Internet Engineering Task Force (IETF)
Request for Comments: 6436
Category: Informational
ISSN: 2070-1721

S. Amante
Level 3
B. Carpenter
Univ. of Auckland
S. Jiang
Huawei

November 2011

Rationale for Update to the IPv6 Flow Label Specification

Abstract

Various published proposals for use of the IPv6 flow label are incompatible with its original specification in RFC 3697. Furthermore, very little practical use is made of the flow label, partly due to some uncertainties about the correct interpretation of the specification. This document discusses and motivates changes to the specification in order to clarify it and to introduce some additional flexibility.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc6436.

Table of Contents

1. Introduction

The flow label field in the IPv6 header was reserved but left Experimental by [RFC2460], which mandates only that "Hosts or routers that do not support the functions of the Flow Label field are required to set the field to zero when originating a packet, pass the field on unchanged when forwarding a packet, and ignore the field when receiving a packet."

The flow label field was normatively specified by [RFC3697]. In particular, we quote three rules from that RFC:

a. "The Flow Label value set by the source MUST be delivered unchanged to the destination node(s)."

b. "IPv6 nodes MUST NOT assume any mathematical or other properties of the Flow Label values assigned by source nodes."

c. "Router performance SHOULD NOT be dependent on the distribution of the Flow Label values. Especially, the Flow Label bits alone make poor material for a hash key."

Additionally, RFC 3697 does not define the method a host should adopt by default to choose the value of the flow label, if no specific method is in use. It was expected that various signaling methods might be defined for agreeing on values of the flow label, but no such methods have been standardized, except a pre-existing option in RSVP [RFC2205].

The flow label is hardly used in practice in widespread IPv6 implementations, although some operating systems do set it [McGann05]. To some extent, this is due to the main focus being on basic deployment of IPv6, but the absence of a default method of choosing the flow label value means that most host implementations simply set it to zero. There is also anecdotal evidence that the rules quoted above have led to uncertainty about exactly what is possible. Furthermore, various use cases have been proposed that infringe one or another of the rules. None of these proposals has been accepted as a standard and in practice there is no significant deployment of any mechanism to set the flow label.

The intention of this document is to explain this situation in more detail and to motivate changes to RFC 3697 intended to remove the uncertainties and encourage active usage of the flow label. It does not formally update RFC 3697, but it serves as background material for [RFC6437].

2. Impact of Current Specification

Rule (a) makes it impossible for the routing system to use the flow label as any form of dynamic routing tag. This was a conscious choice in the early design of IPv6, and there appears to be no practical possibility of revisiting this decision at this stage in the deployment of IPv6, which uses conventional routing mechanisms like those used for IPv4. However, this rule also makes it impossible to make any use at all of the flow label unless hosts choose to set it. It also forbids clearing the flow label for security reasons.

This last point highlights the security properties, or rather the lack thereof, with regards to the flow label. The flow label field is always unprotected as it travels through the network, because there is no IPv6 header checksum, and the flow label is not included in transport pseudo-header checksums, nor in IPsec checksums. As a result, intentional and malicious changes to its value cannot be detected. Also, it could be used as a covert data channel, since apparently pseudo-random flow label values could in fact consist of covert data [NIST]. If the flow label were to carry quality-of-service semantics, then like the diffserv code point [RFC2474], it would not be intrinsically trustworthy across domain boundaries. As

a result, some security specialists believe that flow labels should be cleared for safety [LABEL-SEC] [NSA]. These points must be considered when discussing the immutability of the flow label across domain boundaries. In fact, the adjective "immutable" is confusing, since it implies a property that the flow label field does not actually possess. It has therefore been abandoned as a descriptive term in [RFC6437]. It is only used in the present document to explain why it has been abandoned.

Rule (b) appears to forbid any usage in which the bits of the flow label are encoded with a specific semantic meaning. However, the words "MUST NOT assume" are to be interpreted precisely - if a router knows by configuration or by signaling that the flow label has been assigned in a certain way, it can make use of that knowledge. It is not made clear by the rule that there is an implied distinction between stateless models (in which there is no signaling, so no specific assumption about the meaning of the flow label value can be made) and stateful models (in which there is signaling and the router has explicit knowledge about the label).

