

# IPv6 Extension Headers Review and Considerations

Last updated: October 2006

**This document reviews the concepts of IPv6 Extension Header (EH). It focuses on the implications that the presence of extension headers have on the native IPv6 traffic forwarding performance of network devices.**

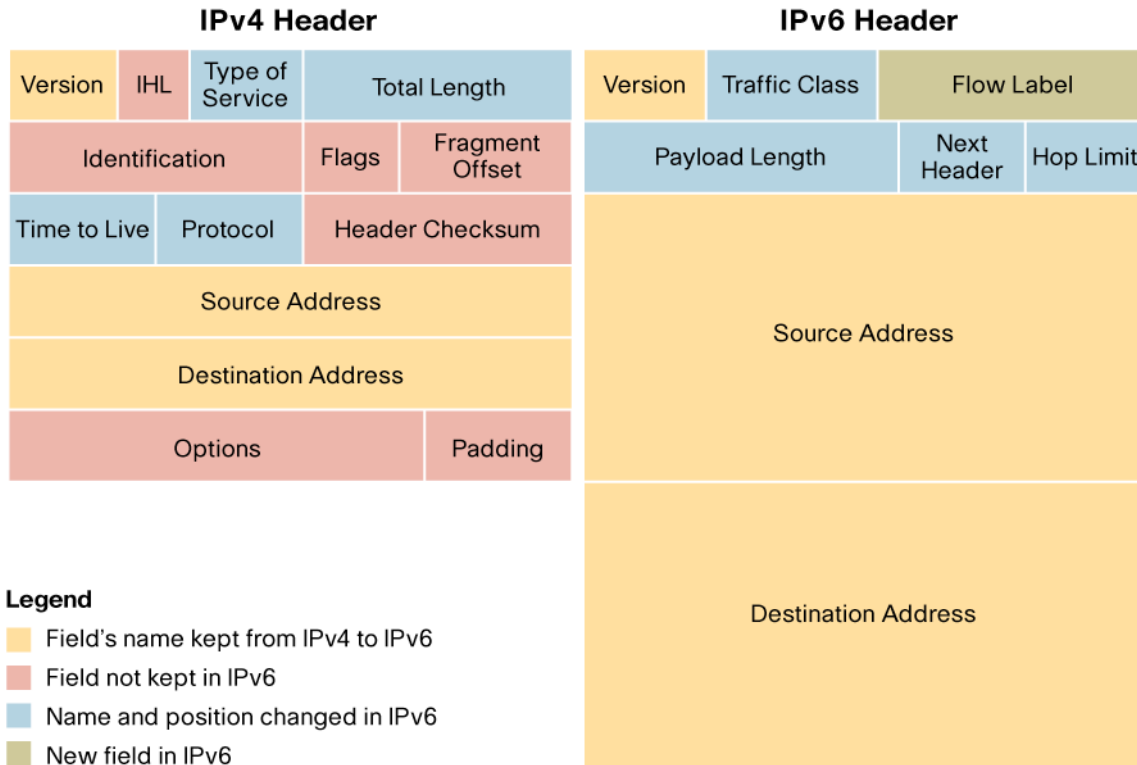
## IPv6 EXTENSION HEADERS

This section reviews the headers used by the IPv6 protocol.

### The Concept

IPv6 is using two distinct types of headers: Main/Regular IPv6 Header and IPv6 Extension Headers. The main IPv6 header is equivalent to the basic IPv4 one despite some field differences that are the result of lessons learned from operating IPv4. Figure 1 presents the IPv4 and IPv6 main headers.

