DNSOP Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 13, 2017

R. Bellis
ISC

January 09, 2017

EDNS X-Proxied-For
draft-bellis-dnsop-xpf-01

Abstract

It is becoming more commonplace to install front end proxy devices in front of DNS servers to provide (for example) load balancing or to perform transport layer conversions.

This document defines an option within the EDNS(0) Extension Mechanism for DNS that allows a DNS server to receive the original client source IP address when supplied by trusted proxies.

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 July 13, 2017.

Table of Contents

1. Introduction

It is becoming more commonplace to install front end proxy devices in front of DNS servers [RFC1035] to provide load balancing or to perform transport layer conversions (e.g. to add DNS over TLS [RFC7858] to a DNS server that lacks native support).

This has the unfortunate side effect of hiding the clients' source IP addresses from the server, making it harder to employ server-side technologies that rely on knowing those address (e.g. ACLs, DNS Response Rate Limiting, etc).

This document defines an option within the EDNS(0) Extension Mechanism for DNS [RFC6891] that allows a DNS server to receive the original client source IP address when supplied by trusted proxies.

Whilst in some circumstances it would be possible to re-use the Client Subnet EDNS Option [RFC7871] to carry this information, a new option is defined to allow both this option and the Client Subnet option to co-exist in the same packet.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

The word "proxy" in this document means a network component that sits on the inbound query path in front of a recursive or authoritative DNS server, receiving DNS queries from clients and dispatching them to local servers. This is to distinguish these from a "forwarder" since that term is usually understood to describe a network component that sits on the oubound query path of a client.

3. Description

3.1. EDNS Option Format

The overall format of an EDNS option is shown for reference below, per [RFC6891], followed by the option specific data:

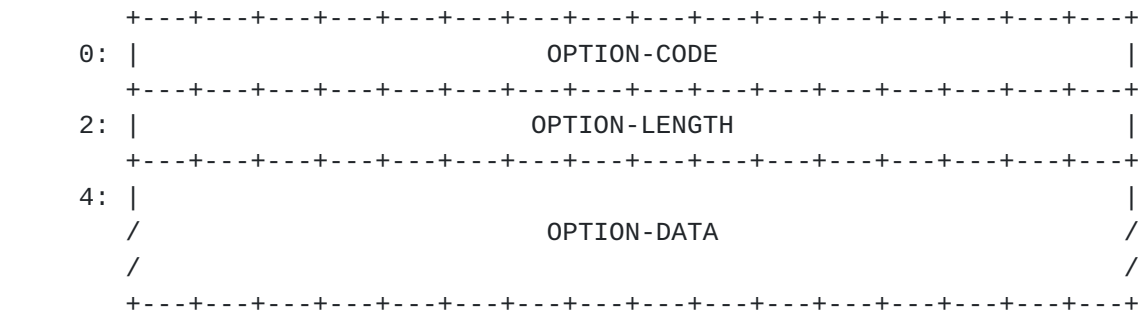

OPTION-CODE: TBD, with mnemonic "XPF".

OPTION-LENGTH: Size (in octets) of OPTION-DATA.

OPTION-DATA: Option specific, as below:

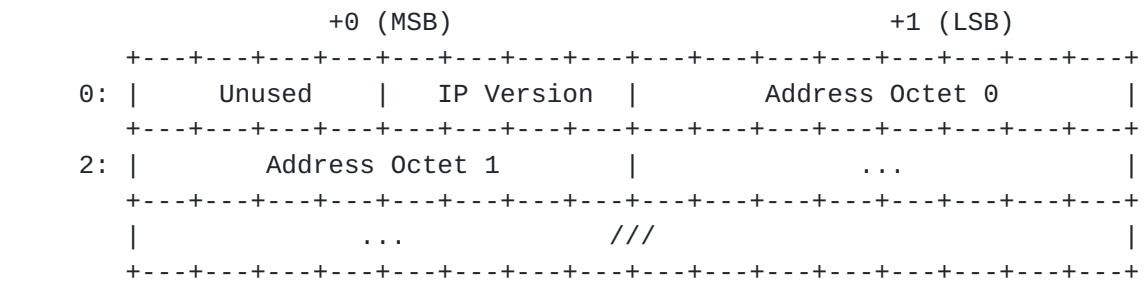

Unused: Currently reserved. These MUST be zero unless redefined in a subsequent specification.

