UDP Checksum for Tunneled Packets

6Man IETF-79

Marshall Eubanks

marshall.eubanks@iformata.com

Phil Chimento

Philip.Chimento@jhuapl.edu

Why do we want to change IPv6 UDP?

- In IPv4, UDP checksums are not required.
- In IPv6, they are. RFC 2460 says
  - `Unlike IPv4, when UDP packets are originated by an IPv6 node, the UDP checksum is not optional.`
- This was done because the IPv6 IP header does not include a checksum.
  - And ***that*** was done to improve router efficiency.
- Why do we want to change this ?
  - To improve router efficiency.

UDP and tunnels

- Consider this from RFC2460.

  `IPv6 receivers must discard UDP packets containing a zero checksum, and should log the error.`

  – Why ? Because this is the only check on the validity of IPv6 UDP packets, including the IP pseudo header (and thus the source and destination IP addresses).

- That's fine, but what about tunneling ?
  – Tunneling is becoming increasingly common to implement changes in Internet Protocols
    - Tunnel protocols increasingly use UDP to get through firewalls
  – This is a real world issue - if you want to make fundamental changes in the way the Internet work, in today's network you are likely to need to tunnel it.

Tunnels have Inner versus Outer checksums

- The difference between UDP tunnel packets and “regular” UDP packets is that tunnel packets encapsulate other packets.
- There is the “outer” packet header and checksum (for the encapsulating UDP packet) but also generally an “inner” packet header which also generally contains a checksum as well.
  - In other words, these packets have 2 headers and 2 checksums.
- We want to take advantage of this to improve the efficiency of tunnel protocols by relaxing RFC 2460 to allow for a zero checksum on the outer packet.
  - This is based on the conclusions of the consideration document
    - draft-ietf-6man-udpzero-02
- In other words, we propose allowing for packets with 2 headers, but only 1 checksum
- **This will be the first fundamental change to IPv6 / RFP2460 since its creation**.
  - We feel it is necessary, will be good for IPv6, and can be done safely.

What Protocols Need this ?

- We first became aware of this issue because of AMT, Automatic Multicast without Explicit Tunnels
  - draft-ietf-mboned-auto-multicast-10
- The Locator/ID Separation Protocol (LISP) has the same issue.
  - `draft-ietf-lisp-09`
- L2TP, Softwires and other efforts also tunnel over UDP.
- The state of the Internet means that there are highly likely to be more.
- Let's look briefly at AMT as an example.

AMT

- AMT uses tunneling to extend the multicast Internet to remote domains.
  - Could be just one node, or an entire network.

