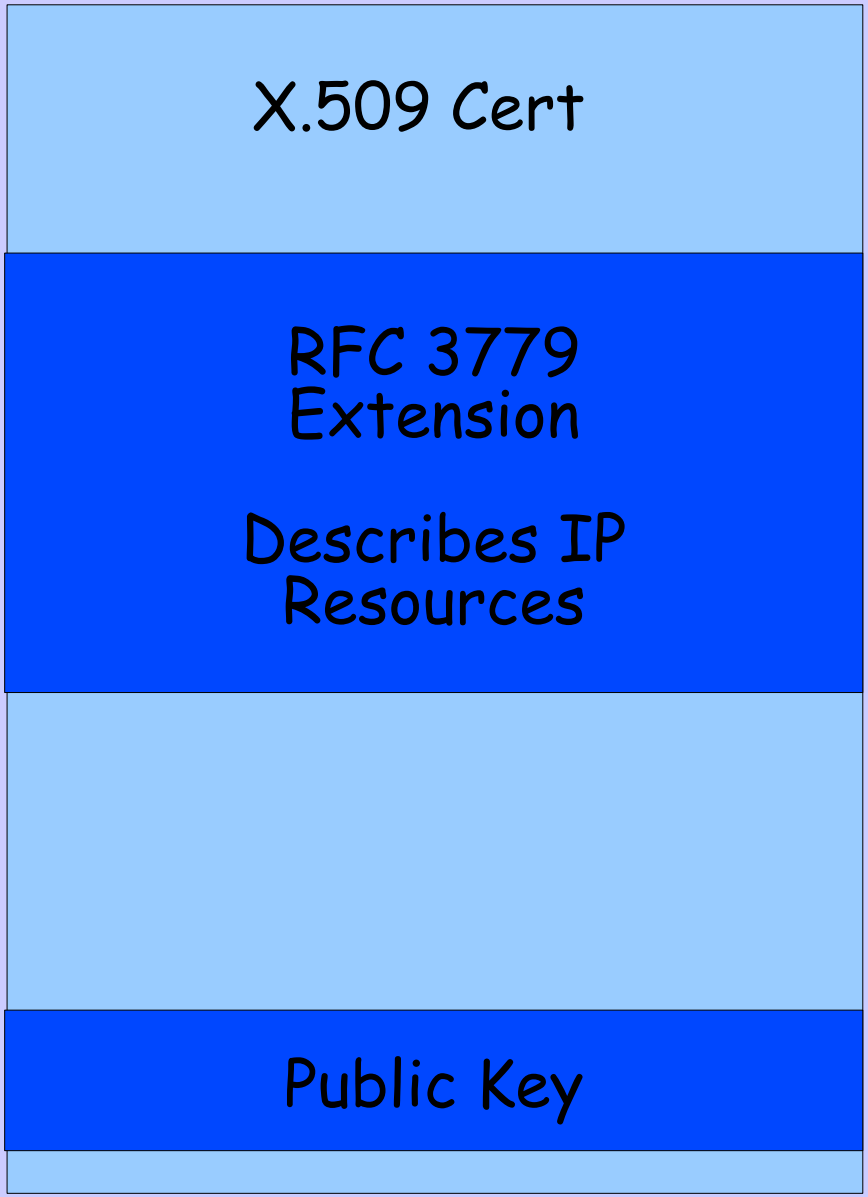

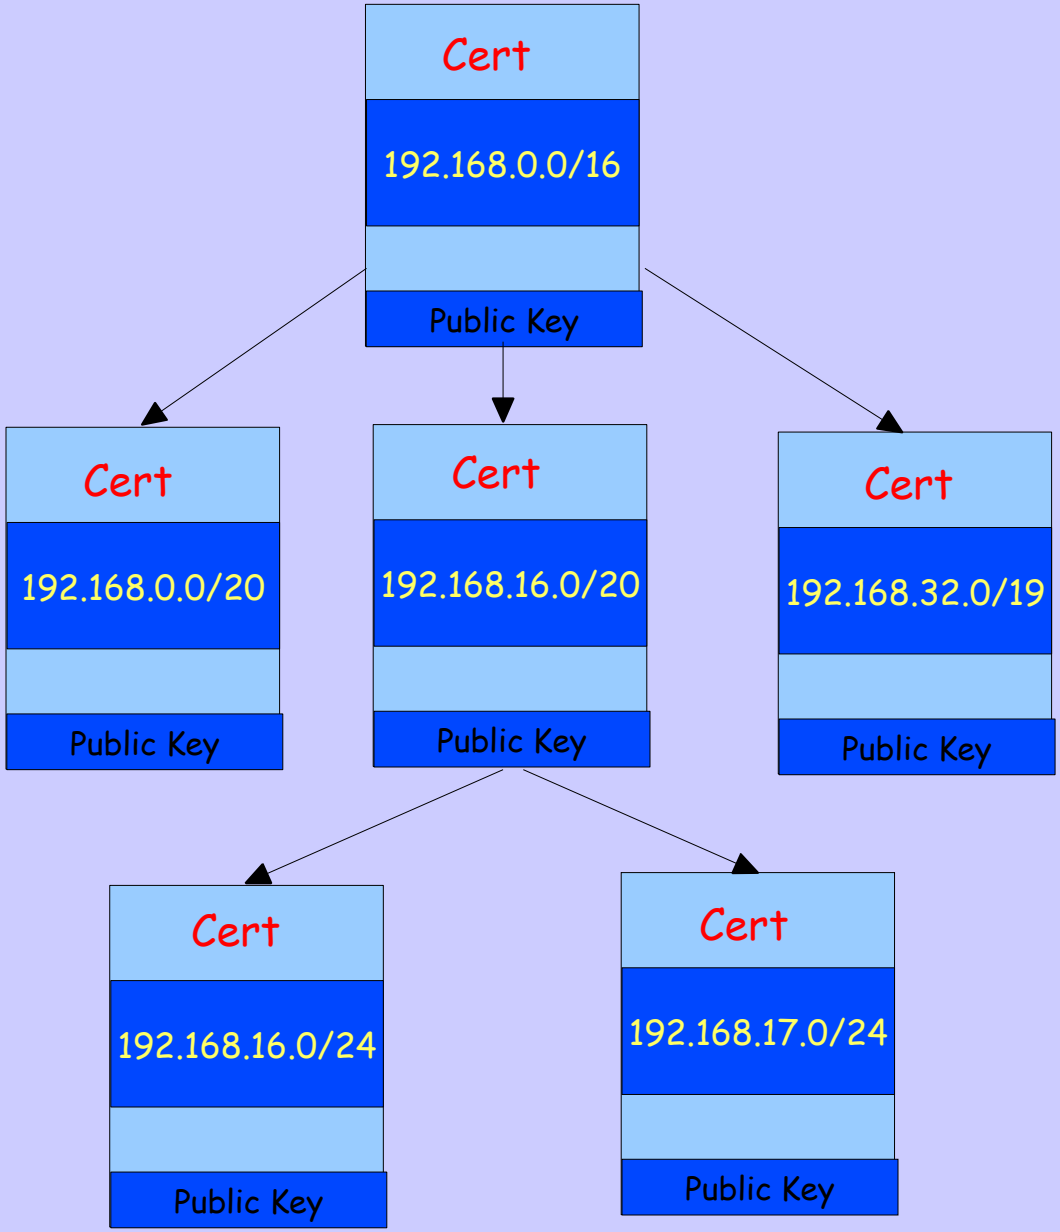
Cert
192.168.0.0/16
Public Key
Cert
192.168.0.0/20
Public Key
Cert
192.168.16.0/20
Public Key
Cert
192.168.32.0/19
Public Key
Cert
192.168.16.0/24
Public Key
Cert
192.168.17.0/24
Public Key

Route Origin Attestation (ROA)