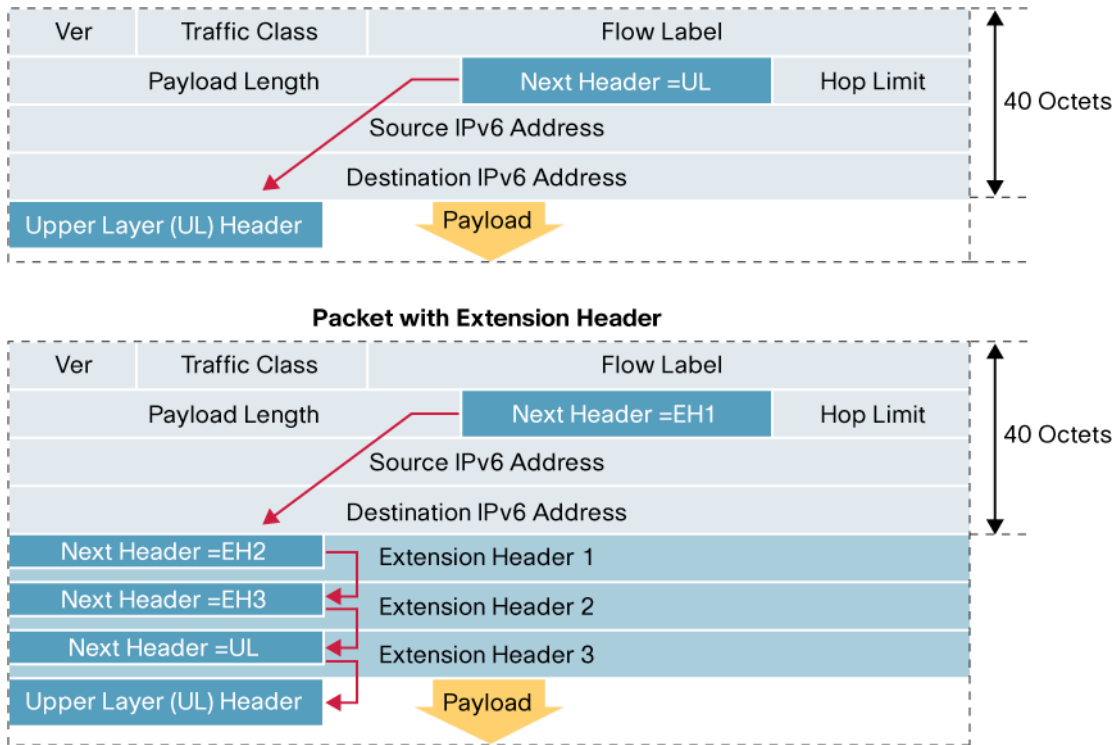
**Figure 1.** IPv4 and IPv6 Headers



The options field in the IPv4 header is used to convey additional information on the packet or on the way it should be processed. Routers, unless instructed otherwise [1], must process the options in the IPv4 header. The processing of most header options pushes the packet into the slow path leading to a forwarding performance hit.

IPv4 Options perform a very important role in the IP protocol operation therefore the capability had to be preserved in IPv6. On the other hand, the impact of IPv4 Options on performance was taken into consideration in the development of IPv6. The functionality of options is removed from the main header and implemented through a set of additional headers called extension headers [2]. The main header remains fixed in size (40 bytes) while customized EHs are added as needed. Figure 2 shows how the headers are linked together in an IPv6 packet.

**Figure 2.** Chaining Extension Headers in IPv6 Packets



RFC2460 defines the extension headers as shown in the following table along with the Next Header values assigned to them:

**Table 1.** IPv6 Extension Headers and their Recommended Order in a Packet

Order	Header Type	Next Header Code
1	Basic IPv6 Header	–
2	Hop-by-Hop Options	0
3	Destination Options (with Routing Options)	60
4	Routing Header	43
5	Fragment Header	44
6	Authentication Header	51
7	Encapsulation Security Payload Header	50
8	Destination Options	60
9	Mobility Header	135
	No next header	59
Upper Layer	TCP	6
Upper Layer	UDP	17
Upper Layer	ICMPv6	58

RFC2460 also recommends the order in which they should be chained in an IPv6 packet:

1. **IPv6 main header**
2. **Hop-by-Hop Options header (if present, it MUST be the first one following the main/regular header)**
3. **Destination Options header**
4. **Routing header**
5. **Fragment header**
6. **Authentication header**
7. **Encapsulating Security Payload header**
8. **Destination Options header**
9. **Upper-layer header**

The only **MUST** requirement is that the Hop-by-Hop EH has to be the first one.

