

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: March 19, 2012

D. Zhang  
S. Jiang  
Huawei Technologies Co., Ltd  
B. Carpenter  
Univ. of Auckland  
September 16, 2011

An Offset Indicating Option for IPv6  
draft-zhang-6man-offset-option-01

## Abstract

This document defines an Offset Indicating option (OI option) encapsulated within an IPv6 Options header. An OI option can provide offset information to locate the end of the IPv6 header chain so that a node receiving an IPv6 packet is able to skip over the IP header chain and access the transport header or other protocol data unit directly.

## Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on March 19, 2012.

## Copyright Notice

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

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 Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Requirements Language . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Format of the Offset Indicating option . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Processing Rules . . . . .	<a href="#">4</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">5</a>
<a href="#">6.</a>	IANA Considerations . . . . .	<a href="#">6</a>
<a href="#">7.</a>	Acknowledgements . . . . .	<a href="#">6</a>
<a href="#">8.</a>	References . . . . .	<a href="#">6</a>
<a href="#">8.1.</a>	Normative References . . . . .	<a href="#">6</a>
<a href="#">8.2.</a>	Informative References . . . . .	<a href="#">6</a>
	Authors' Addresses . . . . .	<a href="#">7</a>

## 1. Introduction

According to [\[RFC2460\]](#), when a node intends to access the payload of an IPv6 packet, it needs to parse the extension headers one by one until it reaches the end of the header chain. This approach may be inefficient for nodes which have no interest in the extension headers and intend to quickly access the payload of IPv6 packets.

A common case is any form of flow classification requiring access to the basic IP header 5-tuple {destination address, source address, protocol, destination port, source port}. The last three elements are only available by following the extension header chain to its end. This could be required for various forms of quality of service support or for flow logging purposes. Another case would be any form of deep packet inspection requiring rapid access to the payload, which also requires skipping over the header chain. If packets must be processed at line speed, this can be a significant performance issue. A method is needed to short-circuit this process.

A brief discussion of this issue from a security standpoint is provided in [Section 2.1.9.2 of \[RFC4942\]](#). In addition, most existing firewall implementations have the capability to verify the correctness of IP headers. Therefore, in some cases, it may be more efficient for the equipment behind a firewall, such as a host or a deep packet inspection device, to skip over the extension headers of the IP packets it receives and access the payload directly.

This document addresses this issue by introducing an Offset Indicating option (OI option for short) which indicates the end of the header chain. The option is transferred in an IPv6 Options header. If there is an existing Hop-by-Hop Options header, the OI option will be in it. Otherwise, it will be in a Destination Options header. According to the recommendations in [\[RFC2460\]](#), this will always place the OI option at the beginning of the header chain. Therefore, if necessary, a node receiving an IPv6 packet can jump over the whole header chain in a single step to directly access the

transport header or other protocol data unit.

This option is an optimization option for certain forwarding nodes. It may be safely ignored by nodes that have no interest in the header chain. Hence, it does not create any performance degradation. In particular, unless there is a Hop-by-Hop Options header for some other reason, it does not create any overhead for simple forwarding nodes.

## 2. 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 [\[RFC2119\]](#).

### 3. Format of the Offset Indicating option

The format of the Offset Indicating option (OI) option is described in Figure 1.

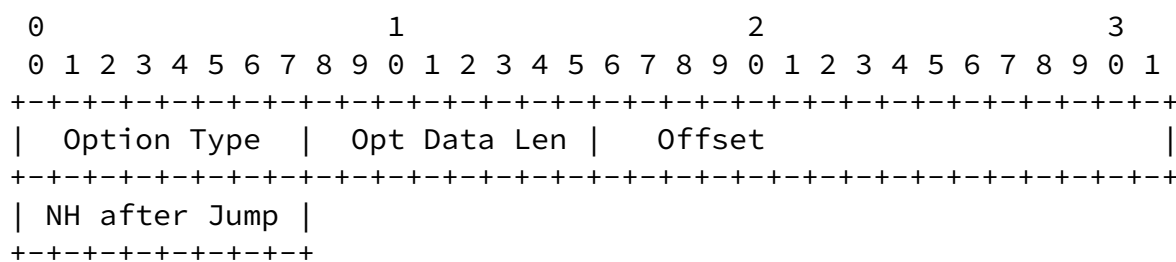


Figure 1. Option Format

Option Type: 8 bits. The value is TBD1.

Note to RFC Editor: please replace TBD1 with the value assigned by IANA and delete this note.

Opt Data Len: as defined in [\[RFC2460\]](#).

