Some Thoughts on Integrity in Routing

Geoff Huston
Chief Scientist, APNIC

What we want…

- We want the routing system to advertise the **correct** reachability information for “**legitimately connected** prefixes at all times
- That means that we want to **avoid**:
  - promulgating reachability for bogus address prefixes
  - promulgating incorrect paths for reachable prefixes
  - blocking paths for legitimately connected prefixes

What do we do today?

I ask you to route my address prefix

You look for these addresses on **whois***

If it all seems to match then accept the request and add it to the network filters for this customer

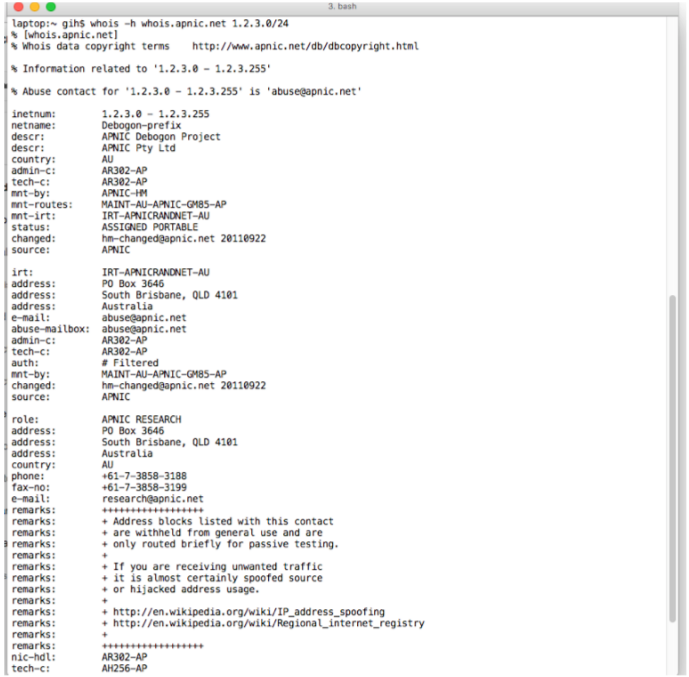

```
laptop:~ gih$ whois -h whois.apnic.net 1.2.3.0/24
% [whois.apnic.net]
% Whois data copyright terms    http://www.apnic.net/db/dbcopyright.html

% Information related to '1.2.3.0 - 1.2.3.255'

% Abuse contact for '1.2.3.0 - 1.2.3.255' is 'abuse@apnic.net'

inetnum:        1.2.3.0 - 1.2.3.255
netname:        Debogon-prefix
descr:          APNIC Debogon Project
descr:          APNIC Pty Ltd
country:        AU
admin-c:        AR302-AP
tech-c:         AR302-AP
mnt-by:         APNIC-HM
mnt-routes:     MAINT-AU-APNIC-GM85-AP
mnt-irt:        IRT-APNICRANDNET-AU
status:         ASSIGNED PORTABLE
changed:        hm-changed@apnic.net 20110922
source:         APNIC

irt:            IRT-APNICRANDNET-AU
address:        PO Box 3646
address:        South Brisbane, QLD 4101
address:        Australia
e-mail:         abuse@apnic.net
abuse-mailbox:  abuse@apnic.net
admin-c:        AR302-AP
tech-c:         AR302-AP
auth:           # Filtered
mnt-by:         MAINT-AU-APNIC-GM85-AP
changed:        hm-changed@apnic.net 20110922
source:         APNIC

role:           APNIC RESEARCH
address:        PO Box 3646
address:        South Brisbane, QLD 4101
address:        Australia
country:        AU
phone:          +61-7-3858-3188
fax-no:         +61-7-3858-3199
e-mail:         research@apnic.net
remarks:        +++++++++++++++++
remarks:        + Address blocks listed with this contact
remarks:        + are withheld from general use and are
remarks:        + only routed briefly for passive testing.
remarks:        +
remarks:        + If you are receiving unwanted traffic
remarks:        + it is almost certainly spoofed source
remarks:        + or hijacked address usage.
remarks:        +
remarks:        + http://en.wikipedia.org/wiki/IP_address_spoofing
remarks:        + http://en.wikipedia.org/wiki/Regional_internet_registry
remarks:        +
remarks:        +++++++++++++++++
nic-hdl:        AR302-AP
tech-c:         AH256-AP
```

* As usual, its not as simple as that, as there are a number of whois servers, and you probably have to negotiate across a number of them to get what you are after, or to be assured that the entry is not in any of the registry data collections

What do we do today?