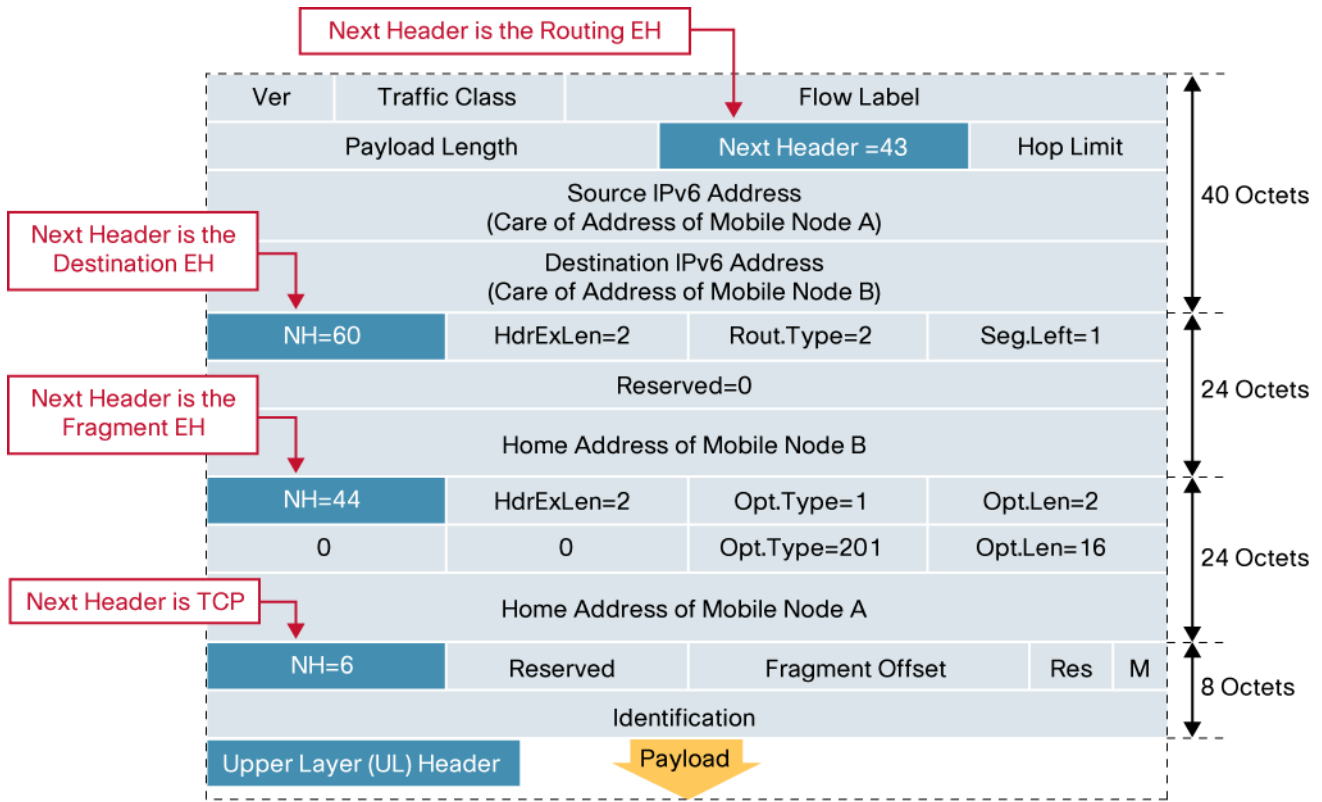
### Commonly Used Extension Headers

The Extension Header should not be viewed as an esoteric feature of IPv6 that would be encountered only at later stages of the network and service deployment. Extension headers are an intrinsic part of the IPv6 protocol and they support some basic functions and certain services. The following is a list of circumstances where EHs are commonly used:

- **Hop-by-Hop EH** is used for the support of Jumbo-grams or, with the Router Alert option, it is an integral part in the operation of MLD. Router Alert [3] is an integral part in the operations of IPv6 Multicast through Multicast Listener Discovery (MLD) and RSVP for IPv6.
- **Destination EH** is used in IPv6 Mobility as well as support of certain applications.
- Routing EH is used in IPv6 Mobility and in Source Routing. It may be necessary to disable “IPv6 source routing” on routers to protect against DDoS.
- **Fragmentation EH** is critical in support of communication using fragmented packets (in IPv6, the traffic source must do fragmentation—routers do not perform fragmentation of the packets they forward)
- **Mobility EH** is used in support of Mobile IPv6 service
- **Authentication EH** is similar in format and use to the IPv4 authentication header defined in RFC2402 [4].
- **Encapsulating Security Payload EH** is similar in format and use to the IPv4 ESP header defined in RFC2406 [5]. All information following the Encapsulating Security Header (ESH) is encrypted and for that reason, it is inaccessible to intermediary network devices. The ESH can be followed by an additional Destination Options EH and the upper layer datagram.

Figure 3 presents the structure of IPv6 data plane packets using extension headers. In this example, the packet is sent from Mobile Node A to Mobile Node B over the route optimized path [6], hence the use of the Routing EH (43) and the Destination Options EH (60). It is sent over a path that has an Maximum Transmission Unit (MTU) smaller than that of Mobile Nodes (MNs) access link, hence the use of the Fragmentation EH (44).

