Routing and Forwarding Integrity: Defenses to Common Challenges

John Kristoff
jtk@cymru.com

Systems and policy

- Sane and secure defaults
- Automated configuration management, connected to
- Automated monitoring and measurement capabilities
- Policy, type and consistency checking
- Control plane isolation and protection
- Authentication and cryptography
- Neighbor collaboration and cooperation

Goodput

- Peer/IX transit/forwarding theft mitigation
- Minimizing CPU punts and packet processing
- BCP 38/84, uRPF, SAVI
  - WARNING: all spoofing-related discussion is tabled
- RTBH, flow-spec, rate limiting, filters
- Queuing and active traffic management (RED, CoS)
- Darknets, quarantines and sinkholes
- Redirects (e.g. fabricated GFW DNS answers)

Route Integrity

- RPKI-Based Origin Validation / BGPsec
- IRR-Based ACLs
- Route history monitoring and alerting
- Route flap dampening
- Prefix allocation boundary filtering
- Prefix announcement count limits
- “Golden Routes” protection

Thank you

- John Kristoff
- <jtk@cymru.com> - https://www.cymru.com/jtk/

Locking Down the Control Plane

At a New Zealand Exchange

Randy Bush <randy@psg.com>

Cristel Pelsser <cristel@iij.ad.jp>

Dean Pemberton <dean@deanpemberton.com>

Josh Bailey <joshb@google.com>

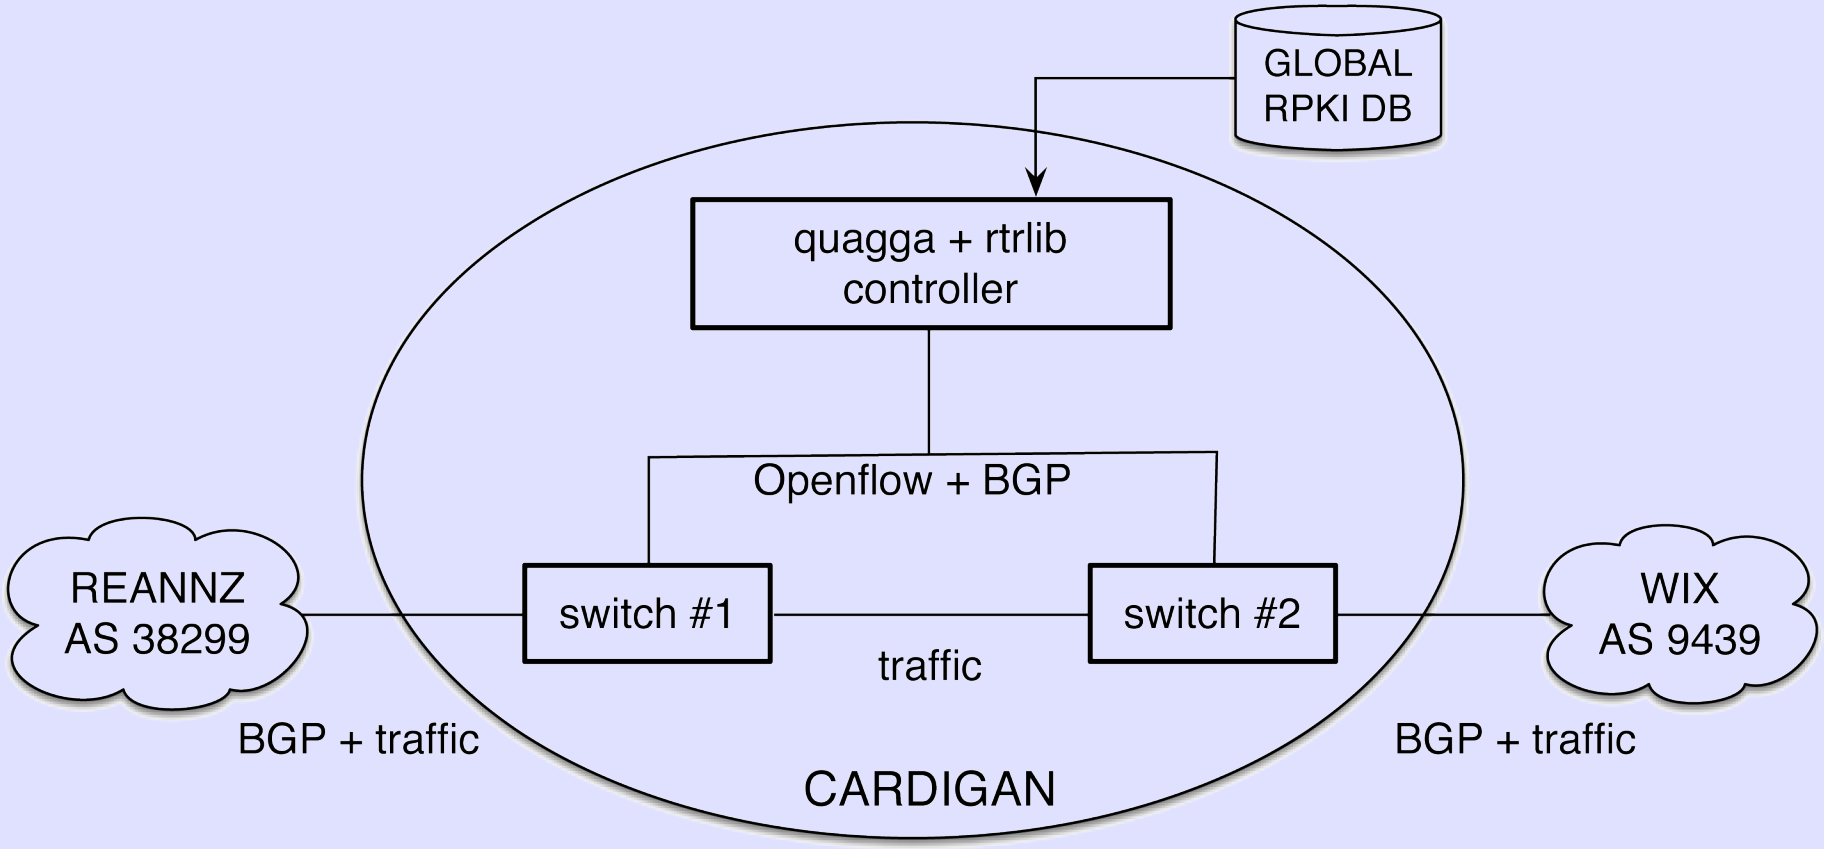

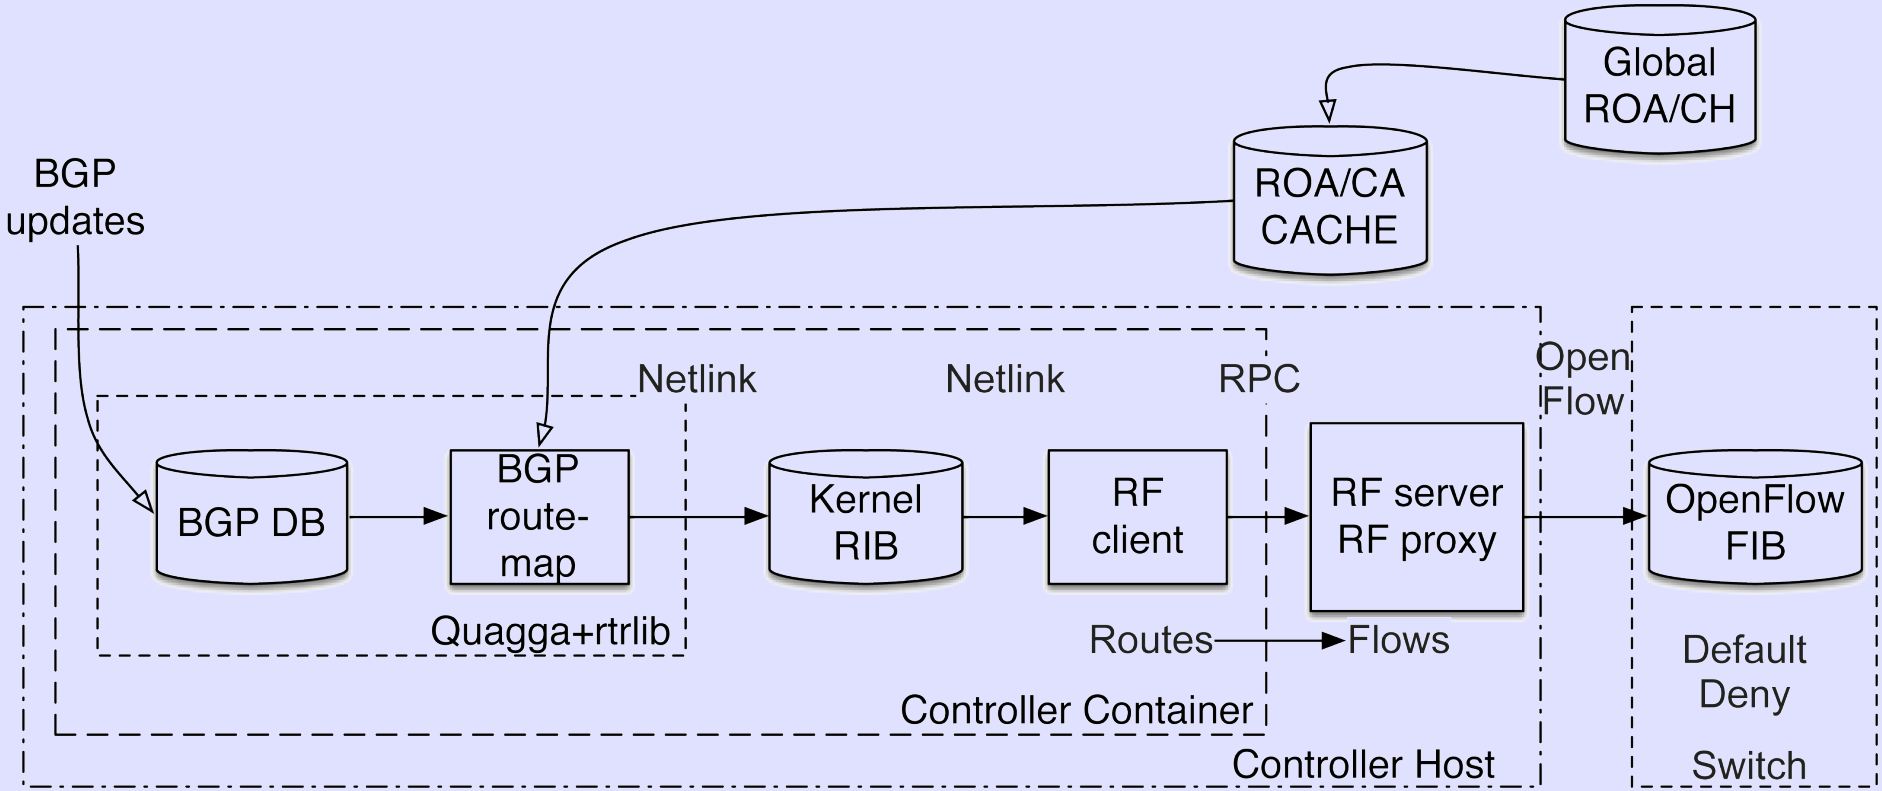

Making the best of IRR

Job Snijders
<job@ntt.net>

IRR Lockdown

NTT evaluates ignoring route objects that cover RIPE space that don’t come from RIPE itself.

In other words: Register route objects for RIPE space in the RIPE registry.

```
inetnum:        193.0.0.0 - 195.255.255.255
netname:        EU-ZZ-193-194-195
descr:          European Regional Registry
```

Good:

```
route:   193.0.0.0/21
descr:   RIPE-NCC
origin:  AS3333
mnt-by:  RIPE-NCC-MNT
source:  RIPE
```

BAD!

```
route:   193.0.0.0/21
descr:   RIPE-NCC
origin:  AS666
mnt-by:  RIPE-NCC-MNT
source:  RADB
```

Why would we ever honor the bad route object?!

