IPv6 Extension Headers on the Internet

Eric Vyncke, evyncke@cisco.com
Mehdi Kouhen, Polytechnique Paris,

Extension Header Security Policy for End User

- White list approach for your traffic
  - Only allow the REQUIRED extension headers (and types), for example:
    - Fragmentation header
    - Routing header type 2 & destination option (when using mobile IPv6)
    - IPsec ☺ AH and ESP
    - And layer 4: ICMPv6, UDP, TCP, GRE, ...

Source: Tony Webster, Flickr

Extension Header Loss over the Internet

- End users SHOULD filter packets with extension headers
- But, what are your ISP and its transit providers doing to your packets?

Source: Paul Townsend, Flickr

Previous Extension Headers Research by Others

- IETF-88, Nov-2013, fgont-iepg-ietf88-ipv6-frag-and-eh.pdf
  - *"Fragmentation and Extension Header Support in the IPv6 Internet"*
  - Single origin, destination = Alexa top web sites (883 unique addr)
  - Ext header size: 8 bytes and 1024 bytes; Failure rate: 45%
- IETF-89, with Tim Chown: 60% packet drops
- IETF-90, Jul-2014, iepg-ietf90-ipv6-ehs-in-the-real-world-v2.0.pdf
  - *"IPv6 Extension Headers in the Real World v2.0"*
  - Origin: RIPE Atlas probes, destination = Alexa again
  - Ext header size: 8, 256, 512 and 1024 bytes
  - Failure rate: between 60% and 90%
- December 2015, draft-ietf-v6ops-ipv6-ehs-in-real-world-02
  - Campaign in June 2015

Issues with Previous Experiments

- Destination: big web sites (Alexa)
- Destination drops are to be expected
- Not testing about Routing Header (for segment routing)

Methodology of our study

1. Determine a set of IPv6 addresses to test :
   - From Alexa's Top 1 Million list
   - From IPv6 BGP-advertised prefixes
2. TCP Traceroute without EHs :
   - Send v6 packets with TCP payload to port 80 of the destination with varying TTL => Routers in the path answer with ICMPv6 Time Exceeded
3. TCP Traceroute with EHs:
   - Same thing but adding an Extension Header before the TCP payload
4. Analysing the traceroutes

Step 1) Determining a set of IPv6 addresses to test

- From Alexa's Top 1 Million list :
  - Take those that have a AAAA record
  - … with a reachable IPv6 address in the AAAA record

- From BGP-advertised IPv6 prefixes
  - Address = [prefix]::1
  - Doesn't exist ? No problem, we are supposed to reach the AS -> Enough

2) TCP Traceroute with EHs

First, normal TCP traceroute without EH, then with EH

EH set :

- Destination Option Header
  16, 256, 512 bytes
- Hop-by-Hop Header
  16 bytes
- DO 16B + HbH 16B
- Routing Header type 4 (expected for SR)
- Fragment Header
  Normal and Atomic

EHs blocked by our ISP (so no result) :

- Hop-by-Hop Header
  256, 512 bytes
- Routing Header type 0 (deprecated)

Methodology of our study : Analysing the traceroutes

- Is it a problem ? Depends where it was dropped !
  - If dropped by the destination organization (host or same AS): Not a problem !
  - If dropped in transit: not cool…

- Where is the dropping node ?
  - If IP corresponds to some major IXPs, we look up the corresponding ASN by knowing the addressing logic, or in a database
  - Otherwise, normal Maxmind GeoIP ASN lookup

Question: publicity of active testing?

- Those test campaigns involve
  - Huge amount of DNS request
  - Active probing (sending more than 300 IPv6 packets per destination)
  - About one week of run
- Should we make those sweeping public?
  - To prevent being marked as 'hostile'
  - To prevent triggering some alerts in some places?