**Figure 3.** Data Traffic Between Two Mobile Nodes over the Route Optimized Path

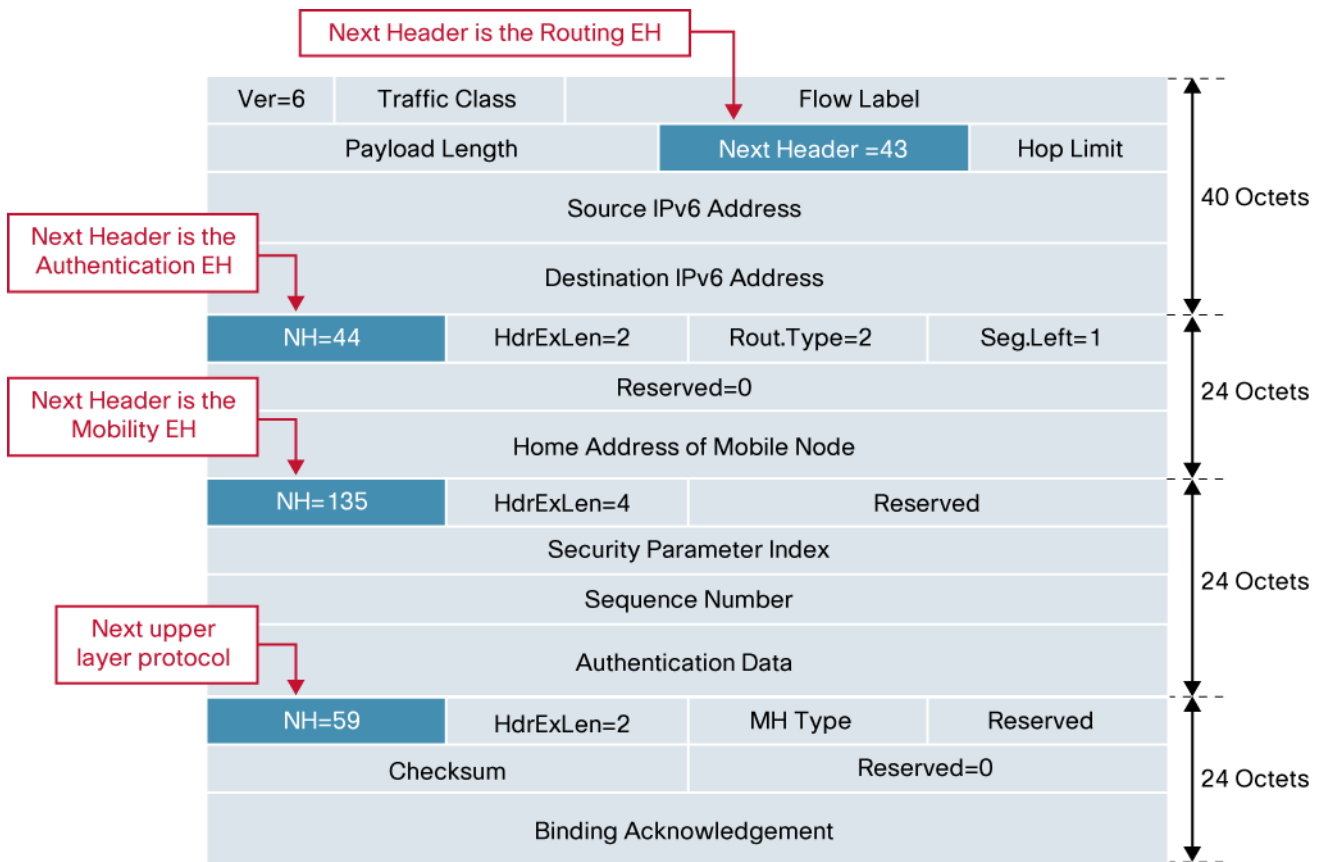


The length of the Extension Header chain in this case is 56 bytes. By comparison, the IPv6 packet between a Mobile Node and Correspondent Node over the route optimized path would have either the Routing or the Destination Options EH (not both) leading to a shorter EH chain.

To exemplify the case of an IPv6 packet with a longer EH chain, Figure 4 depicts the structure of a Binding Acknowledgement sent from the Home Agent to the Mobile Node. The first header in the chain is the Routing EH (43), the second is the Authentication EH (51) and finally the Mobility EH (135).

**Note:** Mobility standards are not allowing for data to be piggy-backed to the binding maintenance messages (control plane). This is the reason why the packet in Figure 4 does not have a payload.

**Figure 4.** Binding Acknowledgment Sent from a Correspondent Node to a Mobile Node



This packet, built to emphasize special cases where multiple extension headers might be used, has an EH chain length of 72 bytes. Since the length of most individual extension headers is variable, the length of EH chains can be even larger. Note however that this size is typically driven by the need to carry certain information in addition to that in the main header. The larger EH chains (whether due to many EH or to long individual EH) are used for control plane traffic. Packets carrying user data such as the one showed in Figure 3 generally have shorter EH chains.

**Note:** The control plane traffic generally has a very low rate compared with the data plane traffic, from a forwarding performance perspective, packets with shorter EH chains (Figure 3) are of more interest.

The common use of IPv6 EH makes it important to analyze and understand the way in which network devices (routers, layer 3 switches and generally devices that forward traffic based on layer 3 information) process the extension headers.

**Note:** The EHs are considered a powerful tool in extending IPv6 to adapt to future protocol requirements and service needs. It is expected the other uses will be identified for the existent EHs and that new EHs will be defined.

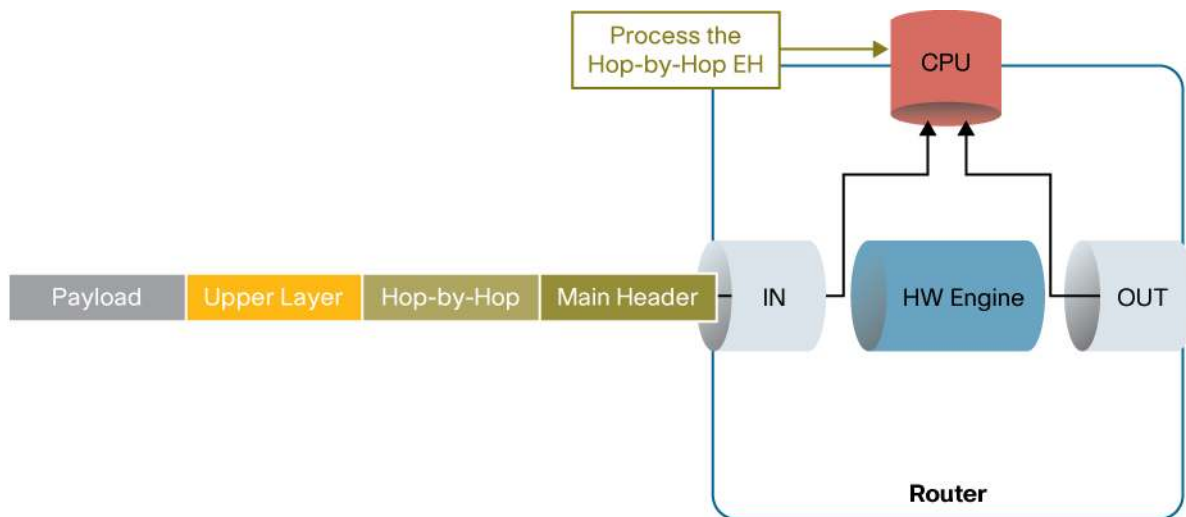
## IPV6 EXTENSION HEADERS PROCESSING

This section describes the way in which various Extension Header types must be processed by network devices under basic forwarding conditions or in the context of advanced features such as Access Lists. It identifies the protocol requirements that must be observed.

