I E T F®

IPv6 Hop-by-Hop Options Processing Procedures

<draft-ietf-6man-hbh-processing-01>

Bob Hinden
Gorry Fairhurst

March 2022
IETF113

Introduction

- Hop-by-Hop Options are not working in the Internet:
  - Very common for routers on a path to drop packets with HBH Option headers.
  - RFC8200 documented current practice by saying only required if router configured.
  - We need to do something different if we expect to use HBH Options in the future.
- This is a proposal to modify Hop-by-Hop Option Processing.
- Document adopted by 6MAN w.g. 22 January 2022.

Changes from -01 individual draft

- Reworked text to talk about processing HBH options at full forwarding rates, instead of "fast path"
- Revised Section 6 "New Hop-by-Hop Options" to allow variable sized HBH options, remove specific length requirements, and other clarifications.
- Editorial changes.

Issue Tracker Setup

- Captured Issues raised in Adoption Call

  https://github.com/ietf-6man/hbh-processing/issues

- Remainder of presentation is to discuss issues

Issues (short)

- Issue #3 - Does the WG have ASIC experience with EH?
  - Most issue text is about <draft-ietf-6man-enhanced-vpn-vtn-id>
- Issue # 7 – Cite RFC7827
  - Fix in next version
- Issue # 14 - Are HBH options of size 8B? or EH of size 8B?
  - Fixed in -00 draft
- Issue #16 - Better to encode HBH processing in Destination Address
  - Propose to close, major change to IPv6 architecture and waste of address space
- Issue #19 - Section 5.2 last para s/patch/path/
  - Fix in next version

Issue #2 - Leading edge line-speed routers might ignore HBH Extension Headers

- I can't see any vendor of really high speed routers implementing that as a default behavior. Why would they even look beyond the Destination Address? They wouldn't even have an option to process or not process HBH.
- Editors: Much might depend on the set of HBH options we discuss and how important these are for these routers. The meaning of "MUST" or "SHOULD" needs carefull thought with respect to this PS and how this will relate to the IS for IPv6.

Issue #4 - RFC9098 notes that nodes do need to process payload

- Misses the discussion in RFC9098: quite often forwarding nodes do need to process IPv6 payloads.

- Are HBH Options considered payloads?

- Editors: Add text about RFC9098

Issue #5 - Definition of Fast Path and Slow Path?

- Comment Proposed:
  - Fast path: A path through a router that is optimized for forwarding packets only based on the network layer headers. The current text "without processing their payloads" may not be clear about which layer information would be used.
  - Slow path A path through a router that can process both the network layer headers and the payloads.
- Similar issue raised in Issue #10
- Editors: Current draft moves toward discussing full forwarding rate, the distinction proposed above is helpful.

Issue #6 - Did RFC8200 make a difference to deployed practice?

- Has enough time elapsed to be able to tell if the RFC8200 made a difference or not?

- Given the constraints on moving IPv6 to “Standard”, it was not intended to change behavior.
- Editors: Text in Section 4 will be improved to reflect this.

Issue #8 - Lack of graceful handling of malformed EHs

- One of the main security problems associated with IPv6 EHs is that too many implementations seem fail to gracefully handle malformed/malicious Ehs.

- Is this different from other headers (IPv6, EH, Transport)?
- Editors: No change.

Issue #9 - Should we deprecate Router Alert?

- There was debate on whether to deprecate Router Alert Option at the level of standards. I'm on the side of supporting deprecation, because IMHO the primary purpose of HBH drafts is to push forward the implementation and deployment of HBH Options.

- Editors: In scope for this draft?
  - See: draft-bonica-6man-deprecate-router-alert-00.txt

Issue #10 - distinction between the Slow and Fast paths is platform-specific

- Overlap with Issue #5
- The distinction between the Slow and Fast paths is platform-specific and is an ever-moving target.

- Current draft talks about Full Forwarding rates
- Editors: Need to agree a way to discuss this

Issue #12 - The difference between slow/fast path may become moot

- I shared Brian Carpenter's view that the difference between slow/fast path may become moot with VM being routers or with the generalization of Network processors

- Editors: As noted in Issues # 5 & 10, need to find a good way to discuss this

Issue #15 - Are routers required to process (RAO)?