Results and analysis

- Drop rates depend on the Extension Header

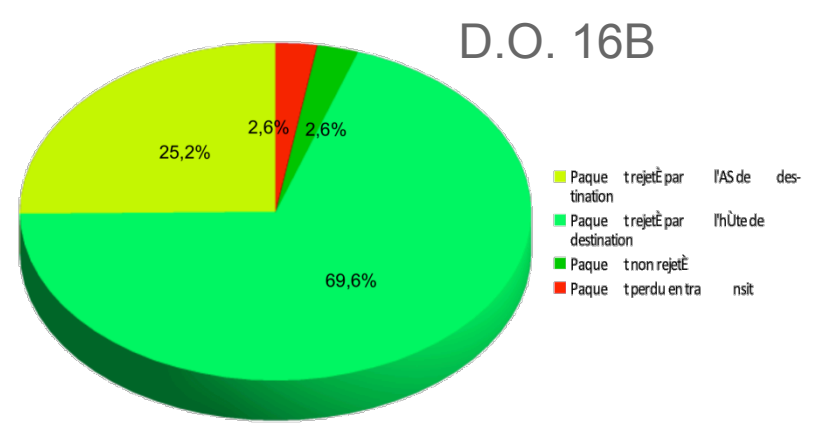

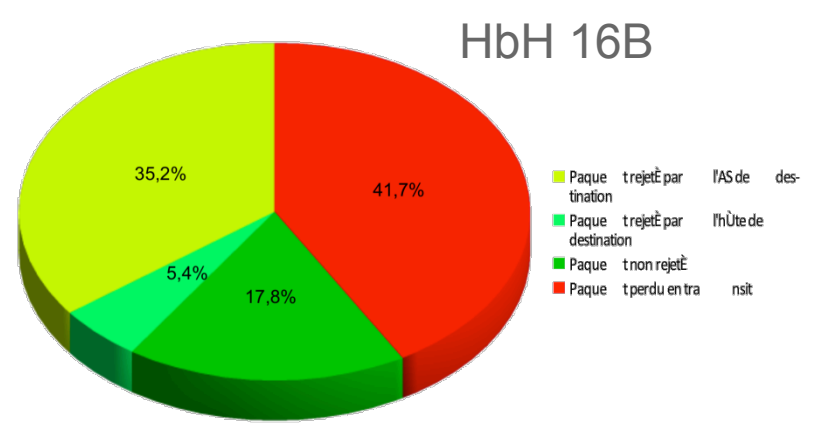

For Alexa, Spring 2015

Transit Drop rates in Spring 2015

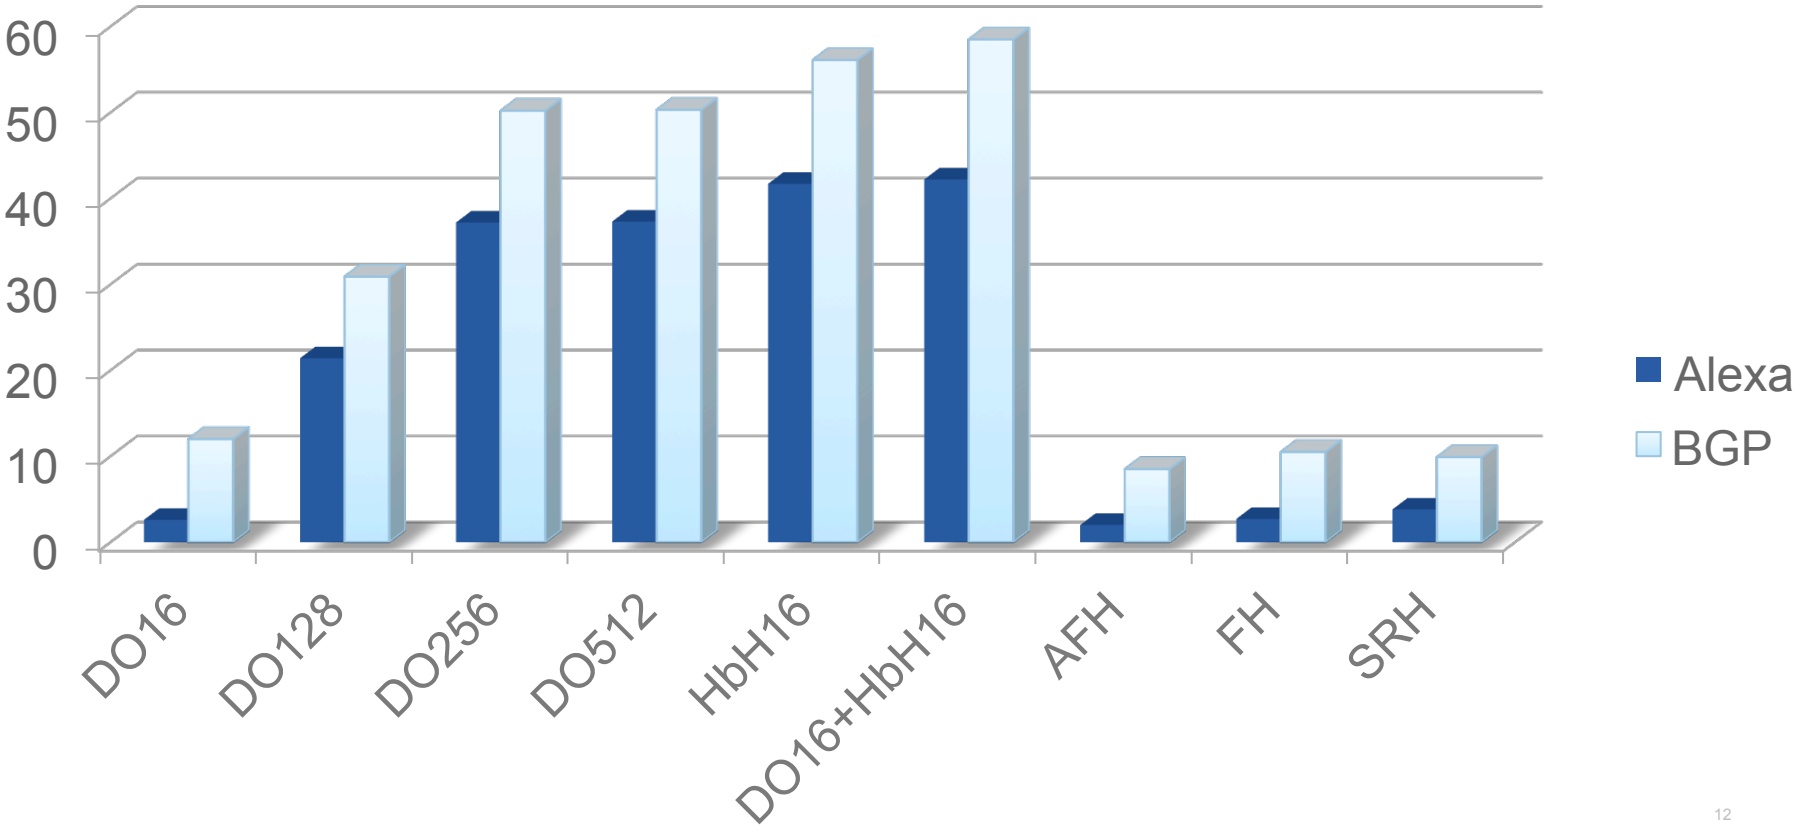

Things Keeps Improving Though

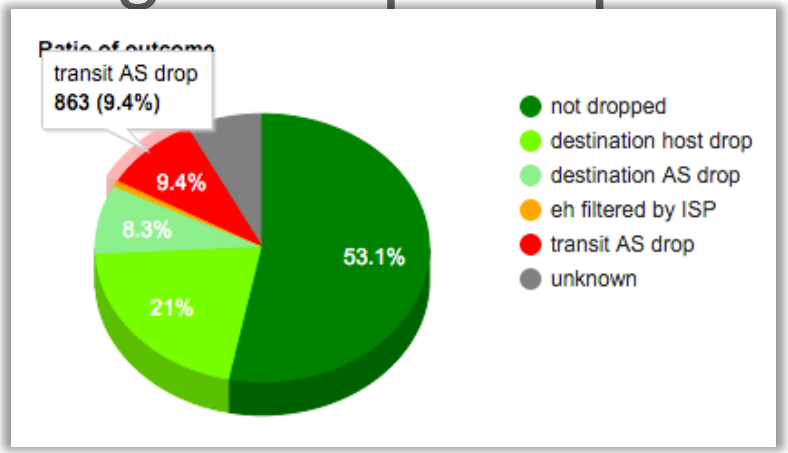

BGP in Spring 2015

