IETF81 Secure IDR Rollup – TREX Workshop 2011

David Freedman, Claranet

claranet

Introduction to Secure IDR (SIDR)

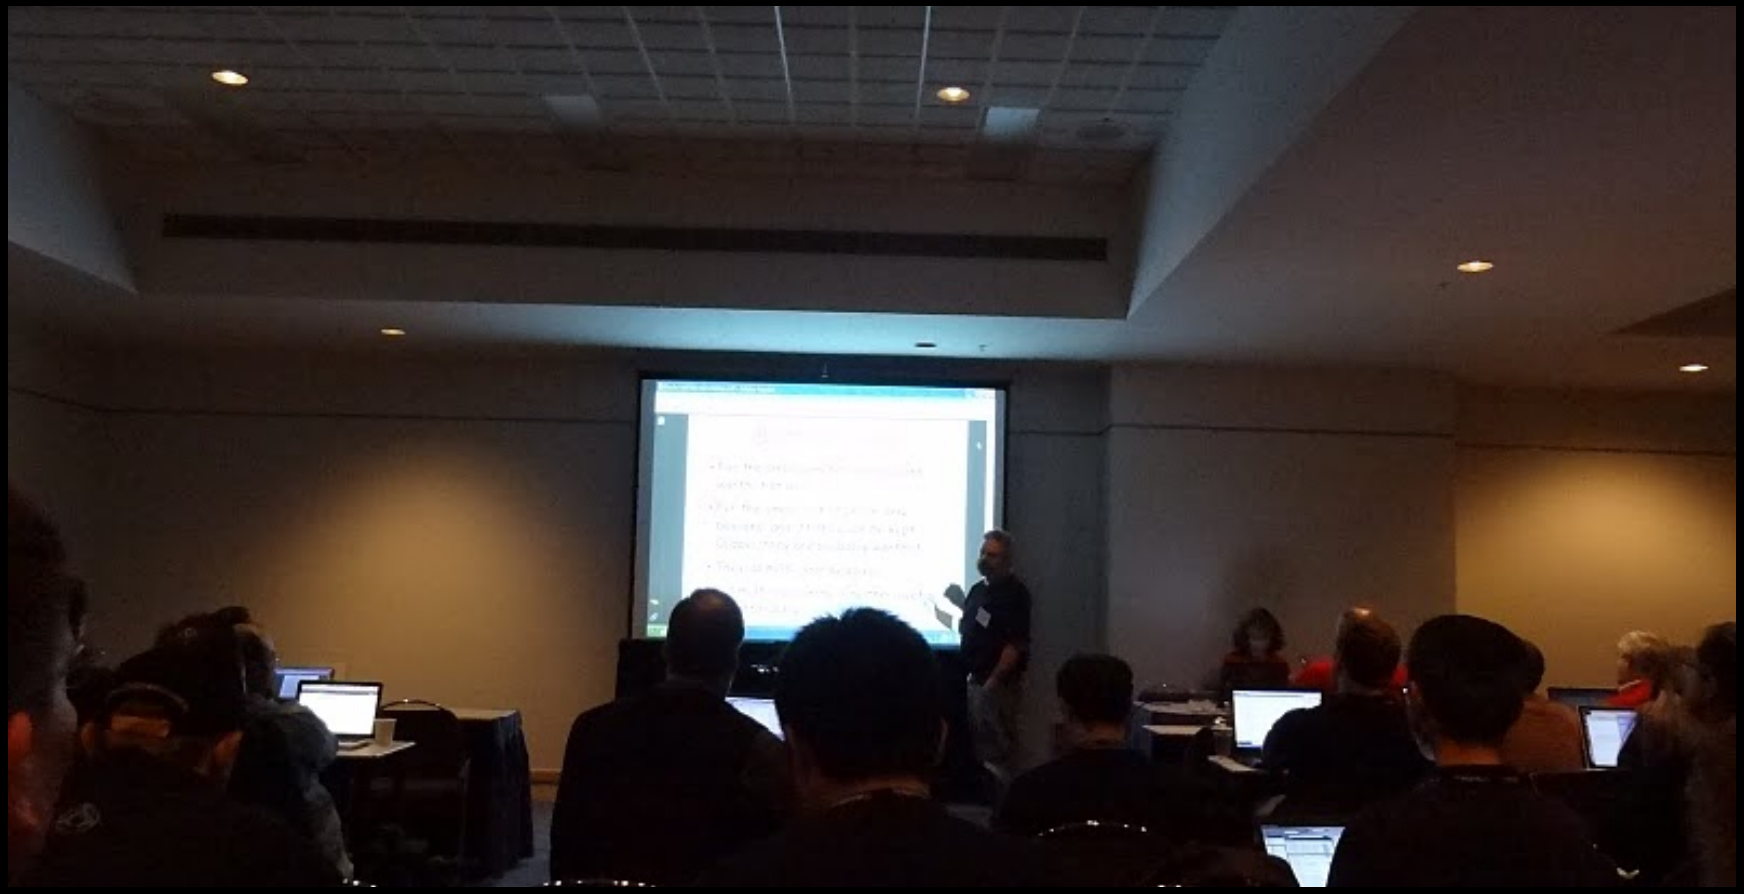

You are in a darkened room at the IETF.
You are surrounded by vendors.
A lone operator stands quietly in the corner.....

A reminder of concepts.. RPKI

The current situation

IANA are the numbering authority, coordinating allocation of numbering resources to Regional Internet Registries

RIPE NCC are the Regional Internet Registry, providing numbering allocations and database services to Local Internet Registries.

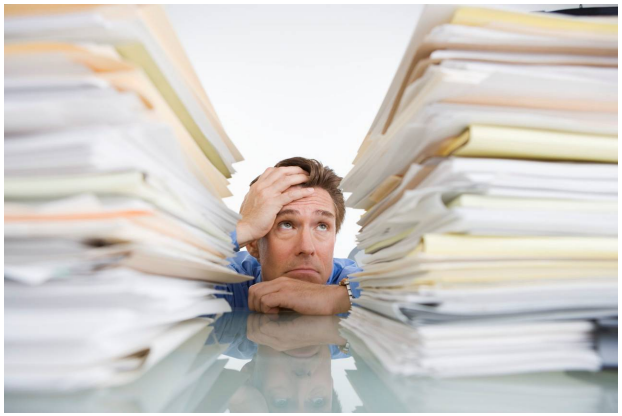

Local Internet Registries are usually network operators, they consume allocated resources and use the database to build customer (and sometimes peer) filters.

Simple filtering from the database

- **Q. Which routes should I accept from my customer AS8272?**
- **A. Only 193.221.118.0/24**

```
$ whois -r -Troute -i origin AS8272
% This is the RIPE Database query service.
% The objects are in RPSL format.
%
% The RIPE Database is subject to Terms and Conditions.
% See http://www.ripe.net/db/support/db-terms-conditions.pdf

% Note: this output has been filtered.
%       To receive output for a database update, use the "-B" flag.

% Information related to '193.221.118.0/24AS8272'

route:          193.221.118.0/24
descr:          CONVERGENCE
origin:         AS8272
mnt-by:         CONVERGENCE-MNT
org:            ORG-CNV1-RIPE
source:         RIPE # Filtered
```

claranet

I have it too easy….

- **My customer has no downstreams**
  - This means no macro evaluation
- **I'm only filtering my customer**
  - Somebody I'm able to (somewhat) control the quality of the data held in the database before I build my filters (try doing this to peers, not much fun).
- **I'm an LIR of the RIPE NCC**
  - I actually have a *great* database from which to build such filters, other service regions are not as stringent with regards to this and as such, independent databases (such as Merit's RADB and ALTDB) exist to bridge this gap (but being decoupled from the RIR, they can suffer from authority problems) – Nobody else has this great IP<->ASN binding.

claranet

Why filter anyway?

We should just trust each other, right?

