Securing Internet Routing

Tashi Phuntsho (tashi@apnic.net)
Senior Network Janitor/Technical Trainer

Why should we bother?

- As a Manager
  - I don’t want to be front page news of a IT paper, or an actual newspaper for routing errors

Headlines

BGPmon.net
@bgpmon

Following

looking into BGP leak incident involving @google prefixes, AS37282 out of Niger and China Telecom.

3:40 AM - 13 Nov 2018

54 Retweets **48** Likes

MainOne
@Mainoneservice

Follow

Replying to @bgpmon @Google

We have investigated the advertisement of @Google prefixes through one of our upstream partners. This was an error during a planned network upgrade due to a misconfiguration on our BGP filters. The error was corrected within 74mins & processes put in place to avoid reoccurrence

5:29 PM - 13 Nov 2018

38 Retweets **50** Likes

ThousandEyes
@thousandeyes

Following

BREAKING: Potential hijack underway. ThousandEyes detected intermittent availability issues to Google services from some locations. Traffic to certain Google destinations appears to be routed through an ISP in Russia & black-holed at a China Telecom gateway router.

2:57 AM - 13 Nov 2018

609 Retweets **525** Likes

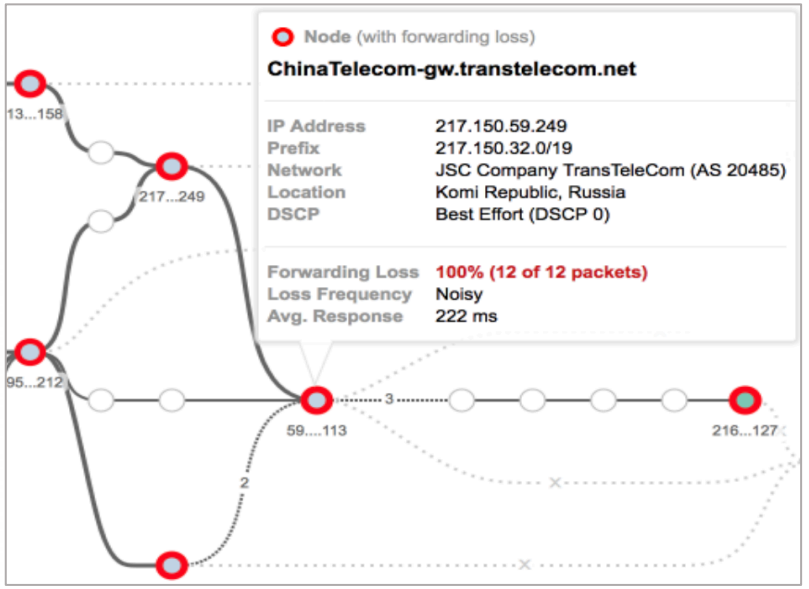

https://blog.thousandeyes.com/internet-vulnerability-takes-down-google/

Headlines

ars TECHNICA BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE

BORDER GATEWAY PROTOCOL ATTACK —

Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency

Almost 1,300 addresses for Amazon Route 53 rerouted for two hours.

DAN GOODIN - 4/25/2018, 1:30 AM

InternetIntelligence
@InternetIntel
Follow

BGP hijack this morning affected Amazon DNS. eNet (AS10297) of Columbus, OH announced the following more-specifics of Amazon routes from 11:05 to 13:03 UTC today:
205.251.192.0/24
205.251.193.0/24
205.251.195.0/24
205.251.197.0/24
205.251.199.0/24

7:52 AM - 24 Apr 2018

Kevin Beaumont @GossiTheDog · Apr 24, 2018
MyEtherWallet subject to a DNS hijack. DNS was redirected via AWS DNS to a server in Russia, Ether stolen. Server is https only so users clicked through certificate errors.

Doug Madory
@DougMadory

Maybe related to this: twitter.com/InternetIntel/…

> **InternetIntelligence** @InternetIntel
> BGP hijack this morning affected Amazon DNS. eNet (AS10297) of Columbus, OH announced the following more-specifics of Amazon routes from 11:05 to 13:03 UTC today:
> 205.251.192.0/24
> 205.251.193.0/24
> 205.251.195.0/24
> 205.251.197.0/24
> 205.251.199.0/24

