

Software
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

D. Hankins
Google
T. Mrugalski
Gdansk University of Technology
March 7, 2011

**Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-
Stack Lite
draft-ietf-software-ds-lite-tunnel-option-10**

Abstract

This document specifies a DHCPv6 option which is meant to be used by a Dual-Stack Lite Basic Bridging Broadband (B4) element to discover the IPv6 address of its corresponding Address Family Transition Router (AFTR).

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 September 8, 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.	Requirements Language	3
2.	Introduction	3
3.	The AFTR-Name DHCPv6 Option	3
4.	DHCPv6 Server Behavior	5
5.	DHCPv6 Client Behavior	5
6.	Security Considerations	6
7.	IANA Considerations	7
8.	Acknowledgements	7
9.	Normative References	7
	Authors' Addresses	8

1. Requirements Language

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](#) [[RFC2119](#)].

2. Introduction

Dual-Stack Lite [[I-D.softwire-ds-lite](#)] is a solution to offer both IPv4 and IPv6 connectivity to customers which are addressed only with an IPv6 prefix (no IPv4 address is assigned to the attachment device). One of its key components is an IPv4-over-IPv6 tunnel, commonly referred to as a Softwire. A DS-Lite "Basic Bridging BroadBand" (B4) device will not know if the network it is attached to offers Dual-Stack Lite service, and if it did would not know the remote endpoint of the tunnel to establish a softwire.

To inform the B4 of the Address Family Transition Router's (AFTR) location, a DNS [[RFC1035](#)] hostname may be used. Once this information is conveyed, the presence of the configuration indicating the AFTR's location also informs a host to initiate Dual-Stack Lite (DS-Lite) service and become a Softwire Initiator.

To provide the conveyance of the configuration information, a single DHCPv6 [[RFC3315](#)] option is used, expressing the AFTR's Fully Qualified Domain Name (FQDN) to the B4 element.


The details of how the B4 establishes an IPv4-in-IPv6 softwire to the AFTR are out of scope for this document.

3. The AFTR-Name DHCPv6 Option

The AFTR-Name option consists of option-code and option-len fields (as all DHCPv6 options have), and a variable length tunnel-endpoint-name field containing a fully qualified domain name that refers to the AFTR which the client MAY connect to.

The AFTR-Name option SHOULD NOT appear in any other than the following DHCPv6 messages: Solicit, Advertise, Request, Renew, Rebind, Information-Request and Reply.

The format of the AFTR-Name option is shown in the following figure:



OPTION_AFTR_NAME: (TBD)

option-len: Length of the tunnel-endpoint-name field in octets.

tunnel-endpoint-name: A fully qualified domain name of the AFTR tunnel endpoint.

Figure 1: AFTR-Name DHCPv6 Option Format

The tunnel-endpoint-name field is formatted as required in DHCPv6 [\[RFC3315\] Section 8](#) ("Representation and Use of Domain Names"). Briefly, the format described is using a single octet noting the length of one DNS label (limited to at most 63 octets), followed by the label contents. This repeats until all labels in the FQDN are exhausted, including a terminating zero-length label. Any updates to [Section 8](#) of DHCPv6 [\[RFC3315\]](#) also apply to encoding of this field. An example format for this option is shown in Figure 2, which conveys the FQDN "aftr.example.com".




Figure 2: Example tunnel-endpoint-name.

Note that in the specific case of the example tunnel-endpoint-name (Figure 2), the length of the tunnel-endpoint-name is 18 octets, and so an option-len field value of 18 would be used.

The option is validated by confirming that all of the following conditions are met:

1. the option-len is greater than 3;

2. the option data can be contained by the option length (that the option length does not run off the end of the packet);
3. the individual label lengths do not exceed the option length;
4. the tunnel-endpoint-name is of valid format as described in DHCPv6 [Section 8 \[RFC3315\]](#);
5. there are no compression tags;
6. there is at least one label of nonzero length.

4. DHCPv6 Server Behavior

A DHCPv6 server SHOULD NOT send more than one AFTR-Name option. It SHOULD NOT permit the configuration of multiple names within one AFTR-Name option. Both of these conditions are handled exceptionally by the client, so an operator using software that does not perform these validations should be careful not to configure multiple domain names.

