IPv6/UDP Zero-Checksum

Magnus Westerlund

Gorry Fairhurst

draft-fairhurst-tsvwg-6man-udpzero-00

Overview

- Why is this being discussed?
- IPv6 UDP checksum
  - Delivery verification
- Tunnel scenario impact
- End-host impact
- Guessing the future
- Impact vs benefit analysis summary

NOPQR

rstuvwx
°±²³
ÇÈÉÊËÌÍÎÏÐ
Þßàáâã
÷øùúûü
đĒēĖėĘę
ĽľĿŀŁłŃńŅ
śŞşŠšŢţ
ŹźŻżŽžƒ

≥fifl

The proposal

- The fundamental proposal is to allow turning off the UDP checksum, i.e. set it to 0, when using IPv6:
  - At least for outer header in tunnels.
- Intended only for specific applications, especially tunneling usages.
- The proposal is a result of two IETF protocols under development:
  - **Automatic IP Multicast Without Explicit Tunnels (AMT)** (draft-ietf-mboned-auto-multicast)
  - **Locator/ID Separation Protocol (LISP)** draft-ietf-lisp
- Checksum change proposed in:
  - **draft-eubanks-chimento-6man-00**

AMT

- › Uses UDP tunnels between an AMT relay router and an AMT gateway
  - – AMT Gateway is either a site gateway router or host
- › UDP chosen for FW traversal
- › The issue is the encapsulated multicast data in UDP + AMT header
  - – Substantial amounts of data
  - – Some routers can't calculate a UDP checksum over a complete packet
    - › Don't have access to the complete packet when encapsulating

LISP

- Encapsulates any IP packet in an IP/UDP/LISP packet between the Ingress Tunnel Router (ITR) and Egress Tunnel Router (ETR).
- The ITR and ETR can be at different locations from site boundary to last hop routers.
- Reasons for using UDP :
  - To allow deployment on routers that can’t access the whole packet when doing encapsulation
  - Equal Cost Multi-Path (ECMP) operations
    - IPv6 Flow label is seen as difficult to use for this purpose
    - UDP ports are a part of the hash

NOPQR

rstuvwx
°±^{23}
ÇÈÉËÌÍÎÏÐ
Þßàáâã
÷øùúûü
đĒĒėĘę
ĽľŁłŃńŅ
śŞşŠšŢţ
ŹźŻżŽžƒ

≥ﬁﬂ

Commonality of LISP and AMT

- LISP and AMT are both tunneling mechanisms
  - Don’t require the UDP checksum to verify data corruption of inner packet, because that will be verified at delivery after de-capsulation

- IP in IP tunneling would work if not for the additional requirements:
  - ECMP
  - Firewall traversal
- Can generate high volumes of traffic

NOPQR

rstuvwx
°±²³
ÇÈËÌÍÎÏÐ
Þßàáâã
÷øùúûü
đĒĖėĘę
ĿŀŁłŃńŅ
śŞşŠšŢţ
ŹźŻżŽžƒ

≥ﬁﬂ

IPv4 vs IPv6

- RFC 2460, section 8 says:
  - Unlike IPv4, when UDP packets are originated by an IPv6 node, the UDP checksum **is not optional**. That is, whenever originating a UDP packet, an IPv6 node must compute a UDP checksum over the packet and the pseudo-header, and, if that computation yields a result of zero, it must be changed to hex FFFF for placement in the UDP header. IPv6 receivers must discard UDP packets containing a zero checksum, and should log the error.
- Using zero-checksum is allowed in v4, but not in v6:
  - The removed IP header checksum resulted in loss of
    - **delivery protection**, i.e. ensuring that it is delivered to the correct right destination address and with correct source address
    - **verification of next header field**
  - In v6, the above are verifed through the transport checksum pseudo header at the end of the delivery, rather than for each hop.
- Remember this is designed this way for a reason