♡ 2 9:23 PM - Apr 24, 2018

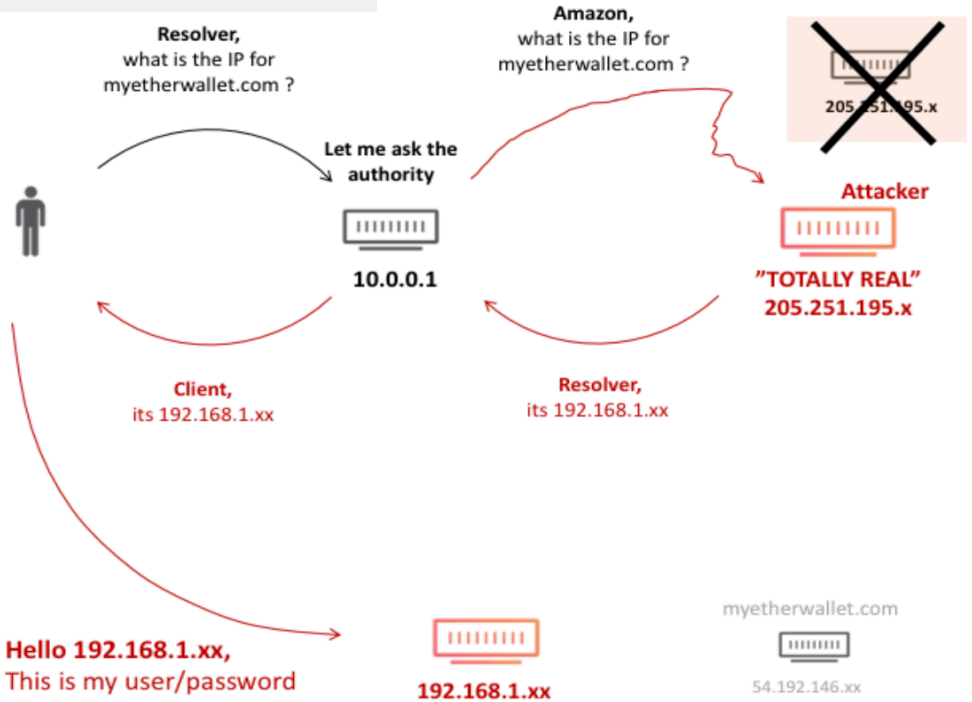

https://blog.cloudflare.com/bgp-leaks-and-crypto-currencies

Headlines

Large BGP Leak by Google Disrupts Internet in Japan

Research // Aug 28, 2017 // Doug Madory

```
trace from Tokyo, Japan to Inuyama, Japan at 03:28 Aug 25, 2017
1  *
2  183.177.32.145   Equinix Asia Pacific              Tokyo       Japan           0.249
3  210.130.154.37   IIJ IPv4 BLOCK ( AS2497 )        Tokyo       Japan           0.618
4  58.138.102.109   tky001bb11.IIJ.Net               Tokyo       Japan           0.877
5  58.138.88.86     sjc002bb12.IIJ.Net               San Jose    United States   97.797
6  152.179.48.117   TenGigE0-3-0-8.GW6.SJC7.ALTER.NET San Jose   United States   97.869
7  *
8  152.179.105.110  google-gw.customer.alter.net     Chicago     United States   337.19
9  108.170.243.197  Google Inc.                      Chicago     United States   246.325
10 *
11 209.85.241.43    Google Inc.                                  United States   256.188
12 72.14.238.38     Google Inc.                      Vancouver   Canada          247.849
13 209.85.245.110   Google Inc.                      Vancouver   Canada          249.291
14 *
15 108.170.242.138  Google Inc.                      Tokyo       Japan           246.267
16 211.0.193.21     OCN (AS4713) CIDR BLOCK 21       Tokyo       Japan           246.351
17 122.1.245.65     OCN (AS4713) CIDR BLOCK 81       Tokyo       Japan           246.426
18 *
19 153.149.218.10   OCN (AS4713) CIDR BLOCK 93       Ōsaka-shi   Japan           256.027
20 125.170.96.38    OCN (AS4713) CIDR BLOCK 77                   Japan           255.683
21 *
22 60.37.32.250     OCN (AS4713) CIDR BLOCK 70                   Japan           254.989
23 118.23.141.202   OCN (AS4713) CIDR BLOCK 86                   Japan           254.526
24 *
25 211.11.83.160    OCN (AS4713) CIDR BLOCK 23       Inuyama     Japan           256.212
```

After leak (JP->JP)

```
trace from London, England to Nürnberg, Germany at 03:30 Aug 25, 2017
1  *
2  195.66.248.190   fe0-2.tr2.linx.net               London      United Kingdom  0.327
3  195.66.249.10    ge0-2-502.tr5.linx.net           London      United Kingdom  0.441
4  195.66.249.13    ge0-2-501.tr4.linx.net           London      United Kingdom  0.477
5  195.66.248.10    uunet-uk-transit.thn.linx.net    London      United Kingdom  0.507
6  158.43.193.245   POS0-0.CR2.LND6.ALTER.NET        London      United Kingdom  0.497
7  140.222.239.41   0.xe-0-0-0.IL1.NYC50.ALTER.NET   New York    United States   108.146
8  146.188.4.197    xe-0-0-1.IL1.NYC41.ALTER.NET     New York    United States   75.719
9  140.222.234.221  0.et-10-1-0.GW7.CHI13.ALTER.NET  Chicago     United States   94.793
10 152.179.105.110  google-gw.customer.alter.net     Chicago     United States   224.352
11 *
12 216.239.40.189   Google Inc.                      Northlake   United States   202.193
13 216.239.58.255   Google Inc.                                                  203.995
14 216.239.58.12    Google Inc.                                                  207.026
15 209.85.253.184   Google Inc.                      Luxembourg  Luxembourg      212.944
16 209.85.252.215   Google Inc.                                                  213.112
17 108.170.252.71   Google Inc.                                                  213.265
18 72.14.222.53     Google Inc.                                  Germany         212.061
19 188.111.165.169  Vodafone GmbH                                Germany         227.077
20 178.7.138.112    Vodafone D2 GmbH                 Nürnberg    Germany         234.226
```

After leak (EU->EU)

https://dyn.com/blog/large-bgp-leak-by-google-disrupts-internet-in-japan/

Headlines

YouTube blames Pakistan network for 2-hour outage

Company appears to confirm reports that Pakistan Telecom was responsible for routing traffic according to erroneous Internet Protocols.

BY GREG SANDOVAL | FEBRUARY 24, 2008 10:15 PM PST

Why should we bother?

- As a Engineer
  - I don’t want to be told at 3AM my routing is broken
  - Or while on a holiday

Why do we keep seeing these?

- Because NO ONE is in charge?
  - No single authority model for the Internet
  - No reference point for what's right in routing

Why do we keep seeing these?

- Routing works by RUMOUR
  - Tell what you know to your neighbors, and Learn what your neighbors know
  - Assume everyone is correct (and *honest*)
    - Is the originating network the rightful owner?

Why do we keep seeing these?

- Routing is VARIABLE
  - The view of the network depends on where you are
    - Different routing outcomes at different locations
  - ~ no reference view to compare the local view ☹

Why do we keep seeing these?

- Routing works in REVERSE
  - Outbound advertisement affects inbound traffic
  - Inbound (*Accepted*) advertisement influence outbound traffic

Why do we keep seeing these?

- And as always, there is no E-bit
  - A bad routing update does not identify itself as BAD
- So tools/techniques try to identify GOOD updates

Why should we worry?

- Because it’s just so easy to do bad in routing!

