

dprive or dnsop  
Internet-Draft  
Intended status: Informational  
Expires: February 14, 2020

K. Henderson  
Verisign  
T. April  
Akamai  
J. Livingood  
Comcast  
August 13, 2019

Authoritative DNS-over-TLS Operational Considerations  
draft-hal-adot-operational-considerations-02

Abstract

DNS over TLS (DoT) has been gaining attention, primarily as a means of communication between stub resolvers and recursive resolvers. There have also been discussions and experiments involving the use of DoT to communicate with authoritative nameservers (Authoritative DNS over TLS or "ADoT"), including communication between recursive and authoritative resolvers. However, we have identified a number of operational concerns with ADoT that have arisen as DNS operators have begun to experiment with and prepare for deploying DoT. These operational concerns need to be addressed prior to ADoT's deployment at scale by DNS operators in order to maintain the stability and resilience of the global DNS. The document also provides some suggested next steps to advance the operator community's understanding of ADoT's operational impact.

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 <https://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 February 14, 2020.

## Copyright Notice

Copyright (c) 2019 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 (<https://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
1.1. Background and Motivation . . . . .	3
1.1.1. Why operational considerations are so important for ADoT . . . . .	3
1.1.2. Other considerations related to ADoT . . . . .	4
2. Terminology . . . . .	5
2.1. Requirements Language . . . . .	5
2.2. Definitions . . . . .	5
3. Key Issues and Questions . . . . .	6
3.1. Signaling Support for ADoT . . . . .	6
3.2. Port number . . . . .	6
3.3. TLS version . . . . .	7
3.4. Resumptions . . . . .	7
3.5. Operational Monitoring . . . . .	8
3.6. Architecture . . . . .	8
3.7. Socket efficiency/tuning considerations . . . . .	8
3.8. Post-Quantum Security . . . . .	9
4. Suggestions for further research and development . . . . .	9
4.1. Required studies and analysis . . . . .	9
4.2. Authoritative DNS over TLS (ADoT) Profile . . . . .	10
5. Security Considerations . . . . .	10
6. References . . . . .	10
6.1. Informative References . . . . .	10
6.2. URIs . . . . .	13
Appendix A. Acknowledgements . . . . .	13
Appendix B. Change Log . . . . .	13
Authors' Addresses . . . . .	14

## 1. Introduction

This is an operational considerations document that focuses on the factors operators need to consider when implementing Authoritative DNS over TLS. An evaluation of the merits of DNS over TLS are beyond the scope and intent of this document.

Typically, DNS communication between stub resolvers, recursive resolvers, and authoritative servers is not encrypted. Some argue that this can pose a privacy challenge for Internet users, because their access to named network resources can potentially be tracked through their DNS communication. In principle, any network element along the path between the user and resolving or authoritative nameservers could observe this unencrypted traffic. Many of these concerns are addressed in [RFC7626].

[RFC8310] proposes using DNS over TLS (DoT) in order to encrypt DNS traffic between hosts.

Historically, much of the work on DNS encryption has focused on the stub-to-recursive path, as the recursive-to-authoritative server path did not leak user specific information, other than if identifying information was encoded in the QName. However, with the increased deployment of EDNS0 Client Subnet [RFC7871], recursive-to-authoritative encryption is becoming an area of interest. Therefore, this document's scope is the recursive-to-authoritative aspect of DoT, or Authoritative DNS over TLS (ADoT), in order to differentiate it from the stub-to-recursive path.

The addition of ADoT, while providing encryption for DNS communication, also introduces other factors that might impact the stability and resiliency of authoritative nameserver operations which may have been optimized for unencrypted DNS, often focusing on UDP transport.

The objective of this document is to try to describe the problem space, make suggestions about solutions, and propose next steps that can help inform both recursive and authoritative operators on how to assess and address the challenges posed by ADoT deployment.

### 1.1. Background and Motivation

#### 1.1.1. Why operational considerations are so important for ADoT

The main concerns for most authoritative operators are the stability, resiliency, scalability, and performance of their platforms. These concerns need to be weighed against the benefits, provided to the end user, by encrypting DNS queries to the authoritative servers.

