

| | |
|--------------------------------|-------------------|
| Independent Submission | Q. Hu |
| Internet-Draft | B. E. Carpenter |
| Intended status: Informational | Univ. of Auckland |
| Expires: August 27, 2011 | February 23, 2011 |

Survey of proposed use cases for the IPv6 flow label
draft-hu-flow-label-cases-03

[Abstract](#)

The IPv6 protocol includes a flow label in every packet header, but this field is not used in practice. This paper describes the flow label standard and discusses the implementation issues that it raises. It then describes various published proposals for using the flow label, and shows that most of them are inconsistent with the standard. Methods to address this problem are briefly reviewed. We also question whether the standard should be revised.

[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 August 27, 2011.

[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](#)

- *1. [Introduction](#)

- *1.1. [A brief history of the flow label](#)
- *1.2. [The flow label and quality of service](#)
- *2. [Flow label definition and issues](#)
 - *2.1. [Flow label properties](#)
 - *2.2. [Dependency prohibition](#)
 - *2.3. [Other issues](#)
- *3. [Documented proposals for the flow label](#)
 - *3.1. [Specify the flow label as a pseudo-random value](#)
 - *3.1.1. [End-to-End QoS Provisioning](#)
 - *3.1.2. [Load balancing](#)
 - *3.1.3. [Security nonce](#)
 - *3.2. [Specify QoS parameters in the flow label](#)
 - *3.3. [Use flow label hop-by-hop to control switching](#)
 - *3.4. [Diffserv use of IPv6 flow label](#)
 - *3.5. [Other uses](#)
- *4. [Conclusion](#)
- *5. [Security Considerations](#)
- *6. [IANA Considerations](#)
- *7. [Acknowledgements](#)
- *8. [Change log](#)
- *9. [References](#)
- *[Authors' Addresses](#)

[1. Introduction](#)

IPv6 is being introduced to overcome the address shortage of the current IPv4 protocol, but it also offers a new feature, i.e., the flow label field in the IPv6 packet header. The flow label is not encrypted by IPsec, and is present in all fragments. However, it is very little used in practice, for reasons discussed below and in [\[I-D.ietf-6man-flow-](#)

[update](#)]. After a short introduction, this document summarizes the current specification of the IPv6 flow label and some open issues about its use in [Section 2](#), and then [Section 3](#) describes and analyses various proposals that have been made for its use. Finally, [Section 4](#) discusses the implications and attempts to draw conclusions.

[1.1. A brief history of the flow label](#)

The original proposal for a flow label has been attributed to Dave Clark [[Deering93](#)], who proposed that it should contain a pseudo-random value. A flow label field was included in the packet header during the preliminary design of IPv6, which followed an intense period of debate about several competing proposals. The final choice was made in 1994 [[RFC1752](#)]. In particular, the IETF rejected a proposal known as CATNIP [[RFC1707](#)], which included so-called 'cache handles' to identify the next hop in high performance routers. Thus CATNIP introduced the notion of a header field that would be shared by all packets belonging to a flow, to control packet forwarding on a hop-by-hop basis. We recognize this today as a precursor of the MPLS label [[RFC3031](#)].

The IETF decided instead to develop a proposal known as SIPP into IP version 6. SIPP included "labeling of packets belonging to particular traffic 'flows' for which the sender requests special handling, such as non-default quality of service or 'real-time' service" [[RFC1710](#)]. In 1994, this used a 28-bit flow label field. In 1995 it was down to 24 bits [[RFC1883](#)] and it was finally reduced to 20 bits [[RFC2460](#)] to accommodate the IPv6 Traffic Class, which is fully compatible with the IPv4 Type of Service byte.

There was considerable debate in the IETF about the very purpose of the flow label. Was it to be a handle for fast switching, as in CATNIP, or was it to be meaningful to applications and used to specify quality of service? Must it be set by the sending host, or could it be set by routers? Could it be modified en route, or must it be delivered with no change?