I ask you to route my address prefix

You look for these addresses on **whois**

If it all seems to m

reque work filters for

this cu

The awesome power of whois!

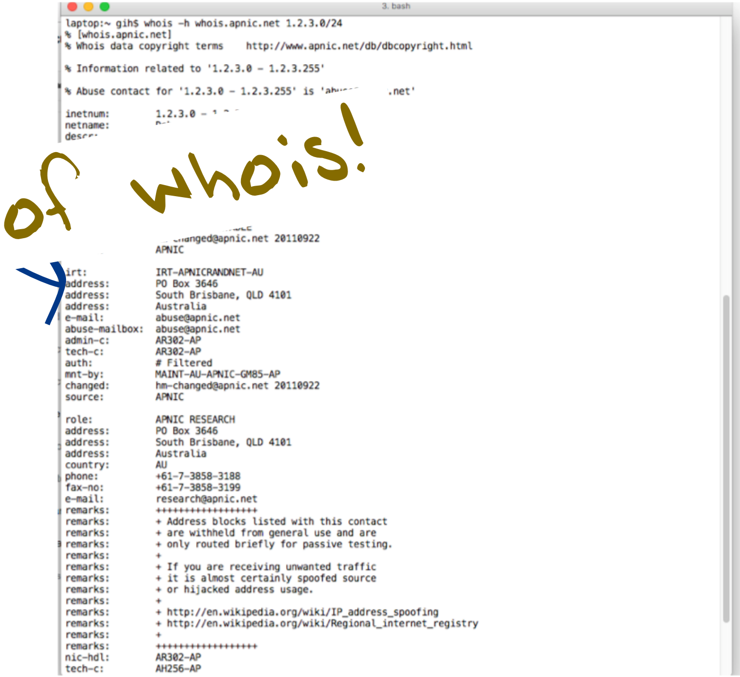

```
laptop:~ gih$ whois -h whois.apnic.net 1.2.3.0/24
% [whois.apnic.net]
% Whois data copyright terms    http://www.apnic.net/db/dbcopyright.html

% Information related to '1.2.3.0 - 1.2.3.255'

% Abuse contact for '1.2.3.0 - 1.2.3.255' is 'ab      .net'

inetnum:        1.2.3.0 -
netname:
descr:

                 changed@apnic.net 20110922
                APNIC

irt:            IRT-APNICRANDNET-AU
address:        PO Box 3646
address:        South Brisbane, QLD 4101
address:        Australia
e-mail:         abuse@apnic.net
abuse-mailbox:  abuse@apnic.net
admin-c:        AR302-AP
tech-c:         AR302-AP
auth:           # Filtered
mnt-by:         MAINT-AU-APNIC-GM85-AP
changed:        hm-changed@apnic.net 20110922
source:         APNIC

role:           APNIC RESEARCH
address:        PO Box 3646
address:        South Brisbane, QLD 4101
address:        Australia
country:        AU
phone:          +61-7-3858-3188
fax-no:         +61-7-3858-3199
e-mail:         research@apnic.net
remarks:        +++++++++++++++++
remarks:        + Address blocks listed with this contact
remarks:        + are withheld from general use and are
remarks:        + only routed briefly for passive testing.
remarks:        +
remarks:        + If you are receiving unwanted traffic
remarks:        + it is almost certainly spoofed source
remarks:        + or hijacked address usage.
remarks:        +
remarks:        + http://en.wikipedia.org/wiki/IP_address_spoofing
remarks:        + http://en.wikipedia.org/wiki/Regional_internet_registry
remarks:        +
remarks:        +++++++++++++++++
nic-hdl:        AR302-AP
tech-c:         AH256-AP
```

* As usual, its not as simple as that, as there are a number of whois servers, and you probably have to negotiate across a number of them to get what you are after, or to be assured that the entry is not in any of the registry data collections

What do we do today?