By Source (WP:NFCC#4), Fair use,
https://en.wikipedia.org/w/index.php?curid=42515224

Current practice

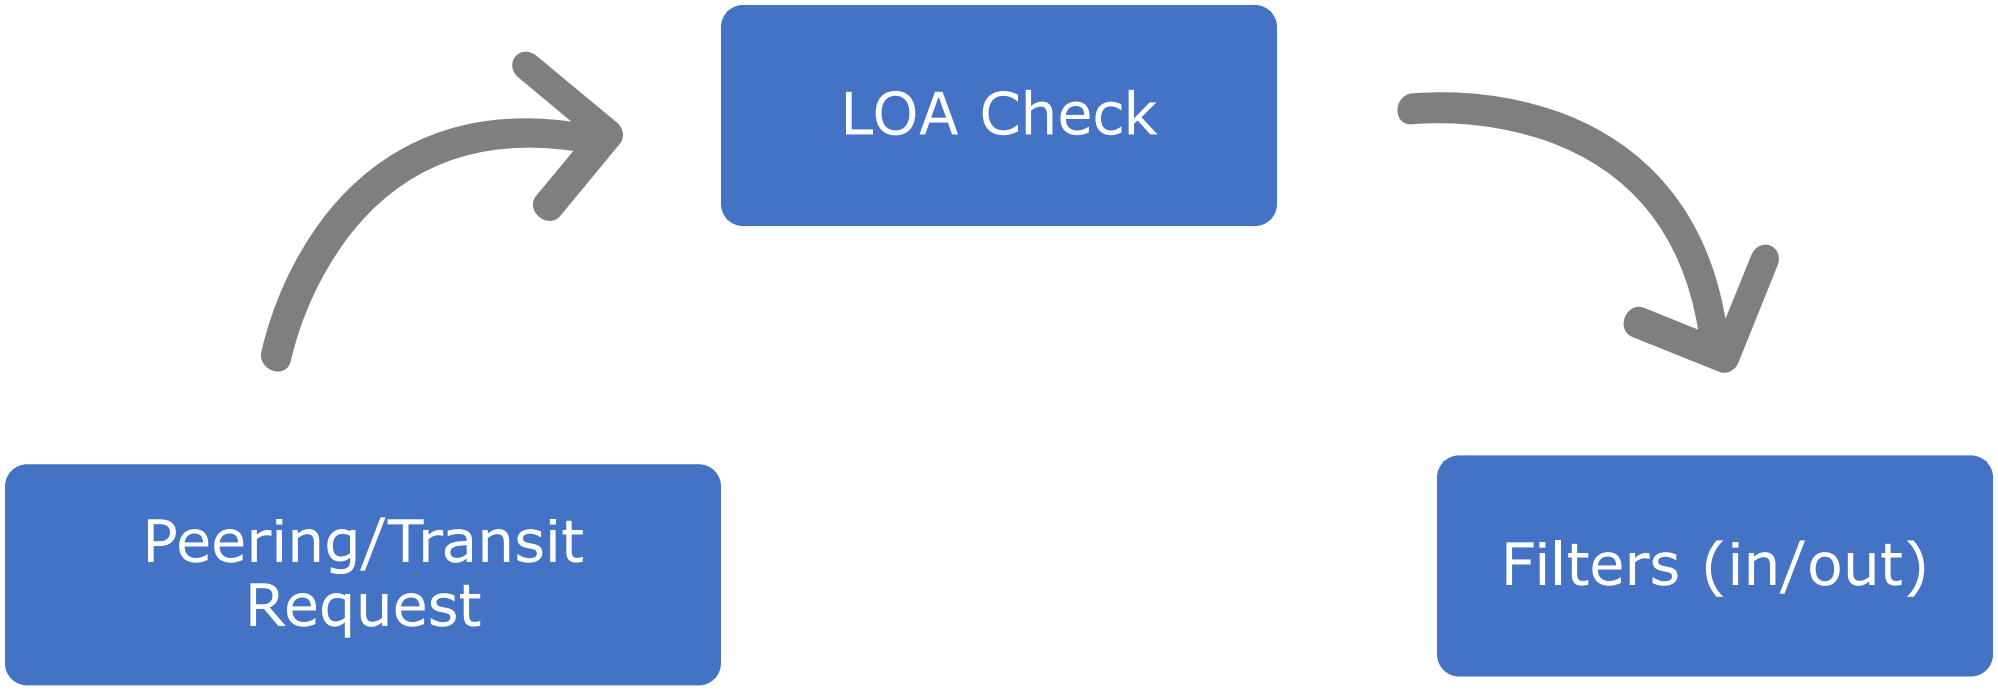

Tools & Techniques

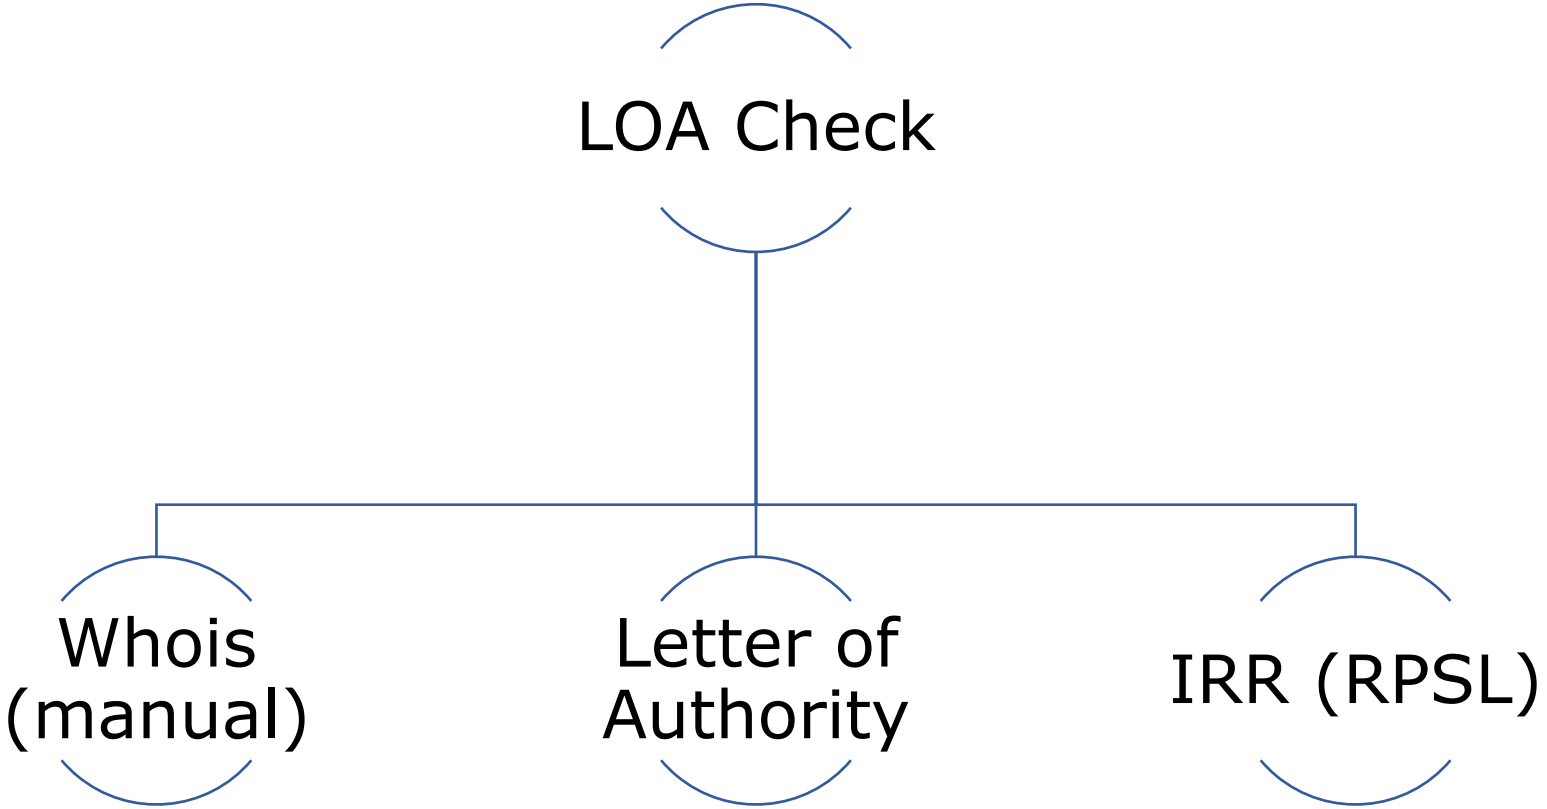

Tools & Techniques

- Look up **whois**
  - verify holder of a resource

```
tashi@tashi ~> whois -h whois.apnic.net 202.125.96.0
% [whois.apnic.net]
% Whois data copyright terms    http://www.apnic.net/db/dbcopyright.html

% Information related to '202.125.96.0 - 202.125.96.255'

% Abuse contact for '202.125.96.0 - 202.125.96.255' is 'training@apnic.net'

inetnum:        202.125.96.0 - 202.125.96.255
netname:        APNICTRAINING-AP
descr:          Prefix for APNICTRAINING LAB DC
country:        AU
admin-c:        AT480-AP
tech-c:         AT480-AP
status:         ALLOCATED NON-PORTABLE
mnt-by:         MAINT-AU-APNICTRAINING
mnt-irt:        IRT-APNICTRAINING-AU
last-modified:  2016-06-17T00:17:28Z
source:         APNIC

irt:            IRT-APNICTRAINING-AU
address:        6 Cordelia Street
address:        South Brisbane
address:        QLD 4101
e-mail:         training@apnic.net
abuse-mailbox:  training@apnic.net
admin-c:        AT480-AP
tech-c:         AT480-AP
auth:           # Filtered
mnt-by:         MAINT-AU-APNICTRAINING
last-modified:  2013-10-31T11:01:10Z
source:         APNIC

role:           APNIC Training
address:        6 Cordelia Street
address:        South Brisbane
address:        QLD 4101
country:        AU
phone:          +61 7 3858 3100
fax-no:         +61 7 3858 3199
e-mail:         training@apnic.net
admin-c:        JW3997-AP
tech-c:         JW3997-AP
nic-hdl:        AT480-AP
mnt-by:         MAINT-AU-APNICTRAINING
last-modified:  2017-08-22T04:59:14Z
source:         APNIC

% Information related to '202.125.96.0/24AS131107'

route:          202.125.96.0/24
descr:          Prefix for APNICTRAINING LAB DC
origin:         AS131107
mnt-by:         MAINT-AU-APNICTRAINING
country:        AU
last-modified:  2016-06-16T23:23:00Z
source:         APNIC
```

Tools & Techniques

- Ask for a **Letter of Authority**
  - Absolve from any liabilities

APNIC

Asia Pacific Network Information Centre
APNIC Pty Ltd
ABN: 42 081 528 010

6 Cordelia Street
PO Box 3646
South Brisbane
QLD 4101 AUSTRALIA

URL www.apnic.net
Enquiries helpdesk@apnic.net
Accounts billing@apnic.net
Phone +61 7 3858 3100
Fax + 61 7 3858 3199

31/03/2018

Letter of Authorization

To whom it may concern,