Finding odd ASpaths

Andree Toonk
<andree@bgpmon.net>

BGP anomaly detection

Right?

BGP anomaly detection

- Expected:
  - 208.67.220.0/24 36692 OpenDNS
  - 558 6461 2914 36692
- Detected:
  - 208.67.220.0/24 4761 Indosat < Hijack
  - 208.67.220.220/32 9121 Turk Telekom < Hijack

  - Detection origin AS changes is pretty simple

BGP anomaly detection

BGP anomaly detection

```
271 6939 35625 6453 3215
```

AS3215 France Telecom *(origin)*

non existing relation

AS6453 Tata

transit

AS35625 Avenir Telematique

peer

AS6939 HE

customer

AS271 BCNET *(BGP feed peer)*

BGP anomaly detection

```
133165 3491 4826 1221 10026 13335
```

13335 CloudFlare *(origin)*

Transit

10026 Pacnet

Transit

1221 Telstra

customer

4826 Vocus

Transit

3491 PCCW

customer

133165 digital ocean *(BGP feed peer)*

Note: 13335 also buys from Vocus, so simple prefix filter caused the 'leak' and interferes with anycast / traffic engineering

Not always as clear…

Routing, Forwarding and IPv6

Jen Linkova
<furry@google.com>

IPv6 is a New Black

- Needs to be secured?
  - Yes
- By copying v4 configs?
  - No
- Test everything again?
  - Yes

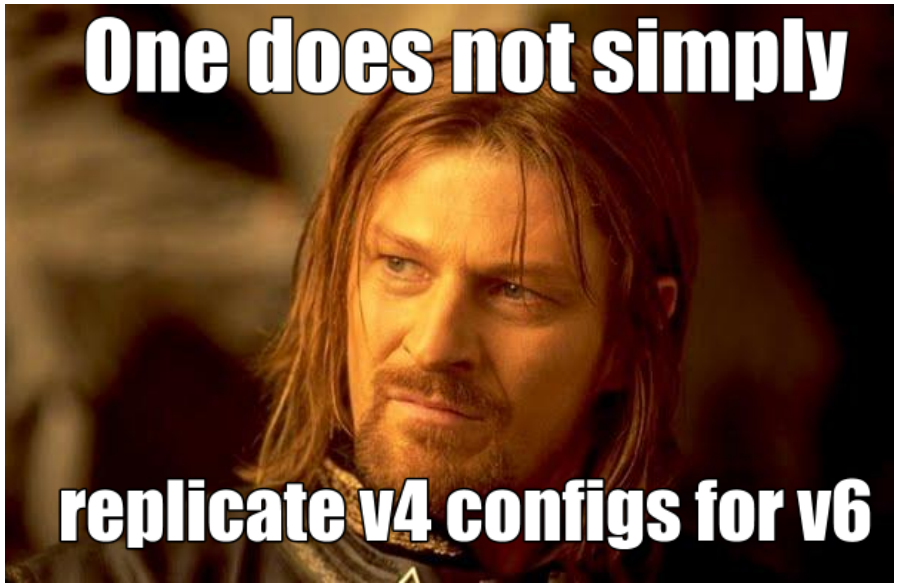

A Few Examples to Think about

- Prefixes longer than /64
  - could your router install it into FIB?
- ACL mismatch due to
  - longer header
  - longer prefixes

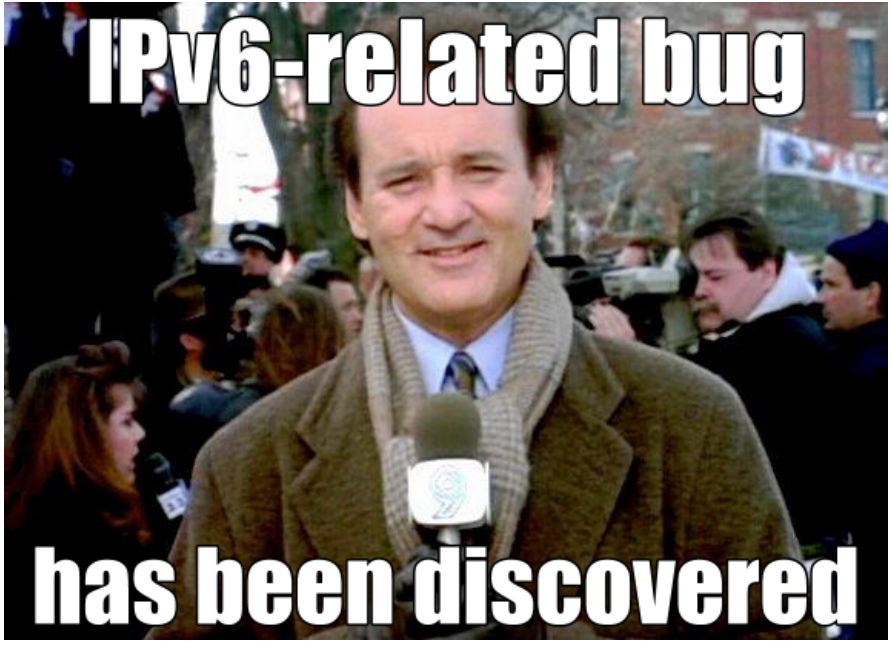

...More Examples...