As a result of caching, the recursive-to-authoritative server communication is less attributable to a particular user than information communicated along the stub-to-recursive path. In cases where the recursive is shared and using QName minimization, some user privacy concerns may be addressed, but cases exist where QName Minimization may not be used, or users may be running their own resolver, defeating the shared service protection.

Initial deployments of ADoT may offer an immediate expansion of the attack surface (additional port, transport protocol, and computationally expensive crypto operations for an attacker to exploit) while, in some cases, providing limited protection to end users.

#### 1.1.2. Other considerations related to ADoT

As resolvers add encryption on the client-to-recursive path, they may also change the way they handle data on the recursive-to-authoritative path. This is expressed in Mozilla Trusted Recursive Resolver (TRR) requirements [1], for example, which require participating resolvers to perform QNAME minimization [RFC7816], and TRR requirement #6, which forbids the EDNS0 Client subnet (ECS) from being propagated unless the recursive-to-authoritative path is encrypted.

The latter requirement may have the possible unintended consequence of reducing the authoritative name servers' ability to provide a best response to DNS queries, until such time as they deploy DNS encryption.

Given that recursive resolvers should be configured to prevent ECS transmission to root, top-level, and effective top-level domain (TLD) servers [RFC7871] section 12.1 [2] - the ECS encryption requirement motivates consideration of authoritative DNS encryption below these levels.

At the higher levels, techniques such as QNAME minimization and Aggressive Use of DNSSEC-Validated Cache [RFC8198] arguably provide an alternate path toward mitigating the risk of disclosure of sensitive information without the operational risk of DNS encryption.

Resolver requirements may change as the understanding of DNS encryption options evolve, but in the meantime, they provide motivation for authoritative name server operators to weigh the risks and benefits of DNS encryption, hence the importance of understanding these operational considerations.

## 2. Terminology

### 2.1. Requirements Language

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

### 2.2. Definitions

Authoritative DoT (ADoT): [I-D.draft-hoffman-dns-terminology-ter-01]

Attack Surface: The sum of attack vectors where an unauthorized user (attacker) can try to enter or extract data from the environment or compromise a service via resource starvation.

Authoritative Operator: An operator of an authoritative DNS server.

CDN: Content Delivery Network - distributed network of servers which proxy traffic between content providers and end users in order to provide high availability and high performance.

ECS: EDNS0 Client Subnet [RFC7871] - an extension to EDNS0 where the client's subnet is included in the DNS query, intended to provide a hint to authoritative servers who may wish to provide different answers in an attempt to provide higher performance for end users based on their network location.

EPSK: External Pre-Shared Key - TLS 1.3 [RFC8446] uses the same PSK extension for keys established both during handshake (resumption PSK) and keys established externally. The EPSK acronym was introduced in draft [I-D.draft-wood-tls-external-psk-importer-00] in order to disambiguate External vs Resumption PSKs.

Performance:

- o QRTT: Query Round-Trip Time - the time it takes between sending a query and receiving a response.
- o Best Response - whether or not the authoritative server, if dynamic responses are used as they are in CDNs, are able to determine or infer location and provide the most local response. It is a key part of the end-to-end performance for end users to get not just an answer quickly but to get the best and most local answer.

- o System Performance - the cost in system resources such as CPU/IO/Memory.
- o Queries Per Second (QPS) - the maximum number of simultaneous queries that a DNS server can handle, on a per second basis.

Authoritative Server - see [RFC8499]

Recursive Resolver - see [RFC8499]

Stub Resolver - see [RFC8499]

TLS: - see [RFC8446]

SUDN: Special-Use Domain Names - see [RFC6761]

### 3. Key Issues and Questions

#### 3.1. Signaling Support for ADoT

[RFC8310] does not define a method for a nameserver to advertise its support of DoT other than to have the client make a connection attempt to the default port of 853. The extra round-trip to check for ADoT support imposes a penalty for clients and resolvers that either do not remember the nameserver or have not communicated with that nameserver before. The extra round-trip required may lead some implementers down a similar path to happy eyeballs [RFC8305] which, in the case of DNS, would send the same query over both encrypted and un-encrypted channels at the same time. A happy eyeballs type approach, which we'll call "leaky resolvers", would defeat the purpose of the encryption protection for the testing query, but may enable subsequent queries to be sent over a private channel with the first query being subject to on-path adversaries. An implementation could use some constant query string as a test query. However any query included in the set of queries comprising the iterative resolution for a QNAME first sent over an encrypted channel that leaks the original stub QNAME, SHOULD NOT be used.