I ask you to route my address pr

You lo

If it a

reque

this cu

data collections

This is a manual process that relies on ascii pattern matching

it is error prone, does not scale, and provides no ongoing assurance

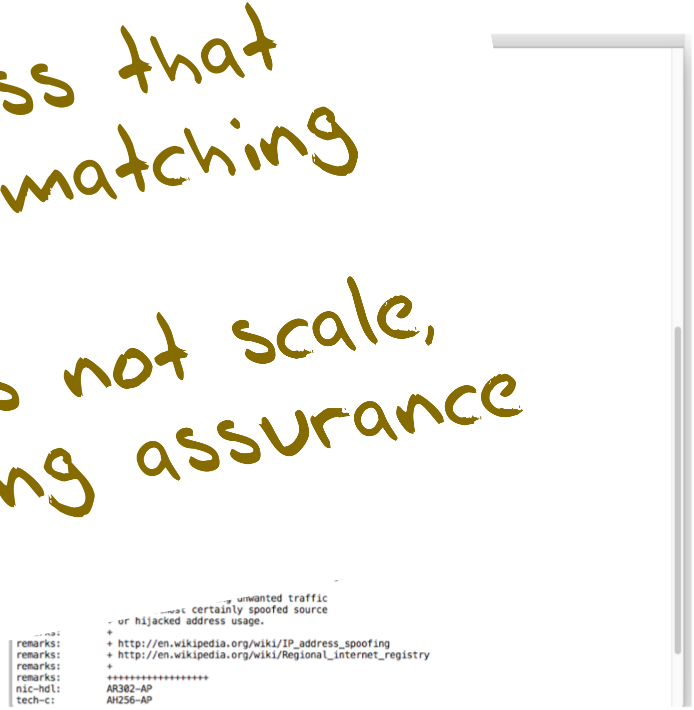

```
                     unwanted traffic
                  certainly spoofed source
                - or hijacked address usage.
              +
remarks:      + http://en.wikipedia.org/wiki/IP_address_spoofing
remarks:      + http://en.wikipedia.org/wiki/Regional_internet_registry
remarks:      +
remarks:      ++++++++++++++++++
nic-hdl:      AR302-AP
tech-c:       AH256-AP
```

What do we do today?

I ask you to route my net

You ask for me to provide a **“Letter of Authority”**

- Which is an effort to absolve you of all liability that may arise from announcing this route

You then add the to the network filters for this customer

What do we (

I ask you to route m

You ask for me to p ”