Because of these uncertainties, and more urgent work, the flow label was consistently ignored by implementors, and today is set to zero in almost every IPv6 packet. In fact, [[RFC2460](#)] defined it as "experimental and subject to change." There was considerable preliminary work such as [[Metzler00](#)], [[Conta01a](#)], [[I-D.conta-diffserv-ipv6-fl-classifier](#)] and [[Hagino01](#)]. The ensuing proposed standard "IPv6 Flow Label Specification" (RFC 3697) [[RFC3697](#)] intended to clarify this situation by providing precise boundary conditions for use of the flow label. However, this has not proved successful in promoting use of the flow label in practice, as a result of which 20 bits are unused in every IPv6 packet header.

[1.2. The flow label and quality of service](#)

Developments in high speed switch design, and the success of MPLS, have largely obviated consideration of the flow label for high-speed

switching. Thus, although various use cases for the flow label have been proposed, most of them assume that it should be used principally to support the provision of quality of service (QoS). For many years it has been recognized that real-time Internet traffic requires a different QoS from general data traffic, and this remains true in the era of network neutrality. Thus an alternative to uniform best-effort service is needed, requiring packets to be classified as belonging to a particular class of service or flow. Currently, this leads to a layer violation problem, since a 5-tuple is often used to classify each packet. The 5-tuple includes source and destination addresses, port numbers, and the transport protocol type, so when we want to forward or process packets, we need to extract information from the layer above IP. This may be impossible when packets are encrypted such that port numbers are hidden, or when packets are fragmented, so the layer violation is not an academic concern. The flow label, being exempt from IPsec encryption and being replicated in packet fragments, avoids this difficulty. It has therefore attracted attention from the designers of new approaches to QoS.

[2. Flow label definition and issues](#)

[2.1. Flow label properties](#)

RFC 3697 [[RFC3697](#)] standardizes properties of the flow label, including: One effect of the second of these rules is to avoid the layer violation problem mentioned in [Section 1](#). By using the 3-tuple, we only use the IP layer to classify packets, without needing any transport layer information. This may reduce the lookup time if flow-based treatment is required, and will work even with IPsec encryption and fragmentation. Therefore, for traffic needing other than best-effort service, such as real-time applications, the flow label can be set to different values to represent different flows, and each node forwarding or receiving the packets may provide different flow-specific treatments by looking at the flow label value. This is more fine-grained than differentiated services (Diffserv) [[Carpenter02](#)], [[RFC2474](#)] but need not be less efficient.

[2.2. Dependency prohibition](#)

An additional important rule in the standard [[RFC3697](#)] effectively forbids any encoding of meaning in the bits of the flow label. To be exact, the standard states that "IPv6 nodes MUST NOT assume any mathematical or other properties of the flow Label values assigned by source nodes." This rule is aimed at the case where a packet from a source using a particular encoding scheme for the flow label reaches a node that is using a different scheme. If by chance the bit pattern in the flow label is meaningful in both schemes, the receiver would misinterpret the flow label. Therefore, in the absence of other information, the receiver must not assume anything about the meaning of the value of the flow label.

The standard [\[RFC3697\]](#) also states "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." The problem this rule is intended to avoid is that if a source uses one method of choosing flow labels (e.g., counting up from 1), any router that assumes another method (e.g., pseudo-randomness) may not perform as intended. Note that there is no easy escape from the combination of these two prohibitions, which we will call the dependency prohibition. Unlike Diffserv code points, flow labels are not locally significant within a single administrative domain; they must be preserved end-to-end. In general, a router cannot know whether a particular packet originated in a host supporting a specific usage of the flow label. Therefore, any method that breaks one or both of these rules will only work if there is some way for a router to determine which sources use the same scheme as itself.

The interpretation of the dependency rule can be subtle and is not spelled out in [\[RFC3697\]](#). A node must not assume properties of the flow label - but it may know them by construction or by signaling. The bits of the flow label alone are poor material for a hash key - but they may be combined with bits from other sources, to provide uniformly distributed hash outputs.