### Hop-by-Hop Extension Header

The Hop-by-Hop Extension Header is the ONLY EH that MUST be fully processed by all network devices as shown in Figure 5. From this perspective, the Hop-by-Hop EH is similar to the IPv4 options. This explains the reason why this EH MUST be the first in a chain of extension headers.

**Figure 5.** Forwarding IPv6 Packets with the Hop-by-Hop Extension Header



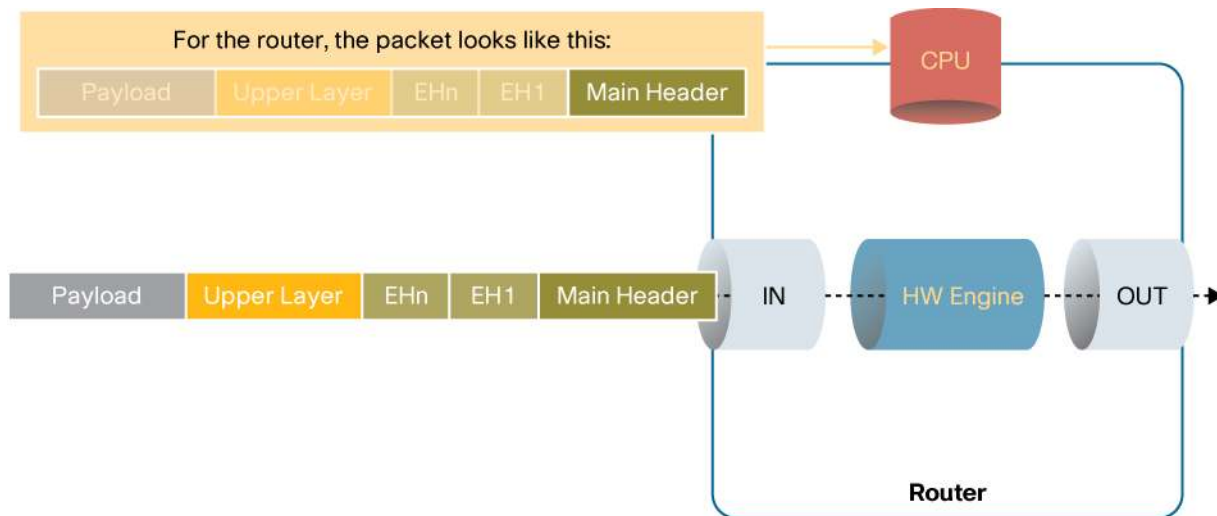
Because the Hop-by-Hop EH must be fully processed, it is handled by the CPU<sup>1</sup> and the IPv6 traffic that contains a Hop-by-Hop EH will go through the slow forwarding path. This rule applies to all vendors. Hardware forwarding is not feasible in this case.

<sup>1</sup> Note: The Hop-by-Hop EH can be processed by the central or Line Card CPU depending on the router architecture.

## Other Extension Headers

Network devices are not required to process any of the other IPv6 extension headers when simply forwarding the traffic<sup>2</sup>. For this reason, IPv6 traffic with one or more EHs other than Hop-by-Hop can be forwarded in hardware as shown in Figure 6.

**Figure 6.** Forwarding IPv6 Packets with Extension Headers other than Hop-by-Hop in the Absence of ACLs



Network devices might however process some EHs if specifically configured to do so while supporting certain services such as IPv6 Mobility.

The extensions headers used to secure the IP communication between two hosts, Authentication and Encapsulating Security Payload Headers, are also ignored by the intermediary network devices while forwarding traffic. These EHs are relevant only to the source<sup>3</sup> and destination of the IP packet. It is important however to remember that all information following the ESH is encrypted and not available for inspection by an intermediary device, if that is required.

## Extension Headers and Access Lists

In the absence of the Hop-by-Hop EH, based on the processing rules described in the previous **Hop-by-Hop Extension Header** and **Other Extension Headers** sections, as long as a router is concerned exclusively with layer 3 information and it is not specifically instructed to process certain EH (for certain services it is supporting), it can forward IPv6 traffic without analyzing the extension headers. An IPv6 packet can have an arbitrary number of EH (other than Hop-by-Hop) and the router would ignore them and simply forward the traffic based on the main header. Under these conditions, routers can forward the IPv6 traffic in hardware despite the EHs. Access Lists (ACL) applied on router interfaces however, can change the router's IPv6 forwarding performance characteristics when extension headers are present.

<sup>2</sup> Note: This document discusses only the behavior of network devices that are natively forwarding IPv6 forwarding. Handling of tunneled IPv6 traffic is not discussed in this document since tunneled traffic is typically process switched.

<sup>3</sup> Note: A router can be the source of an IPv6 IPsec tunnel, adding the AH or the ESH to the packets transmitted over that tunnel. This can be done in software or hardware depending on the platform.