IP Version: The IP protocol version number used by the client, as defined in the IANA IP Version Number Registry [IANA-IP]. Implementations MUST support IPv4 (4) and IPv6 (6).

Address: The source IP address of the client.

3.2. Proxy Processing

Proxies MUST append this option to each request packet received before sending it to the intended DNS server.

If this option is already present in an incoming request it MUST be stripped from the request unless the request was received from an upstream proxy that is itself white-listed by the receiving proxy (i.e. if the proxies are configured in a multi-tier architecture), in which case the original value of the option MUST be preserved.

If the proxy has to create a new OPT RR (because none was present in the original request) it MUST strip any OPT RR subsequently seen in the response for conformance with Section 7 of [RFC6891].

3.3. Server Processing

When this option is received from a white-listed client the DNS server MUST (SHOULD?) use the address from the option contained therein in preference to the client's source IP address for any data processing logic that would otherwise depend on the latter.

If this option is received from a non-white-listed client the server MUST return a REFUSED response.

If the IP version is not understood by the server it MUST return a REFUSED response.

If the length of the client IP address contained in the OPTION-DATA is not consistent with that expected for the given IP version then the server MUST return a FORMERR response.

Servers MUST NOT send this option in DNS responses.

3.4. Secret Key Transaction Authentication for DNS (TSIG)

The considerations for TSIG [RFC2845] from Section 4.5 of "DNS Proxy Implementation Guidelines" [RFC5625] apply here.

A TSIG-signed request MUST either:

1. be forwarded according to RFC 5625 without addition of this option, or

2. be verified using a secret shared between client and proxy, updated with this option, and then re-signed with a (potentially different) shared secret before sending to the server.

In the case of option 1, the server might still be able to uniquely identify and authenticate the client through its shared key, but not by its IP address.

If option 2 is used, there is an operational trade-off to be considered as to whether the two secrets (between client and proxy, and between proxy and server) are actually the same secret. A potential advantage of three-way sharing of the secret is that if the server response requires no modifications it may be returned directly to the client without any further TSIG operations.

Author's note: A third alternative exists, which is to append an additional TSIG signature to the packet based on a secret shared only between the proxy and server. If end-to-end TSIG validation is required alongside TSIG validation between proxy and server, the server would have to 1) validate that second signature, 2) strip it, and then 3) perform further validation on the original signature. Feedback is sought on whether this is worth pursuing.

4. Security Considerations

If the white-list of trusted proxies is implemented as a list of IP addresses, the server administrator MUST have the ability to selectively disable this feature for any transport where there is a possibility of the proxy's source address being spoofed.

This does not mean to imply that use over UDP is impossible - if for example the network architecture keeps all proxy-to-server traffic on a dedicated network and clients have no direct access to the servers then the proxies' source addresses can be considered unspoofable.

5. Privacy Considerations

Used incorrectly, this option could expose internal network information, however it is not intended for use on proxy / forwarder devices that sit on the client-side of a DNS request.

This specification is only intended for use on server-side proxy devices that are under the same administrative control as the DNS servers themselves. As such there is no change in the scope within which any private information might be shared.

6. IANA Considerations

IANA are directed to assign the value TBD for the XPF option in the DNS EDNS0 Option Codes Registry.

7. Acknowledgements

8. References

8.1. Normative References

[IANA-IP] IANA, "IANA IP Version Registry", November 2016, <http://www.iana.org/assignments/version-numbers/>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <http://www.rfc-editor.org/info/rfc1035>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2845] Vixie, P., Gudmundsson, O., Eastlake 3rd, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, DOI 10.17487/RFC2845, May 2000, <http://www.rfc-editor.org/info/rfc2845>.

[RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP 152, RFC 5625, DOI 10.17487/RFC5625, August 2009, <http://www.rfc-editor.org/info/rfc5625>.

[RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <http://www.rfc-editor.org/info/rfc6891>.

8.2. Informative References

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <http://www.rfc-editor.org/info/rfc7858>.

[RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", RFC 7871, DOI 10.17487/RFC7871, May 2016, <http://www.rfc-editor.org/info/rfc7871>.

Author's Address

Ray Bellis
Internet Systems Consortium, Inc.
950 Charter Street
Redwood City CA 94063
USA

Phone: +1 650 423 1200
Email: ray@isc.org