- Which is an effort to y arise from announcing this rout

You then add the tc ustomer

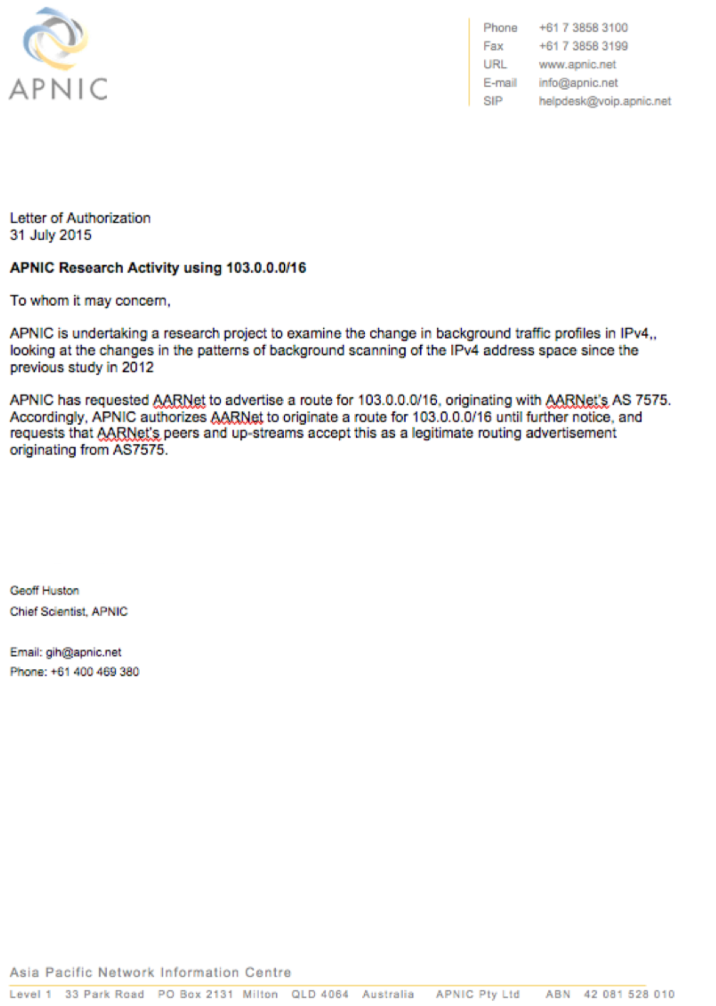

APNIC

Phone +61 7 3858 3100
Fax +61 7 3858 3199
URL www.apnic.net
E-mail info@apnic.net
SIP helpdesk@voip.apnic.net

Letter of Authorization
31 July 2015

APNIC Research Activity using 103.0.0.0/16

To whom it may concern,

APNIC is undertaking a research project to examine the change in background traffic profiles in IPv4,, looking at the changes in the patterns of background scanning of the IPv4 address space since the previous study in 2012

APNIC has requested AARNet to advertise a route for 103.0.0.0/16, originating with AARNet's AS 7575. Accordingly, APNIC authorizes AARNet to originate a route for 103.0.0.0/16 until further notice, and requests that AARNet's peers and up-streams accept this as a legitimate routing advertisement originating from AS7575.

Geoff Huston
Chief Scientist, APNIC

Email: gih@apnic.net
Phone: +61 400 469 380

Asia Pacific Network Information Centre
Level 1 33 Park Road PO Box 2131 Milton QLD 4064 Australia APNIC Pty Ltd ABN 42 081 528 010

What do we (

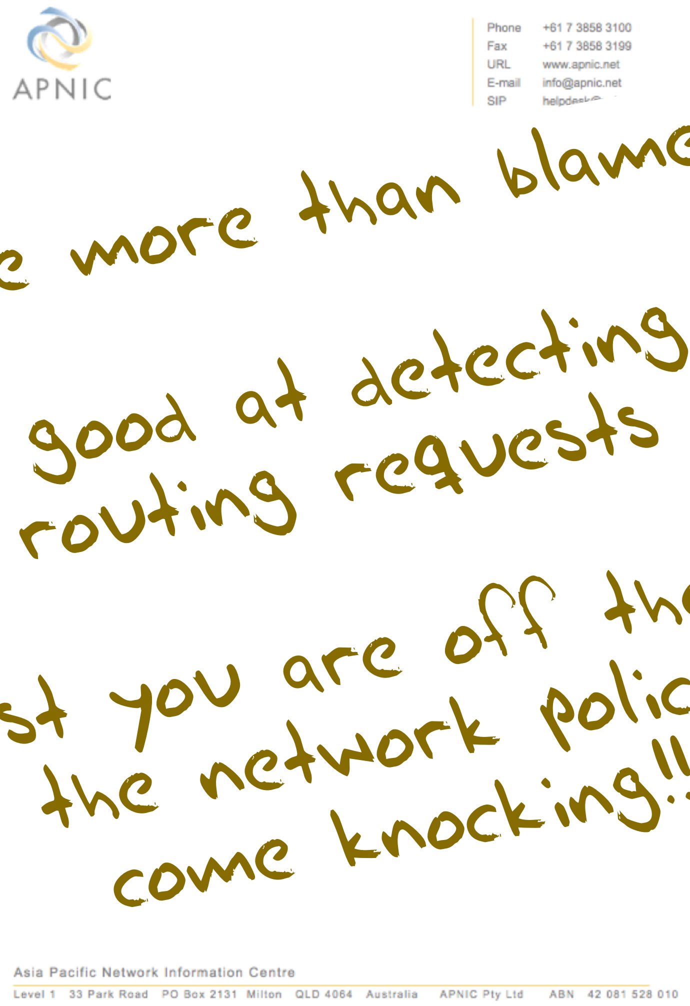

APNIC

Phone	+61 7 3858 3100
Fax	+61 7 3858 3199
URL	www.apnic.net
E-mail	info@apnic.net
SIP	helpdesk@

Asia Pacific Network Information Centre

Level 1 33 Park Road PO Box 2131 Milton QLD 4064 Australia APNIC Pty Ltd ABN 42 081 528 010

I ask you to r

Y

This is little more than blame shifting

This is no good at detecting incorrect routing requests

Yo

But at least you are off the hook when the network police come knocking!!