#### 3.2. Port number

[RFC7858] section 3 [3] indicates that port 853 MUST be used for session establishment unless otherwise negotiated and configured by both the client and server. In the stub-to-recursive connection, changing the port is something that can be done at stub configuration time however, managing this negotiation between the recursive-to-authoritative server is not scalable or standardized. The scalability problem is due to the fact that recursive resolvers

communicate with thousands of authoritative servers, therefore port/service discovery for each of these authoritatives becomes difficult.

Static use of a pre-defined port provides on-path adversaries the ability to more easily drop or manipulate traffic intended for that port, possibly triggering resolvers to downgrade a connection back to a traditional DNS query, eliminating the encryption protections. This attack is more likely to happen on the stub-to-recursive connection but is also a possible threat for recursive-to-authoritative connections.

### 3.3. TLS version

Implementers of ADoT should read, understand, and follow the guidance provided in BCP195 [4], also known as [RFC7525], when deploying DoT on their platforms. At the time of writing, [RFC7525] did not include coverage for TLS 1.3. However, TLS 1.3 should be included in the document that obsoletes this BCP. Until this happens, TLS 1.3 SHOULD be preferred over TLS 1.2, as 1.3 offers both security and performance enhancements. Additionally, operators should monitor TLS version issues and cipher suite vulnerabilities for the version of TLS that their platforms offer.

In the absence of any widespread ADoT deployments, it is easier to limit TLS version 1.3 or greater. The absence of widespread adoption also allows the IETF to create and enforce standards/policies that ensure TLS versions are kept current going forward.

### 3.4. Resumptions

TLS resumption allows clients and servers to use information from a previously established session in order to bootstrap the cryptographic state while avoiding a full handshake. The resumption mechanism is redesigned in TLS 1.3 [RFC8446] section 2.2 [5] and section 2.3 [6], eliminating both [RFC5077] session tickets and session ID resumption.

Resumption improves both connection and resource (socket and CPU) efficiency, therefore operators SHOULD allow for TLS resumption. However, special consideration should be given to 0-RTT resumption as it is vulnerable to replay attacks [RFC3552] see Section 3.3.1 [7]. The replay attack may not be as important for DNS, as DNS queries are generally idempotent. However consideration should be given to possible side-channel attacks [8].

### 3.5. Operational Monitoring

Many operators use external passive monitors in order to understand the health and performance of their infrastructure. Infrastructure monitoring is also often done to retain a copy of traffic for forensic purposes – such as the BIND "packet of death" [9] scenario. These legacy monitoring systems may break with the use of TLS 1.3. Therefore alternatives may need to be deployed/developed in order to maintain effective operational performance and security monitoring functionality.

A number of solutions have been suggested:

- o TLS Security and Data Center Monitoring: Searching for a Path Forward [10]
- o TLS 1.3: Will Your Network Monitoring Go Blind? [11]

### 3.6. Architecture

Operators often reconfigure their architectural designs to best deliver a new product offering or service. Operators should consider the following design alternatives for the new ADoT service:

- o Operators should consider segregating ADoT addresses from traditional DNS over UDP/TCP to enable better attack mitigation, better service monitoring, less service interference, and more stability.
- o Operators should weigh the pros/cons of using a TLS proxy vs direct client-to-host connection. In case of ADoT, the client is most likely a recursive resolver and the host is the authoritative host server.

### 3.7. Socket efficiency/tuning considerations

Operators can realize substantial gains in client session establishment and improve overall RTT by tuning sockets setting for best use-case efficiency.

For the ADoT use case, operators should consult [RFC7766] section 6.2 [12] and minimally consider the following:

- o Optimal number of persistent connections – consideration should be given to the number of persistent connections maintained for both the recursive resolvers and authoritative servers
- o Optimal read/write buffer size



- o Optimal session timeout
- o Optimal close wait state time
- o Optimal connection time/timeout

### 3.8. Post-Quantum Security

Given that ADoT deployments will likely have a long lifetime and are being introduced in an era where post-quantum security is now an important design consideration, it is prudent to consider how protections against quantum computers might be integrated into the deployments.

