

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: March 14, 2012

H. Pfeifer, Ed.
Protocol Labs
September 11, 2011

IPv6 Extension Headers Reserved Space
draft-pfeifer-6man-exthdr-res-01.txt

Abstract

This document specifies a mechanism to allow IPv6 stack implementation to parse upcoming IPv6 extension header by reserving a small range of protocol numbers for well defined IPv6 extension headers. IPv6 stack implementors can code these range into their network stack a priori. These reserved range provide an guarantee that a given next header is a well formed extension header and not a next layer protocol. Operators, companies, vendors et cetera are in the ability to deploy and test not standardized extension headers without modifying every middle box parsing the complete extension header chain in between.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 14, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

Abbreviated Title

September 2011

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	3
2.	Overview	3
3.	Extension Header Encoding	4
4.	Acknowledgements	5
5.	IANA Considerations	5
6.	Security Considerations	5
7.	Normative References	6
Appendix A.	Implementation Example	6
	Author's Address	7

Internet-Draft

Abbreviated Title

September 2011

1. Introduction

Optional IPv6 protocol information are encoded in extension header between the basic IPv6 header and the next layer protocol. Next layer protocols such as TCP or UDP are encoded in a non uniform manner. Network stacks that process the extension header chain are required to know the exact encoding of all staged headers until the required next layer protocol is detected. If a unknown protocol or extension header lies between the processing must be stopped.

Extension header are normally not processed on intermediate nodes [[RFC2460](#)]. Intermediate nodes are urged to not examine or process any extension header. One exception is is the hop-by-hop options header which must be processed. Sometimes nodes along a packet's delivery path are required to parse the complete extension header chain to analyze the transport layer protocol. Middleboxes which analyze the transport protocol header are the most prominent example. Firewalls or network proxies are the most prominent examples. Middle boxes have no possibility to resolve this circumstance cleanly. Firewalls are likely to drop the packet.

End-boxed following [[RFC2460](#)] generate a ICMP Parameter Problem message if an unknown extension header is detected. Providing the sender enough information to resend the packet without the extension header.

[I-D.ietf-6man-exthdr] defines a standard format to unify the layout of IPv6 extension header and advised to use the IPv6 Destination Options Header. Thus providing a way that future extensions are encoded in a consistent manner. This is the first part to address this unlucky fact. The second half of the problem lies in the fact that the next header space is shared with the protocol number. Network stack cannot differentiate if an unknown code-point declaring an extension header or an protocol.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

[2.](#) Overview

This mechanism provide middle boxes the chance to process up to now new extension header and simultaneous do not constrain the possibility to introduce new extension header in future by reserving a small fraction of protocol number space.

Pfeifer

Expires March 14, 2012

[Page 3]

Internet-Draft

Abbreviated Title

September 2011

Future IPv6 extensions which cannot encoded in one of the existing extension header MUST follow the encoding specification in [[I-D.ietf-6man-exthdr](#)]. Upcoming standardized extension header following the proposal MUST be reserved from this space. Network stacks are provided with a guarantee that these extension headers follow the specification.

Currently 60% of the space is already allocated. Thus to be conservative a good compromise is to reserve the range between 245 and 252 (7). This leave space for 98 new next layer protocols, 98 extension header not following the proposal or any combination of both. Upcoming standardized extension header following the proposal MUST be reserved from this space.

Providing an guarantee of the encoding makes it possible that vendors and other parties implement proprietary extension header without standardization and wait years of enrollment of new middle-boxes supporting the new extension header.