If the word "alone" is overlooked, rule (c) has sometimes been interpreted as forbidding the use of the flow label as part of a hash used by load distribution mechanisms. In this case too, the word "alone" needs to be taken into account - a router is allowed to combine the flow label value with other data in order to produce a uniformly distributed hash.

Both before and after these rules were laid down, a considerable number of proposals for use of the flow label were published that seem incompatible with them. Numerous examples and an analysis are presented in [RFC6294]. Those examples propose use cases in which some or all of the following apply:

- The flow label may be changed by intermediate systems.
- It doesn’t matter if the flow label is changed, because the receiver doesn’t use it.
- Some or all bits of the flow label are encoded: they have specific meanings understood by routers and switches along the path.
- The encoding is related to the required quality of service, as well as identifying a flow.
- The flow label is used to control forwarding or switching in some way.

These proposals require either some form of semantics encoding in the bits of the flow label, or the ability for routers to modify the flow label, or both. Thus, they appear to infringe the rules from RFC 3697 quoted above.

We can conclude that a considerable number of researchers and designers have been stymied by RFC 3697. On the other hand, some other proposals discussed in [RFC6294] appear to be compatible with RFC 3697. Several are based on the originator of a packet choosing a pseudo-random flow label for each flow, which is one option suggested in RFC 3697. Thus, we can also conclude that there is a useful role for this approach.

If our goal is for the flow label to be used in practice, the conflict between the various approaches creates a dilemma. There appear to be two major options:

1. Discourage locally defined and/or stateful use of the flow label. Strengthen RFC 3697 to say that hosts should set a label value, without necessarily creating state, which would clarify and limit its possible uses. In particular, its use for load distribution and balancing would be encouraged.

2. Relax the rules to encourage locally defined and/or stateful use of the flow label. This approach would make the flow label completely mutable and would exclude use cases depending on strict end-to-end immutability. It would encourage applications of a pseudo-random flow label, such as load distribution, on a local basis, but it would exclude end-to-end applications.

There was considerable debate about these options and their variants during 2010 - 2011, with a variety of proposals in previous versions of this document and in mailing list discussions. After these discussions, there appears to be a view that simplicity should prevail, and that complicated proposals such as defining quality-of-service semantics in the flow label, or sub-dividing the flow label field into smaller sub-fields, will not prove efficient or deployable, especially in high-speed routers. There is also a clearly expressed view that using the flow label for various forms of stateless load distribution is the best simple application for it. At the same time, it is necessary to recognize that the strict immutability rule has drawbacks as noted above.

Even under the rules of RFC 3697, the flow label is intrinsically untrustworthy, because modifications en route cannot be detected. For this reason, even with the current strict immutability rule, downstream nodes cannot rely mathematically on the value being unchanged. In this sense, any use of the flow label must be viewed

as an optimization on a best-effort basis; a packet with a changed (or zero) flow label value should never cause a hard failure.

The remainder of this document discusses specific modifications to the standard, which are defined normatively in a companion document [RFC6437].

3. Changes to the Specification

Although RFC 3697 requires that the flow label be delivered unchanged, as noted above, it is not included in any transport-layer pseudo-header checksums nor in IPsec authentication [RFC4302]. Both RFC 2460 and RFC 3697 define the default flow label to be zero. At the time of writing, this is the observed value in an overwhelming proportion of IPv6 packets; the most widespread operating systems and applications do not set it, and routers do not rely on it. Thus, there is no reason to expect operational difficulties if a careful change is made to the rules of RFC 3697.

In particular, the facts that the label is not checksummed and rarely used mean that the "immutability" of the label can be moderated without serious operational consequences.

The purposes of the proposed changes are to remove the uncertainties left by RFC 3697, in order to encourage setting of the flow label by default, and to enable its generic use. The proposed generic use is to encourage uniformly distributed flow labels that can be used to assist load distribution or balancing. There should be no impact on existing IETF specifications other than RFC 3697 and no impact on currently operational software and hardware.