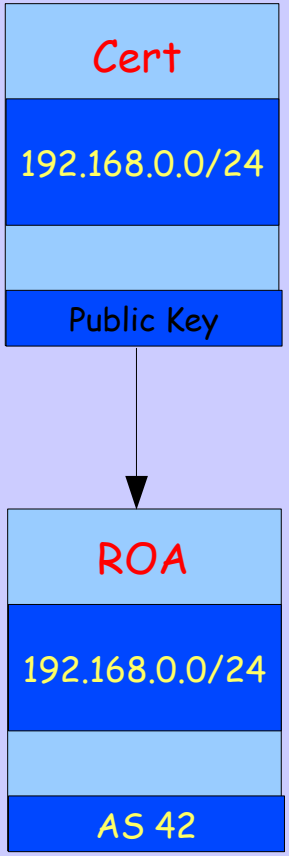

Resource Public Key Infrastructure

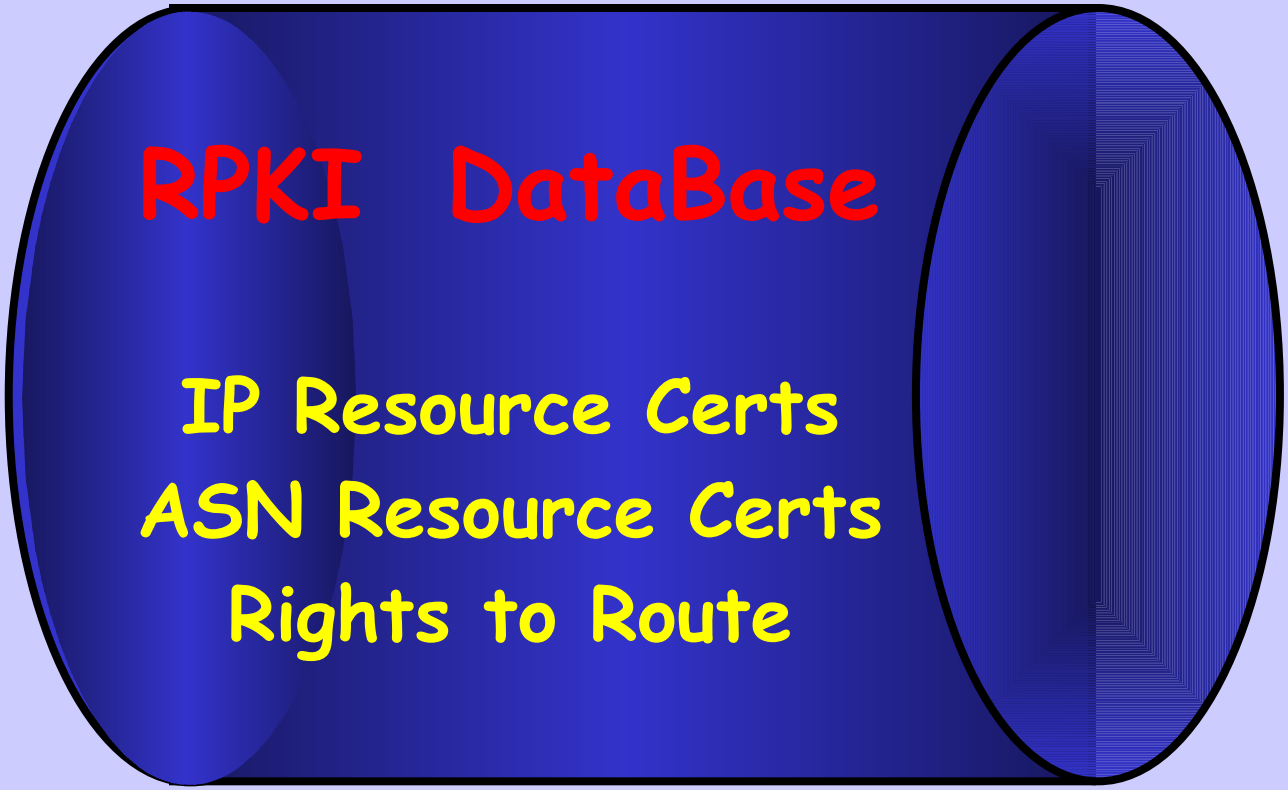

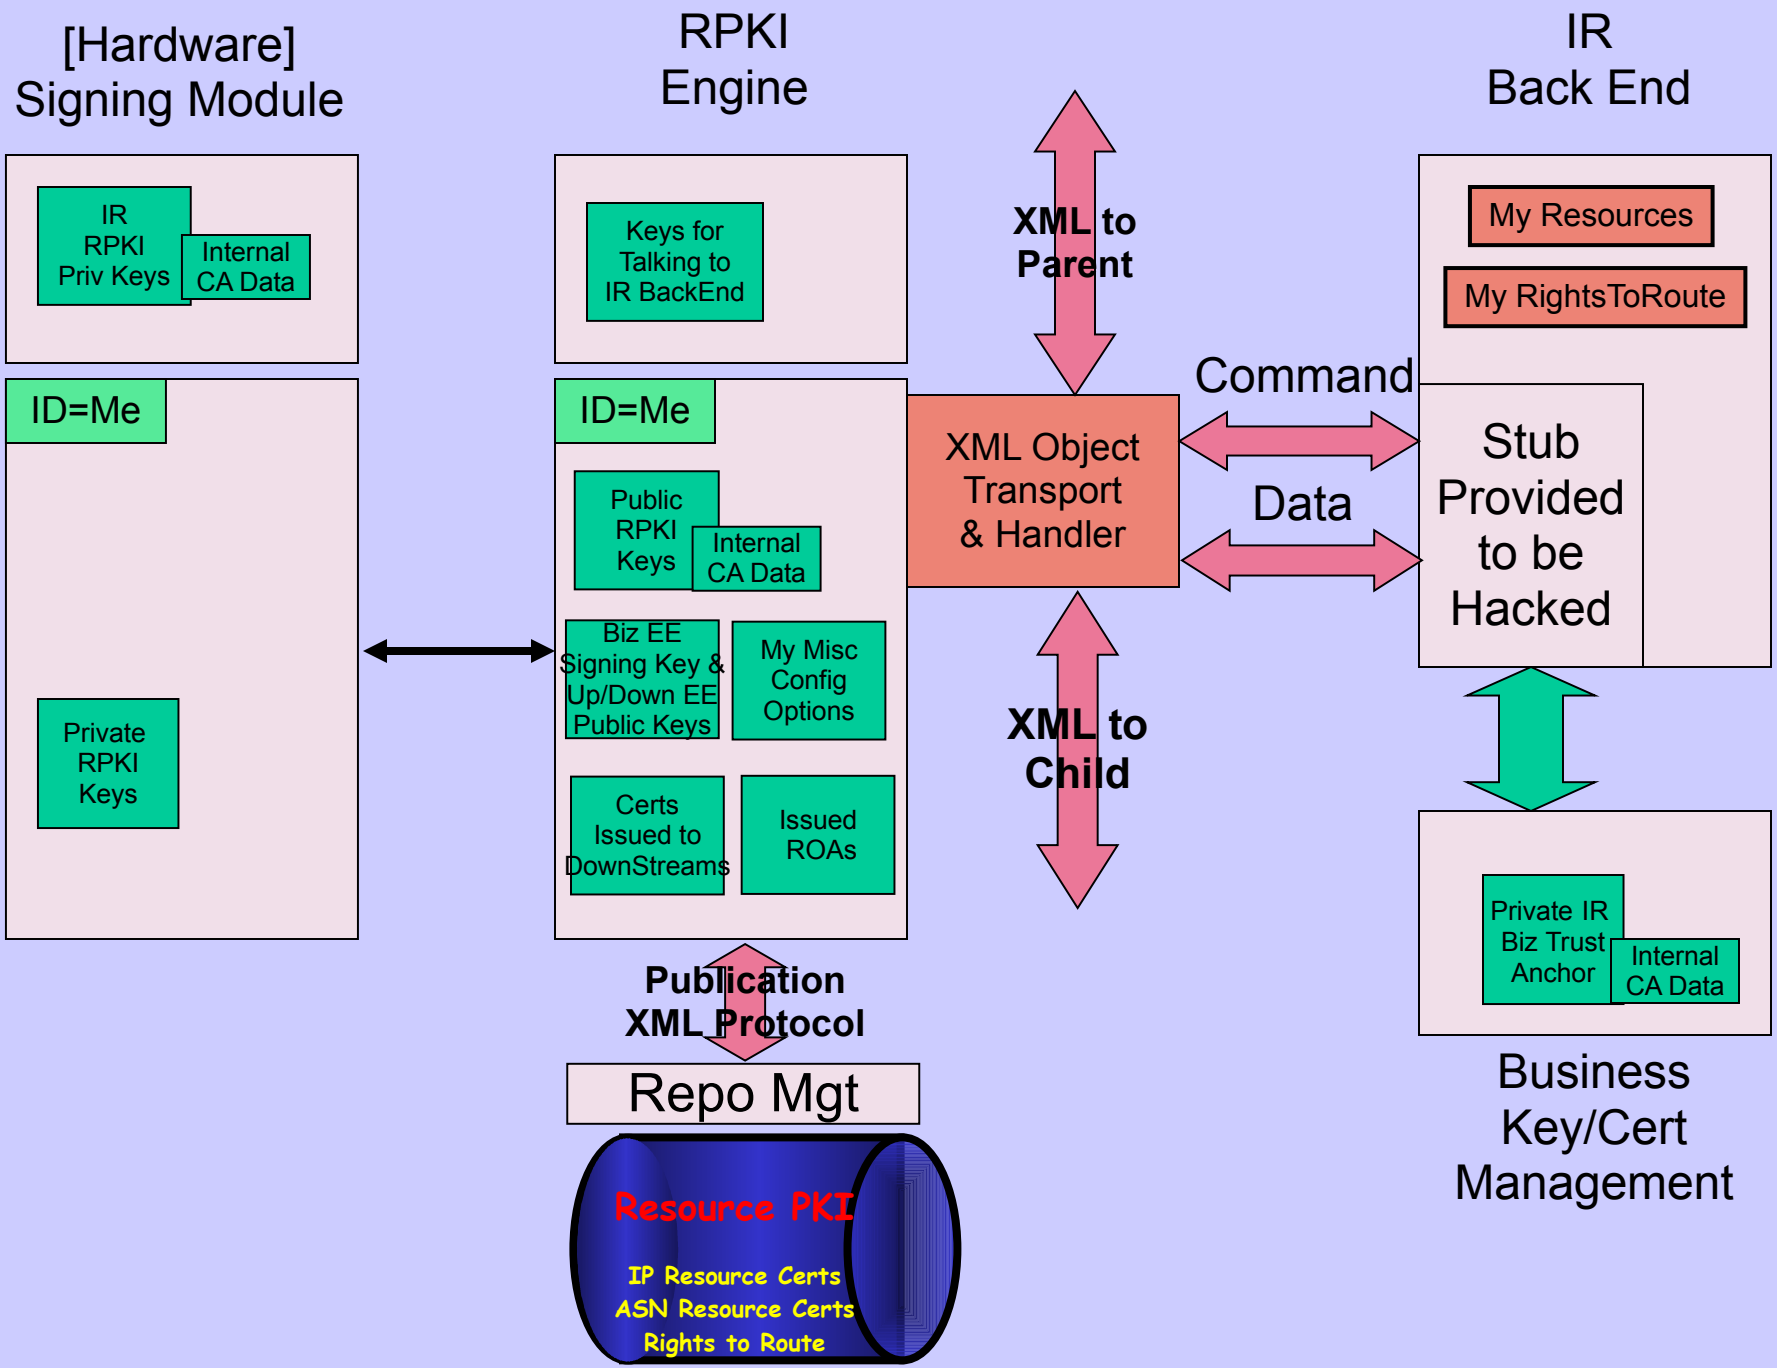
[Hardware]
Signing Module
RPKI
Engine
IR
Back End
IR
RPKI
Priv Keys
Internal
CA Data
Keys for
Talking to
IR BackEnd
My Resources
My RightsToRoute
ID=Me
ID=Me
Public
RPKI
Keys
Internal
CA Data
XML to
Parent
Command
XML Object
Transport
& Handler
Data
Stub
Provided
to be
Hacked
Biz EE
Signing Key &
Up/Down EE
Public Keys
My Misc
Config
Options
Private
RPKI
Keys
XML to
Child
Certs
Issued to
DownStreams
Issued
ROAs
Private IR
Biz Trust
Anchor
Internal
CA Data
Publication
XML Protocol
Repo Mgt
Resource PKI
IP Resource Certs
ASN Resource Certs
Rights to Route
Business
Key/Cert
Management

