



IPv6 Hop-by-Hop Options Processing Procedures

<draft-hinden-6man-hbh-processing-00>

Bob Hinden
Gorry Fairhurst

March 2021
IETF110

Introduction



- Hop-by-Hop Options are not working End-to-End in the Internet
 - Very common to drop packets with HBH Option headers
 - Issue is more severe at the edge
- We need to do something different if we expect to use HBH Options in the future
 - This is a proposal to try something different with Hop-by-Hop Options

Terminology



- Fast Path
 - Hardware, NP, or ASIC packet processing
 - Usual router processing for most packets
 - Also called the “Forwarding Plane”
- Slow Path
 - Software packet processing
 - Router path for special processing
 - Also called “Control Plane”

Background



- HBH Processing in first IPv6 specification was required for all nodes, issues were:
 - Inability to process at wire speed in hardware
 - Packets with HBH options sent to "Slow Path" would degrade routers performance and could be used as a DOS attack
 - Packets could contain multiple HBH options making the problem worse

Background (continued)



- Current IPv6 Specification (RFC8200)
 - Changed requirement to only require HBH processing if router configured to do so.
 - This essentially documented current operational behavior.
- The issues discussed here is focused on routers with specialized hardware forwarding.

Motivation



- HBH Options not practical to be used widely
 - Common to drop all packets with HBH options
 - Multiple HBH options in same packet make problem worse
 - Any mechanism that can be used externally to force packets into the “Slow Path” can be exploited as a DOS attack
- Goal is to redefine HBH procedures to make HBH options practical

Proposed HBH Processing Procedures (1)



- RFC8200 requires HBH Option identified by Next Header value of 0 and must be immediately after IPv6 header
- Change to only allow one HBH option in HBH Header
 - Requires that all HBH options be 8-octet aligned
- Nodes **MUST** discard packets with more than one HBH option

Proposed HBH Processing Procedures (2)



- Change that IPv6 Nodes MUST only process HBH options header if it can be done in the “Fast Path”
 - *Otherwise skip the HBH Options header*
- RFC8200 Section 4.2 defines option type field that controls how the option is processed if not recognized. Change to
 - 00 - skip over this option and continue processing the header.
 - 01 - discard the packet.
 - 10 - discard the packet.
 - 11 - discard the packet.
- *No longer requires sending ICMP error messages*

Configuration



- RFC8200 allows router to control HBH processing by local configuration
- Suggest that routers maintain a lookup table on Option Types that are supported in the "Fast Path"
- Would allow for quick determination if an options is supported
 - If not supported, follow bits in Option Type field

Header Alignment



- HBH Options that are 8-octet aligned
 - Jumbo Payload [RFC2675]
 - Path MTU Record Option [I-D.ietf-6man-mtu-option]
 - RPL Option [I-D.ietf-roll-useofrplinfo]
 - Quick-Start [RFC4782]
 - CALIPSO [RFC5570]
 - SMF_DPD [RFC6621]
 - ILNP Nonce [RFC6744]
 - MPL Option [RFC7731]
- HBH Options that must be deprecated or modified
 - Router Alert [RFC2711]
 - IP_DFF [RFC6971]
 - IOAM [I-D.ietf-ippm-ioam-ipv6-options]

New HBH Options Requirements



- Any new HBH Options defined in the future MUST have 8-octet alignment
- Good to include when they should be sent

Issue Raised - General



- Allow HBH Header elsewhere in the packet
 - This would make it harder to process on “Fast Path”
- Reduces flexibility with only one HBH Option and “Fast Path” only
 - Current flexibility makes it unusable in the Internet
- Restricting HBH Options isn't the right approach
 - We need to do something different if we expect it to work in the future
- Must process in the “Fast Path” too restrictive
 - Could be made a SHOULD
 - How to avoid “Slow Path” being used as a DOS Attack?

Issues Raised

Number of HBH Options



- Only one HBH Option per packet too restrictive
 - It does restrict usage, but different HBH options could be in different packets. For example, the Path MTU option doesn't need to be in all packets in a transport connection.
- Would a limit of 2, 3, 4, etc. be better?
 - Given it doesn't work E-E now, how much flexibility do we need?
 - How would the hosts learn how many to send?
 - Supporting more than one may require the PAD options, could be avoided by 8-octet alignment requirement
- If more than one, should number of HBH Options be global or per node limit?
 - Hosts need to know what the limits are
 - Global may be better

Issues Raised - Other



- Better to Keep ICMP Error messages?
 - Goal is to simplify, is this practical at speed?
- Relationship to SRH
 - SRH is a routing header, does not have HBH properties, processed by node in Destination Address
- Will forwarding packets with unknown HBH Options work at the edge?
- Is operational deployment possible, requires changes in all nodes to be useful?
 - Clearly will take time, will require support from router vendors and operators

Next Steps



- Thanks for all the feedback and editorial comments!
- We need to hear from router vendors and operators
- Ready to Adopt as 6MAN w.g. draft?



QUESTIONS / COMMENTS?