- ~~Deaggregation~~ Traffic Engineering
  - huge number of prefixes
- Using /64 on p2p links
- Using LLA as security feature

UDP Amplification update

Jared Mauch

2014-Feb-03

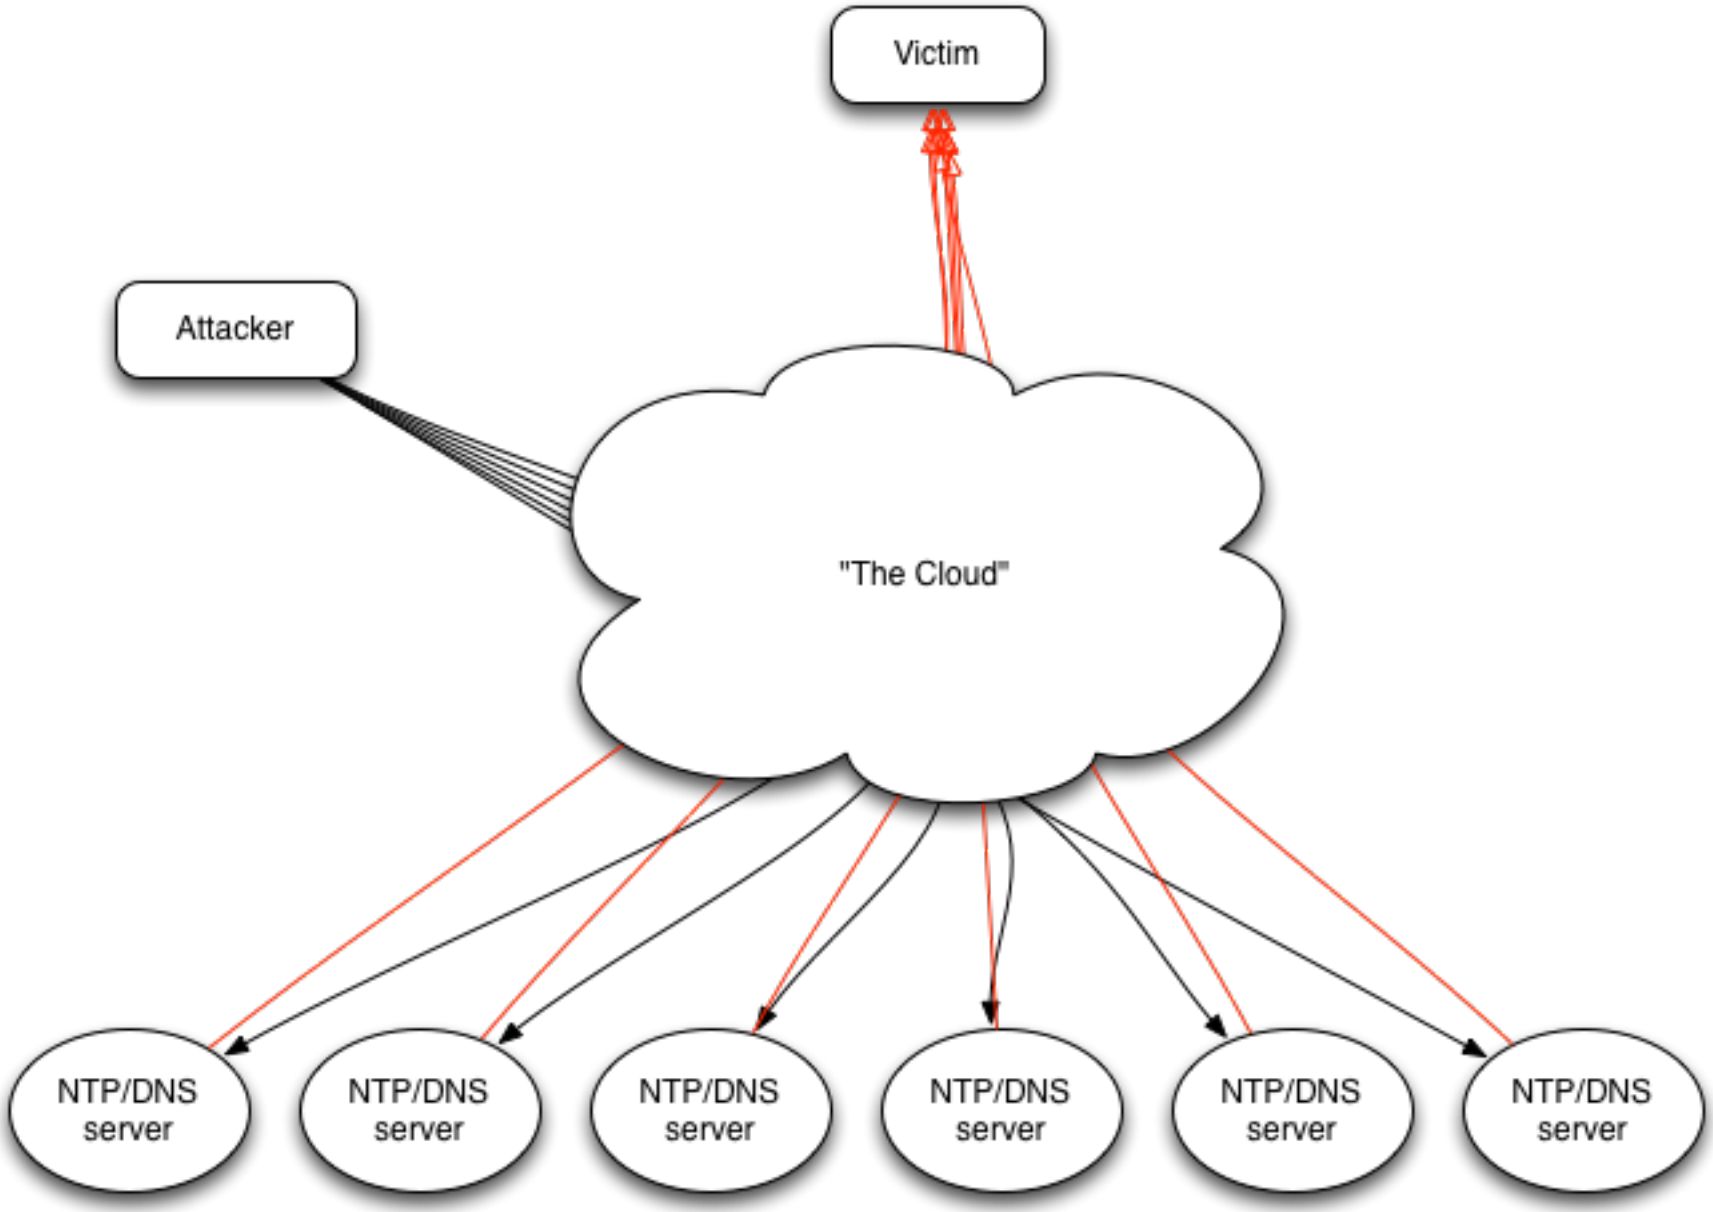
Victim
Attacker
"The Cloud"
NTP/DNS server
NTP/DNS server
NTP/DNS server
NTP/DNS server
NTP/DNS server
NTP/DNS server