[2.3. Other issues](#)

[\[RFC3697\]](#) does not discuss how to use flow label most effectively. This remains the major open issue, but some authors propose that the label should be used with reserved bandwidth to achieve customized QoS provision. Coupled with admission control at the edge router, this could limit congestion. However, as we will see below, this is not the only proposed use.

We now introduce some other open issues.

[3. Documented proposals for the flow label](#)

In the following, we do not intend to recommend or criticise various proposals. This section shows the variety of proposals that have been published, and whether they are compatible with the existing standard. These proposals almost all assume that the flow label's main purpose is to support QoS, and their flow label mechanisms are entangled with QoS mechanisms. We describe the proposals in five broad, and somewhat overlapping, categories, i.e.,

Among the proposals described in the following five sections, various methods are proposed to set up the flow label value. It should be noted that some of these proposals embody novel and perhaps controversial approaches to QoS provision, and these cannot readily be separated from their use of the flow label. We give a reasonable amount of technical detail for some of the proposals, to show the extent to which they propose detailed semantics for the flow label value.

3.1. Specify the flow label as a pseudo-random value

3.1.1. End-to-End QoS Provisioning

As our first example, [\[Lin06\]](#) specifies a 17-bit pseudo-random value.

The figure below shows the proposed flow label structure.

[\[Lin06\]](#) also describes a signalling process between source, routing and destination nodes based on this label structure and on the IPv6 Traffic Class byte, in order to reserve and release router resources for a given flow within a given class of traffic. The pseudo-random LN value is used to uniquely identify a given flow.

Flow Label Specification (figure simplified from [\[Lin06\]](#))



- LF 0 Disable
- 1 Enable
- LT 00 Flow label requested by source
- 01 Flow label returned by destination
- 10 Flow label for data delivery
- 11 Flow label terminates connection
- LN Random number created by source

3.1.2. Load balancing

There have been numerous informal discussions of using pseudo-random flow labels to allow load-balancing or at least load-sharing. This would be achieved by including the flow label value among the fields in each packet header used as input to a modulo(N) hash used to select among N alternative paths. However, concerns about the interpretation of the dependency prohibition have generally prevented such proposals from being written up until recently [\[I-D.ietf-6man-flow-ecmp\]](#).

3.1.3. Security nonce

Another proposal for a pseudo-random flow label value is [\[I-D.blake-ipv6-flow-label-nonce\]](#). This states that off-path spoofing attacks have become a big issue for TCP and other transport-layer applications, and proposes that in IPv6 we should set a random value in the flow label to make the packet header more complex and less easy for the attacker to guess. The two ends of the session will agree on flow label values during the SYN/ACK exchange, but off-path attackers will be unlikely to guess the agreed value. Naturally, on-path attackers who can observe the flow labels in use can trivially defeat this protection. This proposal does not involve using the flow label value to retrieve routing state.

3.2. Specify QoS parameters in the flow label

[\[Prakash04\]](#) proposes to utilize the flow label to indicate required QoS parameters in detail. It uses the first few bits of the flow label field as codes to support different approaches, as summarized in following table. Again, this is incompatible with the dependency prohibition in [\[RFC3697\]](#), since a source that does not use this method may also set the first two bits to non-zero.

Classification for various approaches (from [\[Prakash04\]](#))

| Bit Pattern | Approach |
|-------------|--|
| 00 | No QoS requirement (Default QoS value) |
| 01 | Pseudo-Random value used for the value of Flow-Label |
| 10 | Support for Direct Parametric Representation |
| 1100 | Support for the DiffServ Model |
| 1101 | Reserved for future use |
| 111 | Reserved for future use |

This method allows a pseudo-random option, but also adds options for a direct QoS request and for Diffserv. In the direct QoS parameters approach, 18 bits are used to encode requirements for one way delay, IP delay variation, bandwidth and one way packet loss. The proposal appears to assume that RSVP [\[RFC2205\]](#) mechanisms are used to actually implement these QoS parameters.