RPKI Interfaces/Users

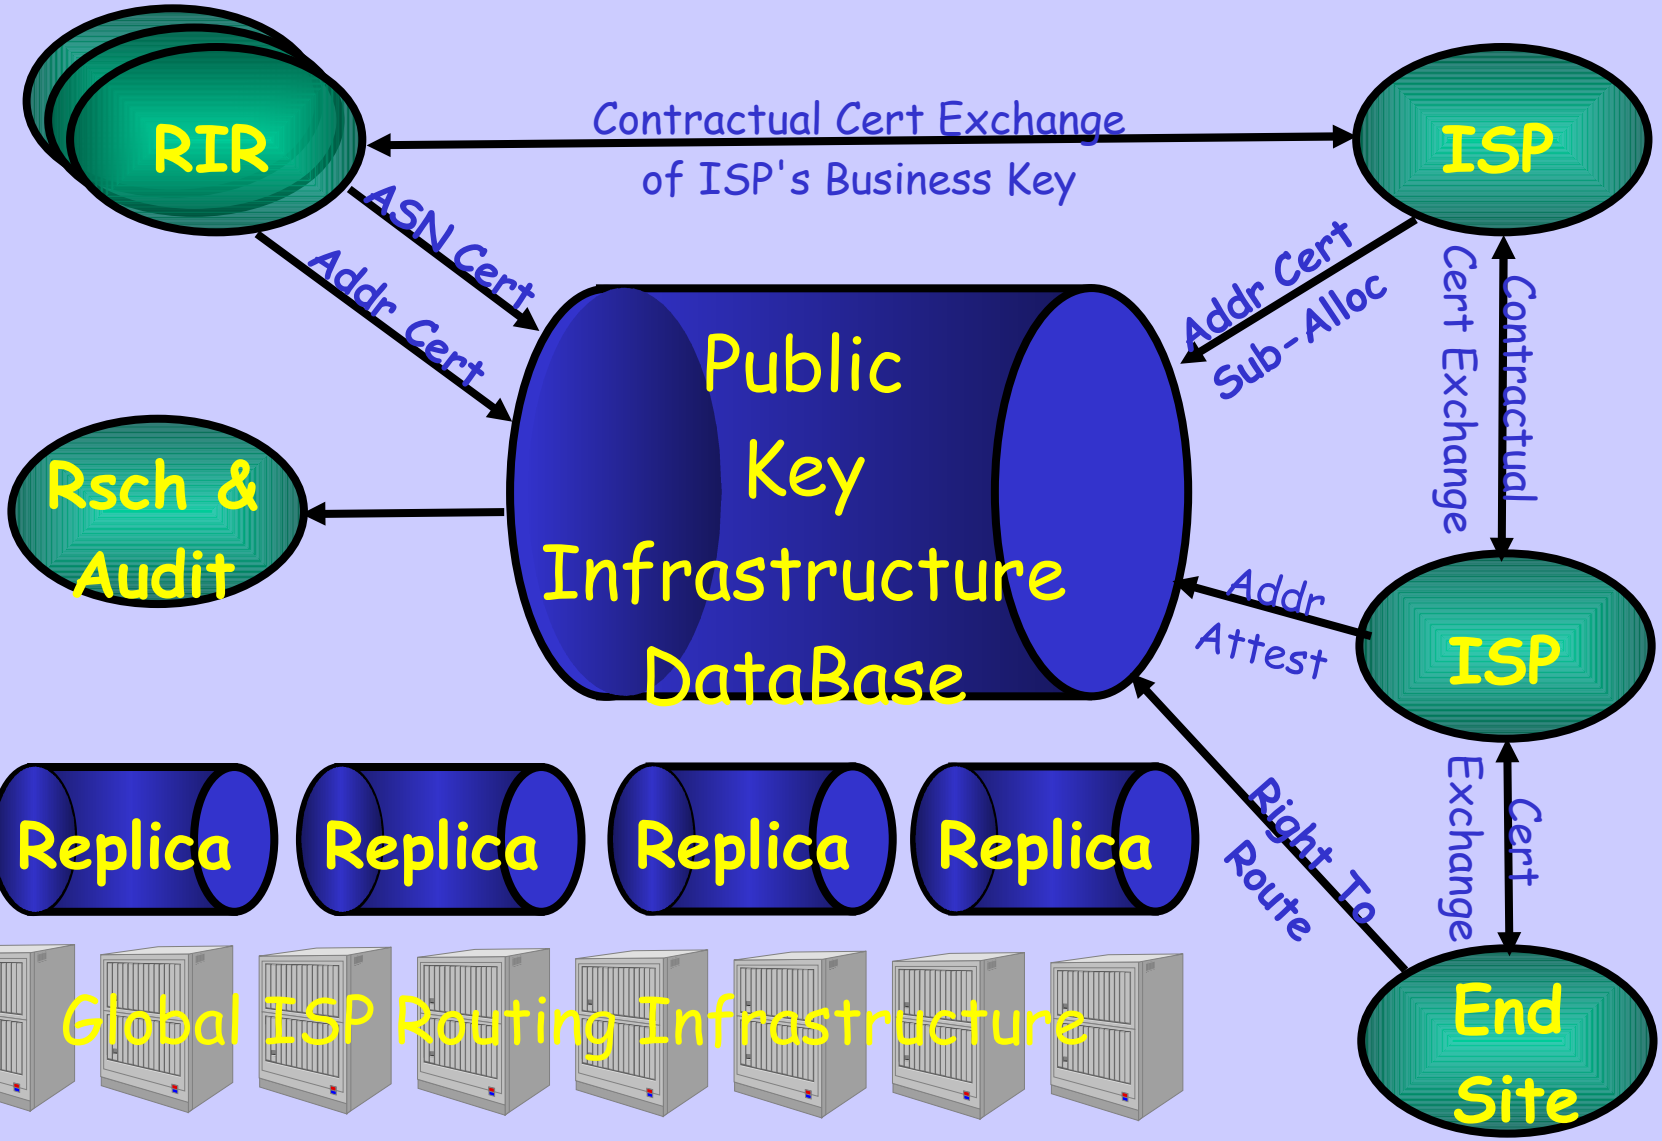

Layer 9 War

- RIRs do not want IANA to sign their certs!
- They want to each be their own root trust anchor
- OTOH, they each want to 'own' their customer ISPs
- It is all about power, not technology

Why Do I Care?

- Formal validation of who can ask me to route what prefixes
- Automation of route filters
- Real routing security in the long term
- Fairness in address trading

Cheap Filter Automation

- This is Ruediger's hack, not mine
- Use ROAs to generate a fake IRR of Route: objects
- Put this ersatz-IRR in front of the other IRRs when running peval()
- A lot of benefit at zero RPSL or software change!

But where I am really going in the long term is

BGP Routing Security

Diversion Attack

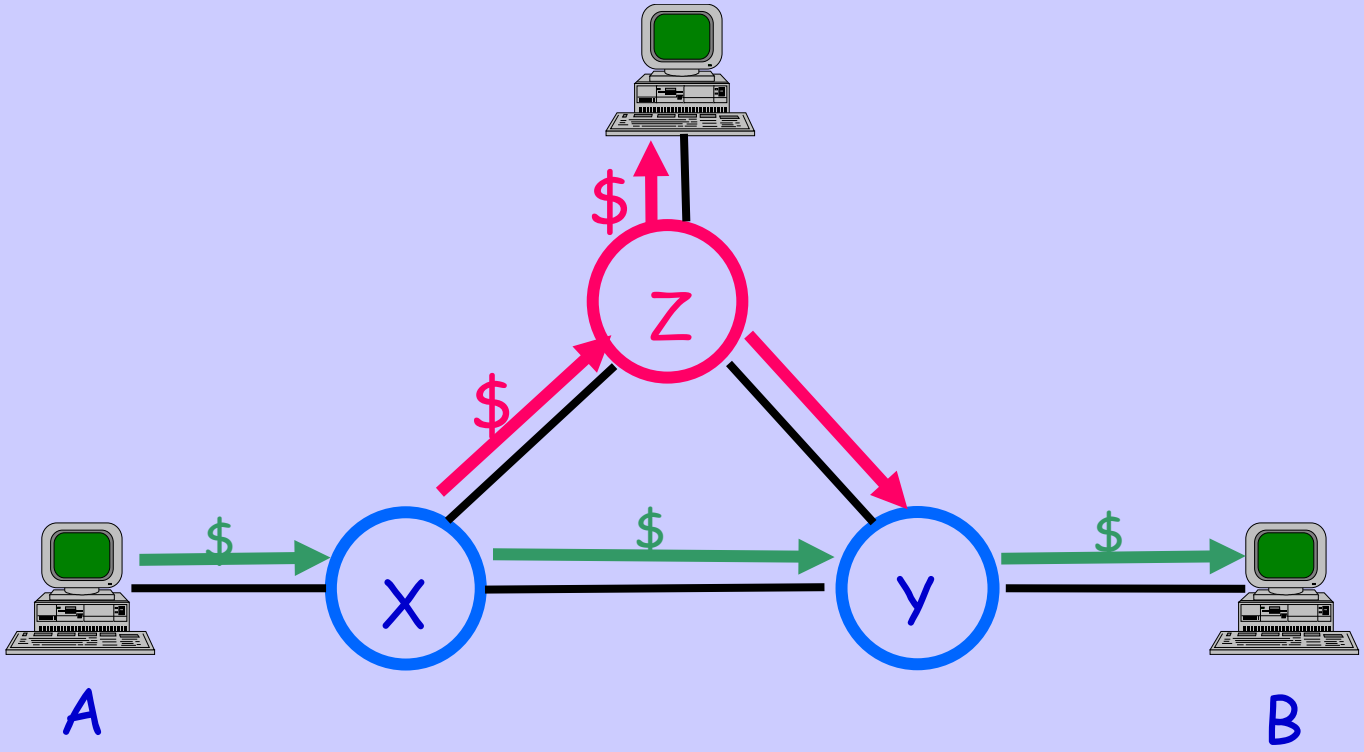

Expected Path - A->X->Y->B

Diverted Path - A->X->Z->Y->B

How Does Z Do It?

