IPv6 Working Group
INTERNET-DRAFT
<draft-ietf-ipv6-flow-label-02.txt>

J. Rajahalme
Nokia
A. Conta
Transwitch
B. Carpenter
IBM
S. Deering
Cisco

Expires: December 2002

June 2002

IPv6 Flow Label Specification
draft-ietf-ipv6-flow-label-02.txt

Status of this memo

Abstract

This document specifies the usage of the IPv6 Flow Label field, the requirements for IPv6 source nodes labeling flows, and the requirements for flow state establishment methods.

1. Terminology and Definitions

Classifier	An IP layer entity that selects packets based on the content of packet headers according to defined rules.
Flow	A sequence of related packets sent from a source to a unicast, anycast, or multicast destination(s). A flow could consist of all packets in a specific transport connection, or a media stream. However, a flow is not necessarily 1:1 mapped to a transport connection. This definition should not be confused with the more restrictive definitions for "flow" and "microflow" in [RSVP] and [DiffServ], respectively. This definition includes, but is not limited to them.
Flow state	The information stored in an IP node driving the flow classification and the flow-specific treatment. The required information is specified by the method defining the flow-specific treatment.
Flow state establishment method	A control mechanism used to set up the flow state. A flow state establishment method can be either - Dynamic, under source node control (e.g. RSVP), - Quasi-dynamic, under network management control, or - Static, through manual configuration. - Algorithmic (e.g. load-spreading)

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

2. Introduction

A flow is a sequence of related packets sent from a source to a unicast, anycast, or multicast destination(s). To enable specific processing for the flow, flow state needs to be established on the nodes providing the flow-specific treatment. The flow state defines what kind of treatment should be provided, and how to classify the packets to the flow.

Traditionally, flow classifiers have been based on the 5-tuple of the source and destination addresses, ports and the transport protocol type (e.g. the "microflow" definition in [DiffServ]). However, these fields may be unavailable due to either fragmentation or encryption, or locating them past a chain of IPv6 option headers may be inefficient. Additionally, dependence on higher layer headers by the IP layer represents a layer violation, possibly hindering the introduction of new transport protocols.

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

The minimum level of IPv6 flow support consists of labeling the flows. IPv6 source nodes can label known flows (e.g. TCP connections, RTP streams), even if the node itself would not require any flow-specific treatment. Doing this enables receiver oriented resource reservations, e.g. [RSVP]. Requirements for flow labeling are given in section 4.

Specific flow state establishment methods and the related service models are out of scope for this specification, but the generic requirements enabling co-existence of different methods in IPv6 nodes are set forth in section 5.

3. IPv6 Flow Label Specification

The 20-bit Flow Label field in the IPv6 header SHOULD be used by a source to label sequences of related packets sent to a specific unicast, anycast, or multicast destination(s). A non-zero Flow Label indicates that the IPv6 packet is labeled. IPv6 nodes receiving a labeled IPv6 packet can use the Flow Label, and Source and Destination Address fields to classify the packet to a certain flow. The packet MAY be given some flow-specific treatment based on the flow state established on a set of IPv6 nodes. The nature of the specific treatment and the methods for the flow state establishment are out of scope for this specification.

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

IPv6 nodes MUST NOT assume mathematical or other non-standardized properties of the Flow Label values assigned by source nodes. 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.

If an IPv6 node is not providing flow-specific treatment, it MUST ignore the field when receiving or forwarding a packet.

4. Flow Labeling Requirements

To enable Flow Label based classification, the source MUST label all packets belonging to a flow with the Flow Label value assigned to the flow.

The assignment of a packet to a flow takes various forms, presented below:

(1) The source MAY take part in a signaling protocol that results in assigning certain transport connection(s) or application data stream(s) to specific flow(s).

(2) The source MAY be configured to assign certain transport connection(s) or application data stream(s) to specific flow(s).

(3) The source SHOULD assign each new application data stream (e.g.

RTP streams) to a new flow.

(4) The source SHOULD assign each new transport connection (e.g. TCP, SCTP) to a new flow.

It is necessary that flow classifiers downstream from the source can classify packets unambiguously, i.e. that all packets which the source has chosen to label as a single flow can be efficiently identified as such.

To enable this, the source node MUST keep track of the Flow Label values it is currently using or has recently used. When a new flow is instantiated, a unique Flow Label MUST be assigned to it. A Flow Label value is considered unique if it is not currently in use with the same Source and Destination addresses. In the case of flows with multiple addresses (e.g., SCTP flows) this requirement for uniqueness extends to all possible (Source, Destination) address pairs.

The IPv6 source node MUST provide a facility for verifying and assigning new Flow Label values, and for storing the Flow Label, and the associated Source and Destination Addresses currently in use. The facility MUST be used whenever a label needs to be assigned for a new flow. The facility SHOULD provide a programming interface with at least the following functionality:

(1) to assign any Flow Label value for a new flow

(2) to assign a specific Flow Label for a new flow, and

(3) to delete a flow, i.e. to free a Flow Label no longer in use.