This proposal allows use of flow label for various important QoS models, so the end user and service provider can choose the most suitable model for their situation; [\[Prakash04\]](#) claims that "this proposal is simple, scalable, modular and generic implementation to provide for QoS using the IPv6 flow label field".

Similarly, [\[Lee04\]](#) defines the flow label field in five parts, with the first 3 bits used as an approach type. The authors define two approaches: a "random" scheme and a "hybrid" scheme. If the first 3 bits equal "001", the flow label will be used as the random identifier of the flow, but if they equal "101", the remaining bits will include a hybrid QoS requirement for this packet, subdivided into traffic type (stringent or best effort), bandwidth, buffer, and delay requirements. Once again the dependency prohibition in [\[RFC3697\]](#) is broken. This proposal also includes throughput monitoring and dynamic capacity allocation.

Effectively this proposal uses the flow label both to signal Intserv-like QoS requirements and to classify traffic into Diffserv-like virtual label-switched paths. Packets with a "random" flow label are mapped into a generic (best effort) virtual path.

3.3. Use flow label hop-by-hop to control switching

[\[I-D.chakravorty-6lsa\]](#) and [\[Chakravorty08a\]](#) describe an architectural framework called "IPv6 Label Switching Architecture" (6LSA). In 6LSA, network components identify a flow by looking at the flow label field in

the IPv6 packet header; all packets with the same flow label must receive the same treatment and be sent to the same next hop. However, 6LSA resembles MPLS by considering that a label only has meaning between 6LSA routers, and setting the flow label at each hop. If the original source sets a non-zero flow label, there is no mechanism to restore it before delivery, a fundamental breach of [\[RFC3697\]](#). The authors of [\[I-D.chakravorty-6lsa\]](#) did at one stage discuss using an IPv6 hop-by-hop option to correct this problem, but this has not been documented. This is a more serious incompatibility than simply breaking the dependency prohibition

Unlike traditional routing algorithms, but like MPLS, 6LSA packets are classified into a Forwarding Equivalence Class (FEC), and routers forward packets on different paths by looking at the FEC. Like previous solutions, the authors divide the flow label field into three parts. The first 3 bits identify the FEC, which will help the router or 6LSA nodes to group the IP packets that receive the same forwarding treatment and forwarding them on the same virtual path. The following 4 bits describe the application type, and the final 13 bits (defined by each node or a group of nodes) specify the hop-specific label. From the table below, we can see the FEC has 6 major categories, each with up to 16 subcategories.

Flow Label Specification (shortened from [\[I-D.chakravorty-6lsa\]](#))

| FEC (First 3 Bits) | Next 4 Bits | Purpose |
|--------------------------|-------------|--------------------------|
| No FEC (000) | 0000 | |
| Domain Specific (000) | 0001 - 1111 | |
| ----- | | |
| VPN (001) | 0001 | (IPSec - Tunnel Mode) |
| | 0010 | (IPSec - Transport Mode) |
| | 0011 | (Special Encryption) |
| | 0100 | (VRF) |
| | 0101 | (End Network Specific) |
| | 0110 - 1111 | (Reserved) |
| ----- | | |
| TE Subset/ | 0001 | (DiffServ) |
| QoS Enhancement (010) | 0010 | (RSVP) |
| . . . | | |
| | 1111 | (Reserved) |
| ----- | | |
| Encapsulation (011) | 0001 | (IPv6 in IPv6) |
| | 0010 | (IPv4 in IPv6) |
| | 0011 | (Other in IPv6) |
| | 0100 | (Enterprise Specific) |
| | 0101 - 1111 | (Reserved) |
| ----- | | |
| Enterprise Specific(111) | 0000 - 1111 | (Reserved) |

The authors claim that fast switching using 20-bit labels instead of 128-bit IPv6 addresses will provide memory and processing savings, as well as network management advantages. "It also allows a network management entity updating available label tables, across the network to reduce man-in-the-middle attacks [sic]" [\[I-D.chakravorty-6lsa\]](#).