[I-D.draft-hoffman-c2pq-05] outlines the threat quantum computing presents to classical cryptographic algorithms.

External Pre-Shared Keys (EPSKs) may be less vulnerable to quantum attacks. A proposed approach to combining EPSKs and certificates in TLS is described in

[I-D.draft-housley-tls-tls13-cert-with-extern-psk-03].

## 4. Suggestions for further research and development

### 4.1. Required studies and analysis

Unlike stub-to-recursive DNS communication, authoritative nameservers affect users in ways that end users cannot avoid or work around. In the event that all authoritative servers for a zone are unreachable, the zone becomes globally unavailable. Hence, in order to preserve stability and resiliency of authoritative nameservers when deploying ADoT, more empirical studies and analysis **MUST** be conducted. The following list is a minimal set of studies and considerations that need to be conducted/addressed in order to maintain authoritative stability and resilience.

- o Attack vectors and mitigation: consider the new adversarial powers enabled by ADoT – types of attacks and denial of service, or other security challenges that are created with the addition of ADoT to authoritative nameservers.
- o Traffic: consider how traffic patterns to authoritative nameservers change with the introduction of ADoT and how these traffic patterns change when the parameters of the service are changed; e.g. persistent connection lifetime, TLS connection parameters, use of TLS session tickets [RFC5077] or Pre-Shared Key extension in TLS 1.3 [RFC8446] section 2.2 [13]. Consider how

these traffic pattern changes will affect the architecture and infrastructure for authoritative operators.

- o ADoT capacity and footprint expansion: consider how common scaling techniques impact authoritative operators; e.g. anycast, load balancing, custom hardware.
- o DTLS/UDP - consider if there is any reason to implement DTLS given that we lose the benefit of pipelining requests and must drop back to TLS/TCP in the case of fragmentation.

It is critical to conduct large-scale measurements of DNS infrastructure in order to quantify some of the scalability issues. While these tests may be performed initially in a controlled lab environment, the public Internet is fundamentally more variable. Therefore, global testing at scale on the Internet MUST also be conducted in order to understand and measure potential issues which must be overcome before full global deployment can occur.

#### 4.2. Authoritative DNS over TLS (ADoT) Profile

Profiles can be used as a mechanism to help mitigate operational concerns over increased attack surface by restricting features such as computationally expensive processes, insecure ciphers, general starvation vectors, or other features that may limit operational performance.

Therefore, an ADoT application profile draft, taking into account the conclusions of required studies and analysis, may help assuage some of the concerns raised in this document.

### 5. Security Considerations

In addition to the applicable security considerations described in RFCs [RFC7626] and [RFC8310], considerations focused on future deployment of quantum computers are described in Post-Quantum Security (Section 3.8). Additional considerations associated with ADoT are TBD based on working group discussions.

### 6. References

#### 6.1. Informative References

- [I-D.draft-hoffman-c2pq-05]  
Hoffman, P., "The Transition from Classical to Post-Quantum Cryptography", draft-hoffman-c2pq-05 (work in progress), May 2019.

- [I-D.draft-hoffman-dns-terminology-ter-01]  
Hoffman, P., "Terminology for DNS Transports and Location", draft-hoffman-dns-terminology-ter-01 (work in progress), April 2019.
- [I-D.draft-housley-tls-tls13-cert-with-extern-psk-03]  
Housley, R., "TLS 1.3 Extension for Certificate-based Authentication with an External Pre-Shared Key", draft-housley-tls-tls13-cert-with-extern-psk-03 (work in progress), November 2018.
- [I-D.draft-wood-tls-external-psk-importer-00]  
Wood, C., "Importing External PSKs for TLS 1.3", draft-wood-tls-external-psk-importer-00 (work in progress), October 2018.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, DOI 10.17487/RFC5077, January 2008, <<https://www.rfc-editor.org/info/rfc5077>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <<https://www.rfc-editor.org/info/rfc6761>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <<https://www.rfc-editor.org/info/rfc7626>>.

