

IPsecme  
Internet-Draft  
Updates: 4301 (if approved)  
Intended status: Experimental  
Expires: 21 August 2026

D. Migault  
J. Halpern  
S. Preda  
D. Liu  
U. Parkholm  
Ericsson  
17 February 2026

Differentiated Services Field Codepoints Internet Key Exchange version 2  
Notification  
draft-mglt-ipsecme-dscp-np-05

Abstract

This document outlines the DSCP Notification Payload, which, during a CREATE\_CHILD\_SA Exchange, explicitly indicates the DSCP code points that will be encapsulated in the newly established tunnel. This document updates RFC 4301.

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 <https://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 21 August 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

#### Table of Contents

1. Requirements Notation . . . . .	2
2. Introduction . . . . .	3
3. RFC4301 Clarification . . . . .	3
4. Protocol Overview . . . . .	4
5. Protocol Description . . . . .	6
6. Payload Description . . . . .	7
7. IANA Considerations . . . . .	7
8. Security Considerations . . . . .	8
9. Acknowledgements . . . . .	8
10. References . . . . .	8
10.1. Normative References . . . . .	8
10.2. Informative References . . . . .	9
Authors' Addresses . . . . .	9

#### 1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 2. Introduction

In the ESP Header Compression Profile Diet-ESP [I-D.ietf-ipsecme-diet-esp], two communicating peers can reach an agreement on DSCP values as part of a compression context. Within this context, DSCP values serve not only as classifiers but are also compressed by the sending peer and subsequently decompressed by the receiving peer for both incoming and outgoing traffic. This process necessitates a mutual agreement on DSCP values for a specific pair of Security Associations (SAs). The DSCP Notification Payload outlined in this specification facilitates the negotiation of these DSCP values for a pair of SAs during the CREATE\_CHILD\_SA Exchange.

Furthermore, the explicit negotiation of DSCP values enhances the "classifier" mechanism, allowing for the establishment of this mechanism and the agreement on various clusters of DSCP values to be considered for each pair of SAs. [RFC4301], Section 4.1 recognizes that aggregating traffic with multiple DSCP values over a single SA may lead to the inappropriate discarding of lower-priority packets due to the windowing mechanism employed by this feature. To mitigate this issue, [RFC4301], Section 4.1 advises that the sender implement a "classifier" mechanism to distribute traffic across multiple SAs. While [RFC4301], Section 4.4.2.1 refers to the "DSCP values" fields in the Security Association Database (SAD), [RFC7296] does not provide a way for peers to indicate which classification is ongoing nor which "DSCP values" are linked to the created SA. This document addresses that deficiency by specifying the DSCP Notification Payload, which explicitly identifies the DSCP code points that will be tunneled in the newly established tunnel during a CREATE\_CHILD\_SA Exchange.

It is essential to recognize that in a standard "classifier" context, there is no necessity for the same cluster of DSCP values to be linked to both the inbound and outbound Security Associations (SAs) of a specific pair. Typically, one peer may employ one pair of SAs, while the other peer may choose a different pair. This flexibility arises because DSCP values are applied exclusively by the sending node, rather than the receiving node. Although the use of the DSCP Notification Payload does not inhibit such configurations, it is likely to diminish their occurrence. Conversely, we anticipate that the explicit negotiation of DSCP values and pairs of SAs will facilitate the management of these "classifiers."

## 3. RFC4301 Clarification

[RFC4301], Section 4.4.2.1 mentions

- \* DSCP values -- the set of DSCP values allowed for packets carried over this SA. If no values are specified, no DSCP-specific filtering is applied. If one or more values are specified, these are used to select one SA among several that match the traffic selectors for an outbound packet. Note that these values are NOT checked against inbound traffic arriving on the SA.

The text does not clearly specify what happens when the DSCP of a packet does not match any of the corresponding DSCP values. This document proposes the following text:

- \* DSCP values -- the set of DSCP values allowed for packets carried over this SA. If no values are specified, no DSCP-specific filtering is applied. If one or more values are specified, these are used to select one SA among several that match the traffic selectors for an outbound packet. In case of multiple matches a preference to the most selective list DSCP value could be implemented by the peer's policy. If the DSCP value of the packet does not match any of the DSCP values provided by the associated matching SAs and there is at least one SA with no DSCP-specific filtering, then, one of these SA SHOULD be selected. On the other hand, if all SAs have DSCP filtering, then, any of the matching SAs can be selected. Note that these values MUST NOT be checked against inbound traffic arriving on the SA.