We note that a similar proposal for QoS-based switching of IPv6 packets [\[I-D.roberts-inband-qos-ipv6\]](#) is designed to use a hop-by-hop option, which has not so far been allocated by the IETF. Proposals related to this have been discussed by the Telecommunications Industry Association and the ITU-T [\[I-D.adams-tsvwg-flow-signalling-codepoint\]](#).

We also note that router lookup efficiency was a major concern at the time when Clark first proposed a flow label [\[Deering93\]](#), but with the advent of very large scale integrated circuits capable of rapid lookup in a routing table, most vendors no longer express such concern.

[3.4. Diffserv use of IPv6 flow label](#)

[\[I-D.banerjee-flowlabel-ipv6-qos\]](#) uses the flow label field as a replacement for the IPv6 Traffic Class field; this proposal suggests the incoming flow label can indicate the QoS requirement by matching a Diffserv classifier. The authors have used the first three bits to indicate this, and the following 16 bits to indicate a Differentiated Services Per-Hop Behavior Identification code (Diffserv PHB-ID) [\[RFC3140\]](#); the last bit is reserved for future use. This method too breaks the dependency prohibition in [\[RFC3697\]](#).

[\[I-D.martinbeckman-ietf-ipv6-fls-ipv6flowswitching\]](#) blends the flow label as an MPLS-like switching tag with DiffServ. Unlike 6LSA, the method attempts to bypass the dependency prohibition by using one bit in the DiffServ Code Point [\[RFC2474\]](#) to indicate that the flow label is a switching tag. In this way, a router can determine whether the flow label conforms to [\[RFC3697\]](#) or to [\[I-D.martinbeckman-ietf-ipv6-fls-ipv6flowswitching\]](#). In [\[I-D.martinbeckman-ietf-ipv6-amp-ipv6hcamp\]](#), the same author proposes using the flow label as a way of compressing IPv6 headers by hashing the addresses into the flow label, again using the DiffServ Code Point to mark the packets accordingly.

[3.5. Other uses](#)

The Integrated Services QoS architecture [\[RFC1633\]](#) specifies that the flow label may be used as a packet filter [\[RFC2205\]](#). At least one implementation supported this [\[Braden10\]](#).

We are not aware of any proposals combining the flow label with the Next Steps in Signaling (NSIS) [\[RFC4080\]](#) architecture.

[\[I-D.donley-software-dslite-flowlabel\]](#) proposes a use case whereby certain flows encapsulated in a particular type of IPv4-in-IPv6 tunnel would be distinguished at the remote end of the tunnel by a specific flow label value. This would allow a service provider to deliver a tailored quality of service. This usage appears to be completely compatible with [\[RFC3697\]](#).

There has been some discussion of possible flow label use in both the ROLL (Routing Over Low power and Lossy networks) [\[RPL-07\]](#) and MEXT (Mobility EXTensions for IPV6) working groups of the IETF. Such uses tend to encode specific local meanings or routing-related tags in the label, so they appear to infringe the dependency prohibition or the immutability of the flow label field. The ROLL group has indeed most recently opted not to use the flow label field for these reasons, despite having to add the undesirable overhead of an IPV6 hop-by-hop option instead [\[I-D.ietf-roll-rpl\]](#). Similarly, MEXT has defined a new mobility option to support flow bindings [\[I-D.ietf-mext-flow-binding\]](#), rather than using the IPV6 flow label field.