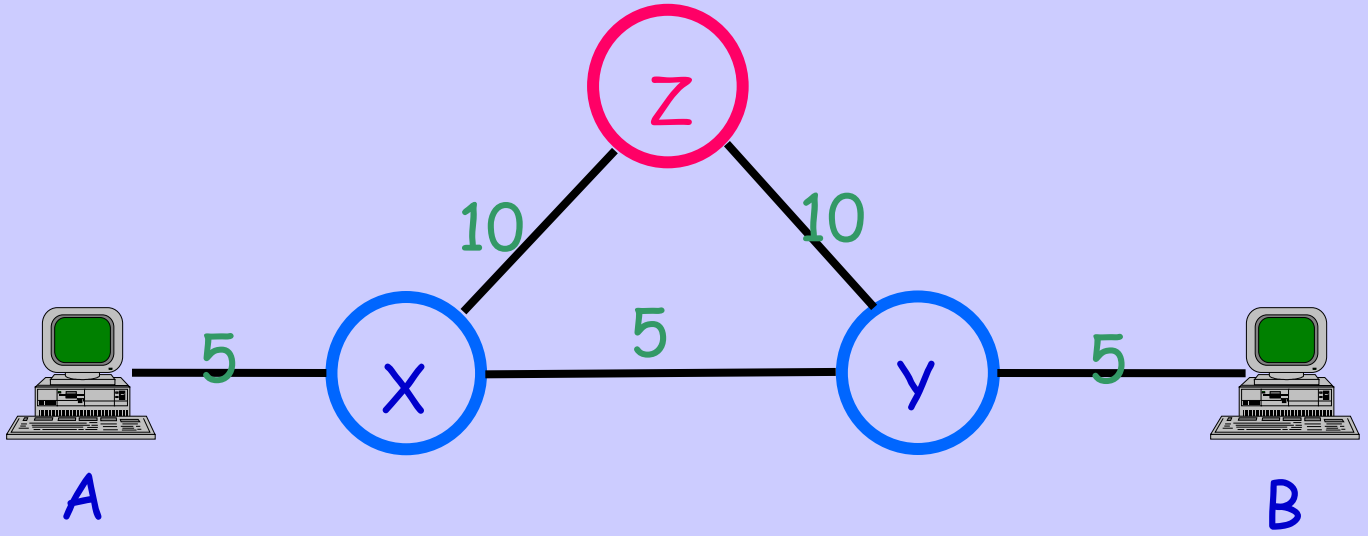

Y tells X and Z that costs are B:5

X tells A and Z that costs are Y:5 B:10

Z tells X that costs are Y:10 B:15

Z tells X that costs are Y:10 **B:4**

X now sends B's traffic to Z!!!

Why is this a Hard Problem?

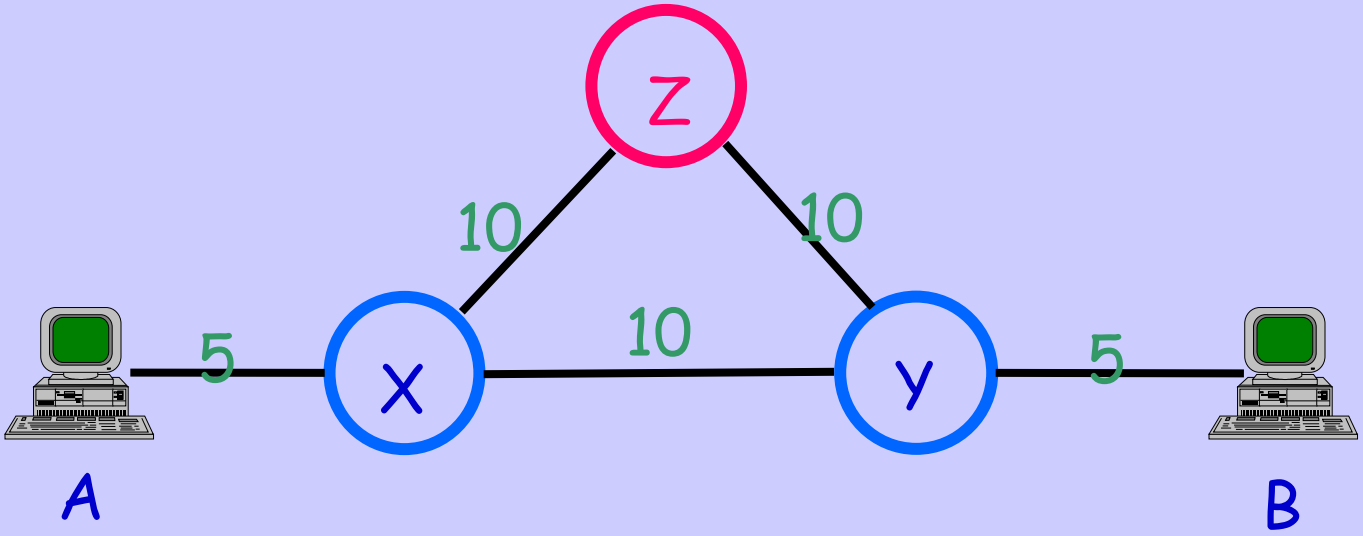

- X does not really know Z's links
- X does not really know Y's links
- They trust each other re costs!

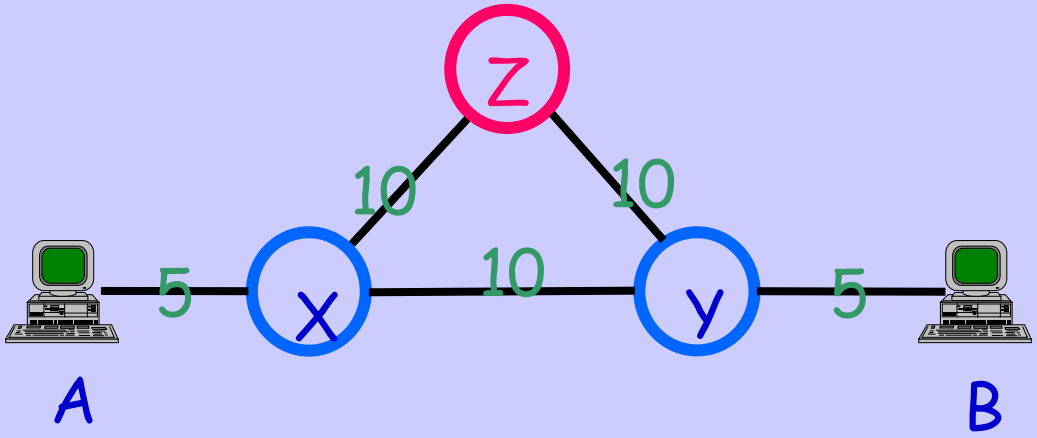

- Validating IP prefix ownership does not help, as Z is not lying about B's owning it
- Using IRR-like peering map does not help, as Z is not lying about who connects to whom

One Approach

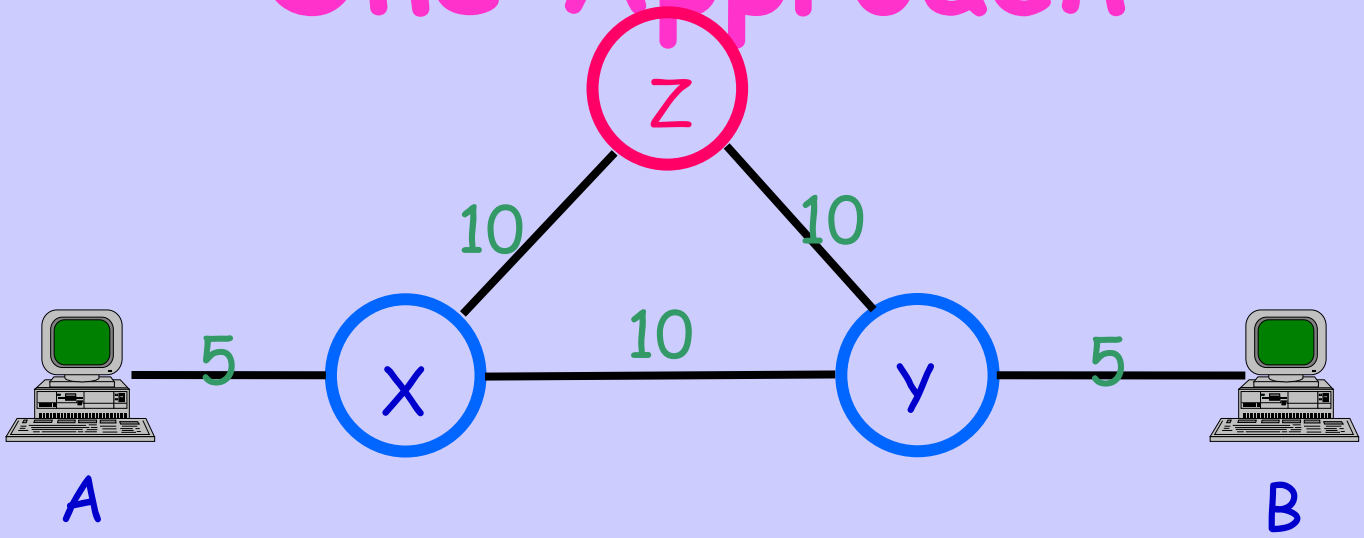

- B cryptographically signs the message to Y Sb(Y->B=5)
- Y signs messages to X and Z encapsulating B's message

