

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: January 7, 2011

C. Donley
CableLabs
K. Erichsen
Time Warner Cable
July 6, 2010

Using the Flow Label with Dual-Stack Lite
draft-donley-softwire-dslite-flowlabel-00

Abstract

This document extends the use of Dual-Stack Lite to identify discrete traffic flows using the IPv6 Flow Label. The identification of discrete traffic flows allows for the application of Quality of Service (QoS) classification and prioritization of traffic traversing Dual-Stack Lite tunnels.

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 January 7, 2011.

Copyright Notice

Copyright (c) 2010 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

Internet-Draft

dslite-flowlabel

July 2010

described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Conventions used in this document	4
3.	Allowing Dual-Stack Lite QoS Using the IPv6 Flow Label	5
3.1.	B4 Considerations	5
3.2.	AFTR Considerations	5
3.3.	Constructing the Flow Label	6
4.	Security Considerations	7
5.	IANA Considerations	8
6.	Normative References	9
Appendix A.	Acknowledgements	10
	Authors' Addresses	11

1. Introduction

Dual-Stack Lite [[I-D.ietf-softwire-dual-stack-lite](#)] describes a method for transitioning to IPv6 by encapsulating IPv4 traffic within an IPv6 tunnel and translating it at the Address Family Translation Router. Through such encapsulation, DS-Lite obfuscates the IPv4 5-tuple behind the IPv6 header, thereby making it difficult to classify IPv4 traffic. Thus, to QoS classifiers, all IPv4 traffic is encapsulated as IP-in-IP and uses the same IPv6 source address, destination address, and protocol. There is no differentiation of IPv4 traffic requiring differentiated quality of service such as Voice over IP.

This lack of differentiation can be problematic for traffic types where prioritization is desired. For example, it is common practice to provide QoS for voice traffic. Such classification and prioritization is typically performed by network elements located between the Basic BroadBand Bridging element and Address Family Transition Router, particularly on WAN links. However, in a Dual-Stack Lite environment, many such network elements are unable to identify characteristics of a particular IPv4 traffic flow and apply QoS accordingly.

This document proposes a method of providing QoS to IPv4 traffic in a DS-Lite environment by identifying individual traffic flows using the IPv6 Flow Label.

2. Conventions used in this document

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 [[RFC2119](#)].

[3.](#) Allowing Dual-Stack Lite QoS Using the IPv6 Flow Label

As described in [[RFC3697](#)], a flow is a sequence of packets sent from a particular source to a particular unicast, anycast, or multicast destination that the sender desires to label as a flow. IPv6 uses the Flow Label in the IPv6 header to identify such flows. In cases where the Flow Label field uniquely identifies such traffic flows, packet classifiers can utilize the triplet of Flow Label, Source Address, and Destination Address fields in order to identify packets belonging to a particular flow.

As described in [[I-D.ietf-softwire-dual-stack-lite](#)], traffic is encapsulated between the Basic BroadBand Bridging (B4) element and Address Family Transition Router (AFTR). Both the B4 and AFTR are aware of the 5-tuple of the encapsulated IPv4 traffic. Thus, both the B4 and AFTR are capable of identifying IPv4 traffic flows and setting the IPv6 Flow Label for the DS-Lite tunnel accordingly. By populating the Flow Label field in DS-Lite tunnels, service providers can use Flow Label classifiers to provide priority treatment to appropriate traffic flows.

[3.1.](#) B4 Considerations

The B4 SHOULD uniquely set the IPv6 Flow Label to a non-zero value per IPv4 traffic flow in accordance with [\[RFC3697\]](#). That is, the B4 SHOULD identify IPv4 traffic flows by the IPv4 5-tuple of IPv4 Source Address, Destination Address, Protocol, Source Port, and Destination Port. The B4 SHOULD construct a unique Flow Label based on the IPv4 5-tuple and apply it to the IPv6 header attached to that flow as it is encapsulated within a DS-Lite tunnel. If the B4 sets the IPv6 Flow Label to a non-zero value, it MUST use the same Flow Label value for other IPv4 packets belonging to the same flow (as determined by the IPv4 5-tuple).

[3.2.](#) AFTR Considerations

The AFTR SHOULD uniquely set the IPv6 Flow Label per IPv4 traffic flow in accordance with [\[RFC3697\]](#). That is, the AFTR SHOULD identify IPv4 traffic flows to be sent to the B4 by the IPv4 5-tuple of IPv4 Source Address, Destination Address, Protocol, Source Port, and Destination Port after completing Network Address Translation. The AFTR SHOULD construct a unique Flow Label based on the IPv4 5-tuple and apply it to the IPv6 header attached to that flow as it is encapsulated within a DS-Lite tunnel. If the AFTR sets the IPv6 Flow Label to a non-zero value, it MUST use the same Flow Label value for other IPv4 packets belonging to the same flow (as determined by the IPv4 5-tuple).

[3.3.](#) Constructing the Flow Label

The exact mechanism for constructing the Flow Label is not specified, except as per [\[RFC3697\]](#). Implementations could use a 20-bit hash of the IPv4 5-tuple such that subsequent IPv4 packets with the same 5-tuple will receive the same Flow Label.

As specified by [\[RFC3697\]](#), Flow Label information is only significant to the B4 or AFTR transmitting the particular DS-Lite flow. Since the Flow Label will be consistently applied to all packets in the flow, however, intermediate devices between the B4 and AFTR can use the Flow Label in packet classifiers to provide quality of service treatment to the flow.

[4.](#) Security Considerations

Security considerations are described in [\[RFC3697\], section 5](#). The Flow Label is not protected, and could be modified by an on-path attacker. However, the impact of any such modification would be limited to the QoS treatment of the modified packet(s).

There are no IANA considerations.

6. Normative References

- [I-D.ietf-softwire-dual-stack-lite]
Durand, A., Droms, R., Haberman, B., Woodyatt, J., Lee, Y., and R. Bush, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [draft-ietf-softwire-dual-stack-lite-04](#) (work in progress), March 2010.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3697] Rajahalme, J., Conta, A., Carpenter, B., and S. Deering, "IPv6 Flow Label Specification", [RFC 3697](#), March 2004.

Internet-Draft

dslite-flowlabel

July 2010

[Appendix A](#). Acknowledgements

Thanks to the following people for their guidance and feedback:

Lee Howard

Andy Shappell

Chris Williams

Marla Azinger

Internet-Draft

dslite-flowlabel

July 2010

Authors' Addresses

Chris Donley
CableLabs
858 Coal Creek Circle
Louisville, CO 80027
USA

Email: c.donley@cablelabs.com

Kirk Erichsen
Time Warner Cable
12101 Airport Way
Broomfield, CO 80021
USA

Email: kirk.erichsen@twcable.com

Donley & Erichsen

Expires January 7, 2011

[Page 11]