APNIC Training (AS45192) runs a lab network to reproduce technical problems faced by members to help troubleshoot specific issues.

This letter serves as an authorization for APNIC Infra (AS4608) to advertise the following address blocks:

202.125.96.0/24

As a representative of APNIC Training team, that is the owner of the subnet and ASN, I hereby declare that I am authorized to sign this LOA.

Tashi Phuntsho
Training Delivery Manager

Email: tashi@apnic.net
Phone: +61 7 3858 3114

Tools & Techniques

- Look up/ask to enter details in internet routing registries (IRR)
  - describes route origination and inter-AS routing policies

```
tashi@tashi ~> whois -h whois.radb.net 61.45.248.0/24
route:          61.45.248.0/24
descr:          APNICTRAINING-DC
origin:         AS135533
mnt-by:         MAINT-AS4826
changed:        noc@vocus.com.au 20160702
source:         RADB

route:          61.45.248.0/24
descr:          Prefix for APNICTRAINING LAB - AS135533
origin:         AS135533
mnt-by:         MAINT-AU-APNICTRAININGLAB
country:        AU
last-modified:  2017-10-19T01:36:37Z
source:         APNIC
```

```
tashi@tashi ~> whois -h whois.radb.net AS17660
aut-num:        AS17660
as-name:        BT-Bhutan
descr:          Divinetworks for BT
admin-c:        DUMY-RIPE
tech-c:         DUMY-RIPE
status:         OTHER
mnt-by:         YP67641-MNT
mnt-by:         ES6436-RIPE
created:        2012-11-29T10:31:33Z
last-modified:  2018-09-04T15:26:24Z
source:         RIPE-NONAUTH
remarks:        ****************************
remarks:        * THIS OBJECT IS MODIFIED
remarks:        * Please note that all data that is generally regarded as personal
remarks:        * data has been removed from this object.
remarks:        * To view the original object, please query the RIPE Database at:
remarks:        * http://www.ripe.net/whois
remarks:        ****************************

aut-num:        AS17660
as-name:        DRUKNET-AS
descr:          DrukNet ISP
descr:          Bhutan Telecom
descr:          Thimphu
country:        BT
org:            ORG-BTL2-AP
import:         from AS6461   action pref=100;     accept ANY
export:         to AS6461     announce AS-DRUKNET-TRANSIT
import:         from AS2914   action pref=150;     accept ANY
export:         to AS2914     announce AS-DRUKNET-TRANSIT
import:         from AS6453   action pref=100;     accept ANY
export:         to AS6453     announce AS-DRUKNET-TRANSIT
```