A secondary purpose is to allow changes to the flow label in a limited way, to allow hosts that do not set the flow label to benefit from it nevertheless. The fact that the flow label may in practice be changed en route is also reflected in the reformulation of the rules.

A general description of the changes follows. The normative text is to be found in [RFC6437].

The definition of a flow is subtly changed from RFC 3697 to allow any node, not just the source node, to set the flow label value. However, it is recommended that sources should set a uniformly distributed flow label value in all flows, replacing the less precise recommendation made in Section 3 of RFC 3697. Both stateful and stateless methods of assigning a uniformly distributed value could be used.

Flow label values should be chosen such that their bits exhibit a high degree of variability, thus 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 have a low probability of guessing 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, an approximation to a discrete uniform distribution is preferable as the source of flow label values. In contrast, an implementation in which flow labels are assigned sequentially is definitely not recommended, to avoid guessability.

In practice, it is expected that a uniform distribution of flow label values will be approximated by use of a hash function or a pseudo-random number generator. Either approach will produce values that will appear pseudo-random to an external observer.

Section 3 of RFC 3697 allows nodes to participate in an unspecified stateful method of flow state establishment. The changes do not remove that option, but clarify that stateless models are also possible and are the recommended default. The specific text on requirements for stateful models has been reduced to a bare minimum requirement that they do not interfere with the stateless model. To enable stateless load distribution at any point in the Internet, a node using a stateful model should never send packets whose flow label values do not conform to a uniform distribution.

The main novelty is that a forwarding node (typically a first-hop or ingress router) may set the flow label value if the source has not done so, according to the same recommendations that apply to the source. This might place a considerable processing load on ingress routers that choose to do so, even if they adopted a stateless method of flow identification and label assignment.

The value of the flow label, once it has been set, must not be changed. However, some qualifications are placed on this rule, to allow for the fact that the flow label is an unprotected field and might be misused. No Internet-wide mechanism can depend mathematically on immutable flow labels. The new rules require that flow labels exported to the Internet should always be either zero or uniformly distributed, but even this cannot be relied on mathematically. Use cases need to be robust against non-conforming

flow label values. This will also enhance compatibility with any legacy hosts that set the flow label according to RFC 2460 and RFC 3697.

A complication that led to much discussion is the possibility that hosts inside a particular network domain might use a stateful method of setting the flow label, and that packets bearing stateful labels might then erroneously escape the domain and be received by nodes performing stateless processing, such as load balancing. This might result in undesirable operational implications (e.g., congestion, reordering) for not only the inappropriately flow-labeled packets, but also well-behaved flow-labeled packets, during forwarding at various intermediate devices. It was suggested that border routers might "correct" this problem by overwriting such labels in packets leaving the domain. However, neither domain border egress routers nor intermediate routers/devices (using a flow label, for example, as a part of an input key for a load-distribution hash) can determine by inspection that a value is not part of a uniform distribution. Thus, there is no way that such values can be detected and "corrected". Therefore, the recommendation to choose flow labels from a uniform distribution also applies to stateful schemes.

4. Discussion

The following are some practical consequences of the above changes:

- Sending hosts that are not updated will in practice continue to send all-zero labels. If there is no label-setting router along the path taken by a packet, the label will be delivered as zero.

- Sending hosts conforming to the new specification will by default choose uniformly distributed labels between 1 and 0xFFFFF.

- Sending hosts may continue to send all-zero labels, in which case an ingress router may set uniformly distributed labels between 1 and 0xFFFFF.

- The flow label is no longer unrealistically asserted to be strictly immutable; it is recognized that it may, incorrectly, be changed en route. In some circumstances, this will break end-to-end usage, e.g., potential detection of third-party spoofing attacks [LABEL-SEC].

- The expected default usage of the flow label is some form of stateless load distribution, such as the ECMP/LAG usage defined in [RFC6438].

- o If the new rules are followed, all IPv6 traffic flows on the Internet should have zero or uniformly distributed flow label values.

