

Behave WG
Internet-Draft
Intended status: Standards Track
Expires: August 21, 2011

J. Korhonen
Nokia Siemens Networks
T. Savolainen, Ed.
Nokia
February 17, 2011

EDNS0 Option for Indicating AAAA Record Synthesis and Format
draft-korhonen-edns0-synthesis-flag-02.txt

Abstract

Advanced hosts and applications benefit of the knowledge of an IPv6 address, AAAA record, synthesis taking place in the network. This draft proposes new EDNS0 option for communicating the synthesis is taking place, used address format, and the IPv6 prefix and suffix used by the DNS64. The communicated information enables hosts to perform local IPv6 address synthesis.

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 21, 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](#) [3](#)
- [2. Requirements and Terminology](#) [3](#)
 - [2.1. Requirements](#) [3](#)
- [3. EDNS0 option for indicating address synthesis](#) [3](#)
- [4. Host behavior](#) [4](#)
- [5. Security Considerations](#) [6](#)
- [6. IANA Considerations](#) [6](#)
- [7. Acknowledgements](#) [6](#)
- [8. References](#) [6](#)
 - [8.1. Normative References](#) [6](#)
 - [8.2. Informative References](#) [7](#)
- [Authors' Addresses](#) [7](#)

1. Introduction

As the networks transition to IPv6, connectivity to IPv4-only domains have to be provided. NAT64 [[I-D.ietf-behave-v6v4-xlate-stateful](#)] and DNS64 [[I-D.ietf-behave-dns64](#)] technologies can be utilized to make IPv4-only peers look like being reachable over IPv6. The DNS64 utilizes IPv6 address synthesis to create local IPv6 presentations of peers having only IPv4 addresses. Applications utilizing DNS for resolving peers' IPv6 addresses can work seamlessly through protocol translation taking place at NAT64.

The DNS64 cannot serve applications not using DNS, such as those receiving IPv4 addresses as referrals. Such applications could nevertheless be able to work through NAT64, provided they are able to create locally valid IPv6 presentations of peers' IPv4 addresses.

This document describes a method for advanced applications to learn the information required to perform local IPv6 address synthesis.

The knowledge of IPv6 address synthesizing taking place may also be useful if DNS64 is present in dual-stack network access. In such cases hosts may choose to use IPv4 addresses instead of synthesized IPv6 addresses, and hence avoid traversal through NAT64.

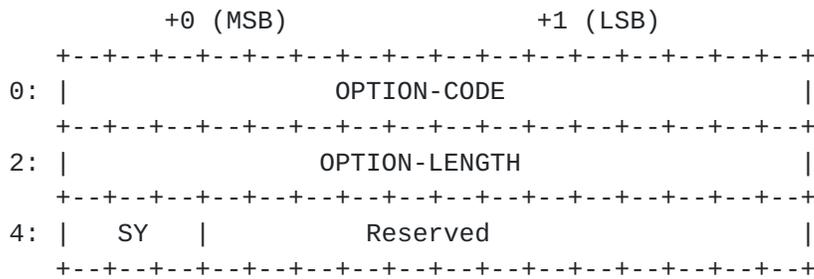
2. Requirements and Terminology

2.1. Requirements

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. EDNS0 option for indicating address synthesis

The mechanism for informing AAAA record synthesis taking place and the used addressing format is communicated in an EDNS0 option in a DNS response. The option has three bits indicating the formats described in [[RFC6052](#)]. The option and bits are structured as follows:



OPTION-CODE (Assigned by IANA)

OPTION-LENGTH 2 (Length of payload in octets)

Possible values for SY-bits are:

- 000 reserved
- 001 prefix length /32
- 010 prefix length /40
- 011 prefix length /48
- 100 prefix length /56
- 101 prefix length /64
- 110 prefix length /96
- 111 address is not synthesized

Reserved Initialized to zero

Figure 1

The prefix length corresponds to the address formats documented in "IPv6 Addressing of IPv4/IPv6 Translators document" [[RFC6052](#)] [section 2.2](#).

Absence of EDNS0 option means that either no synthesis took place or the DNS64 does not support this specification. Either way, when the EDNS0 option is missing, the host cannot conclude for certain whether the AAAA response was synthesized or not. The host may additionally utilize method described in [[I-D.savolainen-heuristic-nat64-discovery](#)].

