

DNSOP Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 11, 2014

W. Wijngaards
NLnet Labs
G. Wiley
VeriSign, Inc.
March 10, 2014

Confidential DNS
draft-wijngaards-dnsop-confidentialdns-01

Abstract

This document offers opportunistic encryption to provide privacy for DNS queries and responses.

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](#)].

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 11, 2014.

Copyright Notice

Copyright (c) 2014 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.	ENCRYPT RR Type	3
3.	Server and Client Algorithm	4
4.	Authenticated Operation	6
5.	Comparison with TLS and DTLS	6
6.	IANA Considerations	7
7.	Security Considerations	7
8.	Acknowledgments	7
9.	Normative References	7

1. Introduction

The privacy of the Question, Answer, Authority and Additional sections in DNS queries and responses is protected by the confidential DNS protocol by encrypting the contents of each section. The goal of this change to the DNS protocol is to make large scale monitoring more expensive, see [[draft-bortzmeyer-dnsop-dns-privacy](#)] and [[draft-koch-perpass-dns-confidentiality](#)]. Authenticity and integrity may be provided by DNSSEC, this protocol does not change DNSSEC and does not offer the means to authenticate responses.

Confidential communication between any pair of DNS servers is supported, both between iterative resolvers and authoritative servers and between stub resolvers and recursive resolvers.

The confidential DNS protocol has minimal impact on the number of packets involved in a typical DNS query/response exchange by leveraging a cacheable ENCRYPT Resource Record and an optionally cacheable shared secret. The protocol supports selectable cryptographic suites and parameters (such as key sizes).

The client fetches an ENCRYPT RR from the server that it wants to contact. The public key retrieved in the ENCRYPT RR is used to encrypt a shared secret or public key that the client uses to encrypt the sections in the DNS query and which the name server uses to encrypt the DNS response. Note that an ENCRYPT RR must be fetched for each name server in order for the entire session to be confidential.

As this is opportunistic encryption, the key is (re-)fetched when the exchange fails. If the key fetch fails or the encrypted query fails, communication in the clear may be performed.

The server advertises which crypto suites and key lengths may be used in the ENCRYPT RR, the client then chooses a crypto suite from this list and includes that selection in subsequent DNS queries.

The key from the server can be cached by the client, using the TTL specified in the ENCRYPT RR, the IP address of the server distinguishes keys in the cache. The server may also cache shared secrets and keys from clients.

2. ENCRYPT RR Type

The RR type for confidential DNS is ENCRYPT TBD (decimal). The presentation format is:

```
. ENCRYPT [flags] [algo] [id] [data]
```


The flags, algo and id are unsigned numbers in decimal and the data is in base-64. The wireformat is: one octet flags, one octet algo, one octet id and the remainder of the rdata is for the data. The type is class independent and it is a hop-by-hop transaction RR type. The domain name of the ENCRYPT record is '.' (the root label) for hop-by-hop exchanges.

In the flags the least two bits are used as usage value. The other flag bits MUST be ignored by receivers and sent as zeroes.

- o PAD (value=0): the ENCRYPT contains padding material. Algo and id are set to 0. Its data length is random (say 1-63 octets), and has some random values. It is a resource record that may be appended to resource records that are encrypted so that identical queries encrypt to different encrypted data of different lengths.
- o KEY (value=1): the ENCRYPT contains a public or symmetric key. The algo field gives the algorithm. The id identifies the key, this id is copied to ENCRYPT type RRS to identify which key to use to decrypt the data. The data contains the key bits.
- o RRS (value=2): encrypted data. The data contains encrypted resource records. The data is encrypted with the selected algorithm and key id. The data contains resource records in DNS wireformat [[RFC1034](#)], with a domain name, type, class, ttl, rdatalength and rdata.
- o SYM (value=3): the ENCRYPT contains an encrypted symmetric key. The contained, encrypted data is rdata of an ENCRYPT of type KEY and has the symmetric key. The data is encrypted with the algorithm and id indicated. The encrypted data encompasses the flags, algo, id, data for the symmetric key.

The ENCRYPT RR type can contain keys. It uses the same format as the DNSKEY record [[RFC4034](#)] for public keys. algo=0 is reserved for future expansion of the algorithm number above 255. algo=1 is RSA, the rdata determines the key size, such as 512 and 768 bits. algo=2 is AES, aes-cbc, size of the rdata determines the size of the key, such as 128 and 192 bits.

3. Server and Client Algorithm

If a clients wants to fetch the keys for the server from the server, it performs a query with query type ENCRYPT and query name '.' (root label). The reply contains the ENCRYPT (or multiple if a choice is offered) in the answer section. These ENCRYPTs have the KEY flag set.

If a client wants to perform an encrypted query, it sends an unencrypted outer packet, with query type ENCRYPT and query name '.' (root label). In the additional section it includes an ENCRYPT record of type RRS. This encrypts a number of records, the first is a query-section style query record, and then zero or more ENCRYPTs of type KEY that the server uses to encrypt the reply. If the client wants to use a symmetric key, there are no ENCRYPTs of type KEY inside the encrypted ENCRYPT data, instead an ENCRYPT of type SYM is positioned in the outer packet, before the ENCRYPT of type RRS and the ENCRYPT of type RRS is encrypted with the symmetric key.