From an operational viewpoint, existing IPv6 hosts that set a default (zero) flow label value and ignore the flow label on receipt will be unaffected by implementations of the new specification. In general, it is assumed that hosts will ignore the value of the flow label on receipt; it cannot be relied on as an end-to-end signal. However, this doesn't apply if a cryptographically generated label is being used to detect attackers [LABEL-SEC].

Similarly, routers that ignore the flow label will be unaffected by implementations of the specification.

Hosts that set a default (zero) flow label but are in a domain where routers set a label as recommended in Section 3 will benefit from whatever flow label handling is used on the path.

Hosts and routers that adopt the recommended mechanism will enhance the performance of any load balancing devices that include the flow label in the hash used to select a particular path or server, even when packets leave the local domain.

5. Security Considerations

See [RFC6437] and [LABEL-SEC] for full discussion. Some useful remarks are in [Partridge].

6. Acknowledgements

The authors are grateful to Qinwen Hu for general discussion about the flow label and for his work in searching the literature. Valuable comments and contributions were made by Ran Atkinson, Fred Baker, Steve Blake, Remi Despres, Alan Ford, Fernando Gont, Brian Haberman, Tony Hain, Joel Halpern, Chris Morrow, Thomas Narten, Pekka Savola, Mark Smith, Pascal Thubert, Iljitsch van Beijnum, and other participants in the 6man working group.

7. Informative References

[FLOWSWITCH] Beckman, M., "IPv6 Dynamic Flow Label Switching (FLS)", Work in Progress, March 2007.

[LABEL-SEC] Gont, F., "Security Assessment of the IPv6 Flow Label", Work in Progress, November 2010.

[McGann05] McGann, O. and D. Malone, "Flow Label Filtering Feasibility", European Conference on Computer Network Defence , 2005.

[NIST] Frankel, S., Graveman, R., Pearce, J., and M. Rooks, "Guidelines for the Secure Deployment of IPv6", National Institute of Standards and Technology Special Publication 800-119, 2010, <http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf>.

[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>.

[Partridge] Partridge, C., Arsenault, A., and S. Kent, "Information Assurance and the Transition to IP Version 6 (IPv6)", Military Communications Conference (MILCOM 2007) , 2007, <http://www.ir.bbn.com/documents/articles/MILCOM_Paper_from_Proceedings.pdf>.

[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.

[RFC2474] 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.

[RFC3697] Rajahalme, J., Conta, A., Carpenter, B., and S. Deering, "IPv6 Flow Label Specification", RFC 3697, March 2004.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.

[RFC6294] Hu, Q. and B. Carpenter, "Survey of Proposed Use Cases for the IPv6 Flow Label", RFC 6294, June 2011.

[RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, 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.

Appendix A. Alternative Approaches

A model was discussed in an earlier version of this document which defined a notion of 'flow label domain' analogous to a differentiated services domain [RFC2474]. This model would have encouraged local usage of the flow label as an alternative to any form of generic use, but it required complex rules for the behavior of domain boundary routers, and proved controversial in discussion.

Two even more complex alternative approaches were also considered and rejected.

The first was to distinguish locally significant flow labels from those conforming to RFC 3697 by setting or clearing the most significant bit (MSB) of the flow label. This led to quite complicated rules, seems impossible to make fully self-consistent, and was not considered practical.

The second was to use a specific differentiated services code point (DSCP) [RFC2474] in the Traffic Class octet instead of the MSB of the flow label itself, to flag a locally defined behavior. A more elaborate version of this was proposed in [FLOWSWITCH]. There are two issues with that approach. One is that DSCP values are themselves only locally significant, inconsistent with the end-to-end nature of the original flow label definition. Secondly, it seems unwise to meld the semantics of differentiated services, which are currently deployed, with the unknown future semantics of flow label usage. However, this approach, while not recommended, does not appear to violate any basic principles if applied strictly within a single differentiated services domain.

Authors' Addresses

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

EMail: shane@level3.net

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

EMail: 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

EMail: jiangsheng@huawei.com