  Sy(X->Y=10 Sb(Y->B=5)) and Sy(Z->Y=10 Sb(Y->B=5))
- Z can only sign Sz(X->Z=10 Sy(Z->Y=10 Sb(Y->B=5)))
- Now X can verify paths and costs
- **Forward path signing** solves the 'simple' case

Costs

- Crypto-CPU-intensive
- Use caching
- Use pre or delayed validation
- Moore's 'Law' is our friend
- Crypto chips are cheap
- Most announcements are boring

Chapter Two

IPv4 Free Pool Run-out,

Best and Fairest Use,

Address 'Trading,'

The Universe, and everything

IPv4 Free-Pool Run-Out

- IPv4 Free Pool will run-out in a few years
- This is not news. See graphs of Frank Solensky over ten years ago; and Geoff's
- IPv4 will go to a **trading model**
- Registries will become title agents, not allocators, of IPv4 space
- RIRs are developing full multi-RIR/LIR open source RPKI software

What Should Have Happened

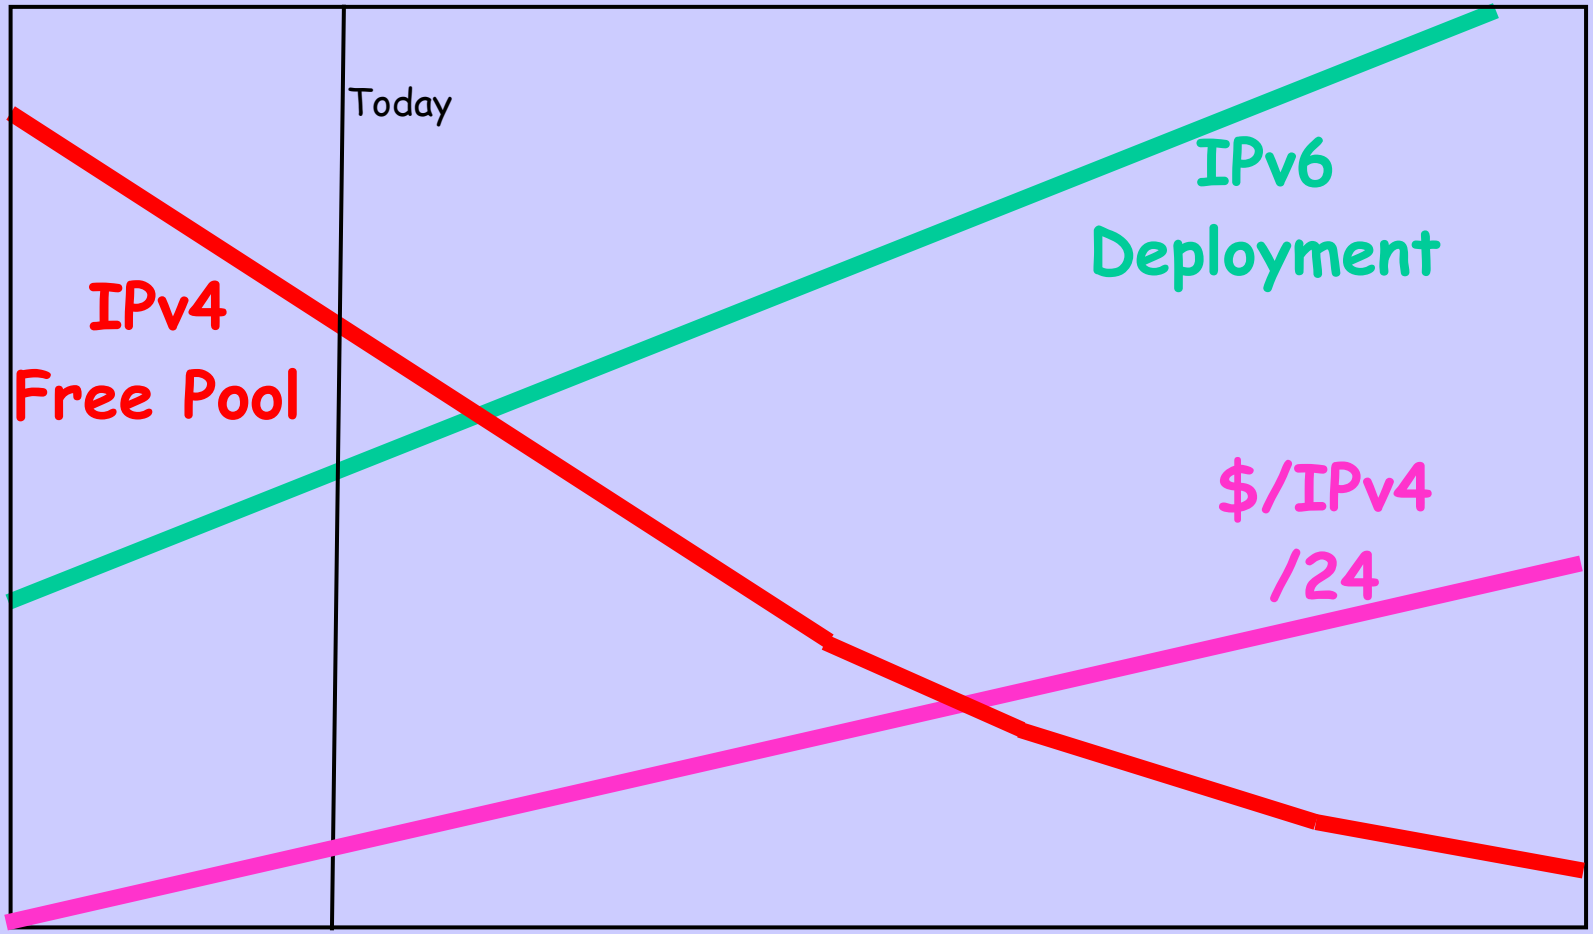

What Is Happening?

