

Internet Draft  
Document: [draft-ietf-ipsec-ikev2-ecnfix-01.txt](#)  
Expires: August 2003

David L. Black  
EMC Corporation  
February 2003

## IKEv2: ECN Requirements for IPsec Tunnels

### Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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>.

### Abstract

IPsec (IP Security) tunnel encapsulation and decapsulation were specified prior to the addition of ECN (Explicit Congestion Notification) to IP, with the potential result that ECN congestion indications could be discarded by IPsec tunnel decapsulation. The current ECN specification includes two ECN operating modes for IPsec tunnels to avoid this situation, plus IKEv1/ISAKMP (Internet Key Exchange/Internet Security Association and Key Management Protocol) negotiation extensions to enable ECN to be used correctly with IPsec tunnels. To simplify this situation, IKEv2 requires changes to tunnel decapsulation that prevent discarding of ECN congestion indication, obviating the need for these multiple ECN operating modes and their associated negotiation support.

## IKEv2: ECN Requirements for IPsec Tunnels February 2003

## Table of Contents

<a href="#">1. Introduction.....</a>	<a href="#">2</a>
<a href="#">2. Conventions used in this document.....</a>	<a href="#">2</a>
<a href="#">3. The ECN and DS Fields in IP headers.....</a>	<a href="#">3</a>
<a href="#">4. ECN vs. IPsec: The Problem.....</a>	<a href="#">3</a>
<a href="#">5. IPsec Changes.....</a>	<a href="#">4</a>
<a href="#">6. Security Considerations.....</a>	<a href="#">5</a>
Normative References.....	<a href="#">5</a>
Informative References.....	<a href="#">6</a>
Author's Address.....	<a href="#">6</a>

[1. Introduction](#)

IPsec tunnel encapsulation and decapsulation were specified [[RFC 2401](#)] prior to the addition of ECN (Explicit Congestion Notification) to IP [[RFC 3168](#)], with the potential result that ECN congestion indications could be discarded by IPsec tunnel decapsulation. The original ECN specification [[RFC 3168](#)] specified two ECN operating modes for IPsec tunnels to avoid this situation, plus IKEv1/ISAKMP negotiation extensions to enable ECN to be used correctly with IPsec tunnels. To simplify this situation, IPsec implementations that support IKEv2 [[IKEv2](#)] MUST implement changes to tunnel decapsulation that prevent discarding of ECN congestion indications, obviating the need for these multiple ECN operating modes and their associated negotiation support.

This document specifies the required changes to IPsec tunnel decapsulation, and updates both [[RFC 2401](#)] (IP Security Architecture) and [[RFC 3168](#)] (The Addition of ECN to IP). In turn, this document is intended to be obsoleted by an updated IP Security Architecture RFC (revision to [RFC 2401](#)) when time permits. This document is necessary at this time to prevent deployment of IKEv2 implementations that discard ECN congestion notifications; such deployment would require perpetuating the two ECN operating modes and the ECN negotiation support for IKEv2.

[2. Conventions used in this document](#)

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [[RFC 2119](#)].


The term "router" is used to refer to all IP network nodes involved in the forwarding of IP traffic between its sender and receiver.

---

IKEv2: ECN Requirements for IPsec Tunnels February 2003

### 3. The ECN and DS Fields in IP headers

Both the IPv4 TOS byte and the IPv6 traffic class octet are divided into a 6-bit DS (Differentiated Services) Field and a 2-bit ECN field [RFC 2474, [RFC 2780](#), [RFC 3168](#)] as follows:



DSCP: differentiated services codepoint

ECN: Explicit Congestion Notification

Figure 2: The Differentiated Services and ECN Fields in IP.

[Section 23.1 of \[RFC 3168\]](#) specifies the ECN field to consist of the two least significant bits of the IPv4 TOS Byte and IPv6 Traffic Class Octet and defines the following four values for that field:

Bits 6-7, ECN Field:

Binary	Keyword	References
-----	-----	-----
00	Not-ECT (Not ECN-Capable Transport)	<a href="#">[RFC 3168]</a>
01	ECT(1) (ECN-Capable Transport(1))	<a href="#">[RFC 3168]</a>
10	ECT(0) (ECN-Capable Transport(0))	<a href="#">[RFC 3168]</a>
11	CE (Congestion Experienced)	<a href="#">[RFC 3168]</a>

Figure 1: The Values of the ECN Field.

The not-ECT codepoint '00' indicates a packet that is not using ECN. The ECN-Capable Transport (ECT) codepoints '10' and '01' are set by the data sender to indicate that the end-points of the transport

protocol are ECN-capable; they are called ECT(0) and ECT(1) respectively. The phrase "the ECT codepoint" in this document refers to either of the two ECT codepoints, which are treated equivalently by routers. Senders are free to use either the ECT(0) or the ECT(1) codepoint to indicate ECT, on a packet-by-packet basis. The CE codepoint '11' is set by a router to indicate congestion to the end nodes. Routers that encounter a packet arriving at a full queue drop the packet, just as they do in the absence of ECN. See [[RFC 3168](#)] for more ECN information.