[RFC 3315 Section 17.2.2 \[RFC3315\]](#) describes how a DHCPv6 client and server negotiate configuration values using the Option Request Option (OPTION_ORO). As a convenience to the reader, we mention here that a server will not reply with a AFTR-Name option if the client has not explicitly enumerated it on its Option Request Option.

5. DHCPv6 Client Behavior

A client that supports the B4 functionality of DS-Lite (defined in [\[I-D.softwire-ds-lite\]](#)) and conforms to this specification MUST include OPTION_AFTR_NAME on its OPTION_ORO.

Because it requires DNS name to address resolution, the client MAY also wish to include the OPTION_DNS_SERVERS [\[RFC3646\]](#) option on its OPTION_ORO.

If the client receives the AFTR-Name option, it MUST verify the option contents as described in [Section 3](#).

Note that in different environments, the B4 element and DHCPv6 client may be integrated, joined, or separated by a third pieces of software. For the purpose of this specification, we refer to the "B4 system" when specifying implementation steps that may be processed at any stage of integration between the DHCPv6 client software and the B4 element it is configuring.

If the B4 system receives more than one AFTR-Name option, it MUST use only the first instance of that option.

If the AFTR-Name option contains more than one FQDN, as distinguished by the presence of multiple root labels, the B4 system MUST use only the first FQDN listed in configuration.

The B4 system performs standard DNS resolution using the provided FQDN to resolve a AAAA Resource Record, as defined in [[RFC3596](#)] and STD 13 [[RFC1034](#)] [[RFC1035](#)].

If any DNS response contains more than one IPv6 address, the B4 system picks only one IPv6 address and uses it as a remote tunnel endpoint for the interface being configured in the current message exchange. The B4 system MUST NOT establish more than one DS-Lite tunnel at the same time per interface. For a redundancy and high availability discussion, see [Section 12.3](#) "High availability" of [[I-D.softwire-ds-lite](#)].

Note that a B4 system may have multiple network interfaces, and these interfaces may be configured differently; some may be connected to networks that call for DS-Lite, and some may be connected to networks that are using normal dual stack or other means. The B4 system should approach this specification on an interface-by-interface basis. For example, if the B4 system is attached to multiple networks that provide the AFTR Name option, then the B4 system MUST configure a tunnel for each interface separately as each DS-Lite tunnel provides IPv4 connectivity for each distinct interface. Means to bind a AFTR Name and DS-Lite tunnel configuration to a given interface in a multiple interfaces device are out of scope of this document.

6. Security Considerations

This document does not present any new security issues, but as with all DHCPv6-derived configuration state, it is completely possible that the configuration is being delivered by a third party (Man In The Middle). As such, there is no basis to trust that the access the DS-Lite Softwire connection represents can be trusted, and it should not therefore bypass any security mechanisms such as IP firewalls.

[RFC 3315](#) [[RFC3315](#)] discusses DHCPv6-related security issues.

[[I-D.softwire-ds-lite](#)] discusses DS-Lite related security issues.

7. IANA Considerations

IANA is requested to allocate single DHCPv6 option code referencing this document, delineating OPTION_AFTR_NAME.

8. Acknowledgements

Authors would like to thank Alain Durand, Rob Austein, Dave Thaler, Paul Selkirk, Ralph Droms, Mohamed Boucadair, Roberta Maglione and Shawn Routhier for their valuable feedback and suggestions.

This work has been partially supported by the Polish Ministry of Science and Higher Education under the European Regional Development Fund, Grant No. POIG.01.01.02-00-045/09-00 (Future Internet Engineering Project).

9. Normative References

- [I-D.softwire-ds-lite]
Durand, A., Ed., "Dual-stack lite broadband deployments post IPv4 exhaustion", [draft-ietf-softwire-dual-stack-lite](#) (work in progress).
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", [RFC 3596](#), October 2003.
- [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3646](#), December 2003.

Authors' Addresses

David W. Hankins
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
USA

Email: dhankins@google.com

Tomasz Mrugalski
Gdansk University of Technology
ul. Storczykowa 22B/12
Gdansk 80-177
Poland

Phone: +48 698 088 272

Email: tomasz.mrugalski@eti.pg.gda.pl

