

Internet Engineering Task Force (IETF)
Request for Comments: 6437
Obsoletes: 3697
Updates: 2205, 2460
Category: Standards Track
ISSN: 2070-1721

S. Amante
Level 3
B. Carpenter
Univ. of Auckland
S. Jiang
Huawei
J. Rajahalme
Nokia Siemens Networks
November 2011

IPv6 Flow Label Specification

Abstract

This document specifies the IPv6 Flow Label field and the minimum requirements for IPv6 nodes labeling flows, IPv6 nodes forwarding labeled packets, and flow state establishment methods. Even when mentioned as examples of possible uses of the flow labeling, more detailed requirements for specific use cases are out of the scope for this document.

The usage of the Flow Label field enables efficient IPv6 flow classification based only on IPv6 main header fields in fixed positions.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6437>.

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.

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. Introduction 3
2. IPv6 Flow Label Specification 4
3. Flow Labeling Requirements in the Stateless Scenario 5
4. Flow State Establishment Requirements 7
5. Essential Correction to RFC 2205 7
6. Security Considerations 7
6.1. Covert Channel Risk 8
6.2. Theft and Denial of Service 8
6.3. IPsec and Tunneling Interactions 10
6.4. Security Filtering Interactions 11
7. Differences from RFC 3697 11
8. Acknowledgements 11
9. References 12
9.1. Normative References 12
9.2. Informative References 12
Appendix A. Example 20-Bit Hash Function 14

1. Introduction

From the viewpoint of the network layer, a flow is a sequence of packets sent from a particular source to a particular unicast, anycast, or multicast destination that a node desires to label as a flow. From an upper-layer viewpoint, a flow could consist of all packets in one direction of a specific transport connection or media stream. However, a flow is not necessarily 1:1 mapped to a transport connection.

Traditionally, flow classifiers have been based on the 5-tuple of the source address, destination address, source port, destination port, and the transport protocol type. However, some of these fields may be unavailable due to either fragmentation or encryption, or locating them past a chain of IPv6 extension headers may be inefficient. Additionally, if classifiers depend only on IP-layer headers, later introduction of alternative transport-layer protocols will be easier.

The usage of the 3-tuple of the Flow Label, Source Address, and Destination Address fields enables efficient IPv6 flow classification, where only IPv6 main header fields in fixed positions are used.

The flow label could be used in both stateless and stateful scenarios. A stateless scenario is one where any node that processes the flow label in any way does not need to store any information about a flow before or after a packet has been processed. A stateful scenario is one where a node that processes the flow label value needs to store information about the flow, including the flow label value. A stateful scenario might also require a signaling mechanism to inform downstream nodes that the flow label is being used in a certain way and to establish flow state in the network. For example, RSVP [RFC2205] and General Internet Signaling Transport (GIST) [RFC5971] can signal flow label values.

The flow label can be used most simply in stateless scenarios. This specification concentrates on the stateless model and how it can be used as a default mechanism. Details of stateful models, signaling, specific flow state establishment methods, and their related service models are out of scope for this specification. The basic requirement for stateful models is set forth in Section 4.

The minimum level of IPv6 flow support consists of labeling the flows. A specific goal is to enable and encourage the use of the flow label for various forms of stateless load distribution, especially across Equal Cost Multi-Path (ECMP) and/or Link Aggregation Group (LAG) paths. ECMP and LAG are methods to bond together multiple physical links used to procure the required

capacity necessary to carry an offered load greater than the bandwidth of an individual physical link. Further details are in a separate document [RFC6438]. IPv6 source nodes SHOULD be able to label known flows (e.g., TCP connections and application streams), even if the node itself does not require any flow-specific treatment. Node requirements for stateless flow labeling are given in Section 3.

This document replaces [RFC3697] and Section 6 and Appendix A of [RFC2460]. A rationale for the changes made is documented in [RFC6436]. The present document also includes a correction to [RFC2205] concerning the flow label.

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 [RFC2119].

2. IPv6 Flow Label Specification

The 20-bit Flow Label field in the IPv6 header [RFC2460] is used by a node to label packets of a flow. A Flow Label of zero is used to indicate packets that have not been labeled. Packet classifiers can use the triplet of Flow Label, Source Address, and Destination Address fields to identify the flow to which a particular packet belongs. Packets are processed in a flow-specific manner by nodes that are able to do so in a stateless manner or that have been set up with flow-specific state. The nature of the specific treatment and the methods for flow state establishment are out of scope for this specification.