If a server wants to encrypt a reply, it also uses the ENCRYPT type. The reply looks like a normal DNS packet, i.e. it has a normal unencrypted outer DNS packet. Because the query name and query type have been encrypted, the outer packet has a query name of '.' and query type of ENCRYPT and the reply has an ENCRYPT record in the answer section with flag RRS. The reply RRs have been encrypted into the data of the ENCRYPT record. The RR counts for every section are stored in the outer (unencrypted) header. Thus, the combination of the original header and the decrypted data from this record results in the decrypted packet.

The client may lookup keys whenever it wants to. It may cache the keys for the server, using the TTL of those ENCRYPT records. It should also cache failures to lookup the ENCRYPT record for some time (eg. the negative TTL if the reply contained one). Errors and also timeouts should also be taken as an indication that the ENCRYPT cannot be looked up, and the client MUST fall back to unencrypted communication (this is the opportunistic encryption case). The result of an encrypted query may also be timeouts, errors or replies with mangled contents, in that case the client MUST fall back to unencrypted communication (this is the opportunistic encryption case). Note that if some middlebox removes the ENCRYPT from the additional section of an encrypted query, likely a reply with ENCRYPTs of type KEY is returned instead of the encrypted reply with an ENCRYPT of type RRS, and again the client does the unencrypted fallback. If the server has changed its keys and does not recognize the keys in an encrypted query, it should return FORMERR, and include its current ENCRYPTs of type KEY in that FORMERR reply. A server may decide it does not (any longer) have the resources for encryption and reply with SERVFAIL to encrypted queries, forcing unencrypted fallback. Keys for unknown algorithms should be ignored by the client, if no usable keys remain, fallback to insecure.

The client may cache the ENCRYPT of type SYM for a server together with the symmetric secret, this is better for performance, as public-key operations can be avoided for repeated queries. The server may also cache the ENCRYPTs of type SYM with the decoded secret,

associating a lookup for the rdata of the SYM record with the decoded secret, avoiding public-key operations for repeated queries.

Key rollover is possible, use different key ids and support the old key for its TTL, while advertising the new key, for the servers. For clients, generate a new public or symmetric key and use it.

4. Authenticated Operation

The previous documented the opportunistic operation, where deployment is easier, but security is weaker. This documents options for authenticated operation. `[[is this out of scope?]]`. The client selects if encryption is authenticated, opportunistic, or disabled.

The ENCRYPT KEYS for authority servers can be signed with DNSSEC. The domain name of the nameserver is used to store the ENCRYPTs of type KEY. `[[Could also be in the reverse tree for the IP address]]`.

For authenticated operation, fallback to insecure should not be performed. However, this will significantly harm deployment as unclean lookup paths result in lookup failure. Keys with unsupported crypto algorithms MUST still be ignored and if no keys are left, fallback to insecure MUST still be performed, also for authenticated operation.

The key for recursive resolvers can be configured into the stub machines, or a domain name can be configured where the keys are looked up and they are signed with DNSSEC.

5. Comparison with TLS and DTLS

An alternative method of accomplishing confidential DNS would be to leverage one of the existing means for establishing a secure transport layer. For example a secure TCP session could be established to the name server over which DNS queries could be sent with no changes to the DNS protocol. The most significant down side to this approach is the burden that it places on high volume name servers. Very large scale DNS operators expect to answer hundreds of thousands of queries per second (possibly even more than a million qps) for each host in their name server footprint. The use of technologies such as IPsec or TLS may have such a severe impact on the largest name server operators as to impede adoption of confidential DNS.

DTLS ([RFC 6347](#)) offers a more interesting approach to securing the connection to a name server that may be implemented in a way that is less abusive to the large scale name servers. It looks as though the overhead imposed by DTLS would probably be significantly higher than

the protocol described in this draft, however if the session established via DTLS is used over a large number of queries then the cost of the handshake could be amortized over the total number of queries.

6. IANA Considerations

An RR type registration for type ENCRYPT with number TBD and it references this document [[to be done when this becomes RFC]].

7. Security Considerations

Opportunistic encryption is the main mode here. Opportunistic encryption has many drawbacks, against active intrusion, but works against passives. The pervasive passive surveillance problem statement and also its security considerations are applicable to this document. Hence the suggested short key sizes and opportunistic encryption.

This technique does not provide perfect forward secrecy. Additionally, timing, traffic analysis (what IP address is contacted), packet size, RR count, header flags and header RCODE can be observed. These could provide almost all the information that was encrypted. Such as: query to IP address for example.com nameservers, size of the packet is similar to a www.example.com lookup and is followed by http packets to www.example.com's IP address.

8. Acknowledgments

Roy Arends

9. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", [RFC 4034](#), March 2005.

Authors' Addresses

Wouter Wijngaards
NLnet Labs
Science Park 140
Amsterdam 1098 XH
The Netherlands

EMail: wouter@nlnetlabs.nl

Glen Wiley
VeriSign, Inc.
Reston, VA
USA

EMail: gwiley@verisign.com