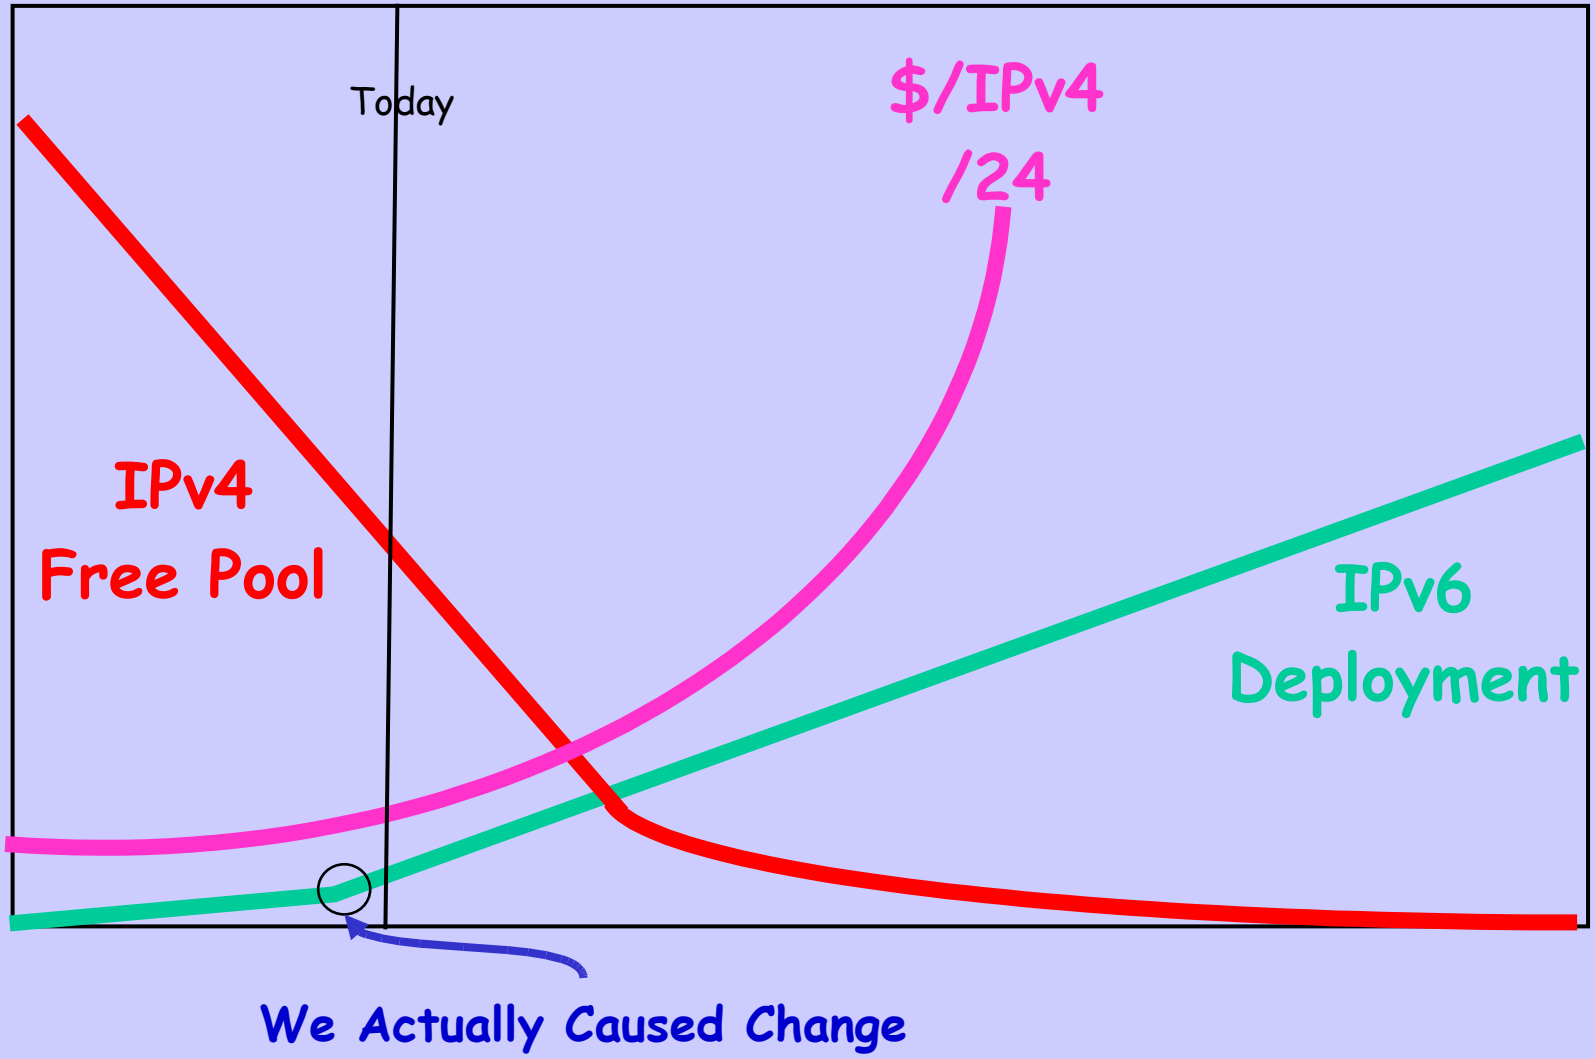

If You Don't Believe It

IPv6 Prefix Allocations

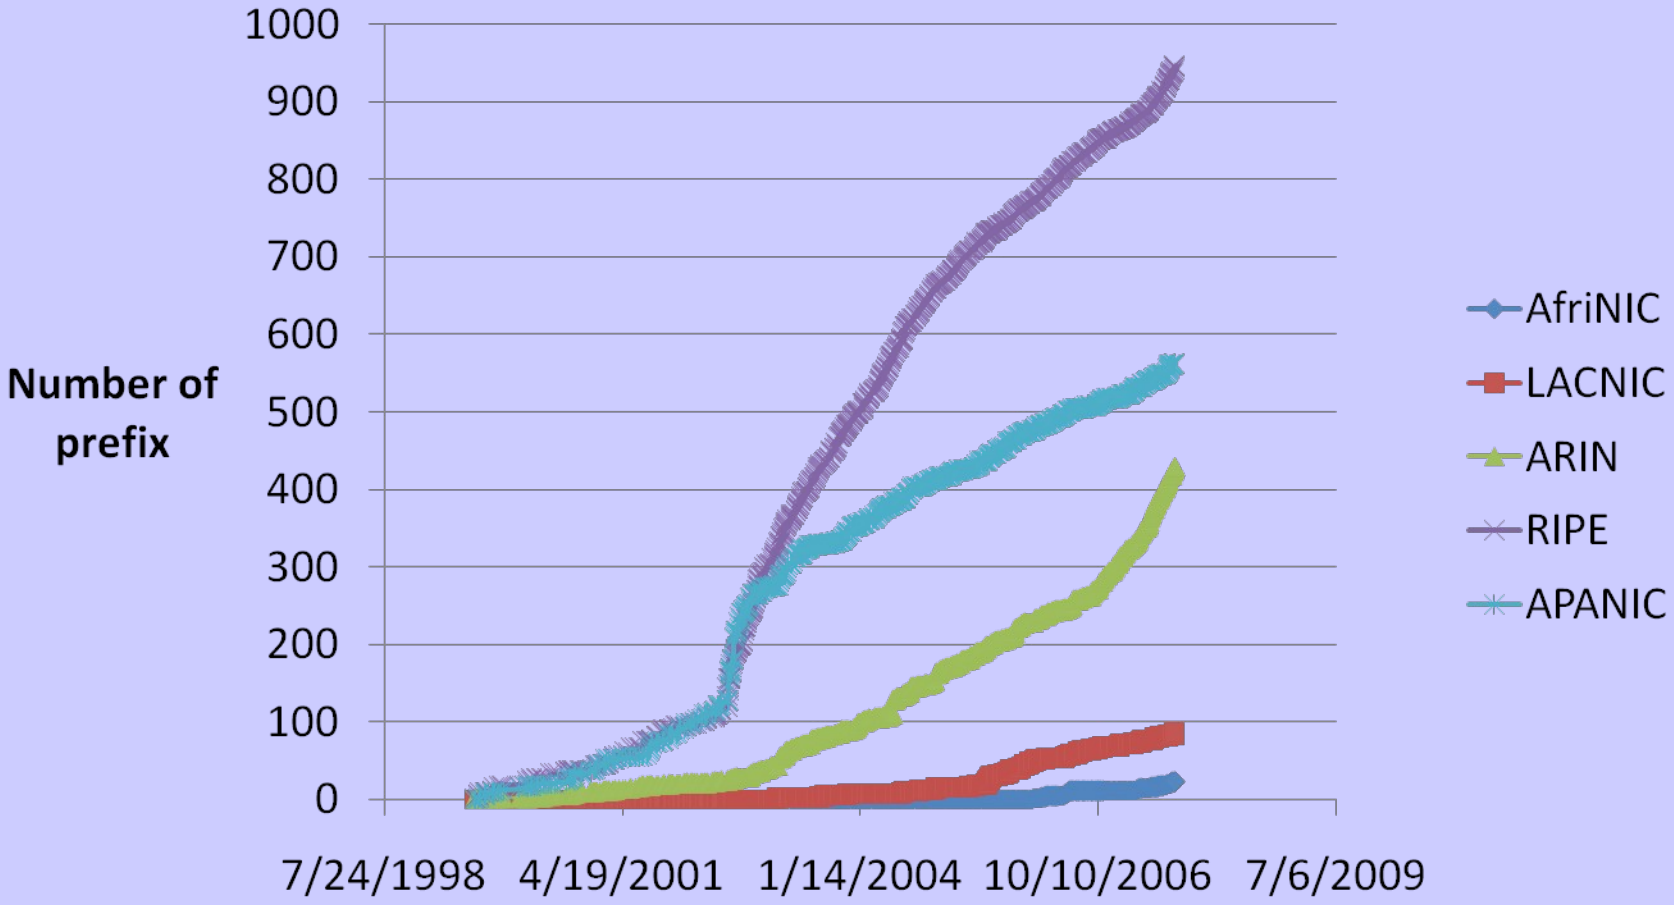

BGP Prefix Announcements

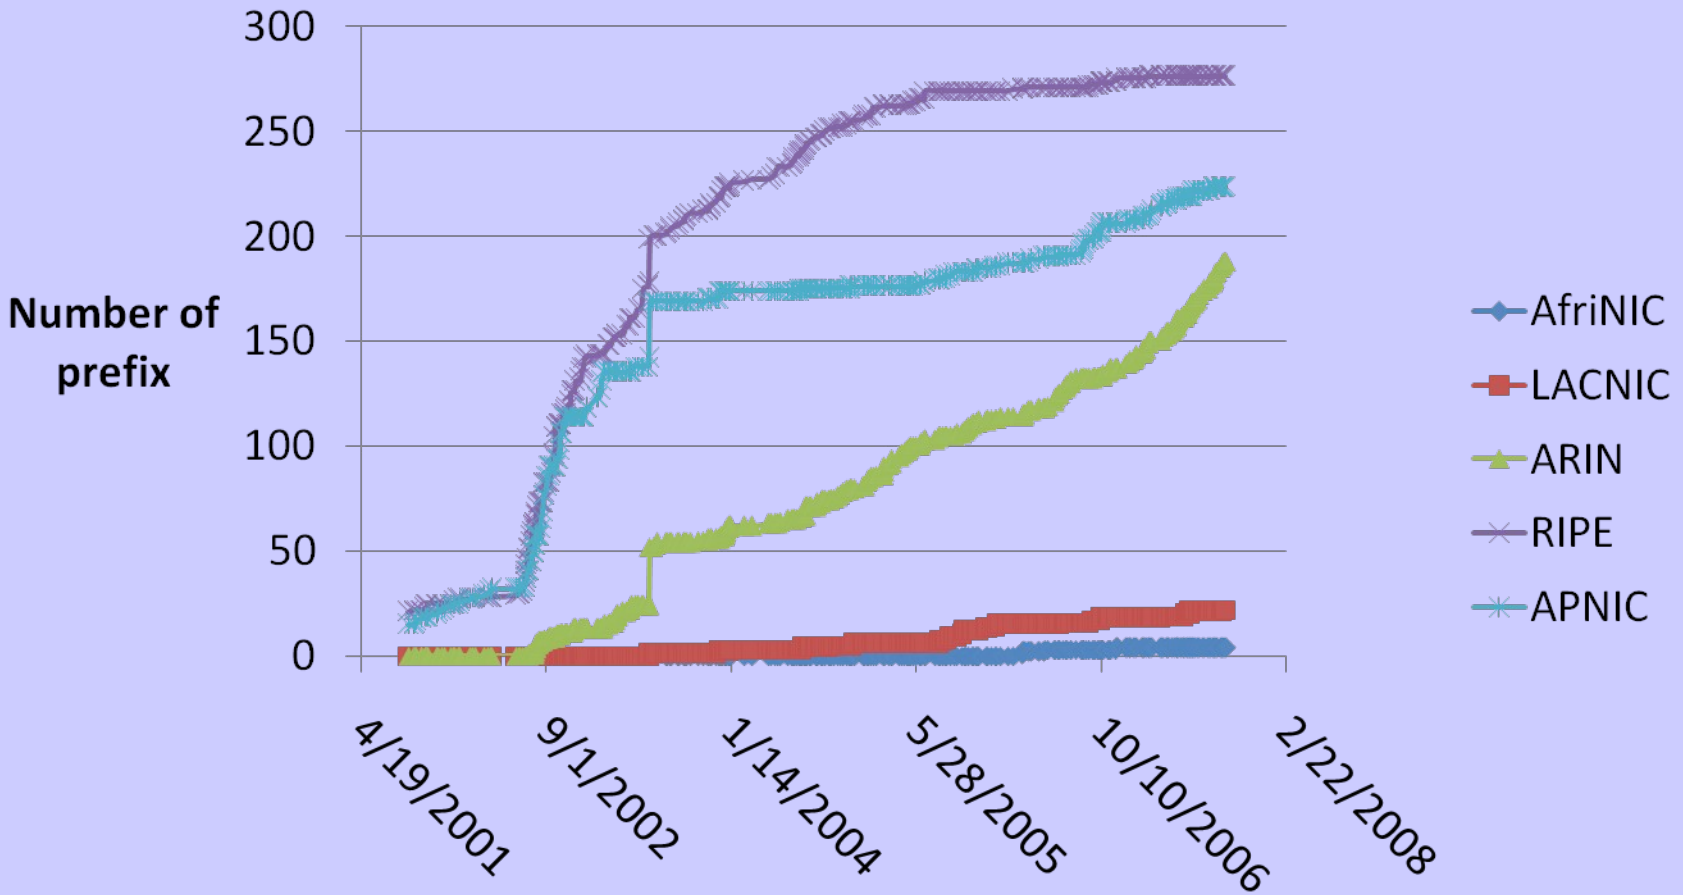

Geoff has more recent measurements and the last year is better!

So How is IPv4 Going to Play Out?

Are current societal and administrative systems fair?

What's 'fair'?

Is This 'Fair'?

In total, 386,590 /24s were issued to 1,866 organizations.

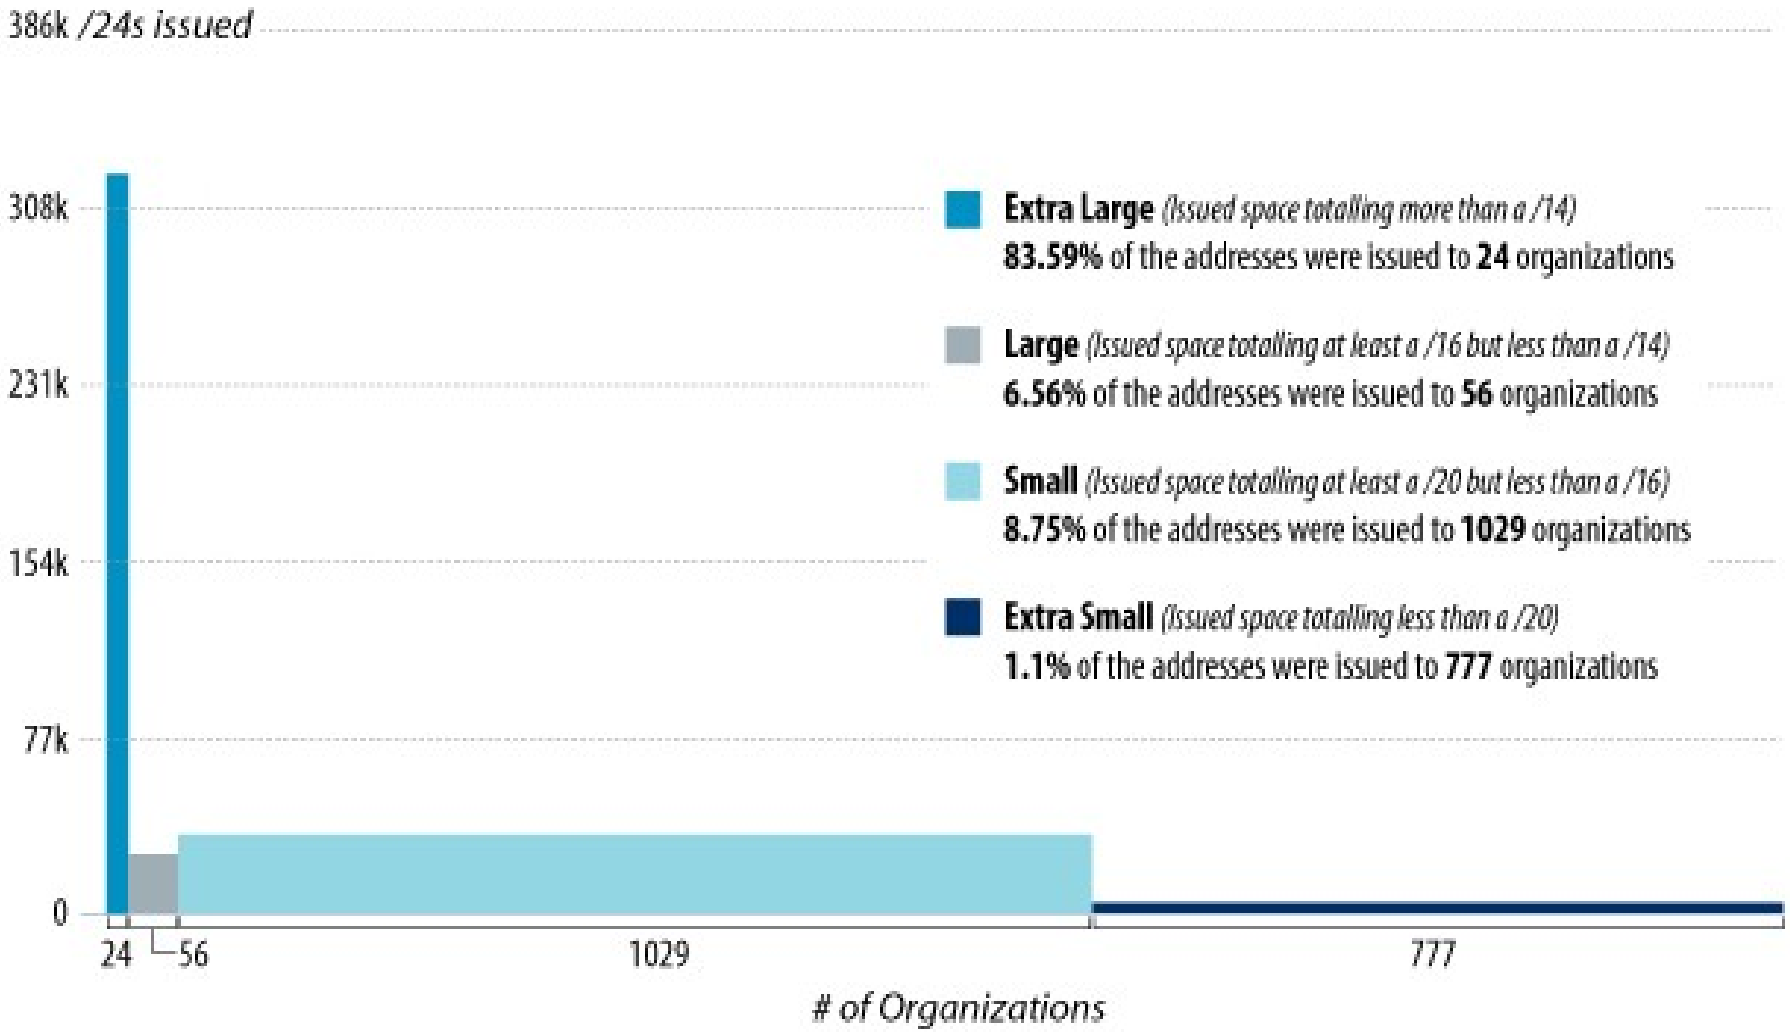

That was ARIN for 2006-7
Other regions have somewhat different distributions.

No one wants to talk about this because grown-ups might be listening.

Yes, it models the market concentration in North America but ...

The RIR communities have placed severe barriers to entry at the low end !

A newcomer may not be able to 'justify' a /20-/24

Why is This?

- We're saving routing table size at the expense of barrier to entry
- Should we be doing this at the end?
- Instead, give me tools to filter out intentional deaggregation
- Note that RPKI certificates are maximally aggregated

Is this how we think the last few /8s should be distributed?

What Might We Do?