- [RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and D. Wessels, "DNS Transport over TCP - Implementation Requirements", RFC 7766, DOI 10.17487/RFC7766, March 2016, <<https://www.rfc-editor.org/info/rfc7766>>.
- [RFC7816] Bortzmeyer, S., "DNS Query Name Minimisation to Improve Privacy", RFC 7816, DOI 10.17487/RFC7816, March 2016, <<https://www.rfc-editor.org/info/rfc7816>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc7871>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8198] Fujiwara, K., Kato, A., and W. Kumari, "Aggressive Use of DNSSEC-Validated Cache", RFC 8198, DOI 10.17487/RFC8198, July 2017, <<https://www.rfc-editor.org/info/rfc8198>>.
- [RFC8305] Schinazi, D. and T. Pauly, "Happy Eyeballs Version 2: Better Connectivity Using Concurrency", RFC 8305, DOI 10.17487/RFC8305, December 2017, <<https://www.rfc-editor.org/info/rfc8305>>.
- [RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <<https://www.rfc-editor.org/info/rfc8310>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.

## 6.2. URIs

- [1] [https://wiki.mozilla.org/Security/DOH-resolver-policy#Privacy\\_Requirements](https://wiki.mozilla.org/Security/DOH-resolver-policy#Privacy_Requirements)
- [2] <https://tools.ietf.org/html/rfc7871#section-12.1>
- [3] <https://tools.ietf.org/html/rfc7858#section-3>
- [4] <https://tools.ietf.org/html/bcp195>
- [5] <https://tools.ietf.org/html/rfc8446#section-2.2>
- [6] <https://tools.ietf.org/html/rfc8446#section-2.3>
- [7] <https://tools.ietf.org/html/rfc3552#section-3.3.1>
- [8] <https://eprint.iacr.org/2005/388.pdf>
- [9] <https://www.nominet.uk/the-packet-of-death/>
- [10] <https://www.rsa.com/en-us/blog/2017-08/tls-security-and-data-center-monitoring-searching-for-a-path-forward>
- [11] <https://www.extrahop.com/company/blog/2018/maintain-visibility-with-tls-1.3/>
- [12] <https://tools.ietf.org/html/rfc7766#section-6.2>
- [13] <https://tools.ietf.org/html/rfc8446#section-2.2>

## Appendix A. Acknowledgements

Thanks to those that provided usage data, reviewed and/or improved this document, including: Piet Barber, Michael Bentkofsky, David Blacka, Florent Guiliani, Scott Hollenbeck, Burt Kaliski, Glen Wiley, and Richard Wilhelm.

## Appendix B. Change Log

RFC EDITOR: PLEASE REMOVE THE THIS SECTION PRIOR TO PUBLICATION.

TODO: Zero this change log out when -00 is submitted to IETF.

From -01 to -02:

- o Updates to the introduction framing

- o Rewording the "Why operational considerations are so important for ADoT" section

From -00 to -01:

- o Removing static-dh reference
- o Further clarifying the scope of the document