Flow label values should be chosen such that their bits exhibit a high degree of variability, making them suitable for use as part of the input to a hash function used in a load distribution scheme. At the same time, third parties should be unlikely to be able to guess the next value that a source of flow labels will choose.

In statistics, a discrete uniform distribution is defined as a probability distribution in which each value in a given range of equally spaced values (such as a sequence of integers) is equally likely to be chosen as the next value. The values in such a distribution exhibit both variability and unguessability. Thus, as specified in Section 3, an approximation to a discrete uniform distribution is preferable as the source of flow label values. Intentionally, there are no precise mathematical requirements placed on the distribution or the method used to achieve such a distribution.

Once set to a non-zero value, the Flow Label is expected to be delivered unchanged to the destination node(s). A forwarding node MUST either leave a non-zero flow label value unchanged or change it only for compelling operational security reasons as described in Section 6.1.

There is no way to verify whether a flow label has been modified en route or whether it belongs to a uniform distribution. Therefore, no Internet-wide mechanism can depend mathematically on unmodified and uniformly distributed flow labels; they have a "best effort" quality. Implementers should be aware that the flow label is an unprotected field that could have been accidentally or intentionally changed en route (see Section 6). This leads to the following formal rule:

- o Forwarding nodes such as routers and load distributors MUST NOT depend only on Flow Label values being uniformly distributed. In any usage such as a hash key for load distribution, the Flow Label bits MUST be combined at least with bits from other sources within the packet, so as to produce a constant hash value for each flow and a suitable distribution of hash values across flows. Typically, the other fields used will be some or all components of the usual 5-tuple. In this way, load distribution will still occur even if the Flow Label values are poorly distributed.

Although uniformly distributed flow label values are recommended below, and will always be helpful for load distribution, it is unsafe to assume their presence in the general case, and the use case needs to work even if the flow label value is zero.

As a general practice, packet flows should not be reordered, and the use of the Flow Label field does not affect this. In particular, a Flow label value of zero does not imply that reordering is acceptable.

3. Flow Labeling Requirements in the Stateless Scenario

This section defines the minimum requirements for methods of setting the flow label value within the stateless scenario of flow label usage.

To enable Flow-Label-based classification, source nodes SHOULD assign each unrelated transport connection and application data stream to a new flow. A typical definition of a flow for this purpose is any set of packets carrying the same 5-tuple {dest addr, source addr, protocol, dest port, source port}. It should be noted that a source node always has convenient and efficient access to this 5-tuple, which is not always the case for nodes that subsequently forward the packet.

It is desirable that flow label values should be uniformly distributed to assist load distribution. It is therefore RECOMMENDED that source hosts support the flow label by setting the flow label field for all packets of a given flow to the same value chosen from an approximation to a discrete uniform distribution. Both stateful and stateless methods of assigning a value could be used, but it is outside the scope of this specification to mandate an algorithm. The algorithm SHOULD ensure that the resulting flow label values are unique with high probability. However, if two simultaneous flows are assigned the same flow label value by chance and have the same source and destination addresses, it simply means that they will receive the same flow label treatment throughout the network. As long as this is a low-probability event, it will not significantly affect load distribution.

A possible stateless algorithm is to use a suitable 20-bit hash of values from the IP packet's 5-tuple. A simple example hash function is described in Appendix A.

An alternative approach is to use a pseudo-random number generator to assign a flow label value for a given transport session; such a method will require minimal local state to be kept at the source node by recording the flow label associated with each transport socket.

Viewed externally, either of these approaches will produce values that appear to be uniformly distributed and pseudo-random.

An implementation in which flow labels are assigned sequentially is NOT RECOMMENDED, as it would then be simple for on-path observers to guess the next value.

A source node that does not otherwise set the flow label MUST set its value to zero.

A node that forwards a flow whose flow label value in arriving packets is zero MAY change the flow label value. In that case, it is RECOMMENDED that the forwarding node sets the flow label field for a flow to a uniformly distributed value as just described for source nodes.