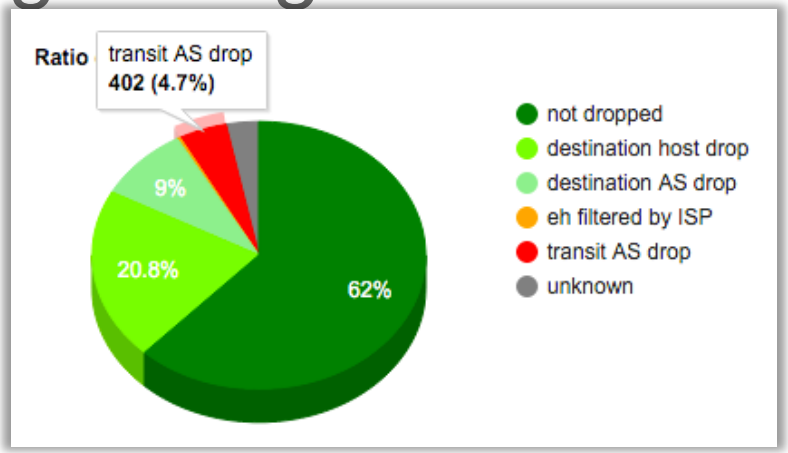

BGP in Spring 2016

- Current research by Polytechnique Paris (Mehdi Kouhen) and Cisco (Eric Vyncke)
  - And VM provided by Sander Steffann and Jan Zorz (Spring 2016)
- https://btv6.vyncke.org/exthdr/index.php?ds=bgp2016&t=fh (work in progress!)
- https://evyncke.go6lab.si/exthdr/index.php