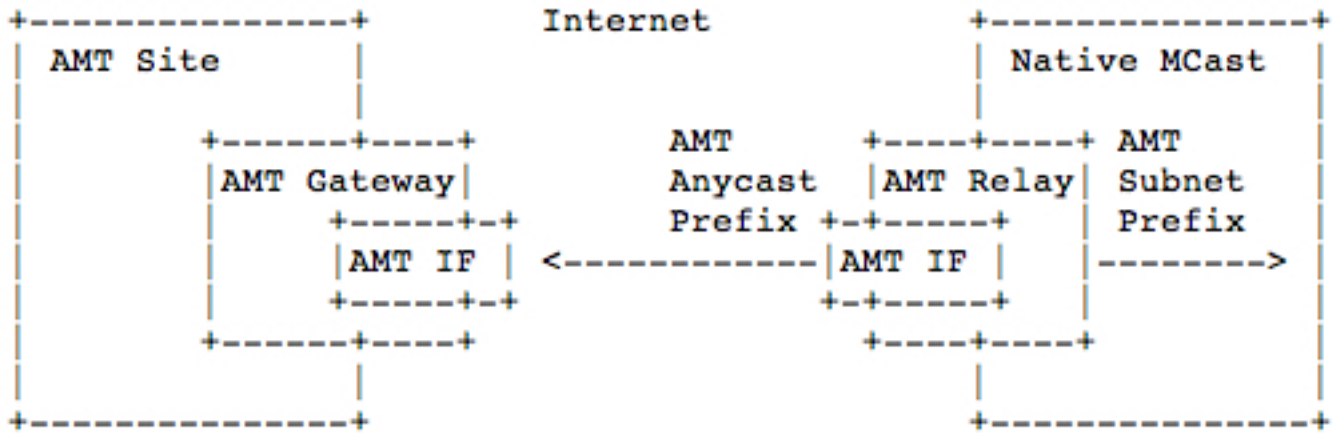

AMT tunnels

- In AMT a relay takes a multicast packet (the “inner” packet), encapsulates it in UDP (creating the “outer” packet), and unicasts it to an AMT gateway, there to be de-encapsulated and placed on the local network.
  - These tunnels will generally not terminate in end-nodes, but inside the network.
  - The relay and gateway are likely to be routers, and the desire is to have these devices handle very high rate video.
  - Having these routers deal with UDP checksums is a big hit in their efficiency.
- We are referring to the data packet - other AMT packets, with checksums, deal with command and control

Potential Problems

- We proposed this back in IETF 74, and it engendered a lot of discussion
  - As is appropriate.
- As a result, Gorry and Magnus wrote a considerations draft, which also engendered a lot of discussion.
- They conclude that this is the best way to do this, and provide a list of considerations, which we have adopted.

Considerations

1. IPv6 protocol stack implementations should not by default allow the new method. The default node receiver behaviour must discard all IPv6 packets carrying UDP packets with a zero checksum.
2. Implementations must provide a way to signal the set of ports that will be enabled to receive UDP datagrams with a zero checksum. An IPv6 node that enables reception of UDP packets with a zero-checksum, must enable this only for a specific port or port-range. This may be implemented via a socket API call, or similar mechanism.
3. RFC 2460 specifies that IPv6 nodes should log UDP datagrams with a zero-checksum. This should remain the case for any datagram received on a port that does not explicitly enable zero-checksum processing. A port for which zero-checksum has been enabled must not log the datagram.
4. A stack may separately identify UDP datagrams that are discarded with a zero checksum. It should not add these to the standard log, since the endpoint has not been verified.
5. Tunnels that encapsulate IP may rely on the inner packet integrity checks provided that the tunnel will not significantly increase the rate of corruption of the inner IP packet. If a significantly increased corruption rate can occur, then the tunnel must provide an additional integrity verification mechanism. An integrity mechanisms is always recommended at the tunnel layer to ensure that corruption rates of the inner most packet are not increased.

Considerations

6. Tunnels that encapsulate Non-IP packets must have a CRC or other mechanism for checking packet integrity, unless the Non-IP packet specifically is designed for transmission over lower layers that do not provide any packet integrity guarantee. In particular, the application must be designed so that corruption of this information does not result in accumulated state or incorrect processing of a tunneled payload.
7. UDP applications that support use of a zero-checksum, should not rely upon correct reception of the IP and UDP protocol information (including the length of the packet) when decoding and processing the packet payload. In particular, the application must be designed so that corruption of this information does not result in accumulated state or incorrect processing of a tunneled payload.
8. If a method proposes recursive tunnels, it needs to provide guidance that is appropriate for all use-cases. Restrictions may be needed to the use of a tunnel encapsulations and the use of recursive tunnels (e.g. Necessary when the endpoint is not verified).
9. IPv6 nodes that receive ICMPv6 messages that refer to packets with a zero UDP checksum must provide appropriate checks concerning the consistency of the reported packet to verify that the reported packet actually originated from the node, before acting upon the information (e.g. validating the address and port numbers in the ICMPv6 message body).

In conclusion, a “modest proposal”

- We propose that the checksum be not required :
  - On the “Outer” UDP packet header of encapsulation protocols with complete “inner packets.”
  - The Inner packet MUST have a checksum.
- As RFC 2460 is a standards track document, this needs to be one as well.
- We would like this to be adopted by the Working Group.

Questions ?
Comments ?
Rotten fruit ?