pre-00

- o Initial draft.

~~~ 0123456789012345678901234567890123456789012345678901234  
5678912

#### Authors' Addresses

Karl Henderson  
Verisign

Email: khenderson@verisign.com

Tim April  
Akamai

Email: ietf@tapril.net

Jason Livingood  
Comcast

Email: jason\_livingood@comcast.com

dprive  
Internet-Draft  
Updates: 1995 (if approved)  
Intended status: Standards Track  
Expires: January 9, 2020

H. Zhang  
P. Aras  
Salesforce  
W. Toorop  
NLnet Labs  
S. Dickinson  
Sinodun IT  
A. Mankin  
Salesforce  
July 8, 2019

DNS Zone Transfer-over-TLS  
draft-hzpa-dprive-xfr-over-tls-02

Abstract

DNS zone transfers are transmitted in clear text, which gives attackers the opportunity to collect the content of a zone by eavesdropping on network connections. The DNS Transaction Signature (TSIG) mechanism is specified to restrict direct zone transfer to authorized clients only, but it does not add confidentiality. This document specifies use of DNS-over-TLS to prevent zone contents collection via passive monitoring of zone transfers.

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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 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. Terminology . . . . .                                          | 4  |
| 3. Use Cases for XFR-over-TLS . . . . .                           | 4  |
| 4. Connection and Data Flows in Existing XFR Mechanisms . . . . . | 5  |
| 4.1. AXFR Mechanism . . . . .                                     | 5  |
| 4.2. IXFR Mechanism . . . . .                                     | 6  |
| 4.3. Data Leakage of NOTIFY and SOA Message Exchanges . . . . .   | 7  |
| 4.3.1. NOTIFY . . . . .                                           | 7  |
| 4.3.2. SOA . . . . .                                              | 8  |
| 5. Connection and Data Flows in XoT . . . . .                     | 8  |
| 5.1. Performance Considerations . . . . .                         | 8  |
| 5.2. AXoT mechanism . . . . .                                     | 8  |
| 5.3. IXoT mechanism . . . . .                                     | 9  |
| 5.3.1. Fallback to AXFR . . . . .                                 | 10 |
| 6. Zone Transfer with DoT - Authentication . . . . .              | 10 |
| 6.1. TSIG . . . . .                                               | 10 |
| 6.2. TLS . . . . .                                                | 10 |
| 6.2.1. Opportunistic . . . . .                                    | 10 |
| 6.2.2. Strict . . . . .                                           | 10 |
| 6.2.3. Mutual . . . . .                                           | 10 |
| 6.3. IP Based ACL on the Primary . . . . .                        | 11 |
| 6.4. ZONEMD . . . . .                                             | 11 |
| 6.5. Comparison of Authentication Methods . . . . .               | 11 |
| 7. Policies for Both AXFR and IXFR . . . . .                      | 12 |
| 8. Multi-primary Configurations . . . . .                         | 13 |
| 9. Implementation Considerations . . . . .                        | 13 |
| 10. Implementation Status . . . . .                               | 14 |
| 11. IANA Considerations . . . . .                                 | 14 |
| 12. Security Considerations . . . . .                             | 14 |
| 13. Acknowledgements . . . . .                                    | 14 |
| 14. Changelog . . . . .                                           | 14 |
| 15. References . . . . .                                          | 15 |
| 15.1. Normative References . . . . .                              | 15 |
| 15.2. Informative References . . . . .                            | 16 |
| 15.3. URIs . . . . .                                              | 17 |
| Authors' Addresses . . . . .                                      | 17 |



## 1. Introduction

DNS has a number of privacy vulnerabilities, as discussed in detail in [I-D.bortzmeyer-dprive-rfc7626-bis]. Stub client to recursive resolver query privacy has received the most attention to date. There are now standards track documents for three encryption capabilities for stub to recursive queries and more work going on to guide deployment of specifically DNS-over-TLS (DoT) [RFC7858] and DNS-over-HTTPS (DoH) [RFC8484].

[I-D.bortzmeyer-dprive-rfc7626-bis] established that stub client DNS query transactions are not public and needed protection, but on zone transfer [RFC1995] [RFC5936] it says only:

"Privacy risks for the holder of a zone (the risk that someone gets the data) are discussed in [RFC5936] and [RFC5155]."

In what way is exposing the full contents of a zone a privacy risk? The contents of the zone could include information such as names of persons used in names of hosts. Best practice is not to use personal information for domain names, but many such domain names exist. There may also be regulatory, policy or other reasons why the zone contents in full must be treated as private.