## Filtering Based on Extension Header Type

Routers can be explicitly instructed to look at and act on traffic that contains certain extension header types. This functionality is available on Cisco platforms and can be configured with the help of IPv6 Access List options:

```
deny protocol {source-ipv6-prefix/prefix-length | any | host source-ipv6-address}
[operator [port-number]] {destination-ipv6-prefix/prefix-length | any | host destination-
ipv6-address} [operator [port-number]] [dest-option-type [doh-number | doh-type]]
[dscp value] [flow-label value] [fragments] [log] [log-input] [mobility] [mobility-type
[mh-number | mh-type]] [routing] [routing-type routing-number] [sequence value]
[time-range name] [undetermined-transport]
```

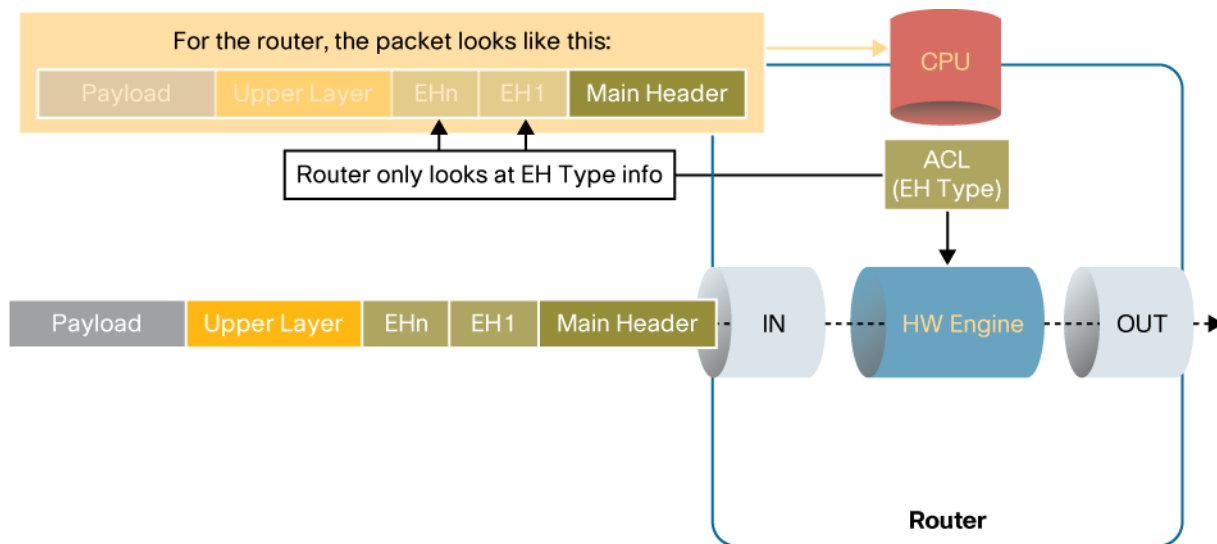
Example 1: Access List Filtering based on Extension Header Type

```
ipv6 access-list EH-type
deny ipv6 any 2001:2B8:1:1::/64 routing
```

In order to permit or deny certain types of extension headers, routers are configured with the ACL features listed above to filter based on the “Header Type” value (see Table 1).

**Note that in this case, the content of the EH is not processed and the router simply makes a decision based on the presence of a certain EH type.** Software platforms can also analyze and filter based on additional EH fields such as the “Type Field”. These filtering capabilities are very useful in implementing, for instance, security policies such as blocking source routing.

**Figure 7.** Forwarding IPv6 Packets with Extension Headers other than Hop-by-Hop with ACLs Filtering Based on EH type



Since this functionality is implemented through ACLs, platforms that support hardware forwarding when ACLs are applied, will be able to handle the IPv6 traffic with EHs in hardware as well (Figure 7).

## Filtering Based on Upper Layer Protocol

Often, routers filter traffic based on the upper layer protocol information (Example 2). In these cases, a router must process the main header of the packet as well as the information in its payload. In the absence of extension headers, routers perform these functions on IPv6 traffic in the same way they do on IPv4 traffic, so the traffic can be forwarded in hardware.

Example 2: Access List Filtering on Upper Layer Protocol Information

```
ipv6 access-list Upper-Layer-Info
deny tcp any 2001:2B8:1:1::/64 eq whois
```

In the presence of extension headers (not Hop-by-Hop) however, the upper layer protocol information is pushed deeper into the payload of the packet, impacting the packet inspection process. In these cases, the router will have to traverse the chain of headers (main plus extension headers), header by header until it reaches the upper layer protocol header and the information it needs for the filter. The extension headers are not processed, the router simply looks at the “Next Header” value and the length of the EH in order to understand what header follows and the offset to its beginning.

Even though a router might be able to process upper layer protocol ACLs or one EH in hardware, if it was not designed while considering all aspects of IPv6, it might not be able to handle filtering when packets contain both EH and Upper Layer data as in the scenario described above.