Usages

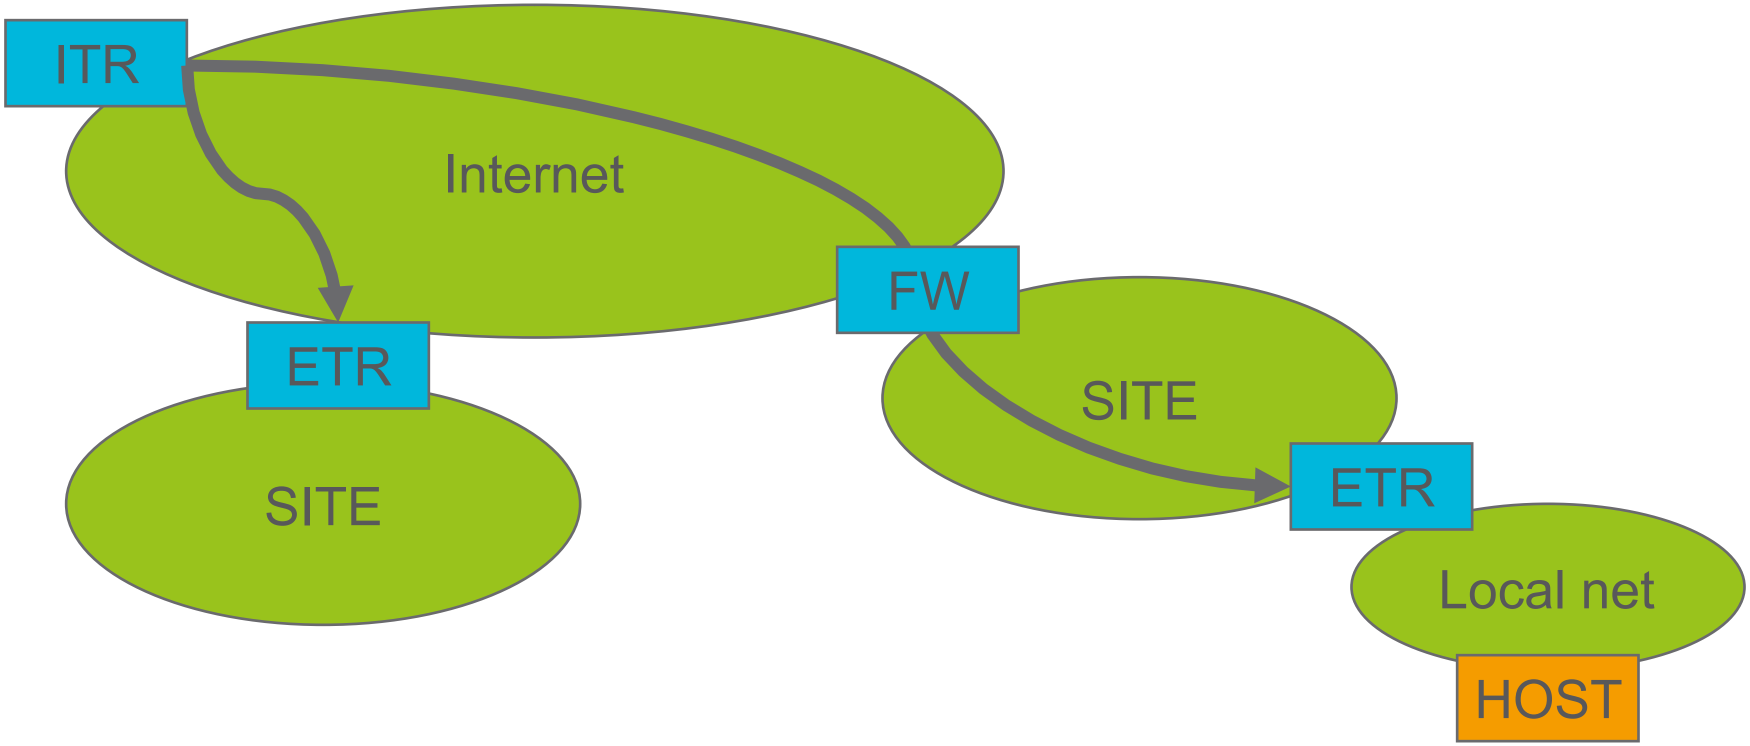

NOPQR

rstuvwx
°±²³
ÇÈËÌÍÎÏÐ
Þßàáâã
÷øùúûü
đĒĔĖėĘę
ĽľŁłŃńŅ
śŞşŠšŢţ
ŹźŻżŽžƒ

≥fifl

Tunnel USAGE Impact

- Uncertain that IPv6/UDP with zero checksum will be passed by firewalls:
  - Packet is not according to RFC2460 and may therefore be considered dangerous or a waste of bandwidth by middlebox
- Turning off the checksum in some host operating systems/ routers/CPEs is not possible or affects the whole system:
  - Margaret Wasserman said on LISP mailing list that this applies to major host operating systems and most checksum offloading hardware in hosts or CPEs.
  - Does not apply to all router cases, but the egress for some use cases may be CPE or end-user hosts