#### [4.](#) ECN vs. IPsec: The Problem

Sections [5.1.2.1](#) and [5.1.2.2](#) of [[RFC 2401](#)] specify that the IPv4 TOS byte and IPv6 traffic class octet are to be copied from the

### IKEv2: ECN Requirements for IPsec Tunnels February 2003

inner header to the outer header by the IPsec tunnel encapsulator and that these fields in the outer header are to be discarded (no change to inner header) by the IPsec tunnel decapsulator. If ECN is in use, ECT codepoints will be copied to the outer header, but when a router within the tunnel changes an ECT codepoint to a CE codepoint to indicate congestion, that indication will be discarded by the decapsulator (the inner header's ECT codepoint will be forwarded). This behavior is highly undesirable, and [Section 9.2 of \[RFC 3168\]](#) specifies changes to IPsec to avoid it. These changes include two ECN operating modes and negotiation support to detect and cope with IPsec decapsulators that discard ECN congestion indications; use of ECN in the outer IP header of IPsec tunnels is not permitted when such discarding is a possibility.

#### [5.](#) IPsec Changes

In order to avoid multiple ECN operating modes and negotiation, IPsec tunnel decapsulators for tunnel-mode Security Associations (SAs) created by IKEv2 MUST implement the following modifications to prevent discarding of ECN congestion indications. IKEv2 tunnel-mode SA negotiation is performed by the USE-TRANSPORT-MODE Notify Message (see Section 5.10.1 of [[IKEv2](#)]). The following IPsec modifications UPDATE [Section 9.2 of \[RFC 3168\]](#) and Sections [5.1.2.1](#) and [5.1.2.2](#) of [[RFC 2401](#)].

Encapsulation and Decapsulation of packets for a tunnel-mode SA created by IKEv2 MUST NOT follow the modifications specified by [Section 9.2 of \[RFC 3168\]](#) and its subsections. Instead, the

following modifications to encapsulation and decapsulation in Sections [5.1.2.1](#) and [5.1.2.2](#) of [[RFC 2401](#)] MUST be performed:

	Outer Hdr at Encapsulator	Inner Hdr at Decapsulator
IPv4		
Header fields:	-----	-----
DS Field	copied from inner hdr (5)	no change
ECN Field	copied from inner hdr	constructed (7)
IPv6		
Header fields:		
DS Field	copied from inner hdr (6)	no change
ECN Field	copied from inner hdr	constructed (7)

(5)(6) If the packet will immediately enter a domain for which the DSCP value in the outer header is not appropriate, that value MUST be mapped to an appropriate value for the domain [[RFC 2474](#)]. See [[RFC 2475](#)] for further information.

(7) If the ECN field in the inner header is set to ECT(0) or ECT(1) and the ECN field in the outer header is set to CE, then

---

## IKEv2: ECN Requirements for IPsec Tunnels February 2003

set the ECN field in the inner header to CE, otherwise make no change to the ECN field in the inner header.

(5) and (6) are identical to match the original usage in [[RFC 2401](#)], where they are different. These actions are not related to ECN, but are part of Differentiated Services support, and are carried over to this document from [[RFC 3168](#)] so that all of [[RFC 3168](#)]'s changes to IPsec can be made inapplicable to SAs created by IKEv2. [Section 9.2 of \[RFC 3168\]](#) continues to apply to IPsec tunnel-mode Security Associations created by IKEv1.

## [6. Security Considerations](#)

[[RFC 3168](#)] contains an extensive discussion of the security considerations of adding ECN to IP, including considerations specific to IPsec. This document is based on those considerations and makes ECN support for IPsec tunnels MANDATORY as opposed to [[RFC 3168](#)]'s treatment of it as a matter of security policy. See [[RFC 3168](#)] for a full discussion of ECN security considerations.

## Normative References

- [IKEv2] Kaufman, C. (ed), "Internet Key Exchange (IKEv2) Protocol", [draft-ietf-ipsec-ikev2-04.txt](#), Work in Progress, January 2003.
- [RFC 2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC 2401] Kent, S. and R. Atkinson, Security Architecture for the Internet Protocol, [RFC 2401](#), November 1998.
- [RFC 2474] Nichols, K., Blake, S., Baker, F. and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), December 1998.
- [RFC 2780] Bradner S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", [BCP 37](#), [RFC 2780](#), March 2000.
- [RFC 3168] Ramakrishnan, K., Floyd, S. and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), September 2001.

Black

Expires - July 2003

[Page 5]

---

IKEv2: ECN Requirements for IPsec Tunnels February 2003

## Informative References

- [RFC 2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Services", [RFC 2475](#), December 1998.

## Author's Address

David L. Black  
EMC Corporation  
176 South Street  
Hopkinton, MA, 01748, USA

Phone: +1 (508) 293-7953  
Email: [black\\_david@emc.com](mailto:black_david@emc.com)

## Acknowledgements

Significant portions of the text of this document were copied or adapted from text in [RFC 3168](#). The contributions of the authors of [RFC 3168](#) are hereby acknowledged.

#### Full Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

#### Intellectual Property Rights Notices

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and

standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.