What do we do today?

I ask you to route my net

You ask for me to enter the details in a **route registry**

Your routers' access filters may be automatically generated from the route registry data that I entered

What do we do today?

I ask you to route my net

You ask for me to enter the details ir

Your routers’ access filters may be a

the route registry data that I enterec

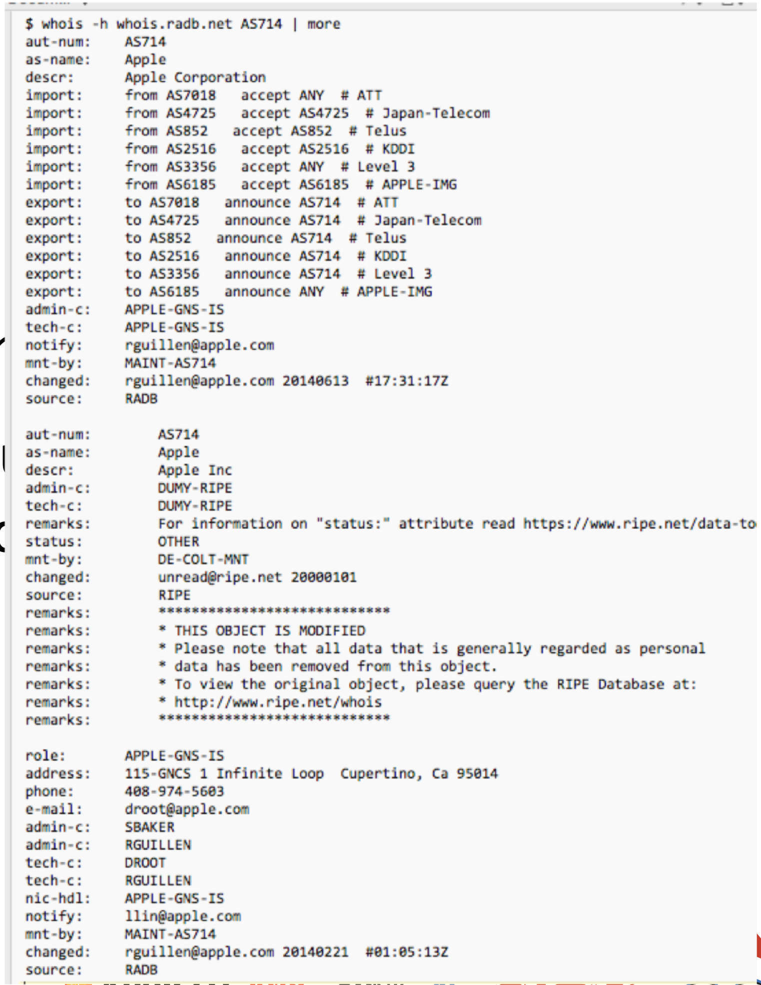

```
$ whois -h whois.radb.net AS714 | more
aut-num:    AS714
as-name:    Apple
descr:      Apple Corporation
import:     from AS7018   accept ANY  # ATT
import:     from AS4725   accept AS4725  # Japan-Telecom
import:     from AS852   accept AS852  # Telus
import:     from AS2516   accept AS2516  # KDDI
import:     from AS3356   accept ANY  # Level 3
import:     from AS6185   accept AS6185  # APPLE-IMG
export:     to AS7018   announce AS714  # ATT
export:     to AS4725   announce AS714  # Japan-Telecom
export:     to AS852   announce AS714  # Telus
export:     to AS2516   announce AS714  # KDDI
export:     to AS3356   announce AS714  # Level 3
export:     to AS6185   announce ANY  # APPLE-IMG
admin-c:    APPLE-GNS-IS
tech-c:     APPLE-GNS-IS
notify:     rguillen@apple.com
mnt-by:     MAINT-AS714
changed:    rguillen@apple.com 20140613  #17:31:17Z
source:     RADB

aut-num:        AS714
as-name:        Apple
descr:          Apple Inc
admin-c:        DUMY-RIPE
tech-c:         DUMY-RIPE
remarks:        For information on "status:" attribute read https://www.ripe.net/data-to
status:         OTHER
mnt-by:         DE-COLT-MNT
changed:        unread@ripe.net 20000101
source:         RIPE
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:        ****************************

role:       APPLE-GNS-IS
address:    115-GNCS 1 Infinite Loop  Cupertino, Ca 95014
phone:      408-974-5603
e-mail:     droot@apple.com
admin-c:    SBAKER
admin-c:    RGUILLEN
tech-c:     DROOT
tech-c:     RGUILLEN
nic-hdl:    APPLE-GNS-IS
notify:     llin@apple.com
mnt-by:     MAINT-AS714
changed:    rguillen@apple.com 20140221  #01:05:13Z
source:     RADB
```