- o The same considerations apply as to source hosts setting the flow label; in particular, the preferred case is that a flow is defined by the 5-tuple. However, there are cases in which the complete 5-tuple for all packets is not readily available to a forwarding node, in particular for fragmented packets. In such cases, a flow can be defined by fewer IPv6 header fields, typically using only the 2-tuple {dest addr, source addr}. There are alternative approaches that implementers could choose, such as:

- * A forwarding node might use the 5-tuple to define a flow whenever possible but use the 2-tuple when the complete 5-tuple is not available. In this case, unfragmented and fragmented packets belonging to the same transport session would receive different flow label values, altering the effect of subsequent load distribution based on the flow label.
- * A forwarding node might use the 2-tuple to define a flow in all cases. In this case, subsequent load distribution would be based only on IP addresses.
- o The option to set the flow label in a forwarding node, if implemented, would presumably be of value in first-hop or ingress routers. It might place a considerable per-packet processing load on them, even if they adopted a stateless method of flow identification and label assignment. However, it will not interfere with host-to-router load sharing [RFC4311]. It needs to be under the control of network managers, to avoid unwanted processing load and any other undesirable effects. For this reason, it MUST be a configurable option, disabled by default.

The preceding rules taken together allow a given network to include routers that set flow labels on behalf of hosts that do not do so. The complications described explain why the principal recommendation is that the source hosts should set the label.

4. Flow State Establishment Requirements

A node that sets the flow label MAY also take part in a flow state establishment method that results in assigning specific treatments to specific flows, possibly including signaling. Any such method MUST NOT disturb nodes taking part in the stateless scenario just described. Thus, any node that sets flow label values according to a stateful scheme MUST choose labels that conform to Section 3 of this specification. Further details are not discussed in this document.

5. Essential Correction to RFC 2205

[RFC2460] reduced the size of the flow label field from 24 to 20 bits. The references to a 24-bit flow label field in Section A.9 of [RFC2205] are updated accordingly.

6. Security Considerations

This section considers security issues raised by the use of the Flow Label, including the potential for denial-of-service attacks and the related potential for theft of service by unauthorized traffic (Section 6.2). Section 6.3 addresses the use of the Flow Label in

the presence of IPsec, including its interaction with IPsec tunnel mode and other tunneling protocols. We also note that inspection of unencrypted Flow Labels may allow some forms of traffic analysis by revealing some structure of the underlying communications. Even if the flow label was encrypted, its presence as a constant value in a fixed position might assist traffic analysis and cryptanalysis.

The flow label is not protected in any way, even if IPsec authentication [RFC4302] is in use, so it can be forged by an on-path attacker. Implementers are advised that any en-route change to the flow label value is undetectable. On the other hand, a uniformly distributed pseudo-random flow label cannot be readily guessed by an attacker; see [LABEL-SEC] for further discussion. If a hash algorithm is used, as suggested in Section 3, it SHOULD include a step that makes the flow label value significantly difficult to predict [RFC4086], even with knowledge of the algorithm being used.

6.1. Covert Channel Risk

The flow label could be used as a covert data channel, since apparently pseudo-random flow label values could, in fact, consist of covert data [NSA]. This could, for example, be achieved using a series of otherwise innocuous UDP packets whose flow label values constitute a covert message, or by co-opting a TCP session to carry a covert message in the flow labels of successive packets. Both of these could be recognized as suspicious -- the first because isolated UDP packets would not normally be expected to have non-zero flow labels, and the second because the flow label values in a given TCP session should all be equal. However, other methods, such as co-opting the flow labels of occasional packets, might be rather hard to detect.

In situations where the covert channel risk is considered significant, the only certain defense is for a firewall to rewrite non-zero flow labels. This would be an exceptional violation of the rule that the flow label, once set to a non-zero value, must not be changed. To preserve load distribution capability, such a firewall SHOULD rewrite labels by following the method described for a forwarding node (see Section 3), as if the incoming label value were zero, and MUST NOT set non-zero flow labels to zero. This behavior is nevertheless undesirable, since (as discussed in Section 3) only source nodes have straightforward access to the complete 5-tuple.

6.2. Theft and Denial of Service

Since the mapping of network traffic to flow-specific treatment is triggered by the IP addresses and Flow Label value of the IPv6 header, an adversary may be able to obtain a class of service that

the network did not intend to provide by modifying the IPv6 header or by injecting packets with false addresses and/or labels. A concrete analysis of this threat is only possible for specific stateful methods of signaling and using the flow label, which are out of scope for this document. Clearly, a full analysis will be required when any such method is specified, but in general, networks SHOULD NOT make resource allocation decisions based on flow labels without some external means of assurance.