Neither of the RFCs mentioned in [I-D.bortzmeyer-dprive-rfc7626-bis] contemplates the risk that someone gets the data through eavesdropping on network connections, only via enumeration or unauthorized transfer as described in the following paragraphs.

[RFC5155] specifies NSEC3 to prevent zone enumeration, which is when queries for the authenticated denial of existences records of DNSSEC allow a client to walk through the entire zone. Note that the need for this protection also motivates NSEC5 [I-D.vcelak-nsec5]; zone walking is now possible with NSEC3 due to crypto-breaking advances, and NSEC5 is a response to this problem.

[RFC5155] does not address data obtained outside zone enumeration (nor does [I-D.vcelak-nsec5]). Preventing eavesdropping of zone transfers (this draft) is orthogonal to preventing zone enumeration, though they aim to protect the same information.

[RFC5936] specifies using TSIG [RFC2845] for authorization of the clients of a zone transfer and for data integrity, but does not express any need for confidentiality, and TSIG does not offer encryption. Some operators use SSH tunneling or IPsec to encrypt the transfer data.

Because the AXFR zone transfer is typically carried out-over-TCP from authoritative DNS protocol implementations, encrypting AXFR using DNS-over-TLS [RFC7858] seems like a simple step forward. This document specifies how to use DoT to prevent zone collection from zone transfers, including discussion of approaches for IXFR, which uses UDP or TCP.

## 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 BCP 14 [RFC2119] and [RFC8174] when, and only when, they appear in all capitals, as shown here.

Privacy terminology is as described in Section 3 of [RFC6973].

Note that in this document we choose to use the terms 'primary' and 'secondary' for two servers engaged in zone transfers.

DNS terminology is as described in [RFC8499].

DoT: DNS-over-TLS as specified in [RFC7858]

DoH: DNS-over-HTTPS as specified in [RFC8484]

XoT: Generic XFR-over-TLS mechanisms as specified in this document

AXoT: AXFR-over-TLS

IXoT: IXFR over-TLS

## 3. Use Cases for XFR-over-TLS

- o Confidentiality. Clearly using an encrypted transport for zone transfers will defeat zone content leakage that can occur via passive surveillance.
- o Authentication. Use of single or mutual TLS authentication (in combination with ACLs) can complement and potentially be an alternative to TSIG.
- o Performance. Existing AXFR and IXFR mechanisms have the burden of backwards compatibility with older implementations based on the original specifications in [RFC1034] and [RFC1035]. For example, some older AXFR servers don't support using a TCP connection for multiple AXFR sessions or XFRs of different zones because they have not been updated to follow the guidance in [RFC5836]. Any

implementation of XFR-over-TLS would obviously be required to implement optimized and interoperable transfers as described in [RFC5936] e.g. transfer of multiple zones-over-one connection.

- o Performance. Current usage of TCP for IXFR is sub-optimal in some cases i.e. connections are frequently closed after a single IXFR.

#### 4. Connection and Data Flows in Existing XFR Mechanisms

The original specification for zone transfers in [RFC1034] and [RFC1035] was based on a polling mechanism: a secondary performed a periodic SOA query (based on the refresh timer) to determine if an AXFR was required.

[RFC1995] and [RFC1996] introduced the concepts of IXFR and NOTIFY respectively, to provide for prompt propagation of zone updates. This has largely replaced AXFR where possible, particularly for dynamically updated zones.

[RFC5936] subsequently redefined the specification of AXFR to improve performance and interoperability.

In this document we use the phrase "XFR mechanism" to describe the entire set of message exchanges between a secondary and a primary that concludes in a successful AXFR or IXFR request/response. This set may or may not include

- o NOTIFY messages
- o SOA queries
- o Fallback from IXFR to AXFR
- o Fallback from IXFR-over-UDP to IXFR-over-TCP

The term is used to encompass the range of permutations that are possible and is useful to distinguish the 'XFR mechanism' from a single XFR request/response exchange.

##### 4.1. AXFR Mechanism

The figure below provides an outline of an AXFR mechanism including NOTIFYS.

Figure 1. AXFR Mechanism [1]

1. An AXFR is often (but not always) preceded by a NOTIFY (over UDP) from the primary to the secondary. A secondary may also initiate

an AXFR based on a refresh timer or scheduled/triggered zone maintenance.

2. The secondary will normally (but not always) make a SOA query to the primary to obtain the serial number of the zone held by the primary.
3. If the primary serial is higher than the secondaries serial (using Serial Number Arithmetic [RFC1982]), the secondary makes an AXFR request (over TCP) to the primary after which the AXFR data flows in one or more AXFR responses on the TCP connection.