Offset: 16 bits. Indicates the distance (in octets) from the end of

the option to the end of the header chain.

NH (Next Header) after Jump: 8 bits. Indicates the type of the transport header or other protocol data unit after the header chain. This MUST equal the Next Header value in the last Extension Header in the packet.

#### 4. Processing Rules

IPv6 source nodes SHOULD insert this option in every packet that contains at least one extension header of any kind, in order to maximise its usefulness. However, it MUST NOT be inserted in packets that include a Fragment Header, to avoid the case where the offset points beyond the end of the first fragment. In any case, performance optimisation is impossible in the case of fragmented packets.

Because the options within a header must be processed strictly in the

order that they appear, the OI option is RECOMMENDED to be the first option within an Options header. This arrangement will maximize the effect of optimization for those routers that use it.

A Hop-by-Hop Options header MUST NOT be created solely for the purpose of carrying the OI option. If and only if the packet contains a Hop-by-Hop Options header for some other reason, the OI option is placed in it. Otherwise it is placed in a Destination Options header.

This option has an alignment requirement of  $4n + 2$ . (See [Section 4.2 of \[RFC2460\]](#) for discussion of option alignment.) If this option is located first within the Options header, the alignment requirement is met naturally; otherwise the host stack that assembles the IPv6 header needs to meet the alignment requirement according to the context by inserting padding options.

The OI option is defined on the basis that the size of extension headers does not change en-route. However, if a future extension header type allows an intermediate device to add additional information in the IP extension header chain, this device MUST also update the value of the Offset field to point to the new position of

the payload header.

If an intermediate device detects that the OI option does not point to a valid transport header, the IPv6 packet MUST be discarded.

## [5.](#) Security Considerations

The OI option provides a method for nodes which have no interest in parsing the header chain to quickly process IP packets. Because transport layer security protocols do not cover extension headers, and the information in the IPv6 header is sufficient to generate the pseudo-header for upper layer protocols, the skipping of extension headers will not impact the security verification performed by transport layer security protocols. However, in IPsec the situation is a little different. Because the ESP header [[RFC4303](#)] or the AH header [[RFC4302](#)] consist of critical information to process the IPsec packet and the extension headers after the ESP or AH header may have to be authenticated or encrypted, these extension headers cannot be skipped over. Therefore, a IPsec implementation MUST NOT skip to the end of the header chain under the instruction of the OI option.

This specification disallows use of the OI option in fragmented packets. In addition to efficiency considerations, this prevents the option from becoming a vector for a buffer overflow attack.

Attackers cannot use the OI option to hide any undesired information in the IPv6 header, because this option is only an optional indication for intermediate devices that do not in any case wish to inspect such information. Security devices may simply ignore this indication and verify every extension header in the chain.

## [6.](#) IANA Considerations

IANA is requested to assign the IPv6 Option Type TBD1 for the Offset Indicating Option and record it in the IPv6 Destination Options and Hop-by-Hop Options registry.

In accordance with [Section 4.2 of \[RFC2460\]](#), this option type has the two most significant bits set to 00 (skip if unrecognized) and the

third-highest-order bit set to 1 (option data may change en-route). This is in case a future IPv6 extension header type may be defined whose size may change en-route, requiring the Offset value to be updated.

Note to RFC Editor: please replace TBD1 with the value assigned by IANA and delete this note.

## [7.](#) Acknowledgements

Valuable comments on this draft were made by Thomas Narten.

## [8.](#) References

### [8.1.](#) Normative References

- [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.

### [8.2.](#) Informative References

- [RFC4302] Kent, S., "IP Authentication Header", [RFC 4302](#), December 2005.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), December 2005.
- [RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/

Zhang, et al. Expires March 19, 2012 [Page 6]

---

Internet-Draft Offset Indicating Option September 2011

Co-existence Security Considerations", [RFC 4942](#),  
September 2007.

## Authors' Addresses

Dacheng Zhang  
Huawei Technologies Co., Ltd

Huawei Building, No.3 Xinxu Rd.,  
Shang-Di Information Industry Base, Hai-Dian District, Beijing  
P.R. China

Email: zhangdacheng@huawei.com

Sheng Jiang  
Huawei Technologies Co., Ltd  
Huawei Building, No.3 Xinxu Rd.,  
Shang-Di Information Industry Base, Hai-Dian District, Beijing  
P.R. China

Email: jiangsheng@huawei.com

Brian Carpenter  
Department of Computer Science  
University of Auckland  
PB 92019  
Auckland, 1142  
New Zealand

Email: brian.e.carpenter@gmail.com