- **Route Hijacking, two modern examples cited:**
  - Pakistan Telecom v Youtube [2008]
    - Lack of filters by Upstreams (PCCW) allowed a leaked /24 for YouTube from Pakistan Telecom to propagate for over two hours, leading to a battle of Longest Prefix Match until the rogue announcement could be pulled.
  - Kapela / Pilosov attack [2008]
    - Hijack a route and use a crafted AS_PATH to blind selective transit networks on the return to the hijacked path (creating a MITM vector).
- **Need to verify the peer against the route and the AS_PATH somehow.**
  - This means having good, authoritative data.
  - The system has to be easy to use and widely adopted
  - The system has to scale.

claranet

Extended X.509 Certification Hierarchy

draft-ietf-sidr-arch-13

- **RFC3779 extensions in Cert carry IP/ASN information.**
- **RFC5280 “Subject Information Access” extension defines publishing URI.**
- **CA certificates issued downstream from the IANA through RIR and LIR participants.**
- **Last CA in the chain issues an EE Certificate.**
- **EE Certificate used to produce a RoA (just a signed blob).**

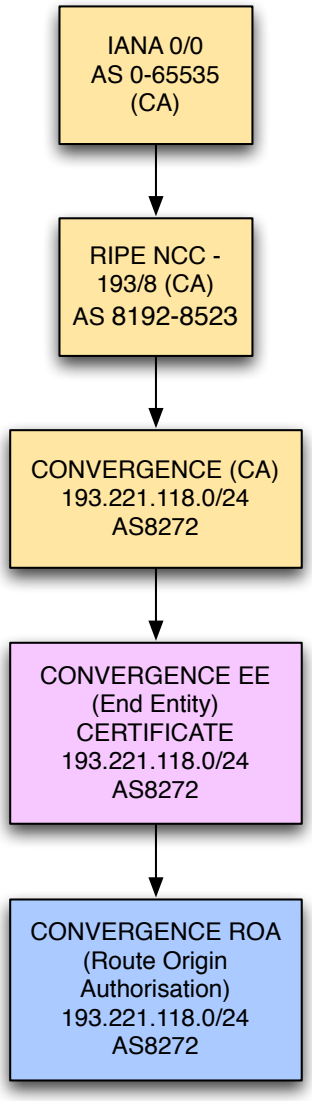

claranet

And we validate this how?

- **draft-ietf-sidr-rpki-rtr-16 explains:**
  - Periodically download entire copy of the RPKI with *rsync* (RFC5781)
  - Offline validation into a *trusted cache*
  - Router peers with cache using rpki-rtr protocol and receives trusted origin data.
- **Sounds futuristic?**
  - Cisco IOS and IOS-XR have PoC code running now
  - Compute load demonstrated to be ~10µsec per update*
- **How much of this do I have to run myself?**
  - Options for RIR Hosted CA, RoA generator and publisher
  - Perhaps your upstream can provide a cache if you are small?

claranet

* Source: http://ripe60.ripe.net/presentations/Bush-The_RPKI_Origin_Validation.pdf

What do we do with this?

draft-ietf-sidr-origin-ops-10

- **Local operator decision**
  - Prefer valid over unvalidated over invalid for instance*

```
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
...
```

claranet

* Example courtesy of : http://ripe60.ripe.net/presentations/Bush-The_RPKI_Origin_Validation.pdf

BGPSec

- **Without AS_PATH filtering everywhere, prefixes and their origin ASNs can be hijacked.**
  - RoA still valid as Origin ASN hasn't changed.
- **Need path validation in the BGP as well**
  - BGPSec being developed by SIDR, *draft-ietf-sidr-bgpsec-overview-00*
- **BGPSec introduces new type of certificate**
  - New "Router" certificate, public key published for an AS, private key held by routers within the AS, *draft-ietf-sidr-bgpsec-overview-00*.
- **Routers use private keys to sign AS_PATH elements**
  - New attribute **BGPSEC_Path_Signatures***

