

IPv6 Packet Truncation

<draft-leddy-6man-truncate-03>



I E T F[®]

John Leddy, Comcast
Ron Bonica, Juniper Networks

IETF102 Montreal

Introduction



- IPv6 allows fragmentation at the source only
- Source nodes can either:
 - Refrain from sending packets whose length is greater than the IPv6 minimum MTU (1280 bytes)
 - Maintain a running estimate of the PMTU and refrain from sending packets whose length exceeds that estimate
- PMTU estimation tools
 - Path MTU Discovery (PMTUD)
 - Packetization Layer PMTUD (PLPMTUD)

How PMTUD Works



- Source produces initial PMTU estimate
 - MTU of first hop along the path to the destination
 - May be larger than actual PMTU
- Source sends packets
 - Never larger than PMTU estimate
 - Possibly larger than the actual PMTU
- When packet is larger than actual PMTU
 - Intermediate node discards the packet
 - Intermediate node sends an ICMP PTB to the source
- Source updates PMTU estimate

How PMTUD Breaks



- Network cannot deliver ICMP PTB message from intermediate node
 - Poorly implemented firewall filters
 - Example: The common home router bug
 - Intermediate node does not have a viable route to source
 - Anycast issues
- Source doesn't update PMTU estimate
- Persistent black holes

New Procedures Required



- Upper-layer protocol marks selected packets as being Truncation Eligible
- When an intermediate node cannot forward a Truncation Eligible packet because of MTU
 - Truncates the packet
 - Marks the packet as being Truncated
 - Forwards the packet to its destination
- Destination node sends ICMP feedback to the source node

New IPv6 Destination Options



- Truncation Eligible option
 - Indicates that a packet is eligible for truncation
- Truncated Packet option
 - Indicates that a packet has been truncated and is eligible for further truncation

Truncation Eligible Option



- Option Type – Value TBD by IANA
 - Act bits specify destination behavior when the option is not recognized
 - Skip over the option and continue processing the packet
 - Value - 00
 - Chg bits indicate whether Option Data can be modified on route
 - Value not consequential. TBD by IANA.
- Opt Data Len **MUST** be equal to 0

Truncated Packet Option



- Option Type – Value TBD by IANA
 - Act bits specify destination behavior when the option is not recognized
 - Discard packet and send ICMP Parameter Problem message to source node
 - Value - 10
 - Chg bits indicate whether Option Data can be modified on route
 - Value not consequential. TBD by IANA.
- Opt Data Len **MUST** be equal to 0

Source Node Procedures



- Upper-layer protocol marks packets as being Truncation Eligible
 - Truncation Eligible option

Intermediate Node Procedures



- Intermediate node attempts to forward the packet and encounters an MTU issue
 - If the packet is eligible for truncation
 - Truncate the packet
 - Update the IPv6 Payload Length
 - Overwrite all instances of the Truncation Eligible with the Packet Truncated option
 - Forward the packet to its destination
 - If the packet is not eligible for truncation
 - Discard Packet and send ICMP PTB to source

Destination Node Procedures



- Destination Node receives a packet with Truncation Eligible option
 - Skip over the Truncation Eligible option and continue to process the packet
 - Do this regardless of whether the destination node recognizes the Truncation Eligible option
 - Deliver the packet to an upper-layer protocol

Destination Node Procedures (continued)



- Destination Node receives packet with Truncated Packet option and recognizes it
 - Send an ICMP PTB to the source node
 - ICMP PTB MTU field reflects packet length
 - By default, discard the packet
 - However, upper-layer protocols can register for delivery of truncated packets

Destination Node Procedures (continued)



- Destination Node receives packet with Truncated Packet option and does not recognize it
 - Send an ICMP Parameter Problem message to the source node
 - ICMP Parameter Problem message contains initial bytes of the original packet, including length
 - Discard the packet
 - Regardless of whether the upper-layer protocol has registered for delivery of truncated packets

Benefits of Truncation



- Use Larger MTU's
 - Fix PMTUD
 - Provide tools that upper-layer protocols can leverage so that they can more intelligently negotiate maximum packet sizes

Fixing PMTUD



- ICMP PTB messages from the destination node are more likely to reach the source than ICMP PTB messages from an intermediate node
- By emitting a single Truncation Eligible packet, the source node can learn the PMTU to a destination
 - This is because a packet can be truncated by multiple nodes along the path to the destination

Tools For Upper-Layers



- Upper-layer protocols can use metadata from truncated packets to infer PMTU information and to negotiate maximum segment sizes
- Upper-layer protocols can use packets that are Truncation Eligible as probes.

Next Steps



- Present to TSVWG
 - Consider requests for additional option data (e.g., pre-truncation length)
- Adopt as 6man WG item



QUESTIONS / COMMENTS?