What do we do today?

I ask you to rout

You ask

Your rou

the route

- How current is this data?
- is it complete?
- Can it be tampered with?
- Who entered this data? With what authority?
- Can i trust it to use the stored information as an automatic filter generator for my network?

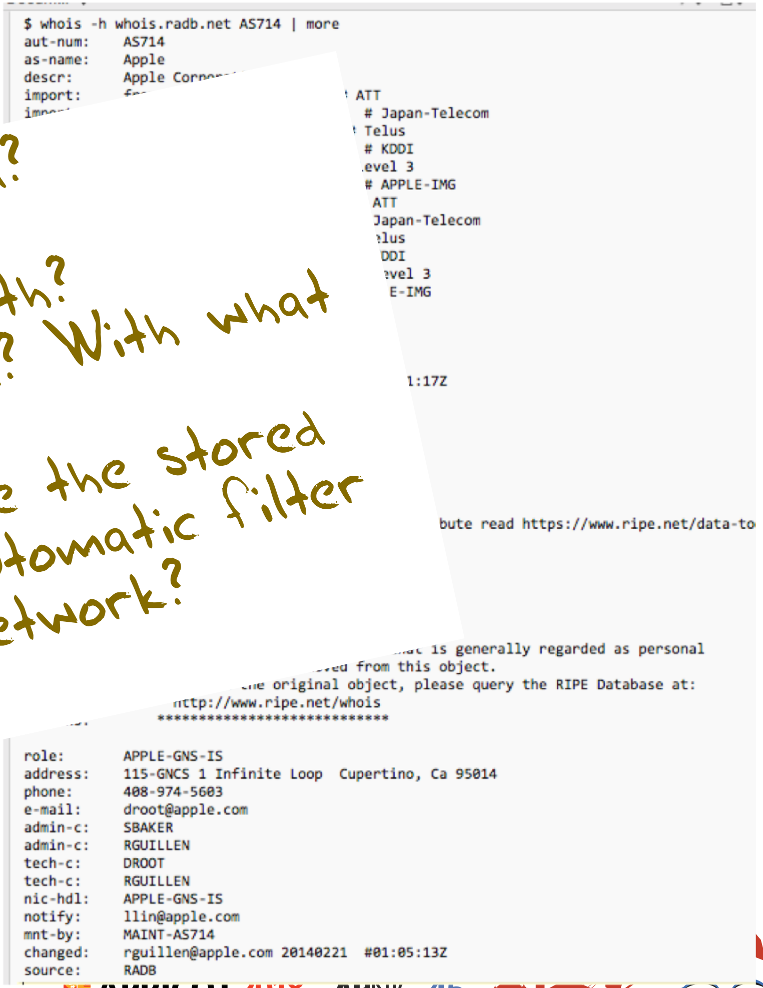

```
$ whois -h whois.radb.net AS714 | more
aut-num:     AS714
as-name:     Apple
descr:       Apple Corpor
import:      fr                          ATT
impor                                     # Japan-Telecom
                                          Telus
                                          # KDDI
                                          evel 3
                                          # APPLE-IMG
                                           ATT
                                           Japan-Telecom
                                           lus
                                           DDI
                                           evel 3
                                            E-IMG

                                            1:17Z

                                              bute read https://www.ripe.net/data-to

                                 at is generally regarded as personal
                            ved from this object.
                       the original object, please query the RIPE Database at:
                    http://www.ripe.net/whois
                   ****************************

role:        APPLE-GNS-IS
address:     115-GNCS 1 Infinite Loop  Cupertino, Ca 95014
phone:       408-974-5603
e-mail:      droot@apple.com
admin-c:     SBAKER
admin-c:     RGUILLEN
tech-c:      DROOT
tech-c:      RGUILLEN
nic-hdl:     APPLE-GNS-IS
notify:      llin@apple.com
mnt-by:      MAINT-AS714
changed:     rguillen@apple.com 20140221  #01:05:13Z
source:      RADB
```