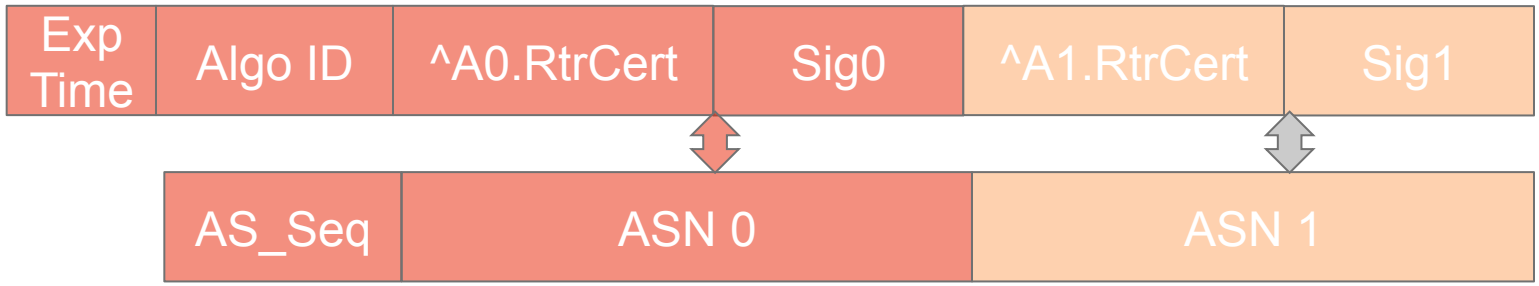

claranet

* Example courtesy of : http://tools.ietf.org/agenda/81/slides/sidr-10.pdf

But what about?

Some as-yet undrafted ideas

- **AS_PATH prepending?**
  - Suggested to add pCNT counter of prepends, decouple policy from path*

