

Christian Hopps  
LabN Consulting, LLC

# IP Traffic Flow Security

## Improving IPsec Traffic Flow Confidentiality

IETF 112 – “draft-ietf-ipsecme-iptfs-11”

# 2021 Recap

- WGLC Competed Feb 2021
  - Doc updated with received comments Feb 22 (-07)
- Post WGLC comments
  - Mar 30: Doc updated (-08) revised language (IPTFS->AGGFRAG) from Valery
  - Apr 5: Draft Write-Up submitted to Shepherd/Chairs
  - July 5: Doc updated (-09) clarifying that reorder window should be small, and should NOT force the replay window to be small as well
  - Sep 3: Doc updated (-10) recommending use of drop timer instead of reorder window to avoid long delays
    - Intending to address important comment from Tero
  - Oct 24: Doc updated (-11) took a guess at text Tero would accept WRT optionally sending immediately out-of-order
  - Oct 31: Text from Tero – one last outstanding issue based on this text

# Last Issue To Resolve

- Update -11 added text saying the receiver **MAY** optionally send whole inner packets on receipt w/o waiting for earlier misordered tunnel packets to arrive.
- Tero's alternate text has same mechanism *but* changes it to "**SHOULD**", restoring the original in-order delivery as a **MAY**.

# New text in -11

“As an optional optimization (e.g., to handle very lossy and/or reordered tunnel paths), the receiver **MAY** transmit any fully formed inner packets contained within the AGGFRAG\_PAYLOADs prior to re-ordering the outer packets.”

# Proposed Tero Text

The receiver **SHOULD** process incoming AGGFRAG\_PAYLOAD payloads as soon as they arrive as much as it can. I.e., if the incoming AGGFRAG\_PAYLOAD packet contains complete inner packet(s), receiver should extract them and forward them immediately. For partial packets the receiver needs to keep the partial packets in the memory until they fall out from the reordering window, or until the missing parts of the packets is received, in which case it will reassemble them and send them out. If AGGFRAG\_PAYLOAD payload contains multiple packets they **SHOULD** be sent out in the order they are in the AGGFRAG\_PAYLOAD (i.e., keep the original order they were received on the other end).

... [reworded original text]

# Counter and Compromise Proposal

- Lou Berger suggested on list, swapping SHOULD/MAYs
  - In-order delivery (which might incur a small delay) remains recommended  
“FWIW I'm basing my comments on my routing area experience where a huge amount of work has been put into maintaining ordering experienced by user traffic at significant implementation expense, i.e., in support of ECMP and other multipath solutions in protocols and hardware.”
  - Out-of-order delivery still allowed
    - I.e., adopt Tero's text *but* keep original as the recommended behavior
- Lou's mail also OK with both **MAYs** with a configuration selection

# Issue with Send Immediately

- Amplifies end-user experienced misordering
- Routers are built to not introduce misordering or bizarre delays in packet flows


# Why reordering outer packets “Just Works”

- Operationally significant delays unlikely from misordering
  - At high send rate (e.g., line rate, no send gap)
    - A reasonable reordering window won't introduce unreasonable delay to correct ordering
  - At lower send rate (wide sending gap)
    - Misordered slots are simply dropped
    - Drop timer limits any delay due to these drops





# Ordered Outer



# Misordered Outer



Send Immediately

many user packets  
misordered



Send In-Order

Small Delay → ←



# For Discussion: MAY vs SHOULD

The receiver **MAY|SHOULD** process incoming AGGFRAG\_PAYLOAD payloads as soon as they arrive as much as it can. I.e., if the incoming AGGFRAG\_PAYLOAD packet contains complete inner packet(s), receiver should extract them and forward them immediately. For partial packets the receiver needs to keep the partial packets in the memory until they fall out from the reordering window, or until the missing parts of the packets is received, in which case it will reassemble them and send them out. If AGGFRAG\_PAYLOAD payload contains multiple packets they SHOULD be sent out in the order they are in the AGGFRAG\_PAYLOAD (i.e., keep the original order they were received on the other end).

Instead of the method described in the previous paragraph the receiver **SHOULD|MAY** reorder out-of-order AGGFRAG\_PAYLOAD payloads received into in-sequence-order AGGFRAG\_PAYLOAD payloads (Section 2.2.3), and only after it has in-order AGGFRAG\_PAYLOAD payload stream, receiver extracts the inner-packets. In this case the receiver considers a packet lost when *the drop timer expires* or its sequence number is abandoned (e.g., pushed out of the re-ordering window, ~~or timed-out~~) by the reordering algorithm. Using this method will make sure the packets are sent in-order, i.e., there is no reordering possible, but the cost is that any lost packet will cause delay of *the drop timer interval full reorder window*, and there will be extra burstiness in the output stream (when lost packet is dropped out from the re-order window, all outer packets received after that are then immediately processed, and sent out back to back).

# Next Steps

- Publish document based on today's discussion/resolution
- No other issues
- Submit to IESG for publication