#### 4. Protocol Overview

The illustrative example of this section considers the following use case:

- \* Expedited Forwarding (EF) with low latency traffic has its own IPsec tunnel,
- \* Assured Forward (AF) classes with different drop precedence (which may take a different route) have their own tunnel,
- \* and all remaining DSCP values are put into a third tunnel.

This section details how a peer uses the DSCP Notify Payload to classify traffic carrying the DSCP values AF11 or AF3 in one tunnel, traffic carrying a DSCP value of EF in another tunnel, and traffic with other DSCP values in a third tunnel. The third SA is designated as the default no-DSCP specific SA. It is RECOMMENDED to configure the Security Policy Data Base (SPD), so that such a default no-DSCP specific SA is created and it is RECOMMENDED its creation happens prior to the SA with specific DSCP values. Note that according to Section 3, there is no specific ordering, but starting with the no-DSCP specific SA ensures compatibility with an IPsec implementation that would for example discard or create a new SA when the DSCP does not match.

Generally, it is recommended that the outer DSCP value matches the inner DSCP value so that the tunneled packet be treated similarly to the inner packet. Such behavior is provided by setting the Bypass DSCP to True. If the initiator prefers for example every tunneled packet being treated similarly, then, an explicit mapping needs to be indicated. Typically, the initiator may be willing to prevent reordered traffic to fall outside the anti-replay windows. Note that such policy is implemented by each peer.

Initiator	Responder
-----	
HDR, SK {IDi, [CERT,] [CERTREQ,] [IDr,] AUTH, SAi2, TSi, TSr} -->	
	<-- HDR, SK {IDr, [CERT,] AUTH, SAr2, TSi, TSr}

Once the no-DSCP specific SA is created, all traffic with any DSCP value is steered to that SA. The initiator then creates the child SA associated with specific DSCP values. In this example, it creates the SA associated with the DSCP value AF11 or AF3, followed by the one associated with value of EF, but this does not follow any specific ordering. The initiator specifies the DSCP values being classified in that SA with a DSCP Notify Payload that carries the DSCP values.

If the responder supports the DSCP Notify Payload, it SHOULD respond with a Notify Payload that indicates the DSCP values selected for that tunnel. By default these values SHOULD be the ones specified by the initiator, but the responder's policy MAY select other values. If the responder does not want to perform DSCP filtering, the responder SHOULD send an empty DSCP Notify Payload, in order to at least indicate support for the DSCP Notify Payload.

As specified in [RFC7296], Section 3.10.1, a Notify Payload with status type MUST be ignored if it is not recognized. The absence of a DSCP Notify Payload by the responder may be due to the responder not supporting the notification, or not advertising the application of DSCP filtering. We do not consider that the absence of classification by the responder prevents the SA from being created. The classification is at least performed for the outbound stream, which is sufficient to justify the creation of the additional SA. Note also that DSCP values are not agreed, and the responder cannot for example narrow down the list of DSCP values being classified. If that would cause a significant issue, the responder can create another SA with the narrowed-down list of DSCP values. The responder may also REKEY\_SA the previous SA to redefine the DSCP values to be considered.

When multiple DSCP values are indicated, and the initiator is mapping the outer DSCP value, the outer DSCP value is expected to be one of these values.

Initiator	Responder
-----	
HDR, SK {SA, Ni, KEi, N(DSCP, AF11, AF3)} -->	
	<-- HDR, SK {SA, Nr, KEr, N(DSCP, AF11, AF3)}

The initiator may then create additional child SAs specifying other DSCP values.

Initiator	Responder
-----	
HDR, SK {SA, Ni, KEi, N(DSCP, EE)} -->	
	<-- HDR, SK {SA, Nr, KEr}

## 5. Protocol Description

During the CREATE\_CHILD\_SA exchange, the initiator or the responder MAY indicate to the other peer the DSCP filtering policy applied to the SA. This is done via the DSCP Notify Payload indicating the DSCP values being considered for that SA.