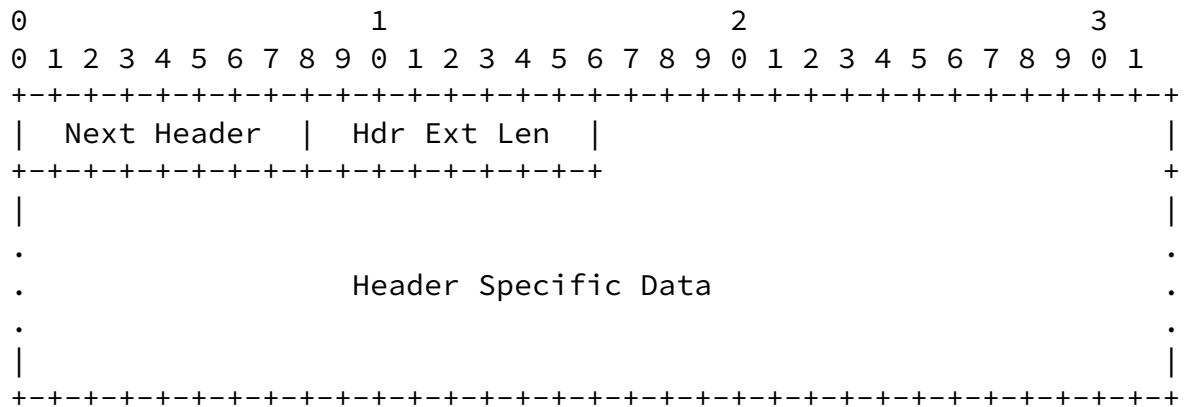
[3.](#) Extension Header Encoding

The following figure illustrate the format of the unified extension header format. It will be removed if [[I-D.ietf-6man-exthdr](#)] is approved.

Internet-Draft

Abbreviated Title

September 2011



Next Header	8-bit selector. Identifies the type of header immediately following the Extension header. Uses the same values as the IPv4 Protocol field.
Hdr Ext Len	8-bit unsigned integer. Length of the Extension header in 8-octet units, not including the first 8 octets.
Header Specific Data	Variable length. Fields specific to the extension header

Figure 1

[4.](#) Acknowledgements

The author thank Suresh Krishnan, James Woodyatt, Erik Kline, James Hoagland, Manav Bhatia and all other people involved in [\[I-D.ietf-6man-exthdr\]](#).

[5.](#) IANA Considerations

This document includes a request to IANA. The editors of this draft request the protocol number 245 till 252 be reserved for uniform formatted IPv6 extension header

[6.](#) Security Considerations

Filled later

Pfeifer

Expires March 14, 2012

[Page 5]

Internet-Draft

Abbreviated Title

September 2011

[7.](#) Normative References

[I-D.ietf-6man-exthdr]

Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "An uniform format for IPv6 extension headers", [draft-ietf-6man-exthdr-04](#) (work in progress), July 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.

[Appendix A.](#) Implementation Example

The following patch applied on net-next-2.6 add support for the code-point reservation.

```
diff --git a/net/ipv6/exthdrs_core.c b/net/ipv6/exthdrs_core.c
index 14ed0a9..d7f2e97 100644
--- a/net/ipv6/exthdrs_core.c
+++ b/net/ipv6/exthdrs_core.c
@@ -4,6 +4,9 @@
     */
     #include <net/ipv6.h>

+#define      IP6_EXTHDR_RESERVED_MIN  245
+#define      IP6_EXTHDR_RESERVED_MAX  252
+
+    /*
+     * find out if nexthdr is a well-known extension header or a protocol
+     */
@@ -18,7 +21,8 @@ int ipv6_ext_hdr(u8 nexthdr)
        (nexthdr == NEXTHDR_FRAGMENT)    ||
        (nexthdr == NEXTHDR_AUTH)        ||
        (nexthdr == NEXTHDR_NONE)        ||
-       (nexthdr == NEXTHDR_DEST);
+       (nexthdr == NEXTHDR_DEST)        ||
+       (nexthdr >= IP6_EXTHDR_RESERVED_MIN &&
+        nexthdr <= IP6_EXTHDR_RESERVED_MAX);
    }
```

Figure 2

Author's Address

Hagen Paul Pfeifer (editor)
Protocol Labs
Hofmannstr. 19
81379, Munich
Germany

Phone: +49 174 5455 209
Email: hagen.pfeifer@protocollabs.com
URI: <http://www.protocollabs.com>