

DNS Extensions Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 2010

G. Barwood
23 October 2009

DNS Transport Signal
draft-barwood-dnsext-dns-transport-signal-02

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 23, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

Describes a DNS resource record that is used to signal support for DNS transport protocols.

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](#)].

1. Introduction

DNS clients are currently unable to efficiently determine which DNS transport protocols a DNS server supports.

DNS clients may try each protocol in turn, but this is an undesirable waste of resources and time, especially as multiple probes have to be sent to take account of packet loss.

Even for the current protocols, TCP and UDP, a client may prefer to use TCP for security reasons, but may not be willing to wait for the TCP connection to fail where the server does not support TCP.

It is expected that new optional protocols may be added in future, for example SCTP or a new special purpose protocol.

Therefore a new resource record type, TPORT is proposed.

Note: throughout this document, unless otherwise qualified, "protocol" means "DNS Transport Protocol", that is the means by which DNS messages are conveyed, not the underlying network protocol. [[RFC1035](#)] uses the term "transmission channel" when discussing truncation. There may be distinct DNS transport protocols operating over the same underlying protocol, for example over UDP. Typically the first DNS transport protocol using a specific underlying protocol will use the name of the underlying protocol, and subsequent DNS transport protocols over the same underlying protocol will be given a different name.

2. TPORT Resource record format

The RDATA wire format is a list of one or more 8-bit numbers that identify DNS transport protocols.

The RDATA presentation format is a list of one or more protocol mnemonics. If the mnemonic is not known, the decimal number for the DNS transport protocol may be used instead, as specified in the IANA considerations, [section 5](#).

Example:

```
NS1.EXAMPLE.NET. 3600 TPORT UDP TCP
```

3. Protocol

The TPORT record for a domain, if it exists, SHOULD be added to the Additional Section of a DNS response whenever an A or AAAA record for the domain is sent.

In particular, a parent zone with a glue A or AAAA record may also have a glue TPORT record. If the parent zone does not support the TPORT

record, or there is no facility for the domain owner to upload a TPORT record to the parent zone, the method described in [Appendix A](#) may be used instead.

The TPORT, A and AAAA records SHOULD have the same TTL, and can be considered to form a single logical, consistent RRset that is divided into distinct RRTypes for historical reasons.

DNS clients SHOULD use the TPORT information to select the most suitable protocol to use. Clients MAY fall back to another protocol if an advertised protocol fails, but SHOULD take account of the security implications, if the fallback protocol is less secure.

The absence of a protocol indicates that clients SHOULD NOT use a protocol for that name server, however if no TPORT record is available, no inferences can be made.

The order in which the protocols are listed has no significance.

To avoid inter-operability problems with old non-conformant resolvers, when the DNS transport protocol is UDP (without EDNS), or according to similar criteria determined by operational experience, TPORT records MAY be omitted unless explicitly requested.

4. Security Considerations

Until server support for a new DNS transport protocol is universal, there is a risk that a server may be downgraded after a protocol has been advertised, resulting in a lame server. The risk is higher where in-zone secondary servers are used that are not under the direct control of the domain owner, and no reliable change notification mechanism is in place. Domain owners may avoid this risk by using out-of-zone name server names where they do not have direct control of the servers, however this is not desirable in some cases. Domain owners should carefully weigh the advantages of a new protocol against this risk. Domain owners may conduct regular checks for lameness to mitigate the risk.

New transport protocols may have different or unforeseen security risks. Otherwise, this specification is not believed to directly cause any new security problems.

5. IANA Considerations

IANA is requested to allocate the TPORT resource record type, and a sub-registry for DNS transport protocols, initialized to

TCP	1	[RFC1035]
UDP	2	[RFC1035]

Numbers 240 to 250 are reserved for private use.

6. Acknowledgments

Thanks to Alex Bligh, Matthew Dempsky, Alfred Hoenes, Shane Kerr,
Olaf Kolkman and Paul Vixie for their comments.

7. Normative References

[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.

Appendix A. Name server name encoding

It may take some time for the TPORT record to be universally supported. In the interim period, TPORT information may be encoded into a name server name.

The convention is that the name server name contains a label starting with 'TPORT-', followed by a list of one or more protocol mnemonics separated by '- '.

For example

EXAMPLE.COM. 3600 NS A.TPORT-TCP-UDP.EXAMPLE.NET.

indicates that the name server for EXAMPLE.COM has support for TCP and UDP.

This method is far from ideal, and is intended mainly for early adopters to experiment with this technology. It does however offer the potential for better security. Operators are reminded that correct procedures need to be followed when changing the name servers for a domain.

Author's Address

George Barwood
33 Sandpiper Close
Gloucester
GL2 4LZ
United Kingdom

Phone: +44 452 722670
EMail: george.barwood@blueyonder.co.uk