4. Host behavior

If a host requires information for local IPv6 address synthesis or NAT64 avoidance, it shall send a DNS query for AAAA record of a well-known IPV4-only fully qualified domain name. This well-known name does not have to be in global DNS system. It is enough that DNS64 recognizes the name and replies to it.

The host may query for well-known IPv4-only name, for example, at the moment the host is configured an IPv6 address of a DNS server. This may also happen at the time first DNS query for AAAA record is initiated.

When sending AAAA query for the known name a host MUST set "Checking Disabled (CD)" bit to zero, as otherwise the DNS64 will not perform IPv6 address synthesis hence does not reveal the IPv6 prefix(es) used for protocol translation.

If a host receives negative reply, it learns there are no NAT64 in the network.

A DNS reply with one or more non-empty AAAA records indicates that the access network is utilizing IPv6 address synthesis. The host reads the flag values on the EDNS0 option to learn the used address format, and with that information fetches from the received IPv6 addresses the information used by the network for IPv6 address synthesis (prefix, suffix, u-bit). The host MUST look through all of the received AAAA records to collect all available prefixes. The prefixes may include Well-Known Prefix or one or more Network-Specific Prefixes.

In the case only one IPv6 prefix was present in the DNS response: a host shall use that IPv6 prefix for both local synthetization and for detecting synthesis done by the DNS64 entity on the network.

In the case multiple IPv6 prefixes were present in the DNS response: a host SHOULD use all received prefixes when determining whether other received IPv6 addresses are synthetic. However, for selecting prefix for the local IPv6 address synthesis host MUST use the following prioritization order, of which purpose is to avoid use of prefixes containing suffixes reserved for the future [[RFC6052](#)]:

1. Use NSP having /96 prefix
2. Use WKP prefix
3. Use longest available NSP prefix

In the case of NXDOMAIN or empty AAAA reply: the DNS64 is not available on the access network, network filtered the well-known AAAA query on purpose, or something went wrong in the DNS resolution. All unsuccessful cases result in unavailability of a host to perform local IPv6 address synthesis. The host MAY periodically resend AAAA query to check if DNS64 has become available or temporary problem cleared. The host MAY also continue monitoring DNS replies with IPv6 addresses constructed from WKP, in which case the host MAY use the

WKP as if it were learned during the query for well-known name.

The information required for local IPv6 address synthesis should be made available for applications to utilize.

Alternatively, the host may learn the required information for the local IPv6 address synthesis or the NAT64 avoidance along with any normal DNS query for an AAAA record. In that case all the above considerations and procedures apply, except for the fact that the fully qualified domain used for the DNS query may or may not be provisioned with an AAAA record. Therefore, if and when the EDNS0 option is absent in the reply, the host cannot reliably determine whether the returned IPv6 address is real or synthesized.

5. Security Considerations

No security considerations have been identified.

6. IANA Considerations

IANA should define a name and an IPv4 address for a well-known IPv4-only name.

IANA should allocate new OPTION-CODE for this option.

7. Acknowledgements

The authors would like to acknowledge Andrew Sullivan for presenting general idea of EDNS0 option and SY-bit in behave WG mailing list.

8. References

8.1. Normative References

[I-D.ietf-behave-dns64]

Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", [draft-ietf-behave-dns64-11](#) (work in progress), October 2010.

[I-D.ietf-behave-v6v4-xlate-stateful]

Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6

Clients to IPv4 Servers",
[draft-ietf-behave-v6v4-xlate-stateful-12](#) (work in progress), July 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), October 2010.

8.2. Informative References

[I-D.savolainen-heuristic-nat64-discovery]
Savolainen, T. and J. Korhonen, "Discovery of a Network-Specific NAT64 Prefix using a Well-Known Name",
[draft-savolainen-heuristic-nat64-discovery-01](#) (work in progress), February 2011.

Authors' Addresses

Jouni Korhonen
Nokia Siemens Networks
Linnoitustie 6
FI-02600 Espoo
Finland

Email: jouni.nospam@gmail.com

Teemu Savolainen (editor)
Nokia
Hermiankatu 12 D
FI-33720 Tampere
Finland

Email: teemu.savolainen@nokia.com