About Open{Resolver,NTP,SSDP,SNMP} Project

- How weekly scanning works
  - DNS since early 2013 (Sundays)
  - NTP since early 2014 (Fridays)
  - SNMP in 2014 (Tuesdays)
  - SSDP in late 2014 (Saturdays)
  - Learned a lot the first weeks
  - Improved the slow-scan methods
- Excluded only 127/8 10/8 and 192.168/16
  - Always room for improvement
  - Few complaints

About the scanning

- What a scan looks like

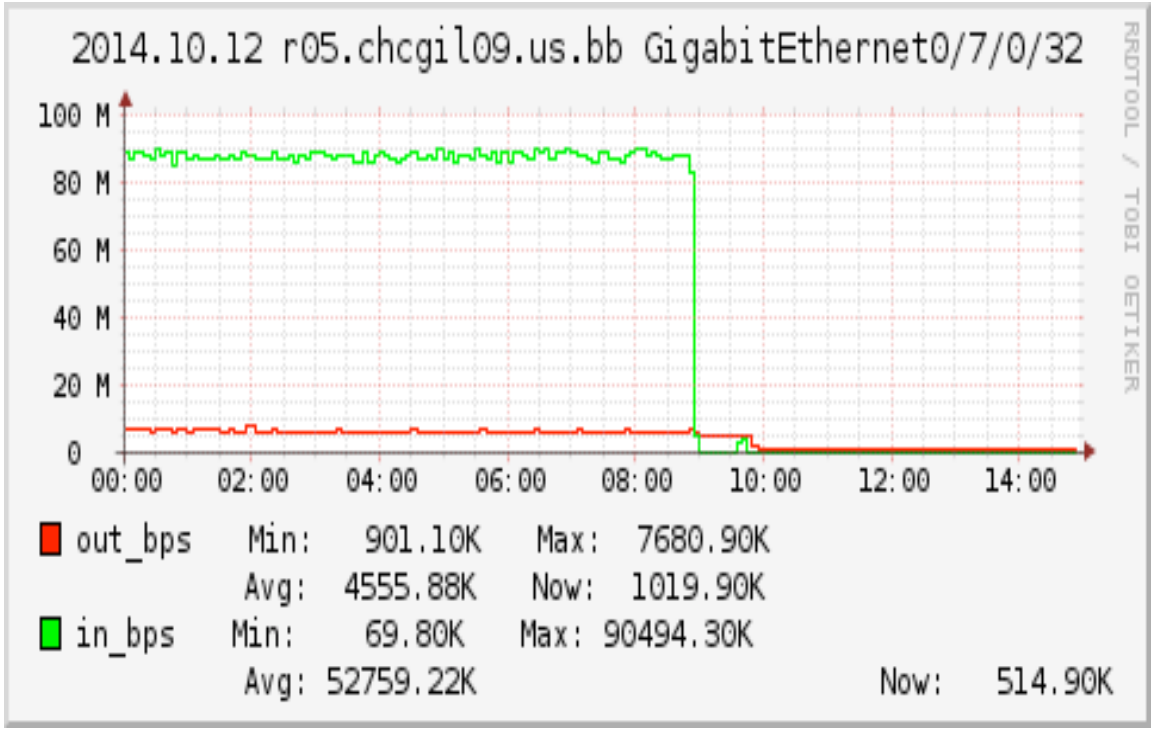

About OpenResolverProject Data

- First data was unusual
  - Took steps to validate results
  - Unexpected mysteries occurred
- DNS uses UDP/53
  - Probes came back from port other than port 53
  - 46-49% of data of this type
- Wrong IP responded
  - 2% from some other IP
  - Can detect and infer spoofing IP networks

About OpenResolverProject Data

- Madness with the Method
  - Unique query sent to each IP
  - Encoded in hex with XOR
- Software has bugs
  - Responds to network, broadcast addresses
  - Responds multiple times
  - Scanned hosts respond for hours, days later

About OpenResolverProject Data

- Misbehaving root causes
  - Many CPE respond on WAN interface
  - Forward query to configured DNS server
    - Alter packet Destination (spoofing scan host IP)
- Remediation
  - Vendors swapped CPE
    - Belkin is amazing to work with
  - Firmware fixes made available