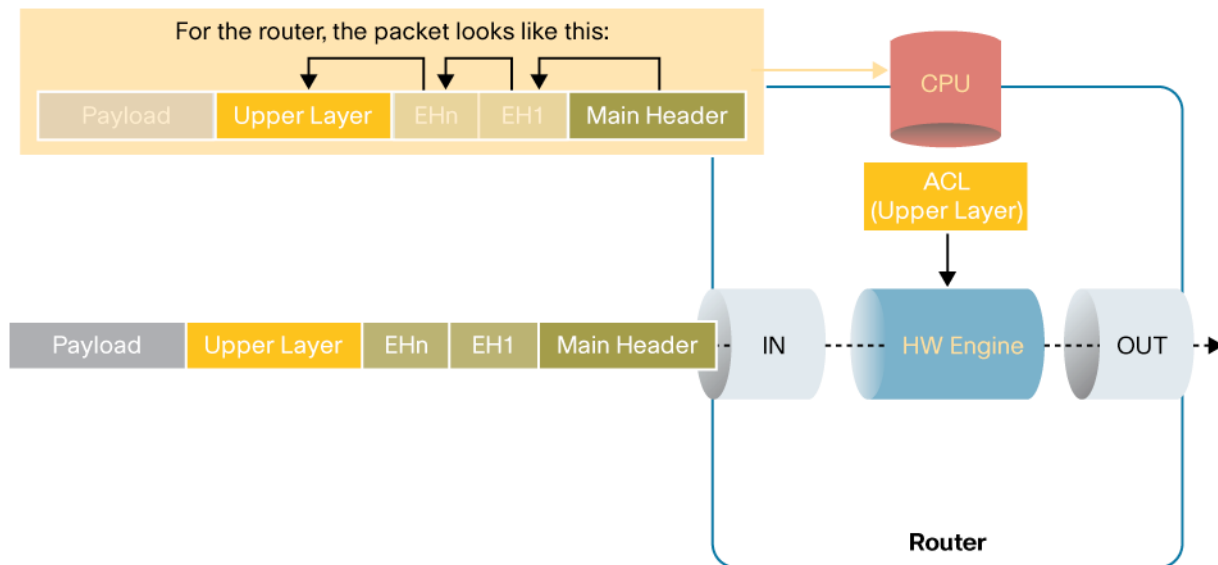
**Since Extension Headers are an important part of IPv6 operation and upper layer protocol filters are commonly used on edge devices (PE routers), it is important to evaluate the impact of EHs (not Hop-by-Hop which always must fully processed) on the IPv6 forwarding performance of the network device.**

Note that if the IPv6 packet contains, along with other EH, the Encapsulating Security Payload EH, the information following it, which includes the upper layer information, is encrypted. In this case, the router ACL described in the scenario discussed in this section will not be able to match on the upper layer information and will drop the packet. When only the Authentication EH is present (no Encapsulating Security Payload EH), the packet payload is not encrypted so the packet will be handled as described in this section.

### FORWARDING PERFORMANCE OF IPV6 TRAFFIC WITH EXTENSION HEADERS ON CISCO ROUTERS

Cisco platforms that leverage hardware forwarding are designed to take into consideration the role and processing requirements of IPv6 extension headers without impacting performance. Figure 8 shows the Cisco products behavior under the conditions described in the previous **Filtering Based on Upper Layer Protocol** section. This is not always the case across the industry.

**Figure 8.** Forwarding IPv6 Packets with Extension Headers other than Hop-by-Hop with ACLs Filtering Based on Upper Layer Information



The following table summarizes the various scenarios, discussed in the previous **IPv6 Extension Headers Processing** section, in which EHs are relevant to packet processing. It also presents the forwarding path used by Cisco hardware based platforms in each of these cases.

**Table 2.** IPv6 Forwarding Path through Cisco Hardware Platforms in the Presence of Extension Headers

	IPv6 Traffic Profile	Applied Advanced Features (ACL)	Forwarding Path <sup>4</sup>
1	IPv6 Main Header Only	No ACL applied	Hardware
2	IPv6 Main Header Only	IP +Upper Layer Info ACL	Hardware
3	IPv6 with Hop-by-Hop EH	With or without ACL	Software (mandatory to fully process the EH)
4	IPv6 with EH other than Hop-by-Hop	No ACL applied	Hardware
5	IPv6 with EH other than Hop-by-Hop	IP + Upper Layer Info ACL	Hardware <sup>5</sup>

The resources allocated for functions performed in hardware are limited and fixed by design. This is a natural constraint driven by implementation costs and it is defined based on the original requirements and specifications that shaped the line card design. The same stands true for the amount of resources available for handling the IPv6 extension headers in hardware. Cisco's hardware accelerated platforms are designed to handle at least 64 bytes of extension headers data. This selected size is considered sufficient to handle the most common chains of EH currently used with various IPv6 traffic types and services. Analyzing the EH chain example provided in Figure 3, the IPv6 packet has a 56 bytes long EH chain, so it would be forwarded in hardware.

Due to the potential combinations and lengths of extension headers, the overall size of the chain will vary. In Cisco routers, when the size of the EH chain exceeds the resources allocated in HW and upper layer protocol filters are applied, the IPv6 traffic will be software switched by the Line-Card CPU. The packet shown in Figure 4 has a 72 bytes long EH chain, which means that a hardware forwarding platform that was designed to handle 64 bytes of EH data will have to software switch this packet. With the current IPv6 applications and services however, this would be an infrequent event. Moreover, packets with very long EH chains are generally control plane packets (such as the one depicted in Figure 4) which have a very low rate, so the impact on the router is not significant.

**Note:** Platforms that do not use hardware acceleration such as the Cisco 7301, 7200 and lower end platforms, will always software switch the IPv6 traffic, independent of the presence or absence of extension headers.