The initiator MAY send an empty DSCP Notify Payload to indicate support of the DSCP Notify Payload as well as an indication the negotiated SA as a no-DSCP specific SA. This SA MAY be followed by the creation of the DSCP-specific SA.

Upon receiving a DSCP Notify Payload, if the responder supports the notification it SHOULD respond with a DSCP Notify Payload. The value indicated SHOULD be the one selected by the initiator.

There is no specific error handling.

## 6. Payload Description

The DSCP Notify Payload is based on the format of the Notify Payload as described in [RFC7296], Section 3.10 and represented in Figure 1.




Figure 1: Notify Payload

The fields Next Payload, Critical Bit, RESERVED, and Payload Length are defined in [RFC7296]. Specific fields defined in this document are:

- \* Protocol ID (1 octet): Set to zero.
- \* Security Parameter Index (SPI) Size (1 octet): Set to zero.
- \* Notify Message Type (2 octets): Specifies the type of notification message. It is set to DSCP\_VALUES (see Section 7).
- \* Notification Data (variable length): lists the DSCP values that are considered for the SA. Each value is encoded over a single byte.

## 7. IANA Considerations

IANA is requested to allocate one value in the "IKEv2 Notify Message Types - Status Types" registry: (available at <https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml#ikev2-parameters-16>) with the following definition:

Value	Notify Messages - Status Types
TBD	DSCP

## 8. Security Considerations

As the DSCP value field is already defined by [RFC4301] in the SA structure, the security considerations of [RFC4301] apply. The DSCP Notification Payload communicates clearly the DSCP value field to the responder.

When the tunnel mode is used, the communication of the DSCP value field could be easily interpreted by monitoring the received DSCP values of the inner traffic when that traffic is encapsulated, and so no secret information is revealed. When the transport mode is used, that value may be changed by the network and eventually, the value of the field could be unknown to the other peer. However, this cannot be considered as a protection mechanism, and the communication of the DSCP value cannot be considered as revealing information that was previously not revealed.

The notification of the set of DSCP values to the other peer does not require additional resources either for the initiator or for the receiver. The SA is created either with DSCP values or without.

Similarly, the notification of the set of DSCP values to the other peer does not introduce additional constraints on the traffic. First, the responder may also ignore the DSCP Notification Payload. Then, when an SA is associated with a set of DSCP values, this does not prevent the other peer from sending traffic with a different DSCP value over that SA. In other words, traffic coming with unexpected DSCP values is not rejected as would have been the case if the DSCP values had been considered as Traffic Selectors.

## 9. Acknowledgements

We would like to thank Scott Fluhrer for his useful comments; Valery Smyslov, Tero Kivinen for their design suggestions we carefully followed. We would also like to thank William Atwood for his careful review and suggestions.

## 10. References

### 10.1. Normative References



- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October 2014, <<https://www.rfc-editor.org/info/rfc7296>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

## 10.2. Informative References

- [I-D.ietf-ipsecme-diet-esp]  
Migault, D., Hatami, M., Cespedes, S., Atwood, J. W., Liu, D., Guggemos, T., Bormann, C., and D. Schinazi, "ESP Header Compression with Diet-ESP", Work in Progress, Internet-Draft, draft-ietf-ipsecme-diet-esp-09, 17 August 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-ipsecme-diet-esp-09>>.

## Authors' Addresses

Daniel Migault  
Ericsson  
Email: [daniel.migault@ericsson.com](mailto:daniel.migault@ericsson.com)

Joel Halpern  
Ericsson  
Email: [joel.halpern@ericsson.com](mailto:joel.halpern@ericsson.com)

Stere Preda  
Ericsson  
Email: [stere.preda@ericsson.com](mailto:stere.preda@ericsson.com)

Daiying Liu  
Ericsson  
Email: [harold.liu@ericsson.com](mailto:harold.liu@ericsson.com)

U. Parkholm  
Ericsson  
Email: [ulf.x.parkholm@ericsson.com](mailto:ulf.x.parkholm@ericsson.com)