Tools & Techniques

- IRR
  - *Helps auto generate network (prefix/as-path) filters using RPSL tools*
    - Filter out route advertisements not described in the registry

```
tashi@tashi ~> bgpq3 -Al PEER-v4IN AS17660
no ip prefix-list PEER-v4IN
ip prefix-list PEER-v4IN permit 45.64.248.0/22
ip prefix-list PEER-v4IN permit 103.7.252.0/22
ip prefix-list PEER-v4IN permit 103.7.254.0/23
ip prefix-list PEER-v4IN permit 103.245.240.0/22
ip prefix-list PEER-v4IN permit 103.245.242.0/23
ip prefix-list PEER-v4IN permit 119.2.96.0/19
ip prefix-list PEER-v4IN permit 119.2.96.0/20
ip prefix-list PEER-v4IN permit 202.89.24.0/21
ip prefix-list PEER-v4IN permit 202.144.128.0/19
ip prefix-list PEER-v4IN permit 202.144.128.0/23
ip prefix-list PEER-v4IN permit 202.144.144.0/20
ip prefix-list PEER-v4IN permit 202.144.148.0/22
tashi@tashi ~> bgpq3 -6Al PEER-v6IN AS17660
no ipv6 prefix-list PEER-v6IN
ipv6 prefix-list PEER-v6IN permit 2405:d000::/32
ipv6 prefix-list PEER-v6IN permit 2405:d000:7000::/36
```

```
tashi@tashi ~> bgpq3 -Abl PEER-v4IN AS17660
PEER-v4IN = [
    45.64.248.0/22,
    103.7.252.0/22,
    103.7.254.0/23,
    103.245.240.0/22,
    103.245.242.0/23,
    119.2.96.0/19,
    119.2.96.0/20,
    202.89.24.0/21,
    202.144.128.0/19,
    202.144.128.0/23,
    202.144.144.0/20,
    202.144.148.0/22
];
tashi@tashi ~> bgpq3 -6Abl PEER-v6IN AS17660
PEER-v6IN = [
    2405:d000::/32,
    2405:d000:7000::/36
];
```

```
tashi@tashi ~> bgpq3 -f 38195 -lSUPERLOOP-IN AS-SUPERLOOP
no ip as-path access-list SUPERLOOP-IN
ip as-path access-list SUPERLOOP-IN permit ^38195(_38195)*$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(681|4647|4749|4785)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(4846|4858|7477|7578)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(7585|7604|7628|7631)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(7699|9290|9297|9336)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(9499|9544|9549|10143)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(10145|11031|12041|15133)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(15967|17462|17498|17766)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(17829|17907|17991|18000)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(18110|18201|18292|23156)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(23456|23677|23858|23935)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(24007|24065|24093|24129)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(24231|24233|24238|24341)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(24459|27232|30215|30762)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(36351|37993|38263|38269)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(38451|38534|38549|38570)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(38595|38716|38719|38790)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(38809|38830|38858|42909)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(44239|45158|45267|45278)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(45570|45577|45638|45671)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(45844|46571|55411|55419)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(55455|55506|55575|55707)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(55752|55766|55803|55845)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(55884|55931|55954|56037)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(56098|56135|56178|56225)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(56271|56287|58422|58443)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(58511|58606|58634|58676)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(58712|58739|58750|58868)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(58914|59256|59330|59339)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(59356|60592|60758|63926)$
ip as-path access-list SUPERLOOP-IN permit ^38195(_[0-9]+)*_(63937|63956)$
```

Tools & Techniques

- Problem(s) with IRR
  - No single authority model
    - How do I know if a RR entry is genuine and correct?
    - How do I differentiate between a current and a lapsed entry?
  - Many RRs
    - If two RRs contain conflicting data, which one do I trust and use?
  - Incomplete data - Not all resources are registered in an IRR
    - If a route is not in a RR, is the route invalid or is the RR just missing data?
  - Scaling
    - How do I apply IRR filters to upstream(s)?

Back to basics – identifying GOOD

- Using digital signatures to convey the "*authority to use*"?
  - A private key to *sign* the *authority*, and
  - the public key to *validate* that *authority*

How about trust in this framework?