What do we do today?

I ask

You

Your

the

A publicly accessible description of every import and export policy to every transit, peer, and customer, has proved to be extremely difficult to maintain

Today, we have many routing registries, not one, and the quality of the data in those registries is close to impossible to ascertain.

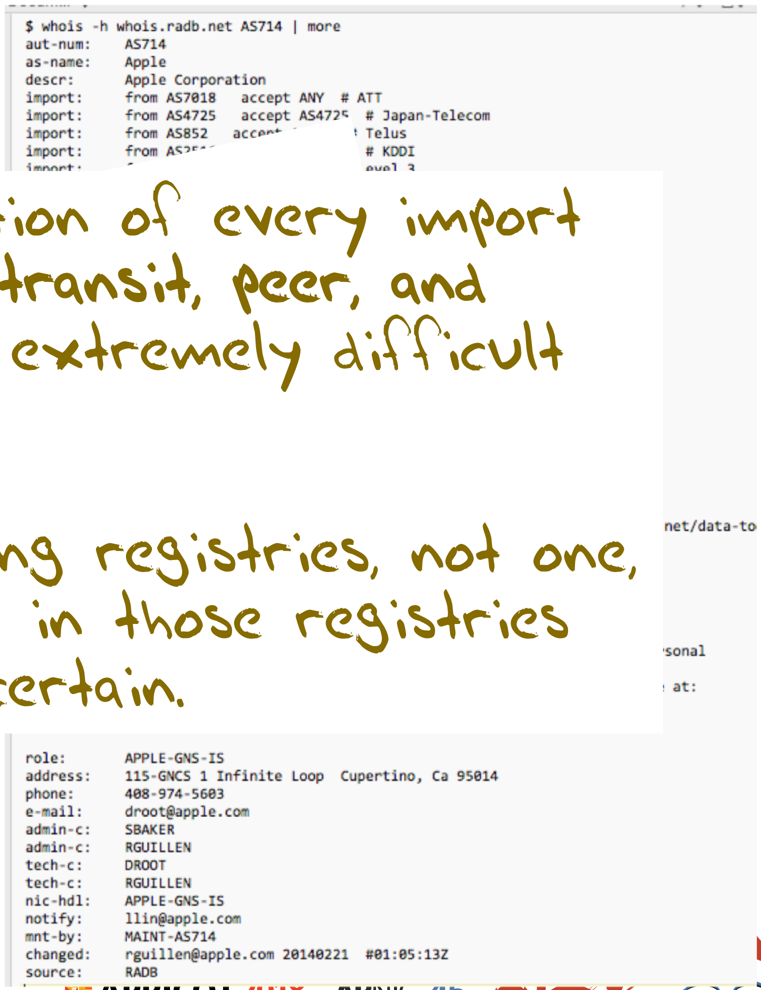

```
$ whois -h whois.radb.net AS714 | more
aut-num:    AS714
as-name:    Apple
descr:      Apple Corporation
import:     from AS7018   accept ANY  # ATT
import:     from AS4725   accept AS4725  # Japan-Telecom
import:     from AS852   accep          # Telus
import:     from AS25                   # KDDI
import:                                 evel 3
```

```
net/data-to

sonal

at:
```

```
role:       APPLE-GNS-IS
address:    115-GNCS 1 Infinite Loop  Cupertino, Ca 95014
phone:      408-974-5603
e-mail:     droot@apple.com
admin-c:    SBAKER
admin-c:    RGUILLEN
tech-c:     DROOT
tech-c:     RGUILLEN
nic-hdl:    APPLE-GNS-IS
notify:     llin@apple.com
mnt-by:     MAINT-AS714
changed:    rguillen@apple.com 20140221  #01:05:13Z
source:     RADB
```

What’s the problem here?