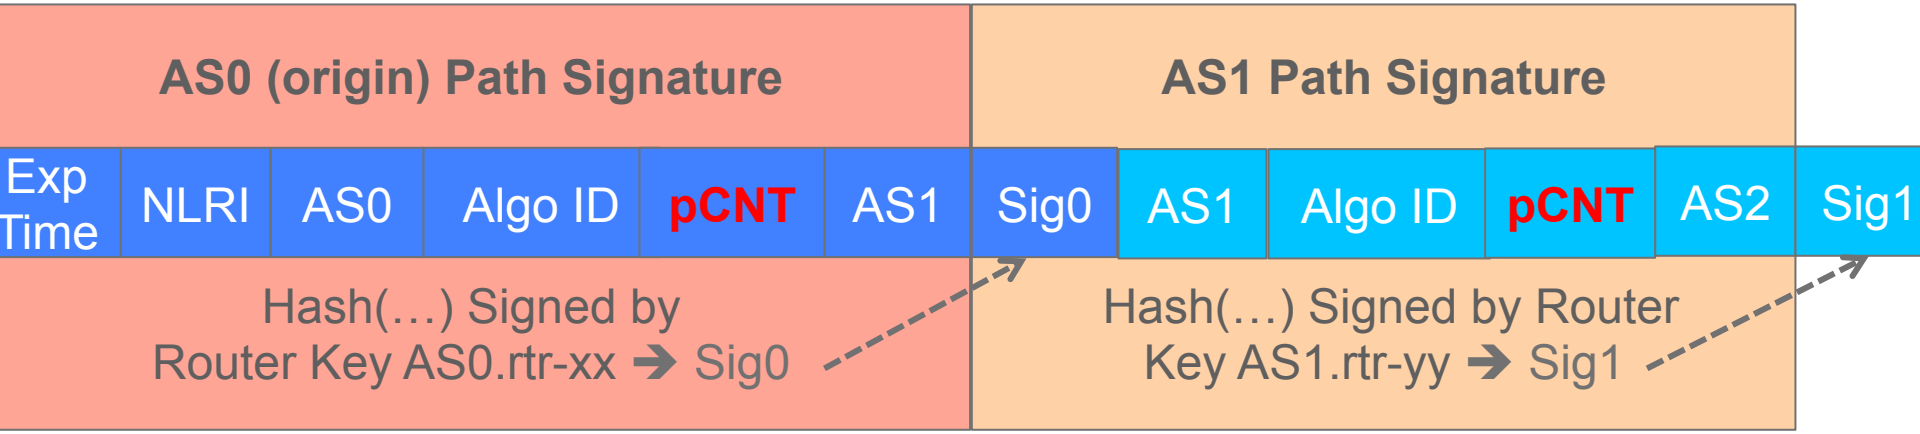

- **Transparent Route Servers (i.e MLP)?**
  - This makes people uneasy as it proposes a bypass mechanism
  - Current thinking to use pCNT=0 as “Special Case” for xref with AS_PATH, meaning “ignore this AS in AS_PATH when bestpathing” , non-BGPSec speakers simply strip the pCNT=0 AS from the AS_PATH completely.

claranet

* Example courtesy of : http://tools.ietf.org/agenda/81/slides/sidr-10.pdf

And what about replay attacks?

draft-ietf-sidr-bgpsec-ops-00

- **Replay attacks are of concern**
  - Provider annoyed at customer who switches
  - Prefix “Stuck in Router” being replayed over and over again.
  - All at Human timescales.
- **Freshness is critical**
  - We could do smaller signing lifetimes
  - Origin could re-announce before this expires (i.e *Beaconing)*
  - Suggested to be days, but can be hours for critical infrastructure)
  - Timing is of course jittered.
- **Beaconing is required to prevent replay attacks**
  - But only origin beaconing is tolerable to internet infrastructure
  - Multi-Beaconing neither useful, nor affordable.

claranet

More SIDR discussed topics

- **draft-ietf-sidr-algorithm-agility-03**
  - How we transition to newer and more secure algorithms without disruption.
  - We already have normal Key Rollovers (i.e old key expires or needs to be revoked)
  - This draft specifies a rollover for reason of a new algorithm.
  - Five Phases (0-4) with six steps:
    - Phase 0: Old algorithm only used.
    - Phase 1: Parent CAs can issue using new algorithm.
    - Phase 2: RPs MAY be able to validate new algorithm.
    - Phase 3: RPs MUST be able to validate new and old algorithms.
    - Phase 4: RPs MAY be able to validate using old algorithm.
    - Phase 0: Return to Phase 0, only now old algorithm is invalid.

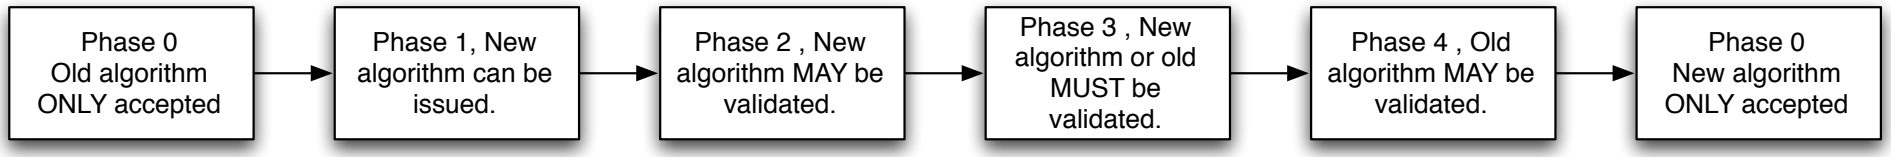

claranet

More SIDR discussed topics

- **draft-ietf-sidr-publication-01**
  - An XMLoHTTP based publication protocol, designed for children publishing back to parents (think outsourced SIA model).
- **draft-ietf-sidr-usecases-02**
  - A brilliant reference document showing various RPKI use cases and their outcomes. I would recommend a read of this, here is an except of the index:

claranet

And finally some fun

- **draft-ietf-sidr-ghostbusters-09**
  - With all these certificates floating around, when there's a problem, who you gonna call?
  - This stuff may not be engineered by neteng/netops folk, perhaps systems/security people, not traditionally ops contacts from external networks!
  - We sign RoAs with our EE certificates, why limit this to just routing data?
  - Use the EE to sign a vCard, for somebody to contact regarding a RoA or certification chain.
  - Not meant to replace WHOIS

```
BEGIN:vCard
 VERSION:4.0
 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
```

Any questions?

claranet