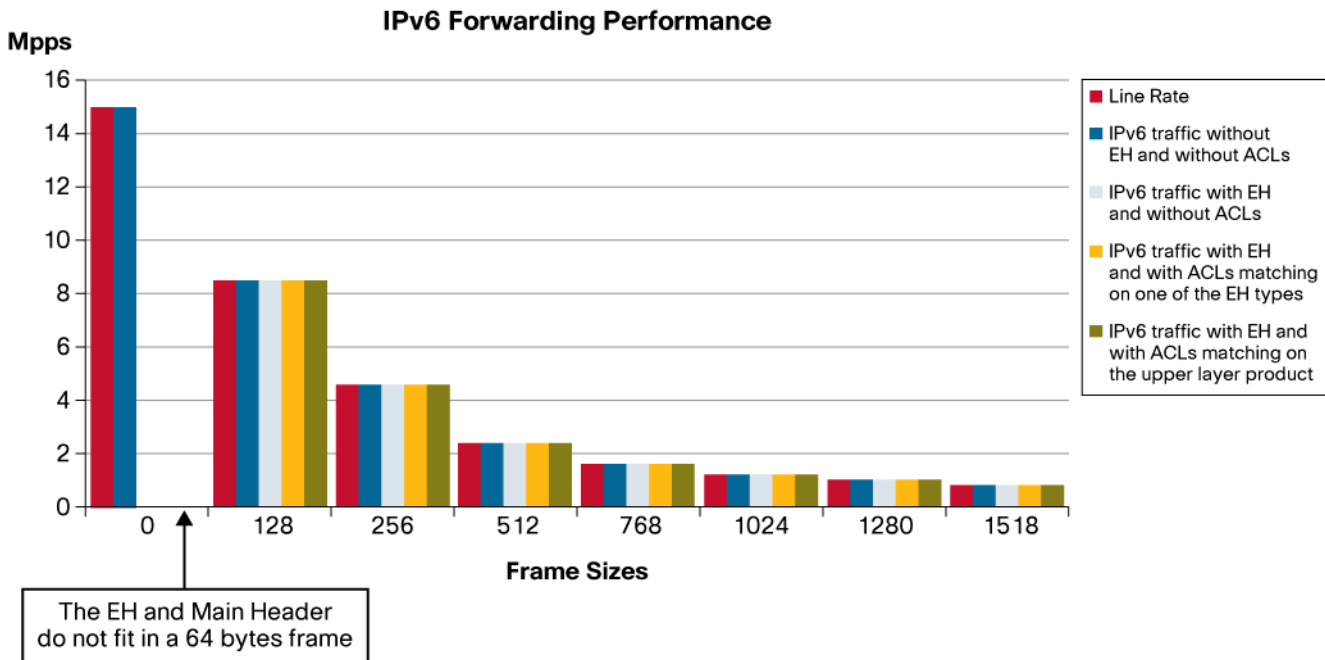
**Security Note:** There is always the possibility that IPv6 traffic with a significant number of extension headers or very large extension headers is sent into the network with the malicious intent of overrunning the HW resources of network devices. Regardless of the platform HW design, this is a possible DDoS type of attack vector. Security features protecting against it must be implemented. To protect the CPU from being overwhelmed by high rates of this type of traffic, Cisco routers implement rate limiting of packets that are diverted from the hardware to software path.

<sup>4</sup> Note: Applicable to Cisco high-end platforms: CRS, GSR (Engine3, Engine5 line cards), Cisco 7600/6500 (SUP720).

<sup>5</sup> Note: If the length of the extension header chain (sum of the lengths of all EH in the chain) is larger than the resources allocated in hardware, the packet is punted to the software path as explained later in this section.

As an example for discussion, the Engine 5 line card of the GSR routers can handle up to 64 bytes of extension headers in hardware. Its IPv6 forwarding performance is not impacted by the presence of EH (other than Hop-by-Hop) in the scenarios discussed in this document as shown in Figure 9 where the tests were executed on a GSR running Cisco IOS Software Release 12.0(32)S with E5 10 Gigabit Ethernet Line Cards.

**Figure 9.** GSR, Engine 5 IPv6 Forwarding Performance in the Presence of Extension Headers



This data exemplifies the behavior of a hardware platform designed with IPv6 in mind from the beginning.

## CONCLUSIONS

Extension headers are an intrinsic building block of IPv6. It is critically important to understand their role and mode of use as well as the processing requirements they have on various network devices. Deployed IPv6 networks must be capable to handle IPv6 traffic containing extension headers. They must forward such IPv6 traffic at optimal, production level performance in the presence of advanced features such as Access Lists.

## REFERENCES

1. [http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products\\_feature\\_guide09186a00801d4a94.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801d4a94.html)
2. RFC2460 "Internet Protocol, Version 6 (IPv6) Specification" S. Deering and R. Hinden, December 1998
3. RFC2711 "IPv6 Router Alert Option" C. Partridge and A. Jackson, October 1999
4. RFC2402 "IP Authentication Header" S. Kent and R. Atkinson, November 1998
5. RFC2406 "IP Encapsulating Security Payload (ESP)" S. Kent and R. Atkinson, November 1998
6. RFC3775 "Mobility Support in IPv6" D. Johnson, C. Perkins and J. Arkko, June 2004



**Corporate Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**  
Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www-europe.cisco.com  
Tel: 31 0 20 357 1000  
Fax: 31 0 20 357 1100

**Americas Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**  
Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco.com Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic  
Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy  
Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal  
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden  
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)