

Incident analysis with RIPE NCC tools

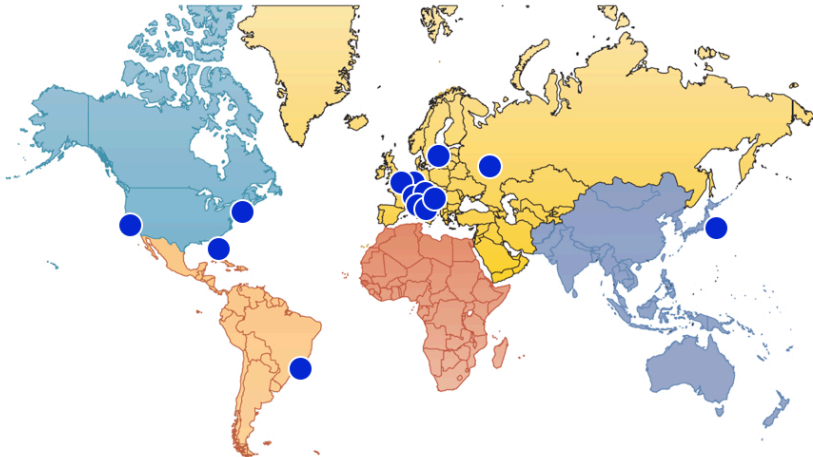
Analysing the RIS/Duke BGP incident

Erik Romijn <eromijn@ripe.net>
Senior Software Engineer



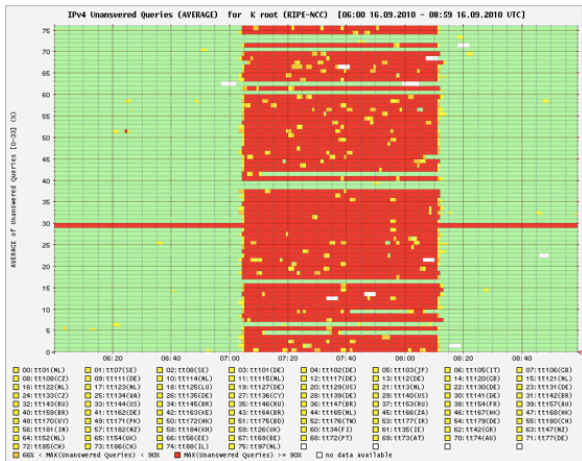
RIPE NCC information collection

- Routing Information Service (RIS)
 - Listens to and stores all BGP updates
 - Receiving data form 600 peers



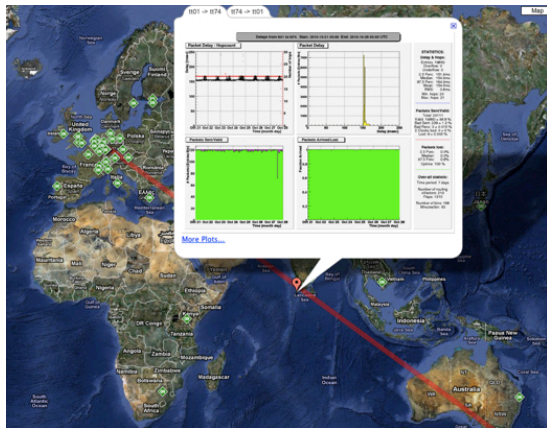
RIPE NCC information collection

- DNS monitoring service (DNSMON)
 - Monitors critical DNS infrastructure
 - About 100 vantage points worldwide



RIPE NCC information collection

- Test Traffic Measurements (TTM)
 - One-way latency/jitter/loss & traceroutes
 - About 100 nodes in full mesh



Case study: RIPE NCC / Duke University BGP experiment

RIS experiments & announcements

- RIS has a long tradition of supporting research
- Second AS in the world to announce 4-byte AS numbers
- Beacon prefixes from RIS available since 2002
 - Also a vital part of debogonizing


Case study: RIPE NCC BGP experiment

- RIPE NCC conducted an experiment on 27-08-2010
 - An optional BGP attribute was announced
 - This was a optional transitive attribute of 3000 bytes
 - The announcement was valid according to RFC4271
- Some routers corrupted the route and sent it
 - Peers who saw this dropped the session
- This caused disruption to some internet traffic