- Follows the resource allocation/delegation hierarchy

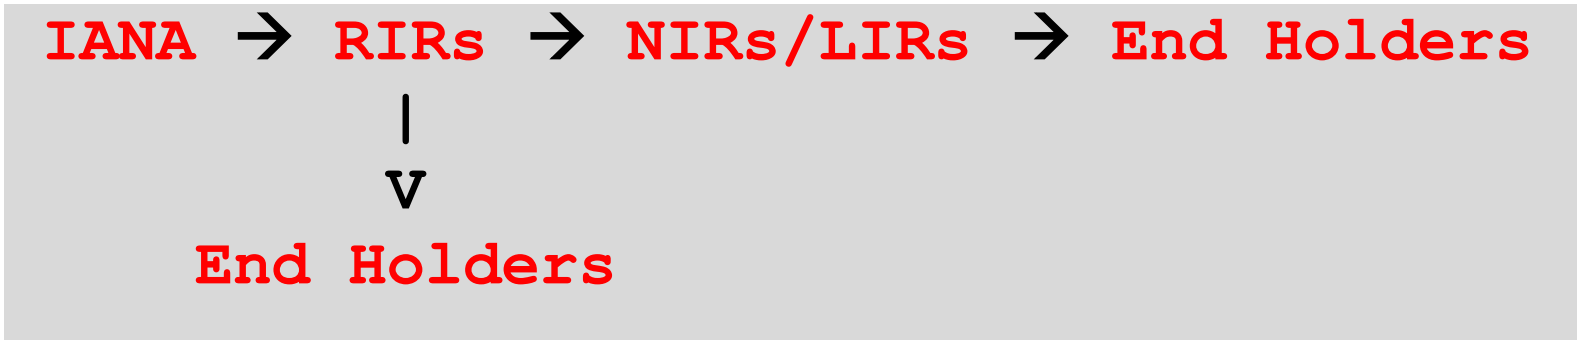

RPKI Chain of Trust

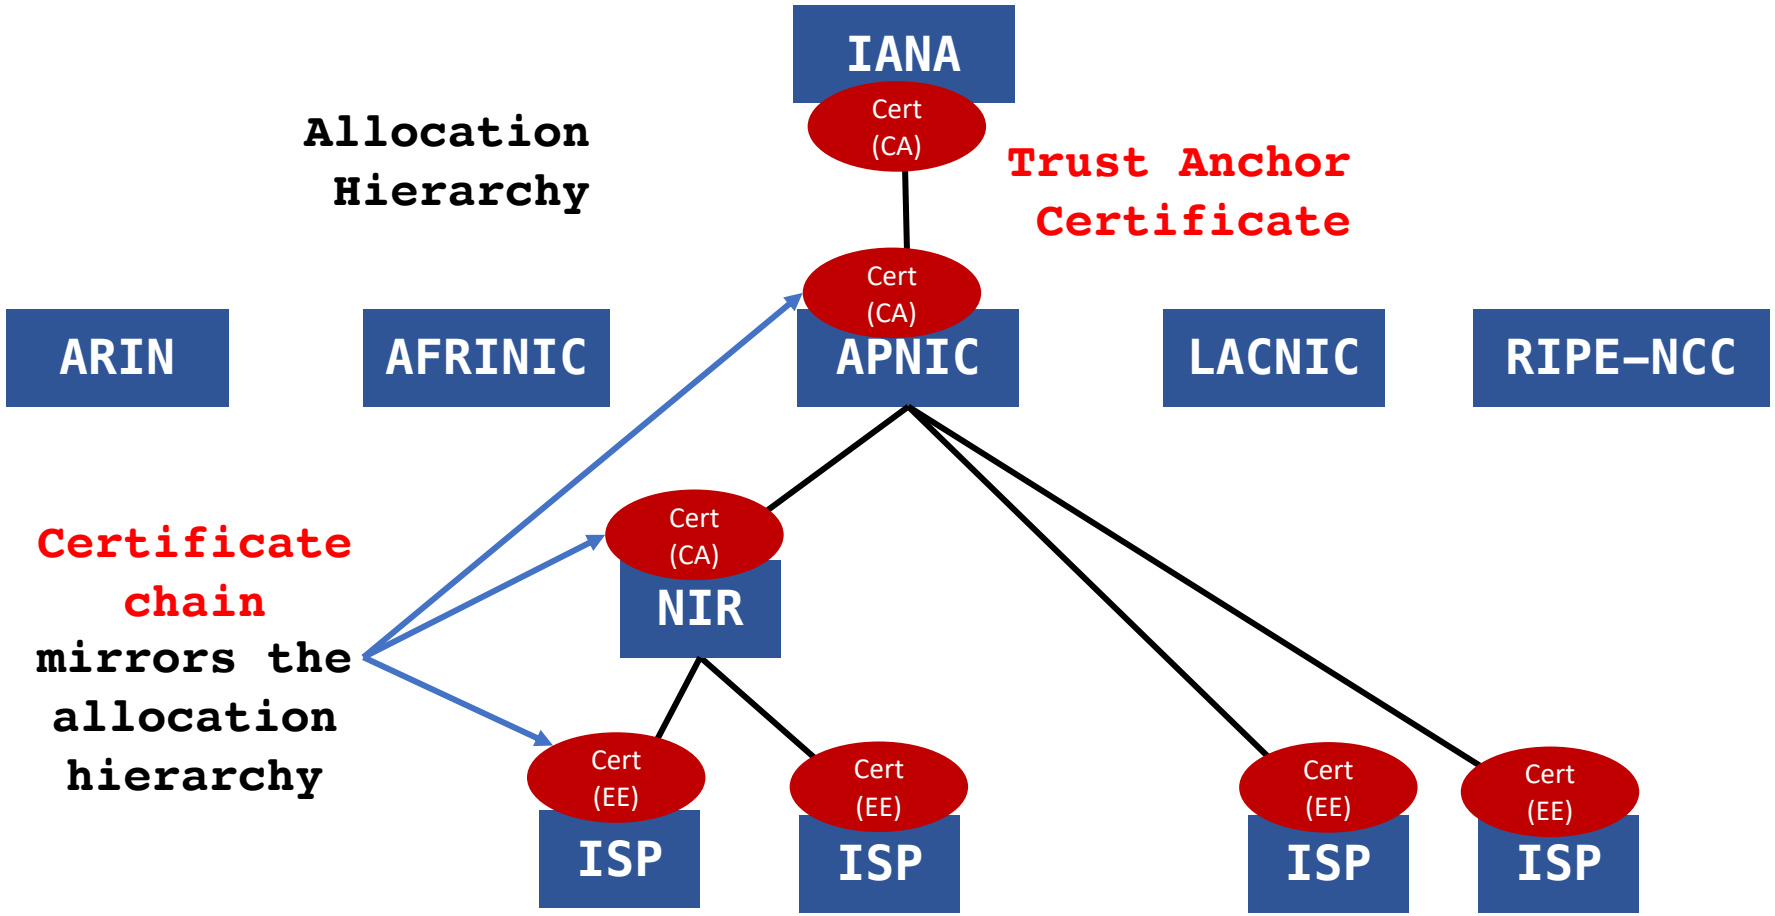

Image 4

Resource Certificates

- When an address holder A (*IRs) allocates resources (*IP address/ASN*) to B (end holders)
  - ***A** issues a resource certificate that binds the allocated address with **B's** public key, all signed by **A's** (CA) private key*
  - *proves the holder of the private key (**B**) is the legitimate holder of the resource!*

Route Origin Authorization (ROA)

- B can now sign *authorities* using its private key, which can be validated by any third party against the TA

- For routing, the address holder can *authorize* a network (ASN) to *originate* a route, and **sign** this permission with its private key (ROA)

Prefix	203.176.32.0/19
Max-length	/24
Origin ASN	**AS17821**

Route Origin Validation (ROV)