About OpenResolverProject Data

- Graphs representing data

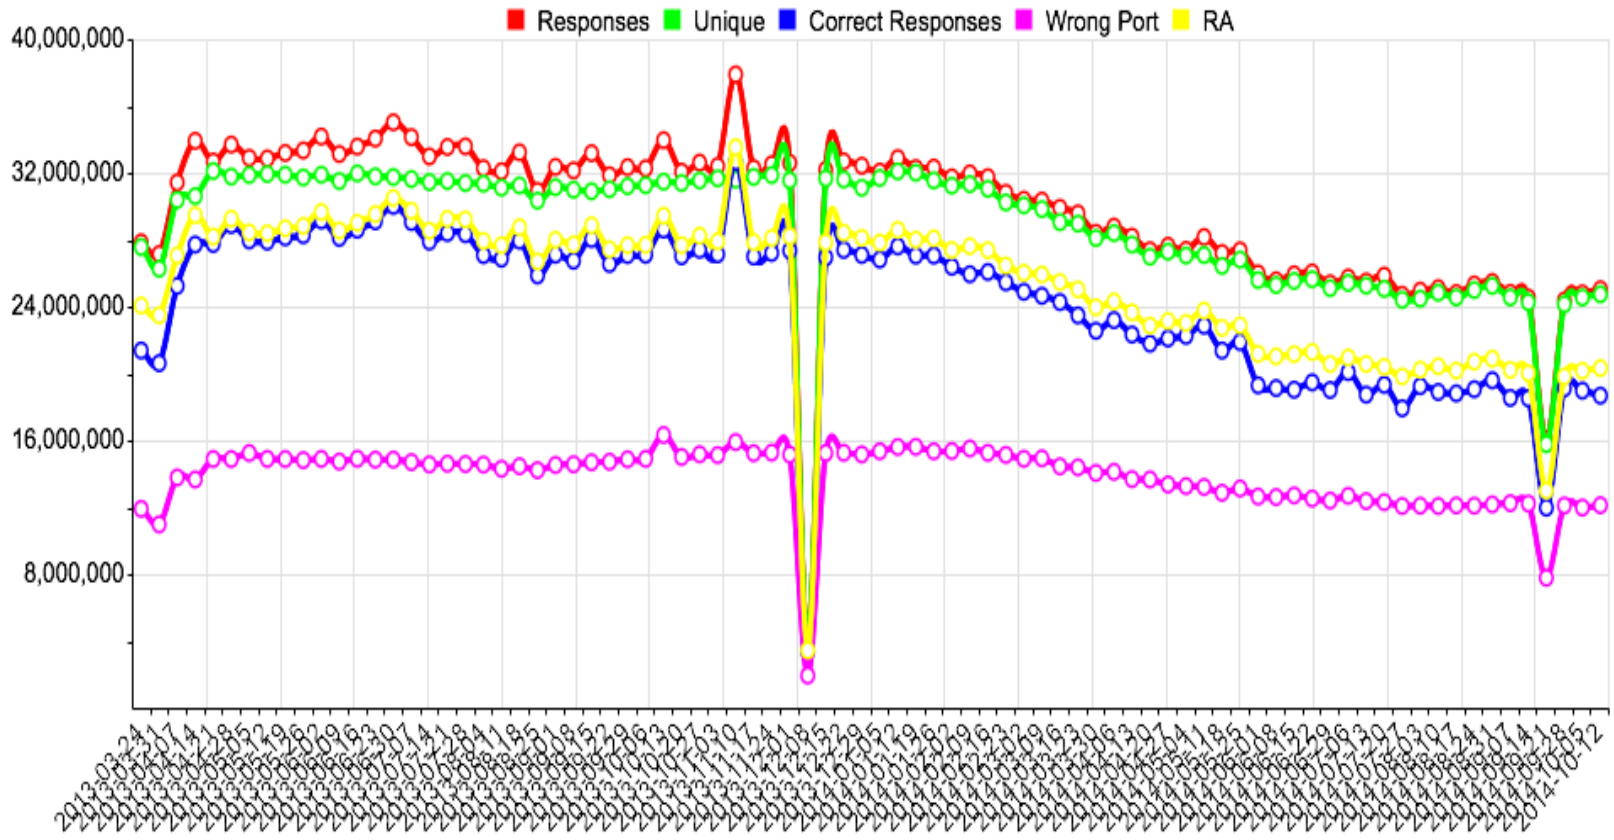

About OpenNTPProject Data

- Default behavior of NTP and Mode 7
  - Monlist provides large amplification effect
- Studied and detailed by researchers
  - Christian Rossow – “Amplification Hell”
  - Jakub Czyz et al “Rise and Decline of NTP DDoS Attacks”
- Provides ~500-1000x bitrate amplification
- Support removed via Bug#1532 in 4.2.7p26 April 26, 2010

About OpenNTPProject Data

- Monlist Amplifier Change Rate
- 1,529,866 2014-01-10
- 1,402,569 2014-01-17
- 803,156 2014-01-24
- 564,027 2014-01-31
- 490,724 2014-02-07
- 349,583 2014-02-14
- ...
- 188,549 2014-10-10