Tunnel Usage IMPACT

- Using zero checksum for the tunnel traffic:
  - – Impact on the encapsulated traffic, either:
    - Protected by inner transport header and thus verified at end-destination
    - Or in the same situtation as before.
- Corruption of outer IPv6 header in a packet in the tunnel has affect
  - – Corrupted destination delivers it to random host
  - – Corrupted source makes it look like it comes from a different source
    - Likely discarded as context will not be existing, but depends on implementation
  - – Corrupted next protocol header will lead this packet to be interpreted as something else
    - Also no difference from using checksum

Corruption

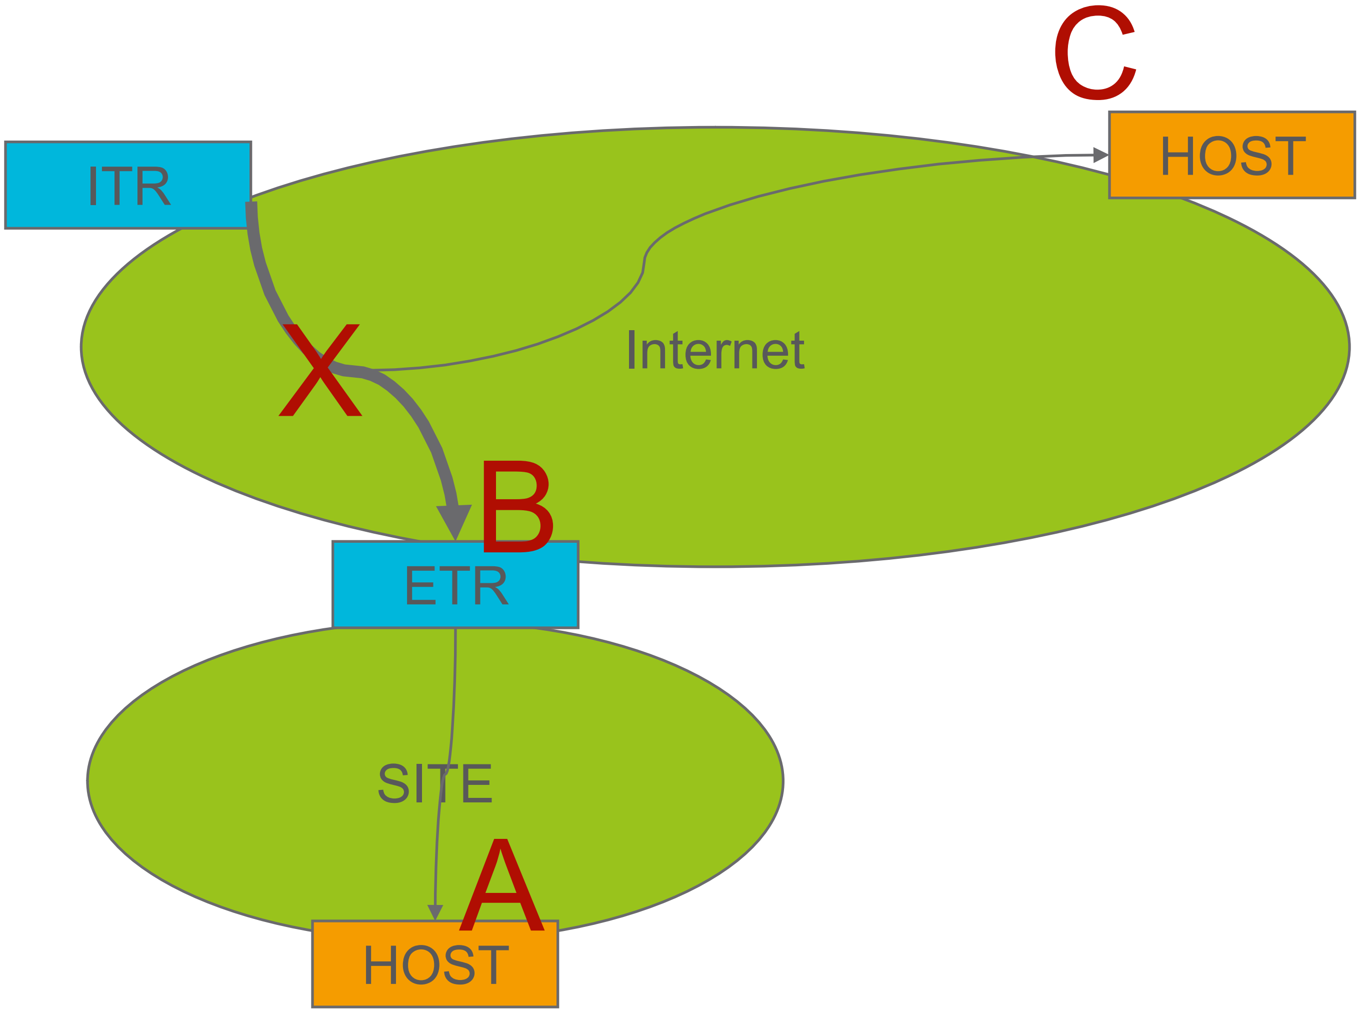

END HOST Impact

- A packet with a corrupted destination arrives at its new target
  - Where it is processed by the UDP stack:
    - This will likely drop it as it has an illegal checksum value
      - Assuming an unchanged host.
    - If the IP and UDP layer isn’t well integrated or the receving host has been changed, it will be forwarded to application
    - Depending on application, possibly may determine this as corrupt data it will (or will not) process.
    - Depending on application, may also modify/create protocol state.
- A host that turns off checksum as a result of allowing this:
  - Has lost its delivery protection
  - Will be 32000 times more likely to get unintended packets delivered to applications

END Host IMPACT

- Without a transport checksum any IPv6 extension header, such as fragmentation will not be protected against corruption
  - Increases the risk of erronous re-assembly of fragmented packets