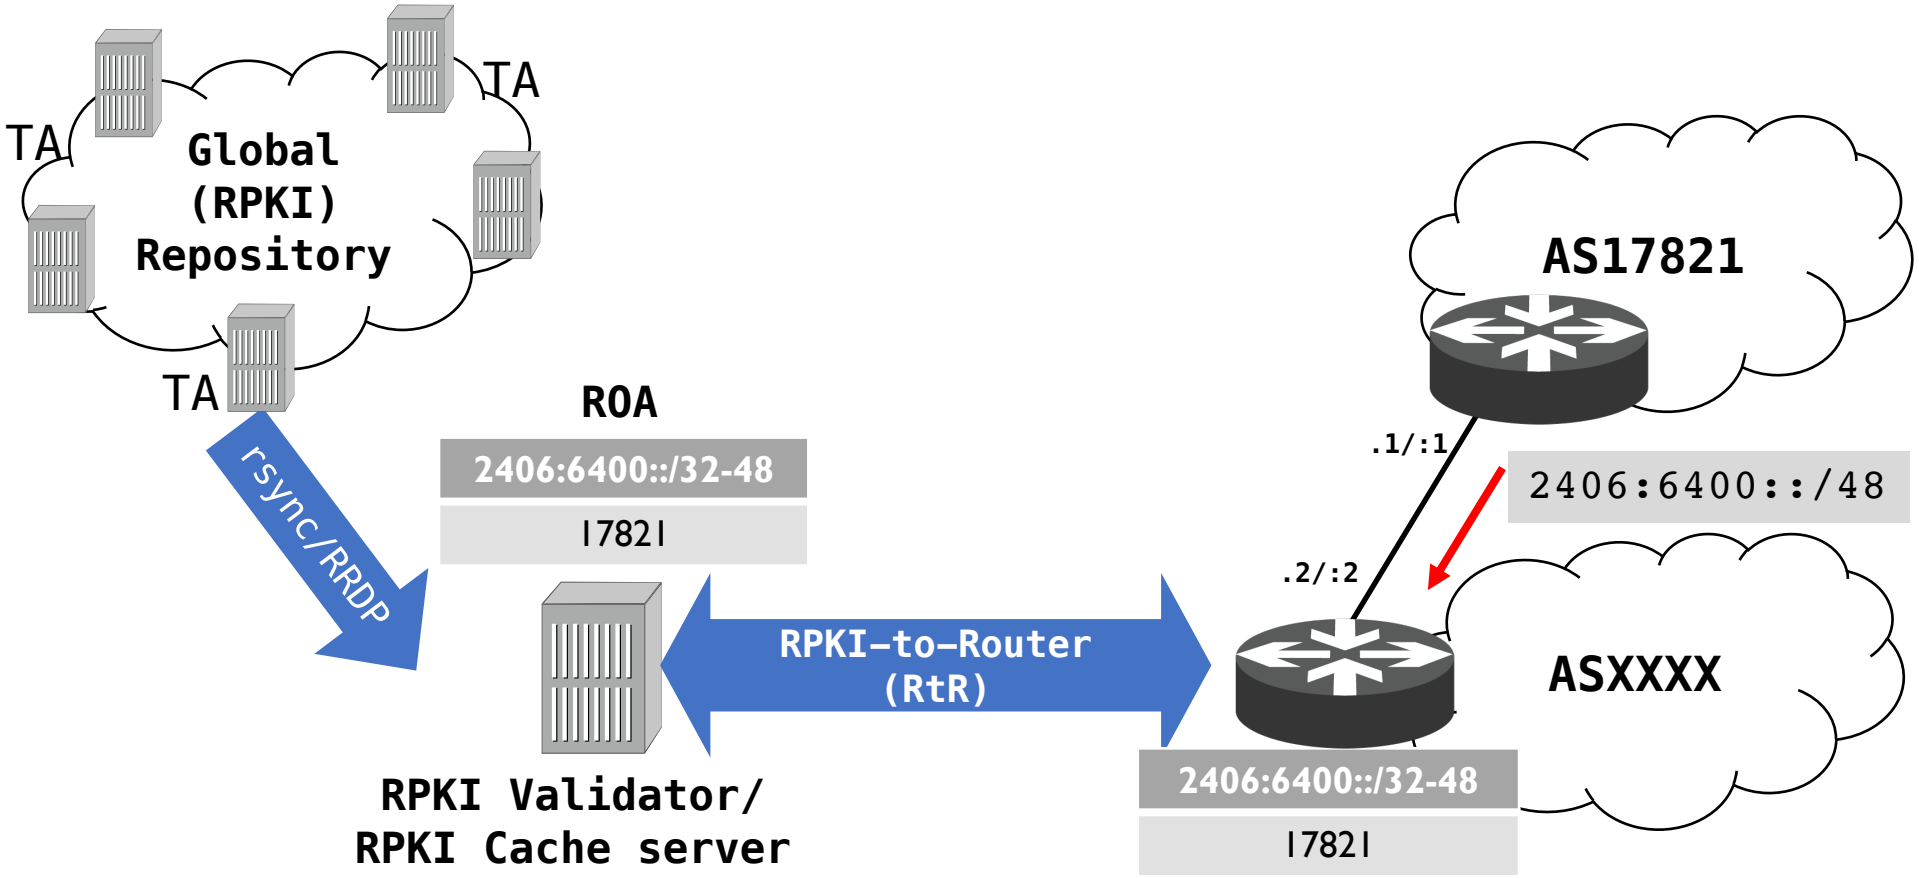

Are ROAs enough?

- What if I forge the origin AS in the AS path?
  - Would be accepted as “good” – pass origin validation!

- Which means, we need to secure the AS path as well
  - Need AS path validation (per-prefix)