About OpenNTPProject Data

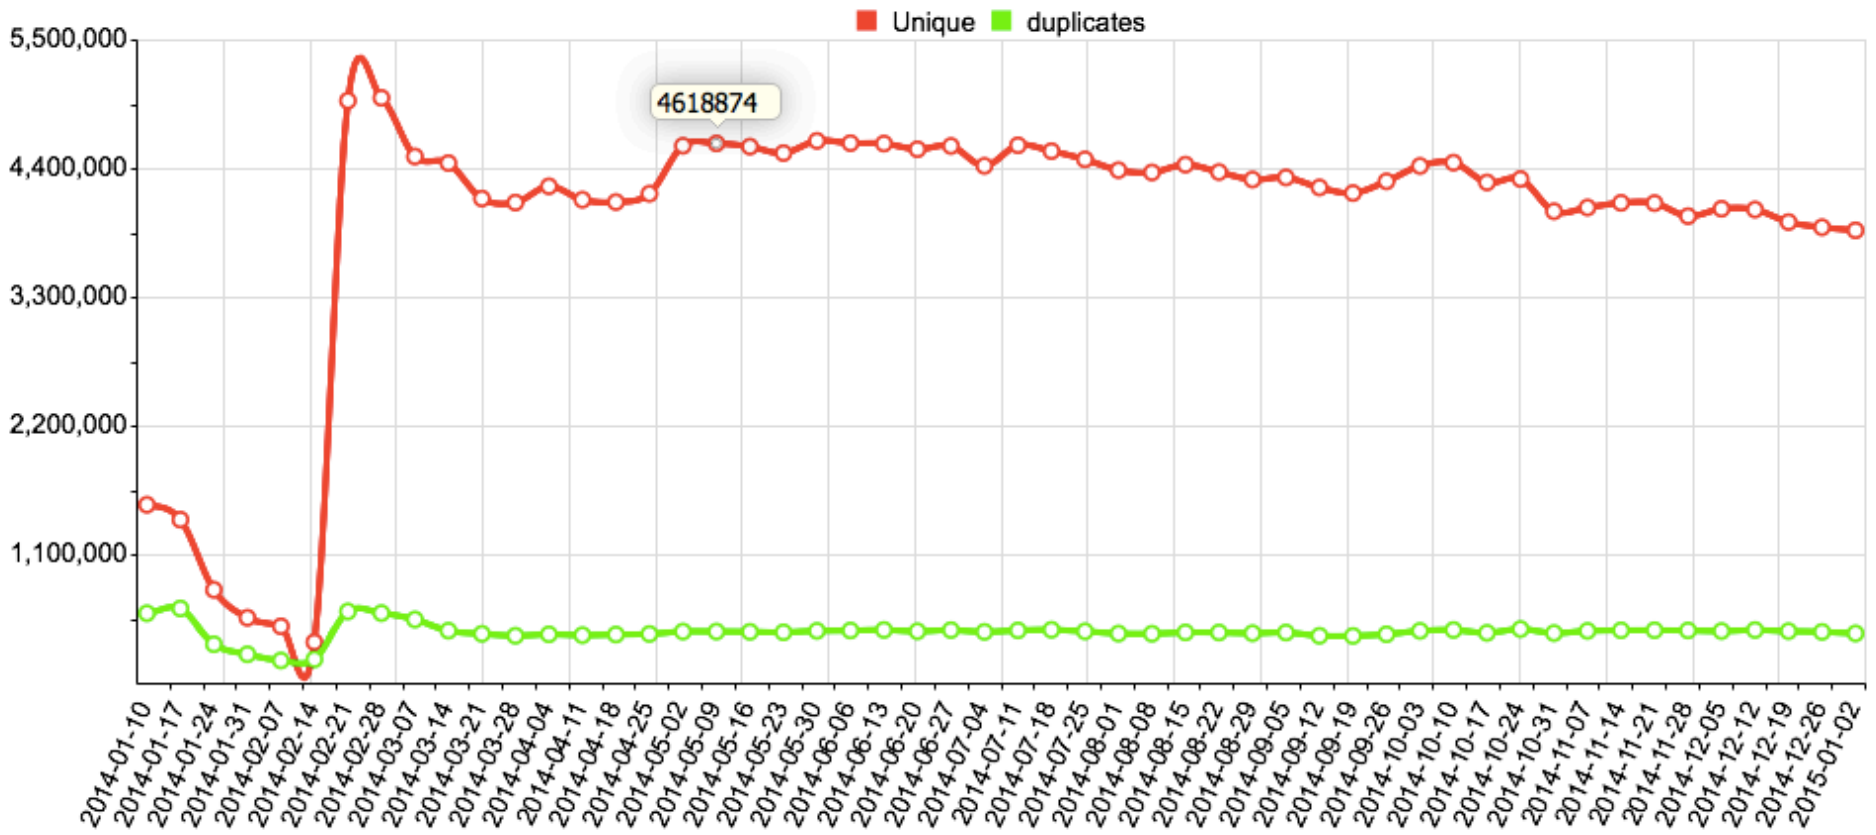

About OpenNTPProject Data

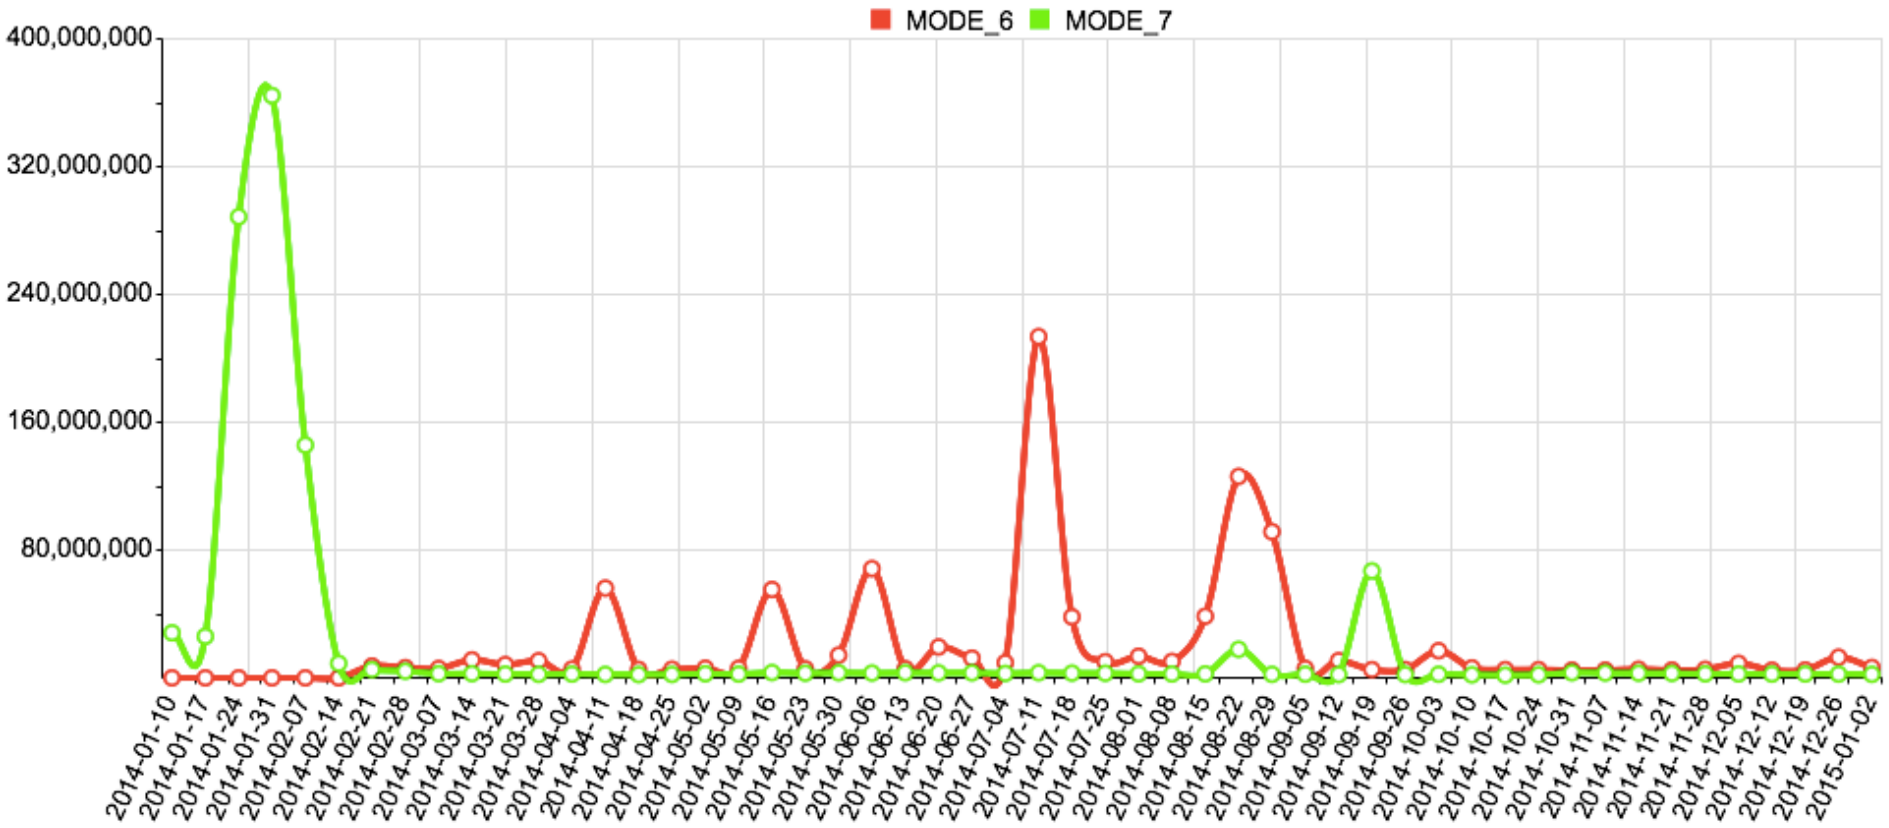

Some graphics are without meaning, like this one

About OpenNTPProject Data

- Version Scanning
  - Gives detailed information about deployed versions
  - Can fingerprint hosts
    - IOS-XR
      - version="4", processor="unknown", system="UNIX”
      - system="cisco",
    - Linux – Broadcom 24xGE + 4x10GE Switch (!)
      - version=“ntpd 4.1.1c-rc1@1.836 Fri Nov 19 10:37:40 KST 2010 (414)", processor="mips", system="Linux2.4.20_mvl31-bcm95836cpci

About OpenNTPProject Data

- Lets get personal details
  - processor="i386", system="JUNOS8.1R4.3”
  - processor="x86_64", system="VMkernel/4.1.0”
  - processor="i386", system="BIG-IPBIG-IP 4.6.2”
  - processor="UltraSparc-IIe", system="sparcv9-wrs-vxworks”
  - system="Linux2.6.18_pro500-p34xx-mips2_fp_le-ubiquoss”
    - FTTX/GPON CPE in Korea

About OpenSNMPProject Data

- Wait, the Management is on the internet?
- SNMP guides online often use the default public/private communities
  - Scans started 2014-06-24
  - 6-10GB of data per week
  - 7,340,773 unique devices responded 2014-10-07
- Similar challenges with embedded solutions and defaults
- Once you talk to a host, some send you their traps