[4. Conclusion](#)

Three aspects of the current standard [\[RFC3697\]](#) have caused problems to many designers:

Taken together, these rules essentially forbid any encoding of the semantics of a flow, or of any information about its path, in the flow label. This was intentional, in accordance with the stateless nature of the Internet architecture and with the end-to-end principle [\[Saltzer84\]](#), [\[RFC3724\]](#). It was also felt that QoS encoding via Diffserv was sufficient, and that the requirement for high-speed switching could be met by MPLS. But this means that the majority of the proposals described above breach the standard and the intent of the standard. The authors often appear to be using the flow label either as an MPLS-like switching handle, or as an encoded QoS signal.

In contrast, a few documents mentioned above do appear to respect the rules of RFC 3697. These are [\[I-D.blake-ipv6-flow-label-nonce\]](#), [\[I-D.donley-softwire-dslite-flowlabel\]](#), [\[I-D.ietf-6man-flow-ecmp\]](#), [\[I-D.martinbeckman-ietf-ipv6-fls-ipv6flowswitching\]](#) and [\[I-D.martinbeckman-ietf-ipv6-amp-ipv6hcamp\]](#). Additionally, [\[Lin06\]](#) would have joined this list if it had not assigned three flag bits in the flow label field. Although predating RFC 3697, the Integrated Services usage [\[RFC2205\]](#) also seems to be compatible.

What would other designers need to do, if they wish to respect RFC 3697? There appear to be two choices. One is to simply accept the existing rules at face value, as in the proposals just listed. This limits the application of the flow label. It can, for example, be used as a nonce or as part of the material for a hash used to share load among alternate paths. It cannot be the only material for such a hash, because of the dependency prohibition. The flow label could also be used consistently with RFC 3697, if an application designer so chose, as a way to associate all packets belonging to a given application session between two hosts, across multiple transport sessions. This, however, would presumably exclude using the flow label to govern routing decisions in any way, and would have widespread implications that have never been explored.

The other choice, for designers who wish to use the flow label to control switching or QoS directly, is to bypass the rules within a given

domain (a set of cooperating nodes) in a way that nodes outside the domain cannot detect. In this case, any deviation from RFC 3697 has no possible effect outside the domain in question.

An example scheme to emulate the immutability of labels is as follows. Within the domain, all hosts set the label to zero, the routers set and interpret the label in any way they wish, and the last hop router always sets the label back to zero. If a packet arrives from outside the domain with a non-zero label, there is a method (such as a special Diffserv code point) to mark this packet so that its label would be ignored and delivered unchanged. An alternative approach would be to define a hop-by-hop option to carry the original flow label across the domain, so that it could be changed within the domain but restored before forwarding the packet beyond the domain.

If a domain allows mutable labels in such a way, it may safely ignore the dependency prohibition, and it may set the bits in the label according to locally defined rules. Within the domain, the label could be used as desired, and most of the proposed designs discussed above could be "rescued."

However, given the considerable number of designers who have proposed solutions incompatible with RFC 3697, the relatively few designs using the standard rules, and the failure of designs such as ROLL and MEXT to make use of the flow label, it seems reasonable to ask whether the RFC 3697 standard has value.

[5. Security Considerations](#)

The flow label is not protected in any way and can be forged by an on-path attacker. Off-path attackers may be able to guess a valid flow label unless a pseudo-random value is used. Specific usage models for the flow label need to allow for these exposures. For further discussion, see [\[RFC3697\]](#).

[6. IANA Considerations](#)

This document requests no action by IANA.

[7. Acknowledgements](#)

An invaluable review of this document was performed by Bob Braden. Helpful comments were made by Sheng Jiang. This document was produced using the xml2rfc tool [\[RFC2629\]](#).