- I am not an expert, but I admit it, which is a differentiator :)
- Even distribution to RIRs of the last /8s
- Within RIRs, damp big request[er]s
- Enable small requests
- Save the last /16 for unknowns and emergencies
- Open market with transparency

ARIN Legacy Prefix Announcements

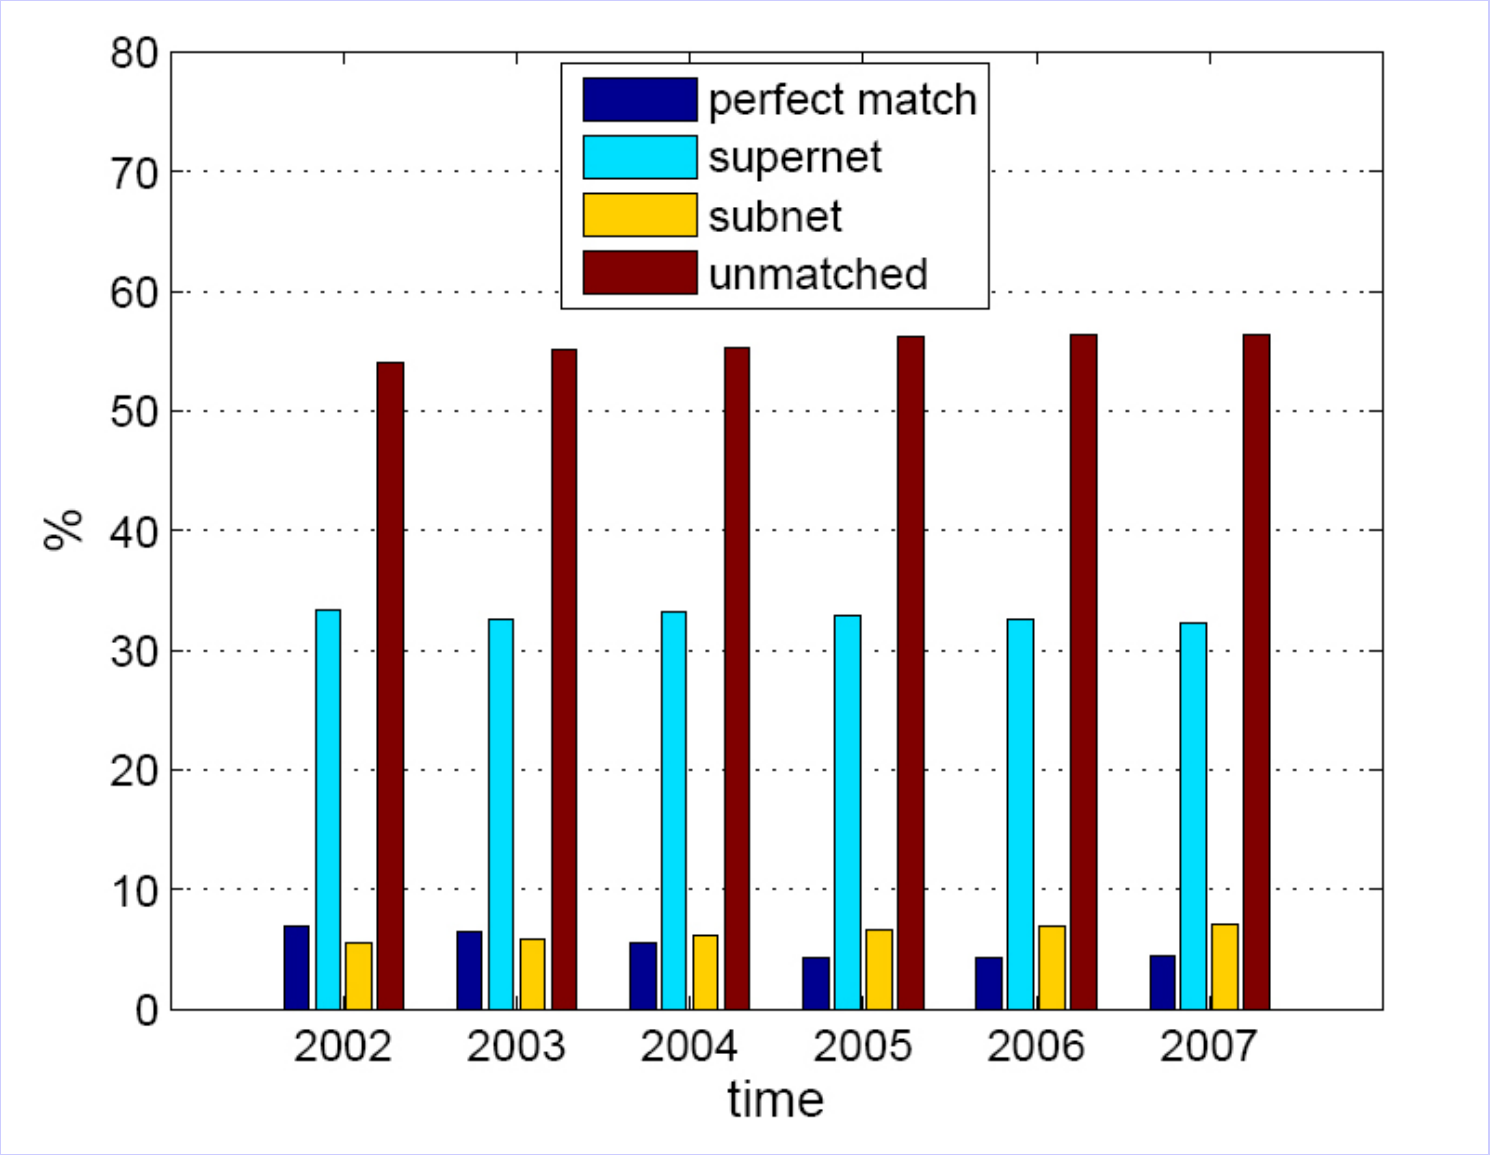

Unannounced /24 Equivalents

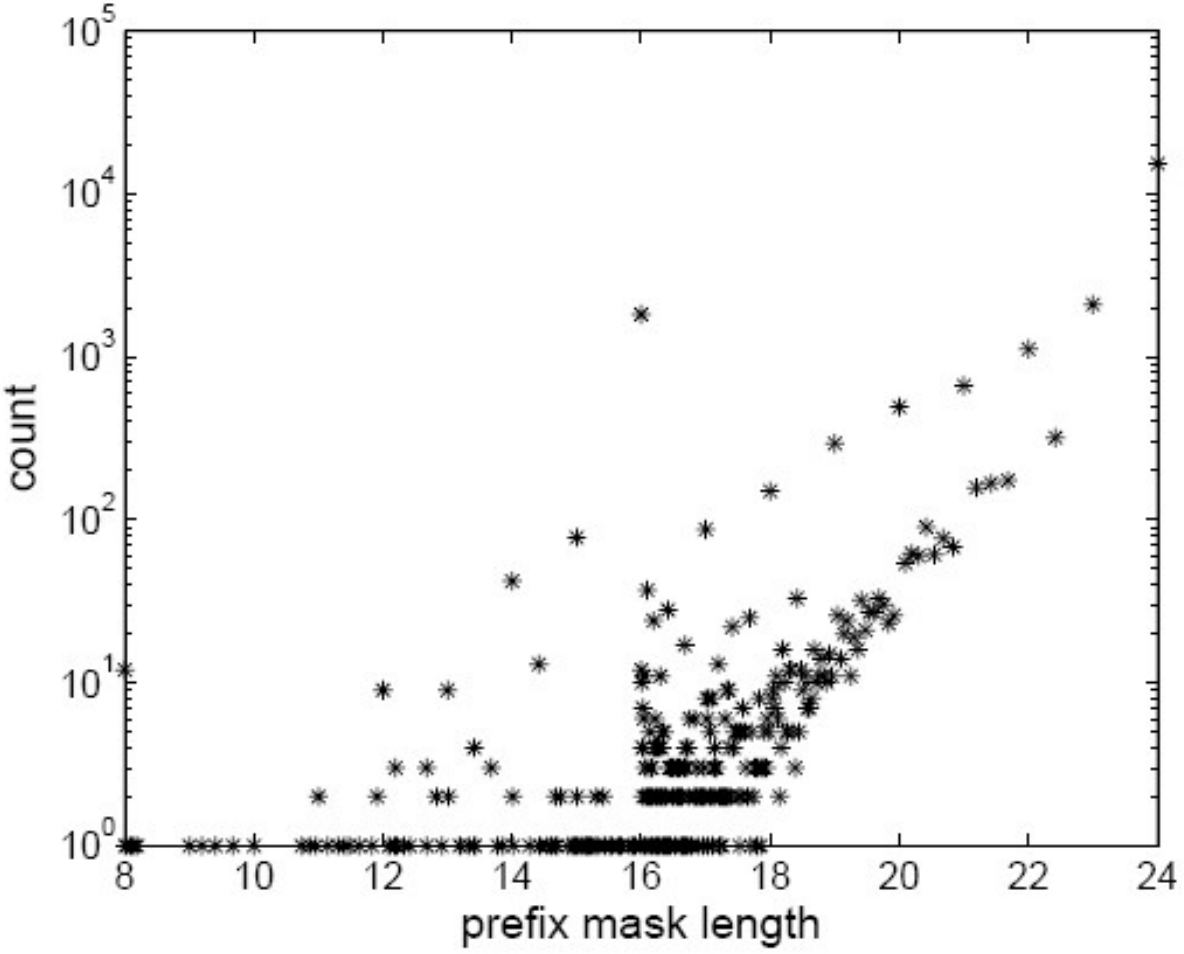

Figure 11: Histograms of the unannounced IP block

That's Legacy Space

There is also a lot of underutilized RIR Space Post-Legacy

How to Put IPv4 Space to Best Use?

Best Use
is Supposed to be
What Markets Do

There Already is a

Black Market in

IPv4 Address Space

Would you Rather Have a Black Market or an Open Market?

I personally prefer a possibly flawed open market to amateur over-regulators

The RPKI certificates are how we make the Market Transparent and Safe

Routing Table Growth

- Same in IPv6 as IPv4
- Proportional to multi-homers
- And traffic engineers
- All the way to the enterprise edge
- 2m+ routes soon, more later
- Multi-vendor is mandatory, I do not want to be owned ever again

Once Again -

Enterprise Scale Routers Must Handle 2m+ Routes Very Soon and More Coming

Routing Improvements

- Where was Clarence 15 Years Ago?
- We have been algorithmically lazy
- We never engaged the maths folk
- Routing is considered uninteresting in today's CS programs
- We have more economists and lawyers in the game than mathematicians

Where I do Not Want to Go

- Complexity
- More devices in my network
- Complexity
- Reliance on more protocols
- Complexity
- Centralization (GENI et alia)
- And did I mention Complexity?

Complexity is the Arch-Enemy of Scalability and Margins

Whose Margins?

"Screw you! I make billions of dollars from selling you complexity."

-- A friend at a vendor

End of my spiel!