The interface definition for the facility is beyond the scope of this document.

When a dynamically instantiated flow terminates, its Flow Label value MUST NOT be reused until it is certain that all associated state has

been deleted from all nodes on the path. With some flow state establishment methods signaling new state may be sufficient. A mechanism with a sufficiently long timeout period before reusing the Flow Label values can also be used.

With [RSVP] or [SDP] either the source or the destination of the flow could have a preference for the Flow Label value to be used. For

example, a destination with multiple sources sending packets to it could require all the sources to use the same Flow Label value in order to collapse the classifier state to a single flow state entry, instead of having separate classifier state for each source (ref. the Wildcard-Filter reservation style in [RSVP]). Therefore the source SHOULD honor the destination's request to mark the packets with the Flow Label value specified.

To enable the peer(s) to know the assigned or requested Flow Label value, the value SHOULD be included along with the Source and Destination addresses as part of any signaling dealing with the flow, e.g. transport layer connection set up, RSVP for resource reservation, or SDP for media session parameters.

5. Flow State Establishment Requirements

To enable flow-specific treatment, flow state needs to be established on all or a subset of the IPv6 nodes on the path from the source to the destination(s). The methods for the state establishment, as well as the models for flow-specific treatment are defined in separate specifications.

To enable co-existence of different methods in IPv6 nodes, the methods MUST meet the following basic requirements:

(1) A packet is classified unambiguously to a flow on the basis of the Flow Label, and the Source and Destination Address fields. Depending on the method semantics, multiple such triplets MAY identify the same flow state (e.g. SCTP flows with multiple addresses at either end-points, or a diffserv classifier with an address range. See also the RSVP Wildcard-Filter example in section 4 above). The flow state establishment method MUST convey this classifying information to the IPv6 nodes that need to perform the classification. Usage of any additional header fields for flow classification is beyond the scope of this specification.

(2) The IPv6 node facility keeping track of the Flow Label, and the associated Source and Destination Addresses MUST be utilized when assigning Flow Label values to new flows (see section 4 above).

(3) The Flow Label value 0 is reserved for non-labeled packets.

(4) The method MUST provide the means for flow state clean-up from the IPv6 nodes providing the flow-specific treatment. Both soft- and hard-state methods are possible.

(5) Flow state establishment methods SHOULD be able to recover from the case where the requested flow state cannot be supported.

(6) Flow state establishment methods SHOULD include the Mobile IP Home Addresses of the source and the destination in the state establishment process in addition to the Care-of Addresses, if available. This enables avoiding state duplication on fixed portions of the path when either end changes its Care-of Address.

Security Considerations

Anything that facilitates flow classification also increases the vulnerability to traffic analysis.

The use of the Flow Label field in general enables flow classification also in the presence of ESP encryption of IPv6 payloads. This allows the transport header values to remain confidential, which may lessen the possibilities for some forms of traffic analysis.

IANA Considerations

This specification does not define any well-known values.

Acknowledgements

The discussion on the topic in the IPv6 WG mailing list has been instrumental for the definition of this specification. The authors want to thank Steve Blake, Jim Bound, Francis Dupont, Robert Elz, Tony Hain, Bob Hinden, Christian Huitema, Frank Kastenholz, Charles Perkins, Hesham Soliman, Michael Thomas, and Margaret Wasserman for their contributions.

Normative References

[IPv6] S. Deering, R. Hinden, "Internet Protocol Version 6 Specification", RFC 2460, December 1998.

Informative References

[DiffServ] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.

[Rajahalme] J. Rajahalme, A. Conta, "An IPv6 Flow Label Specification Proposal", Internet Draft <draft-rajahalme-ipv6-flow-label-00.txt>, November 2001, expires May 2002, Work in progress.

[RFC1809] C. Partridge, "Using the Flow Label Field in IPv6", RFC 1809, June 1995.

[RSVP] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource Reservation Protocol (RSVP) Version 1 Functional Specification", RFC 2205, September 1997.

[SDP] M. Handley, V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

Authors' Addresses

Jarno Rajahalme
Nokia Research Center
P.O. Box 407
FIN-00045 NOKIA GROUP,
Finland
E-mail: jarno.rajahalme@nokia.com

Alex Conta
Transwitch Corporation
3 Enterprise Drive
Shelton, CT 06484
USA
Email: aconta@txc.com

Brian E. Carpenter
IBM Zurich Research Laboratory
Saeumerstrasse 4 / Postfach
8803 Rueschlikon
Switzerland

Email: brian@hursley.ibm.com

Steve Deering
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Email: deering@cisco.com

Expiration Date

This memo is filed as <draft-ietf-ipv6-flow-label-02.txt> and expires in December 2002.