[8. Change log](#)

draft-hu-flow-label-cases-00: first I-D version, 2010-04-19
draft-hu-flow-label-cases-01: updated following review comments, 2010-08-18
draft-hu-flow-label-cases-02: updated following more review comments, 2010-09-30

9. References

| | |
|-----------|---|
| [RFC1752] | Bradner, S. and A. Mankin , " The Recommendation for the IP Next Generation Protocol ", RFC 1752, January 1995. |
| [RFC1883] | Deering, S.E. and R.M. Hinden , " Internet Protocol, Version 6 (IPv6) Specification ", RFC 1883, December 1995. |
| [RFC2460] | Deering, S.E. and R.M. Hinden , " Internet Protocol, Version 6 (IPv6) Specification ", RFC 2460, December 1998. |
| [RFC4080] | Hancock, R., Karagiannis, G., Loughney, J. and S. Van den Bosch, " Next Steps in Signaling (NSIS): Framework ", RFC 4080, June 2005. |
| [RFC1710] | Hinden, R. , " Simple Internet Protocol Plus White Paper ", RFC 1710, October 1994. |
| [RFC1707] | McGovern, M. and R. Ullmann , " CATNIP: Common Architecture for the Internet ", RFC 1707, October 1994. |
| [RFC2474] | Nichols, K. , Blake, S. , Baker, F. and D.L. Black , " Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers ", RFC 2474, December 1998. |
| [RFC1809] | Partridge, C. , " Using the Flow Label Field in IPv6 ", RFC 1809, June 1995. |
| [RFC3697] | Rajahalme, J., Conta, A., Carpenter, B. and S. Deering, " IPv6 Flow Label Specification ", RFC 3697, March 2004. |
| [RFC1633] | Braden, B. , Clark, D. and S. Shenker , " Integrated Services in the Internet Architecture: an Overview ", RFC 1633, June 1994. |
| [RFC3140] | Black, D., Brim, S., Carpenter, B. and F. Le Faucheur, " Per Hop Behavior Identification Codes ", RFC 3140, June 2001. |
| [RFC2629] | Rose, M.T. , " Writing I-Ds and RFCs using XML ", RFC 2629, June 1999. |
| [RFC2205] | Braden, B. , Zhang, L. , Berson, S. , Herzog, S. and S. Jamin , " Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification ", RFC 2205, September 1997. |
| [RFC3724] | Kempf, J., Austein, R., IAB, " The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the |

| | |
|---|---|
| | Internet Architecture ", RFC 3724, March 2004. |
| [RFC3031] | Rosen, E., Viswanathan, A. and R. Callon, " Multiprotocol Label Switching Architecture ", RFC 3031, January 2001. |
| [I-D.conta-diffserv-ipv6-fl-classifier] | Conta, A and J Rajahalme, " A model for Diffserv use of the IPv6 Flow Label Specification ", Internet-Draft draft-conta-diffserv-ipv6-fl-classifier-01, November 2001. |
| [I-D.chakravorty-6lsa] | Chakravorty, S, Bush, J and J Bound, " IPv6 Label Switching Architecture ", Internet-Draft draft-chakravorty-6lsa-03, July 2008. |
| [I-D.banerjee-flowlabel-ipv6-qos] | Banerjee, R, " A Modified Specification for use of the IPv6 Flow Label for providing An efficient Quality of Service using hybrid approach ", Internet-Draft draft-banerjee-flowlabel-ipv6-qos-03, April 2002. |
| [I-D.roberts-inband-qos-ipv6] | Roberts, L and J Harford, " In-Band QoS Signaling for IPv6 ", Internet-Draft draft-roberts-inband-qos-ipv6-00, July 2005. |
| [I-D.blake-ipv6-flow-label-nonce] | Blake, S, " Use of the IPv6 Flow Label as a Transport-Layer Nonce to Defend Against Off-Path Spoofing Attacks ", Internet-Draft draft-blake-ipv6-flow-label-nonce-02, October 2009. |
| [I-D.ietf-6man-flow-ecmp] | Carpenter, B and S Amante, " Using the IPv6 flow label for equal cost multipath routing and link aggregation in tunnels ", Internet-Draft draft-ietf-6man-flow-ecmp-05, July 2011. |
| [I-D.ietf-6man-flow-update] | Amante, S, Carpenter, B and S Jiang, " Rationale for update to the IPv6 flow label specification ", Internet-Draft draft-ietf-6man-flow-update-07, July 2011. |
| [I-D.adams-tsvwg-flow-signalling-codepoint] | song, j, Adams, J and J Joung, " Progress and future development of Flow State Aware standards, and a proposal for alerting nodes or end-systems on data related to a flow ", Internet-Draft draft-adams-tsvwg-flow-signalling-codepoint-00, June 2008. |
| [I-D.donley-softwire-dslite-flowlabel] | Donley, C and K Erichsen, " Using the Flow Label with Dual-Stack Lite ", Internet-Draft draft-donley-softwire-dslite-flowlabel-02, July 2011. |
| [I-D.ietf-mext-flow-binding] | Tsirsis, G, Soliman, H, Montavont, N, Giaretta, G and K Kuladinithi, " Flow Bindings in Mobile IPv6 and NEMO Basic |