- None of these approaches are very satisfactory as a complete solution to this problem
- Let’s take a step back and see if we can use digital signature technology to assist here.
- If we can, then we can construct automated systems that will recognise validly signed attestations about addresses and their use

Registry Role

- The registry plays the role of a neutral third party ‘trust point’ that can provide an impartial record of which entity is the current holder of an IP address
- Which is fine for humans, but of limited use to automated systems
- How can we automate the validation function that allows an entity to validate whether or not a party is the current holder of an IP address?

Crypto to the rescue!

- Public / Private keys can be really useful here
  - I sign <something> using my private key and send it to you
  - Using my public key you can be assured that:
    - I signed this (and no one else)
    - I cannot deny that I signed it
    - What I signed has not been altered on the way between me and you
  - The assurance can be automated, and does not necessarily rely on a manual process of matching ascii text

The RPKI

- If I have the association between a public key and a number block registered by the RIR, then
  - Instead of performing a human match between the registry entry and the party you can get the party to sign an attestation using their local private key
  - If the attestation can be validated by the public key published by the RIR then you have automated the validation function and don’t need eyeballs to read web pages to validate the ‘rights’ of use of IP addresses

The RPKI Certificate Service

- Enhancement to the RIR Registry
  - Offers verifiable proof of the number holdings described in the RIR registry

- Resource Certification is an opt-in service
  - Number Holders choose to request a certificate
    - Derived from registration data

BGPSEC: BGP + RPKI Origination

- One approach is to look at the process of “permissions” that add an advertised address prefix to the routing system:
  - The address holder is authorizing a network to originate a route advertisement into the routing system
- The **ROA** is a digitally signed version of this authority. It contains
  - An address prefix (and range of ‘allowed’ prefix sizes)
  - An originating ASN
- This allows others to check the validity of a BGP route origination:
  If there is a valid ROA, and the origin AS matches the AS in the ROA, and the prefix length is within the bounds of the ROA, then the announcement has been entered into the routing system with the appropriate permissions

BGPSEC: BGP + RPKI Propagation

- In BGP AS Path manipulation is also a problem
- How can a BGPSEC speaker know that the AS Path in a BGP Update is genuine?
- Answering this question in BGPSEC gets very messy very quickly!

In my opinion: It’s highly unlikely that we will see widespread uptake of BGPSEC anytime soon, if ever, largely due to the overheads associated with AS path signing

Errrr – why isn't this being adopted by ISPs?

- Cryptography and Certificate management are operationally challenging:

  which is often seen as one more thing to go wrong!

- Validation of signed data is convoluted – maybe it should've been simpler

- Its not just ROAs – you need AS Path protection as well
  - As long as a hijacker includes your ROA-described originating AS in the faked AS PATH then the hijacker can still inject a false route
  - If ROAs are challenging for operators, then BGPsec is far more so!

The Perfect is the Enemy of the Good

Maybe there are some “Good” things we can do right now instead of just waiting for BGPsec to be sorted out!

More Ideas?

- Waiting for everyone to adopt a complex and challenging technology solution is probably not going to happen anytime soon
- Are that other things we can do that leverage the RPKI in ways that improve upon existing measures?
  - Use ROAs to digitally sign a LOA?
  - Digitally sign whois entries?
  - Digitally sign Routing Policy descriptions in IRRs

  - Signed data could help a user to determine if the information is current and genuine
  - This would not directly impact routing infrastructure, but instead would improve the operators’ route admission process to automatically identify routing requests that do not match signed registry / routing database information

What should we do?

- We could keep on thinking about how to make a routing infrastructure that is impervious to attempts to coerce it into false states
  - But it seems that we are not sure how to do this, and not sure who would pay the cost of trying to do this!

AND/OR

- Perhaps we should undertake some focussed work on open BGP monitoring and alarm services that allow us to detect and identify routing issues as they arise, and assist network operators to respond quickly and effectively

Thanks!