A denial-of-service attack [RFC4732] becomes possible in the stateless model when the modified or injected traffic depletes the resources available to forward it and other traffic streams. If a denial-of-service attack were undertaken against a given Flow Label (or set of Flow Labels), then traffic containing an affected Flow Label might well experience worse-than-best-effort network performance.

Note that since the treatment of IP headers by nodes is typically unverified, there is no guarantee that flow labels sent by a node are set according to the recommendations in this document. A man-in-the-middle or injected-traffic denial-of-service attack specifically directed at flow label handling would involve setting unusual flow labels. For example, an attacker could set all flow labels reaching a given router to the same arbitrary non-zero value or could perform rapid cycling of flow label values such that the packets of a given flow will each have a different value. Either of these attacks would cause a stateless load distribution algorithm to perform badly and would cause a stateful classifier to behave incorrectly. For this reason, stateless classifiers should not use the flow label alone to control load distribution, and stateful classifiers should include explicit methods to detect and ignore suspect flow label values.

Since flows are identified by the 3-tuple of the Flow Label and the Source and Destination Address, the risk of denial of service introduced by the Flow Label is closely related to the risk of denial of service by address spoofing. An adversary who is in a position to forge an address is also likely to be able to forge a label, and vice versa.

There are two issues with different properties: spoofing of the Flow Label only and spoofing of the whole 3-tuple, including Source and Destination Address.

The former can be done inside a node that is using or transmitting the correct source address. The ability to spoof a Flow Label typically implies being in a position to also forge an address, but

in many cases, spoofing an address may not be interesting to the spoofer, especially if the spoofer's goal is theft of service rather than denial of service.

The latter can be done by a host that is not subject to ingress filtering [RFC2827] or by an intermediate router. Due to its properties, this is typically useful only for denial of service. In the absence of ingress filtering, almost any third party could instigate such an attack.

In the presence of ingress filtering, forging a non-zero Flow Label on packets that originated with a zero label, or modifying or clearing a label, could only occur if an intermediate system such as a router was compromised, or through some other form of man-in-the-middle attack.

6.3. IPsec and Tunneling Interactions

The IPsec protocol, as defined in [RFC4301], [RFC4302], and [RFC4303], does not include the IPv6 header's Flow Label in any of its cryptographic calculations (in the case of tunnel mode, it is the outer IPv6 header's Flow Label that is not included). Hence, modification of the Flow Label by a network node has no effect on IPsec end-to-end security, because it cannot cause any IPsec integrity check to fail. As a consequence, IPsec does not provide any defense against an adversary's modification of the Flow Label (i.e., a man-in-the-middle attack).

IPsec tunnel mode provides security for the encapsulated IP header's Flow Label. A tunnel mode IPsec packet contains two IP headers: an outer header supplied by the tunnel ingress node and an encapsulated inner header supplied by the original source of the packet. When an IPsec tunnel is passing through nodes performing flow classification, the intermediate network nodes operate on the Flow Label in the outer header. At the tunnel egress node, IPsec processing includes removing the outer header and forwarding the packet (if required) using the inner header. The IPsec protocol requires that the inner header's Flow Label not be changed by this decapsulation processing to ensure that modifications to the label cannot be used to launch theft- or denial-of-service attacks across an IPsec tunnel endpoint. This document makes no change to that requirement; indeed, it forbids changes to the Flow Label.

When IPsec tunnel egress decapsulation processing includes a sufficiently strong cryptographic integrity check of the encapsulated packet (where sufficiency is determined by local security policy), the tunnel egress node can safely assume that the Flow Label in the inner header has the same value it had at the tunnel ingress node.

This analysis and its implications apply to any tunneling protocol that performs integrity checks. Of course, any Flow Label set in an encapsulating IPv6 header is subject to the risks described in the previous section.

6.4. Security Filtering Interactions

The Flow Label does nothing to eliminate the need for packet filtering based on headers past the IP header if such filtering is deemed necessary for security reasons on nodes such as firewalls or filtering routers.

7. Differences from RFC 3697

The main differences between this specification and its predecessor [RFC3697] are as follows:

- o This specification encourages non-zero flow label values to be used and clearly defines how to set a non-zero value.
- o This specification encourages a stateless model with uniformly distributed flow label values.
- o This specification does not specify any details of a stateful model.
- o This specification retains the rule that the flow label must not be changed en route but allows routers to set the label on behalf of hosts that do not do so.
- o This specification discusses the covert channel risk and its consequences for firewalls.