| | |
|----------------------------|--|
| | Support ", Internet-Draft draft-ietf-mext-flow-binding-11, October 2010. |
| [I-D.ietf-roll-rpl] | Winter, T, Thubert, P, Brandt, A, Clausen, T, Hui, J, Kelsey, R, Levis, P, Pister, K, Struik, R and J Vasseur, " RPL: IPv6 Routing Protocol for Low power and Lossy Networks ", Internet-Draft draft-ietf-roll-rpl-19, March 2011. |
| [RPL-07] | Winter, T. and P. Thubert, " RPL: IPv6 Routing Protocol for Low power and Lossy Networks ", Internet-Draft draft-ietf-roll-rpl-07, March 2010. |
| [Metzler00] | Metzler, J. and S. Hauth, " An end-to-end usage of the IPv6 flow label ", Internet-Draft draft-metzler-ipv6-flowlabel-00, November 2000. |
| [Conta01a] | Conta, A. and B. Carpenter, " A proposal for the IPv6 Flow Label Specification ", Internet-Draft draft-counta-ipv6-flow-label-02, July 2001. |
| [Hagino01] | Hagino, J. , " Socket API for IPv6 flow label field ", Internet-Draft draft-itojun-ipv6-flowlabel-api-01, April 2001. |
| [Deering93] | Deering, S., "SIPP Overview and Issues", Minutes of the Joint Sessions of the SIP and PIP Working Groups , November 1993. |
| [Braden10] | Braden, R., "Private Communication", 2010. |
| [Lee04] | Lee, I. H. and S. J. Kim, "A QoS Improvement Scheme for Real-Time Traffic Using IPv6 Flow Labels", Lecture Notes in Computer Science Vol. 3043, 2004. |
| [Lin06] | Lin, C. N., Tseng, P. C. and W. S. Hwang, "End-to-End QoS Provisioning by Flow Label in IPv6", JCIS , 2006. |
| [Prakash04] | Prakash, B., "Using the 20 bit flow label field in the IPv6 header to indicate desirable quality of service on the internet", University of Colorado (M.Sc. Thesis), 2004. |
| [Chakravorty08a] | Chakravorty, S., "Challenges of IPv6 Flow Label implementation", Proc IEEE MILCOM2008 , 2008. |
| [Carpenter02] | Carpenter, B. and K. Nichols, "Differentiated Services in the Internet", Proc IEEE 90 (9) 1479-1494, 2002. |
| [Saltzer84] | Saltzer, J.H. , Reed, D.P. and D.D. Clark, "End-To-End Arguments in System Design", ACM TOCS 2 (4) 277-288, 1984. |

Authors' Addresses

Qinwen Hu Hu Department of Computer Science University of Auckland PB
92019 Auckland, 1142 New Zealand EMail: qhu009@aucklanduni.ac.nz

Brian Carpenter Carpenter Department of Computer Science University
of Auckland PB 92019 Auckland, 1142 New Zealand EMail:
brian.e.carpenter@gmail.com