  - SNMP can be quite revealing

Internet of Everything, Including…

- NTCIP Signs
  - Eagle EPAC300
  - Skyline NTCIP DMS Sign

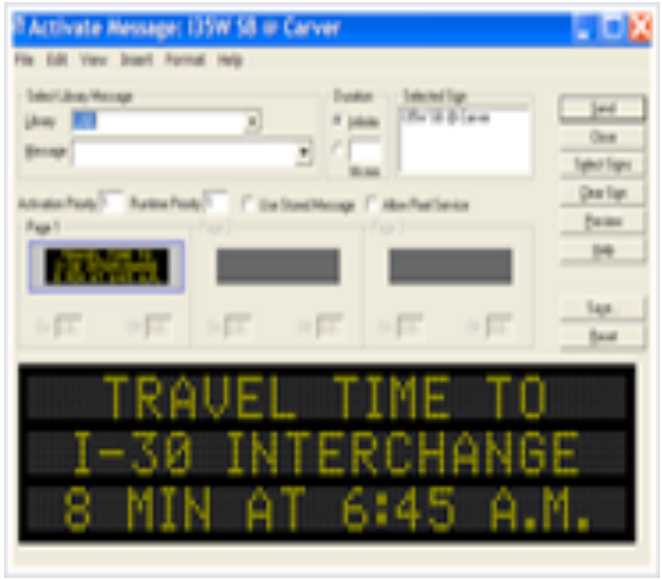

About OpenSSDPProject Data

- SSDP/UPnP is used to establish port forwarding on home routers
  - Think XBox- Live
- This service also exposes details of a home network

- HTTP/1.1 200 OK CACHE-CONTROL: max-age=1800 EXT: LOCATION: http://192.168.0.1:1900/rootDesc.xml SERVER: Ubuntu/7.10 UPnP/1.0 miniupnpd/1.0 ST: urn:schemas-upnp-org:service:Layer3Forwarding:1 USN: uuid:fc4ec57e-b051-11db-88f8-0060085db3f6::urn:schemas-upnp-org:service:Layer3Forwarding:1

- HTTP/1.1 200 OK CACHE-CONTROL: max-age=1800 DATE: Sat, 10 Jan 2015 00:00:02 GMT EXT: LOCATION: http://192.168.1.254:52869/gatedesc.xml SERVER: Linux/2.6.20-Amazon_SE, UPnP/1.0, Intel SDK for UPnP devices /1.2 ST: uuid:973fb8c8-d356-4e02-9093-3687a259f57e USN: uuid:973fb8c8-d356-4e02-9093-3687a259f57e

About Open.*Project

- Started as measurement for internal use
  - Ongoing attack measurement
  - What percentage were from known hosts
- Transformed into public facing data
  - Raw data provided to national CERTs
  - Public access to small data sets
  - ASN based reporting made available
  - Researchers have published papers from derived data

Thank you

Questions?

jmauch@us.ntt.net