Processing of the Hop-by-Hop Options Header

draft-peng-v6ops-hbh-00

Shuping Peng	Huawei
Zhenbin Li	Huawei
Chongfeng Xie	China Telecom
Zhuangzhuang Qin	China Unicom
Gyan Mishra	Verizon

Motivations

- The HBH Options Header is a valuable container for facilitating new services
  - The hop-by-hop processing behavior is very desirable
  - New services: IOAM, Alternate Marking, PMTU, etc.
- The HBH Options Header is rarely utilized in the current operators' networks.
  - Protect the control plane from undesired traffic by dropping or ignoring hbh
- Our main purpose is to
  - Break the endless cycle that resulted in HBH to become a DOS vector.
  - Enable the HBH options header to be utilized in a safe/secure way without impacting the management plane.
  - Ease the deployments of the new HBH based network services in a multi-vendor/operator's scenario that can now be deployed without operational impact.

Modern Router Architecture

- Modern router architecture design maintains a strict separation of its control and forwarding plane
- The control plane
  - realized in software on general-purpose processors
  - vulnerable to the DoS attack
- The forwarding plane
  - realized in hardware on flexible high-performance software programmable NPUs.
  - capable of handling very high packet rates and perform complex tasks such as HBH processing.
- The interface between control and forwarding plane
  - a rate-limit mechanism is always implemented to protect the control plane against DoS attack
    - ✓ cause inconsistent packet drops
    - ✓ impact the normal IP forwarding

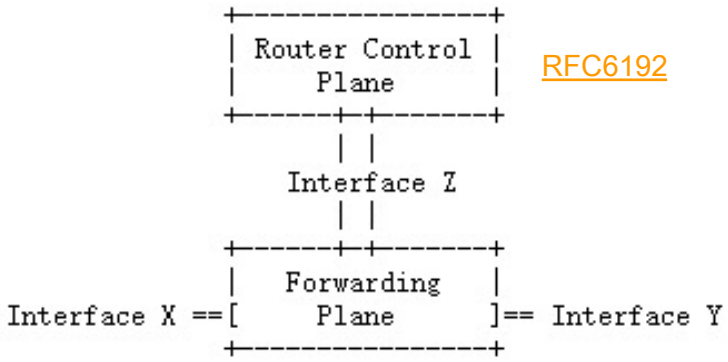

RFC6192

Figure 1. Modern Router Architecture

Common Implementations

- The value of the Next Header field in the IPv6 header
  - the only trigger for the default processing behavior of the HBH
- Common implementations
  - Once the device receives an IPv6 packet with its Next Header field set to 0, it will be directly sent to the slow path.
  - The option type of each option will not be examined before the packet is sent to the slow path.
  - In most of the cases, such processing behavior is the default configuration and cannot be changed.
- Historical Reasons
  - HBH options were not yet well-understood
  - Inflexible Fixed function ASICs performing proprietary functions were not so capable as they are in today's modern software programmable NPU's that can forward line rate in the fast path and still be able to perform complex tasks such as HBH forwarding options processing without having to punt to the slow path.
- Consequences
  - All packets that contain HBH are dispatched to the slow path
  - A risk of a DoS attack on the control plane
  - Congest the slow path, causing other critical functions to fail
  - Rate-limit causes inconsistent packet drops and impact the normal end-to-end IP forwarding of **the new services**

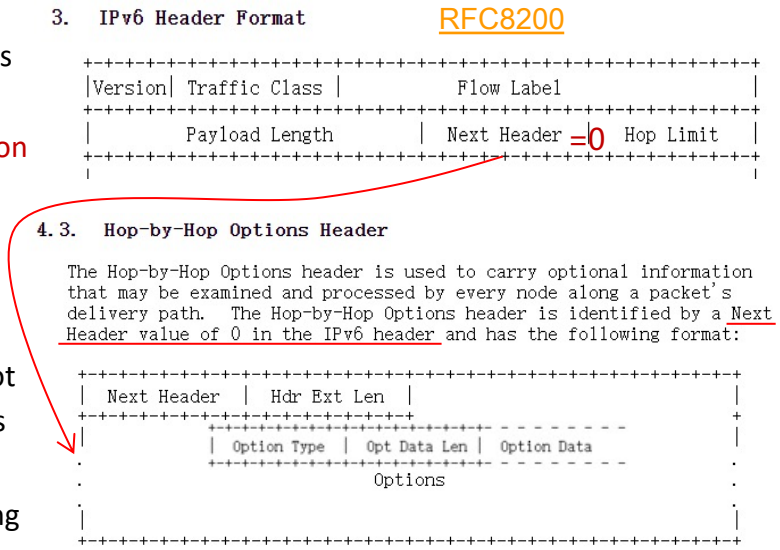

The desired processing behavior

- **Control pane should be protected from undesired traffic**
  - The HBH options that are not supposed to be processed by the control plane should not be sent to the control plane, potentially causing the DoS attack
  - The HBH options header should not be directly sent to the control plane once the packets are received since these options may not aim for the control plane
- **Source Node should not**
  - Encode the HBH options that exceed the maximum length of the HBH options header 2,048 bytes.
  - Encode the number of HBH options that exceeds the lowest processing capability of nodes along the path.
  - Encode the number of HBH options that exceed the maximum overall length of the extended header chain.
- **Better to separate the two types of options that are supposed to be processed by the control plane or by the forwarding plane, respectively**
  - The desired processing behaviors for the two types of options are different
  - The options aimed for the control plane are better not to consume the forwarding plane resources
  - For the current common implementation, all the options are sent to the control plane
  - There is no simple way to differentiate the two types of options except inspecting each option type

- The new deployments should be **compatible** with the existing deployments
  - Since default configuration of some devices cannot be changed or reconfigured
  - The update of the entire network cannot be done within one day

Corresponding migration strategies

- In order to achieve the desired processing behavior of the HBH options header and facilitate the ever-emerging new services to be deployed in operators' networks across multiple vendors' devices, the migration can happen in three parts as described below:
- 1. The source of the HBH options header encapsulation.
  - The information to be carried in the HBH options header needs to be first categorized and encapsulated into either control options or forwarding options, and then encapsulated in different packets.
- 2. The nodes within the network.
  - The nodes are updated to the proposed behavior introduced in the previous section.
- 3. The edge node of the network.
  - The edge node should check whether the packet contains a HBH header with control or forwarding option. Packet with a control option may still be filtered and dropped while packets with forwarding option should be allowed by the ACL.
  - If it is certain that there is no harm that can be introduced by the HBH options to the nodes and the services, they can also be allowed.
- Note: During the migration stage, the nodes that are not yet updated will stay with their existing configurations.

WG Adoption?

Thank you!

One proposal in draft-li-6man-hbh-fwd-hdr

- The HBH Options actually contain information for the use of the forwarding plane and the control plane of the nodes, respectively.
- They can be categorized into HBH Forwarding Options and HBH Control Options [I-D.li-6man-hbh-fwd-hdr].
- It is suggested to separate the two types of HBH options and carry them in different packets since generally they serve for different purposes and require different processing procedures on a node.
- The packets carrying the HBH Forwarding Options are supposed to be maintained in the forwarding plane rather than being directly sent up to the control plane. While the packets carrying the HBH Control Options are supposed to be sent to the control plane.
- If the IPv6 extension header including the HBH options header of a packet cannot be recognized by the node, or the option in the HBH header is unknown to the node, and the node is not the destination of the packet, the packet should not be dropped or sent to the control plane, rather this unrecognized extension header should be skipped and the rest of the packet should be processed.