For further details, see [RFC6436].

8. Acknowledgements

Valuable comments and contributions were made by Jari Arkko, Ran Atkinson, Fred Baker, Richard Barnes, Steve Blake, Tassos Chatzithomaoglou, Remi Despres, Alan Ford, Fernando Gont, Brian Haberman, Tony Hain, Joel Halpern, Qinwen Hu, Chris Morrow, Thomas Narten, Mark Smith, Pascal Thubert, Iljitsch van Beijnum, and other participants in the 6man working group.

Cristian Calude suggested the von Neumann algorithm in Appendix A. David Malone and Donald Eastlake gave additional input about hash algorithms.

Steve Deering and Alex Conta were co-authors of RFC 3697, on which this document is based.

Contributors to the original development of RFC 3697 included Ran Atkinson, Steve Blake, Jim Bound, Francis Dupont, Robert Elz, Tony Hain, Robert Hancock, Bob Hinden, Christian Huitema, Frank Kastenholz, Thomas Narten, Charles Perkins, Pekka Savola, Hesham Soliman, Michael Thomas, Margaret Wasserman, and Alex Zinin.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.

9.2. Informative References

- [LABEL-SEC] Gont, F., "Security Assessment of the IPv6 Flow Label", Work in Progress, November 2010.
- [NSA] Potyraj, C., "Firewall Design Considerations for IPv6", National Security Agency I733-041R-2007, 2007, <http://www.nsa.gov/ia/_files/ipv6/I733-041R-2007.pdf>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC3697] Rajahalme, J., Conta, A., Carpenter, B., and S. Deering, "IPv6 Flow Label Specification", RFC 3697, March 2004.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.

- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.
- [RFC4311] Hinden, R. and D. Thaler, "IPv6 Host-to-Router Load Sharing", RFC 4311, November 2005.
- [RFC4732] Handley, M., Rescorla, E., and IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.
- [RFC5971] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", RFC 5971, October 2010.
- [RFC6436] Amante, S., Carpenter, B., and S. Jiang, "Rationale for Update to the IPv6 Flow Label Specification", RFC 6436, November 2011.
- [RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, November 2011.
- [vonNeumann] von Neumann, J., "Various techniques used in connection with random digits", National Bureau of Standards Applied Math Series 12, 36-38, 1951.

Appendix A. Example 20-Bit Hash Function

As mentioned in Section 3, a stateless hash function may be used to generate a flow label value from an IPv6 packet's 5-tuple. It is not trivial to choose a suitable hash function, and it is expected that extensive practical experience will be required to identify the best choices. An example function, based on an algorithm by von Neumann known to produce an approximately uniform distribution [vonNeumann], follows. For each packet for which a flow label must be generated, execute the following steps:

1. Split the destination and source addresses into two 64-bit values each, thus transforming the 5-tuple into a 7-tuple.
2. Add the following five components together using unsigned 64-bit arithmetic, discarding any carry bits: both parts of the source address, both parts of the destination address, and the protocol number.
3. Apply the von Neumann algorithm to the resulting string of 64 bits:
 1. Starting at the least significant end, select two bits.
 2. If the two bits are 00 or 11, discard them.
 3. If the two bits are 01, output a 0 bit.
 4. If the two bits are 10, output a 1 bit.
 5. Repeat with the next two bits in the input 64-bit string.
 6. Stop when 16 bits have been output (or when the 64-bit string is exhausted).
4. Add the two port numbers to the resulting 16-bit number.
5. Shift the resulting value 4 bits left, and mask with 0xffff.
6. In the highly unlikely event that the result is exactly zero, set the flow label arbitrarily to the value 1.

Note that this simple example does not include a step to prevent predictability, as recommended in Section 6.

Authors' Addresses

Shane Amante
Level 3 Communications, LLC
1025 Eldorado Blvd
Broomfield, CO 80021
USA

EEmail: shane@level3.net

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

EEmail: brian.e.carpenter@gmail.com

Sheng Jiang
Huawei Technologies Co., Ltd
Q14, Huawei Campus
No.156 Beiqing Road
Hai-Dian District, Beijing 100095
P.R. China

EEmail: jiangsheng@huawei.com

Jarno Rajahalme
Nokia Siemens Networks
Linnoitustie 6
02600 Espoo
Finland

EEmail: jarno.rajahalme@nsn.com