AS path validation - BGPsec

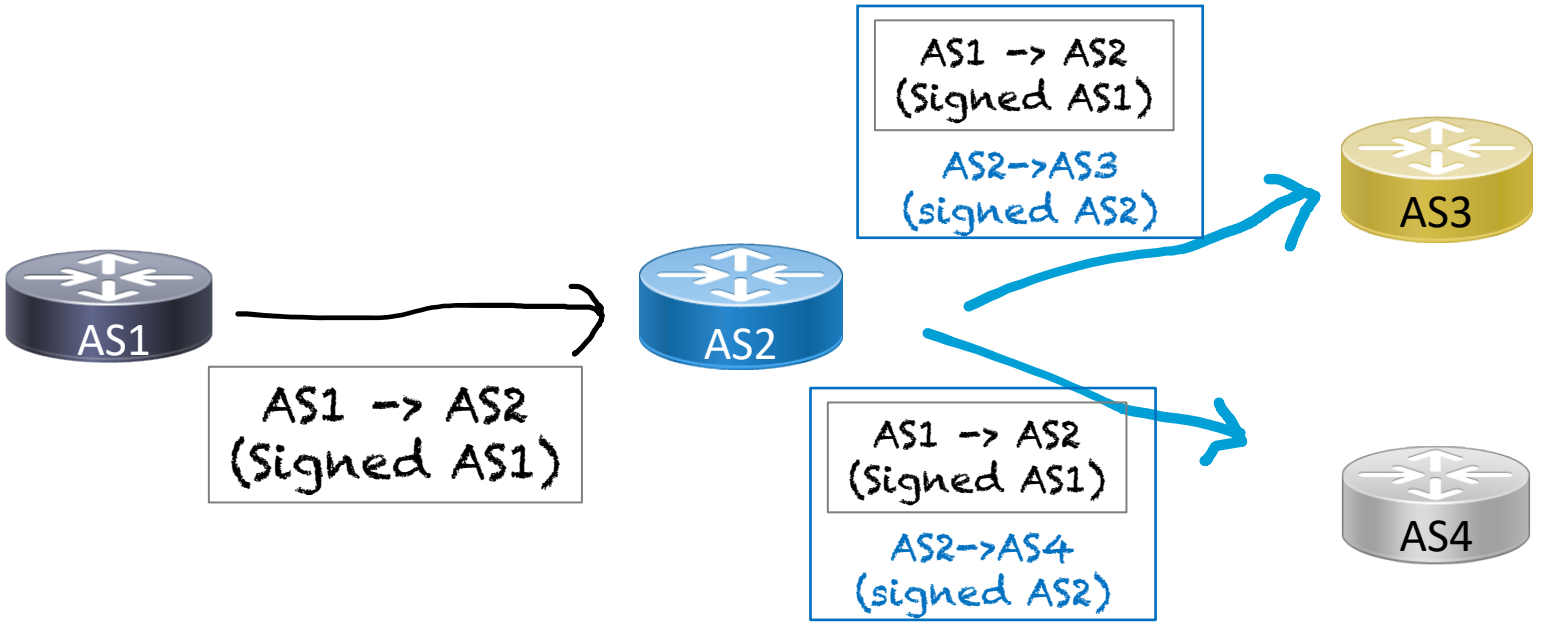

- A BGPsec speaker validates the received update by checking:
  - If there is a ROA that describes the prefix and origin AS
  - If the received AS path can be validated as a chain of signatures (for each AS in the AS path) using the AS keys

AS path validation issues…

- More resources
  - CPU - high crypto overhead to validate signatures, and
  - Memory
    - Updates in BGPsec would be per prefix
    - New attributes carrying signatures and certs/key-id for every AS in the AS path
- How do we distribute the certificates required?
- Can we have partial adoption?
- Given so much overhead, can it do more - Route leaks?

What can we do?

- Basic BGP OpSec hygiene – RFC7454/RFC8212
  - *RFC 8212* – BGP default reject or something similar
  - Filters with your *customers* and *peers*
    - *Prefix filters, Prefix limit*
    - *AS-PATH filters, AS-PATH limit*
    - Use IRR objects (source option) or ROA-to-IRR
  - Filters with your *upstream(s)*
  - Create ROAs for your resources
    - Filter based on ROAs -> ROV
- Join industry initiatives like MANRS
  - https://www.manrs.org/

MANRS

ROV – Industry trends

AT&T/as7018 now drops invalid prefixes from peers

Jay Borkenhagen jayb at braeburn.org
Mon Feb 11 14:53:45 UTC 2019

- Previous message (by thread): BGP topological vs centr
- Next message (by thread): AT&T/as7018 now drops inv
- **Messages sorted by:** [date] [thread] [subject] [aut

```
FYI:

The AT&T/as7018 network is now dropping all RPKI-i
announcements that we receive from our peers.

We continue to accept invalid route announcements
at least for now.  We are communicating with our c
invalid announcements we are propagating, informin
routes will be accepted by fewer and fewer network

Thanks to those of you who are publishing ROAs in
also like to encourage other networks to join us i
to improve the quality of routing information in t

Thanks!

Ja
```

[apops] RPKI ROV & Dropping of Invalids - Africa

- ***To***: apops@apops.net
- ***Subject***: [apops] RPKI ROV & Dropping of Invalids - Africa
- ***From***: Mark Tinka <mark.tinka@seacom.mu>
- ***Date***: Tue, 9 Apr 2019 14:05:03 +0200

Hello all.

In November 2018 during the ZAPF (South Africa Peering Forum) meeting in Cape Town, 3 major ISP's in Africa announced that they would enable RPKI's ROV (Route Origin Validation) and the dropping of Invalid routes as part of an effort to clean up the BGP Internet, on the 1st April, 2019.

On the 1st of April, Workonline Communications (AS37271) enabled ROV and the dropping of Invalid routes. This applies to all eBGP sessions for IPv4 and IPv6.

On the 5th of April, SEACOM (AS37100) enabled ROV and the dropping of Invalid routes. This applies to all eBGP sessions with public peers, private peers and transit providers, both for IPv4 and IPv6. eBGP sessions toward downstream customers will follow in 3 months from now.

We are still standing by for the 3rd ISP to complete their implementation, and we are certain they will communicate with the community accordingly.

Please note that for the legal reasons previously discussed on various fora, neither Workonline Communications nor SEACOM are utilising the ARIN TAL. As a result, any routes covered only by a ROA issued under the ARIN TAL will fall back to a status of Not Found. Unfortunately, this means that ARIN members will not see any improved routing security for their prefixes on our networks until this is resolved. We will each re-evaluate this decision if and when ARIN's policy changes. We are hopeful that this will happen sooner rather than later.

If you interconnect with either of us and may be experiencing any routing issues potentially related to this new policy, please feel free to reach out to:

- noc@workonline.africa
- peering@seacom.mu

Workonline Communications and SEACOM hope that this move encourages the rest of the ISP community around the world to ramp up their deployment of RPKI ROV and dropping of Invalid routes, as we appreciate the work that AT&T have carried out in the same vein.

In the mean time, we are happy to answer any questions you may have about our deployments. Thanks.

Mark Tinka (SEACOM) & Ben Maddison (Workonline Communications).

MMIX is dropping Invalids!

Acknowledgement

- Geoff Huston, APNIC
- Randy Bush, IIJ Labs/Arrcus

Any questions?