- Does it mean we MUST process ANY HBH option in the fast path except RA option, if the device support fast path forwarding?
  - Yes, that is the intent

- Editors: Close item.

Issue #17 - What are the incentives for wider support?

- Several comments
  - We need to agree on a road map of how much HbH space the ASICs should be able to process in 5, 10, 15, and even 20 years
  - Constraining HBH may be the only way to make it deployable. Granted, the constraints that we apply today may not be necessary tomorrow. For that reason, it may be more appropriate to publish the HBH processing draft as a BCP instead of Standards Track document.
- Editors: Needs w.g. discussion

Issue #18

- What status will the recommendations / requirements have?
  - I don't think the IETF's crystal ball is accurate enough to know where the technology will be ten or twenty years from now. I think all we can do is make some practical recommendations (small "r") for the next few years, as a 6man/v6ops collaboration. Should this be standards track
- Editors: Should this be Standards track?

Issue #20 - We cannot expect unrelated routers/devices on the path to apply specific behavior

- Many extension headers will only be processed by nodes (source & dest mainly) trusting each other's (i.e., my own definition of 'limited domain')
  - Extension headers are the obvious way to extend IP features
- Editors:
  - We agree that a specific HBH Option will only be processed by nodes wanting to do so (and don't think we are changing that).
  - Draft is focused on moving HBH process to "fast path" and not dropping packets with HBH Options.

Issue #21 - New options MUST NOT be defined that require Slow Path processing?

- Not supporting "New options MUST NOT be defined that require Slow Path processing"...
  - Doesn’t want to put constraints on new options

- Editors: Could change MUST to SHOULD
  - On the other hand, is this needed, given we have Router Alert?

QUESTIONS / COMMENTS?

https://github.com/ietf-6man/hbh-processing/issues

BACKUP SLIDES

Background

In the first IPv6 specification:

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

In the current IPv6 Specification (RFC8200):

- HBH processing is only required if router configured.
- This essentially documented current operational behavior.
- *Not intended to improve HBH processing*

Motivation

- **Still not practical for HBH Options to be used widely:**
  - Paths commonly *drop all packets with HBH options;*
  - *Multiple HBH options* in a 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.*
- **Our goal is to redefine procedures to make HBH options practical**:
  - This likely won’t work on all paths;
  - Methods can be designed that would still benefit from incremental support where provided.

New Hop-by-Hop Options

- New HBH Options SHOULD be designed for ”Fast Path” processing:
  - Straight forward to process
  - Designed to be the first option in HBH option header
  - Size of an option should not extend beyond what can be reasonably expected to be executed at full forwarding rate
- Any new option that does not meet this needs to explain reason in detail.

Proposal Summary (Changes to RFC8200)

- First HBH option MUST be processed in “Fast Path” **
  - Additional HBH options MAY be processed if configured to do so.
- Nodes creating packets with HBH options SHOULD include a single HBH option;
  - MAY include more based on local configuration.
- If there are more than one HBH options, a node MAY skip the rest without examining them (not processed or verified).
- Nodes unable to process an HBH option in the “Fast Path” MUST treat it as an unrecognized option.

** Router Alert is the exception

Proposal Summary (Changes to RFC8200) Continued…

- If HBH Option not recognized, change processing of high-order 2 bits of Option Type “10” and “11” to:

10 **discard the packet** and, regardless of whether or not the packet's Destination Address was a multicast address, **MAY send an ICMP Parameter Problem**, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

11 **discard the packet** and, only if the packet's Destination Address was not a multicast address, **MAY send an ICMP Parameter Problem**, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

Proposal Summary (Router Alert) Continued….

- **Router Alert**
  - Node SHOULD verify that the Router Alert option contains a supported protocol.
  - Verified packets SHOULD be sent to “Slow Path” for processing.
  - Nodes configured to support Router Alert options MUST protect itself from “Slow Path” infrastructure attacks.