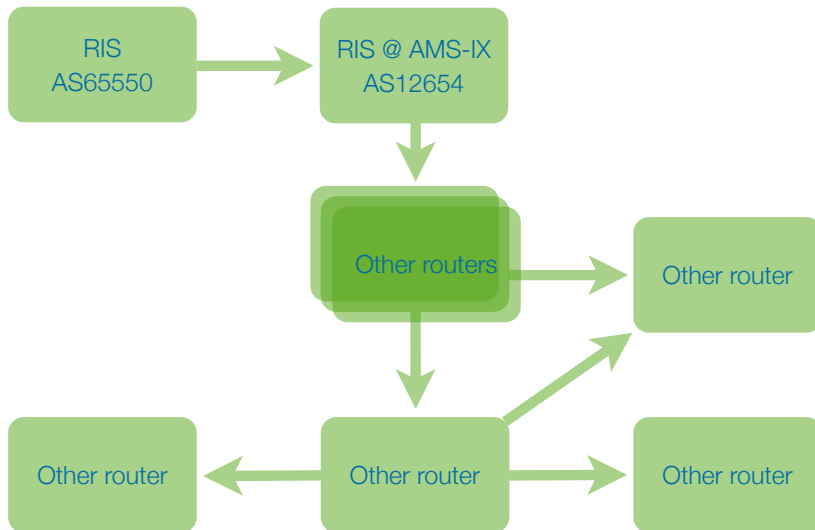
Case study: 27 August 2010

- Announcement active from 08:41 to 09:08 UTC, using 93.175.144.0/24
- We later observed some negative impact
- Immediately started an extensive investigation
 - This pointed towards a Cisco IOS XR bug
 - Sent out a very detailed private announcement
 - Also provided Cisco with all details
- Cisco released cisco-sa-20100827-bgp

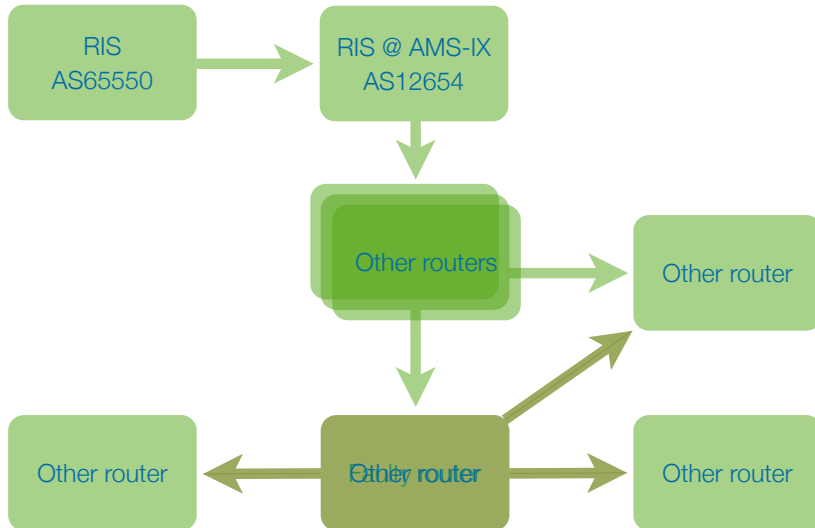
Propagation of the announcement



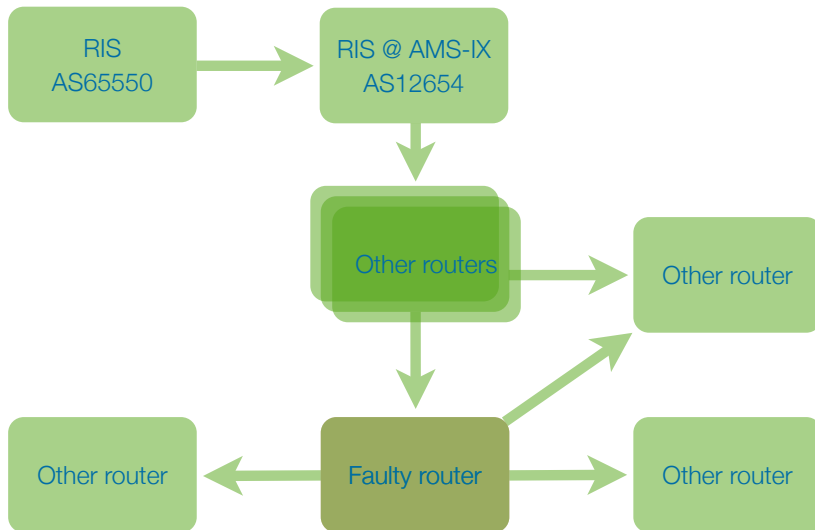
Propagation of the announcement



Propagation of the announcement



Propagation of the announcement



Goal of the experiment


- Research group from Duke University approached RIPE NCC to help
- Their goal was to measure support for long optional transitive attributes
 - Intended to be used for certificates for secure routing
- They did not have an AS number or addresses
- Provided RIPE NCC with a patched Quagga

Expected results


- A. The route propagates with the attribute intact
 - B. The route propagates, with some AS in the path removing the attribute
 - C. The route propagates, but takes a different path because some ASes drop the route
- A and B were seen in 4-byte AS number tests.