NOPQR

rstuvwx
°±²³
ÇÈÉÊËÌÍÎÏÐ
Þßàáâã
÷øùúûü
đĒĒĖėĘę
ĽľŁłŃńŅ
śŞşŠšŢţ
ŹźŻżŽžƒ

≥ﬁﬂ

Summary Pro and CONS

- Using UDP with zero checksum does not always seem to meet goals:
  - Yes, gets ECMP to work
  - May, get you through firewalls
  - Does restrict the deployability to systems that can be changed
- Has impact on other systems and applications
  - Reduced delivery protection capabilities
  - Especially if this gets deployed for other applications
    - Not comparable with IPv4/UDP without checksum usage
- Zero-checksum are not catch all solution for LISP and AMT
  - Consider other solutions and different tradeoff?
    - IP in IP tunnels

What to do

1. Propose to **NOT** perform any checksum rule change in RFC 2460.
2. Propose that someone takes on 6man work to clarify usage of flow label so it can be used in ECMP hashes, etc, to remove one of the reasons for using UDP.

NOPQR

rstuvwx
°±²³
ÇÈÉËÌÍÎÏÐ
Þßàáâã
÷øùúûü
đĒĖėĘę
ĽľŁłŃńŅ
śŞşŠšŢţ
ŹźŻżŽž*ƒ*

≥fifl

Further Reading

- The IPv6 UDP Checksum Considerations [draft-fairhurst-tsvwg-6man-udpzero-00.txt]
- "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998
- "The Lightweight User Datagram Protocol (UDP-Lite)", RFC 3828, July 2004.
- "Unicast UDP Usage Guidelines for Application Designers", BCP 145, RFC 5405, November 2008.
- "Automatic IP Multicast Without Explicit Tunnels (AMT)", [draft-ietf-mboned-auto-multicast-09]
- "Locator/ID Separation Protocol (LISP)", [draft-farinacci-lisp-12.txt]
- "UDP Checksums for Tunneled Packets", [draft-eubanks-chimento-6man-00]
- "The UDP Tunnel Transport mode", [draft-fairhurst-6man-tsvwg-udptt-02]

ERICSSON