Impact of the experiment on the Internet


Unstable prefixes




E-mails per hour - 28-29 August



Unstable prefixes




Length of invisibilities



Critical DNS infrastructure (from DNSMON)



- Root servers unaffected
- 57% of TLDs unaffected
- Minor effects for 38% of the TLDs
 - Some dropped queries for one or two servers
- More significant effects on 5% of the TLDs

Critical DNS infrastructure



Visualization parameters



Prefix /  

Start date  Time  

End date  Time  


View

Cancel


View from a TTM probe in Prague, CZ

Delays from tt108 to tt01. Start: 2010-08-27 08:00 End: 2010-08-27 10:00 UTC

Packet Delay / Hopcount



Packet Delay



STATISTICS:


Delay & Hops:

Entries: 213
Overflow: 0
Underflow: 0
2.5 Perc: 16.0ms
Median: 16.2ms
97.5 Perc: 122.8ms
Mean: 33.6ms
RMS: 39.0ms
Min. hops: 14
Max. hops: 16


Packets Sent/Valid:

Total: 236
Valid: 213 = 90.3 %
Bad Sent: 1 = 0.42 %
Bad Recv: 1 = 0.42 %
2 Clocks bad: 0 = 0 %
Lost: 21 = 8.9 %

Packets Sent/Valid



Packets Arrived/Lost




Packets lost:

2.5 Perc: 0.0%
Median: 0.0%
97.5 Perc: 100.0%
Uptime: 100 %


Over-all statistic:

Time period: 0 day
Number of routing
vVectors: 19
Flaps: 19
Number of bins: 48
Minutes/bin: 2.5

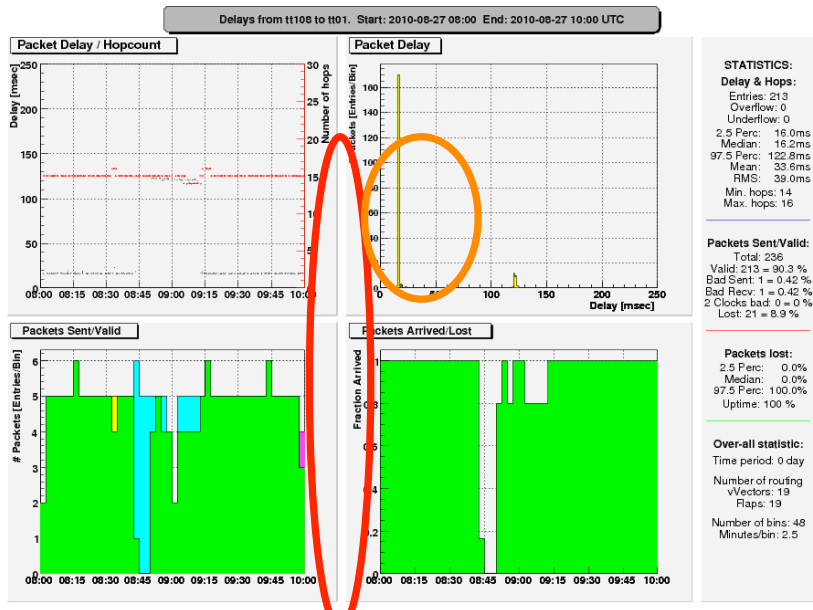
Updates for IPv4 vs IPv6



Unstable prefixes vs number of updates




View from a TTM probe in Prague, CZ



Locality of effects - updates


BGP Updates on all RIS locations (IPv4)



- LINX, London
- CIXP, Geneva
- DIX-IE, Tokyo
- MIX, Milan
- DE-CIX, Frankfurt
- PAIX, Palo Alto
- NOTA, Miami
- AMS-IX/NL-IX/GN-IX, Amsterdam
- VIX, Vienna
- Netnod, Stockholm
- NYIIX, New York
- MSK-IX, Moscow
- PTT, Sao Paulo

Locality of effects - withdrawals


BGP Withdrawals on all RIS locations (IPv4)




- | | |
|---------------------|---------------------------------|
| — LINX, London | — AMS-IX/NL-IX/GN-IX, Amsterdam |
| — CIXP, Geneva | — VIX, Vienna |
| — DIX-IE, Tokyo | — Netnod, Stockholm |
| — MIX, Milan | — NYIIX, New York |
| — DE-CIX, Frankfurt | — MSK-IX, Moscow |
| — PAIX, Palo Alto | — PTT, Sao Paulo |
| — NOTA, Miami | |

Locality of effects - vendors per IX


LINX




NYIIX




AMS-IX



DE-CIX



VIX



Lessons learned


- Future experiments should be pre-announced with sufficient lead time
- Detected vulnerabilities should be handled with more care
- More comprehensive impact assessments are needed
- Your input is welcome: <ris@ripe.net>

Questions?

Erik Romijn <eromijn@ripe.net>



Updates per prefix range



AS path length