[RFC5936] specifies that AXFR must use TCP as the transport protocol but details that there is no restriction in the protocol that a single TCP session must be used only for a single AXFR exchange, or even solely for XFRs. For example, it outlines that the SOA query can also happen on this connection. However, this can cause interoperability problems with older implementations that support only the trivial case of one AXFR per connection.

Further details of the limitations in existing AXFR implementations are outlined in [RFC5936].

It is noted that unless the NOTIFY is sent over a trusted communication channel and/or signed by TSIG it can be spoofed causing unnecessary zone transfer attempts.

Similarly unless the SOA query is sent over a trusted communication channel and/or signed by TSIG the response can, in principle, be spoofed causing a secondary to incorrectly believe its version of the zone is update to date. Repeated successful attacks on the SOA could result in a secondary serving stale zone data.

#### 4.2. IXFR Mechanism

The figure below provides an outline of the IXFR mechanism including NOTIFYs.

Figure 1. IXFR Mechanism [2]

1. An IXFR is normally (but not always) preceded by a NOTIFY (over UDP) from the primary to the secondary. A secondary may also initiate an IXFR based on a refresh timer or scheduled/triggered zone maintenance.
2. The secondary will normally (but not always) make a SOA query to the primary to obtain the serial number of the zone held by the primary.

3. If the primary serial is higher than the secondaries serial (using Serial Number Arithmetic [RFC1982]), the secondary makes an IXFR request to the primary after the primary sends an IXFR response.

[RFC1995] specifies that Incremental Transfer may use UDP if the entire IXFR response can be contained in a single DNS packet, otherwise, TCP is used. In fact it says in non-normative language:

"Thus, a client should first make an IXFR query using UDP."

So there may be a forth step above where the client falls back to IXFR-over-TCP. There may also be a forth step where the secondary must fall back to AXFR because e.g. the primary does not support IXFR.

However it is noted that at least two widely used open source authoritative nameserver implementations (BIND [3] and NSD [4]) do IXFR using TCP by default in their latest releases. For BIND TCP connections are sometimes used for SOA queries but in general they are not used persistently and close after an IXFR is completed.

It is noted that the specification for IXFR was published well before TCP was considered a first class transport for DNS. This document therefore updates [RFC1995] to state that DNS implementations that support IXFR-over-TCP MUST use [RFC7766] to optimise the use of TCP connections and SHOULD use [RFC7858] to manage persistent connections.

#### 4.3. Data Leakage of NOTIFY and SOA Message Exchanges

This section attempts to presents a rationale for also encrypting the other messages in the XFR mechanism.

Since the SOA of the published zone can be trivially discovered by simply querying the publicly available authoritative servers leakage RR of this is not discussed in the following sections.

##### 4.3.1. NOTIFY

Unencrypted NOTIFY messages identify configured secondaries on the primary.

[RFC1996] also states:

"If ANCOUNT>0, then the answer section represents an unsecure hint at the new RRset for this .

But since the only supported QTYPE for NOTIFY is SOA, this does not pose a potential leak.

#### 4.3.2. SOA

For hidden primaries or secondaries the SOA response leaks the degree of lag of any downstream secondary.

### 5. Connection and Data Flows in XoT

#### 5.1. Performance Considerations

The details in [RFC7766], [RFC7858] and [RFC8310] about e.g. using persistent connections and TLS Session Resumption [RFC5077] are fully applicable to XFR-over-TLS as well.

It is RECOMMENDED that clients and servers that support XoT also implement EDNS0 Keepalive [RFC7828].

#### 5.2. AXoT mechanism

The figure below provides an outline of the AXoT mechanism including NOTIFYs.

Figure 3: AXoT mechanism [5]

All implementations that support XoT MUST fully implement [RFC5953] behavior on TLS connections.

Sections 4.1, 4.1.1 and 4.1.2 of [RFC5936] describe guidance for AXFR clients and servers with regard to re-use of sessions for multiple AXFRs, AXFRs of different zones and using TCP session for other queries including SOA.

For clarity we restate here that an AXoT client MAY use an already opened TLS connection to send a AXFR request. Using an existing open connection is RECOMMENDED over opening a new connection. (Non-AXoT session traffic can also use an open connection.)

For clarity we additionally state here that an AXoT client MAY use an already opened TLS connection to send a SOA request. Using an existing open connection is RECOMMENDED over opening a new connection.

The connection for AXFR-over-TLS SHOULD be established using port 853, as specified in [RFC7858], unless there is mutual agreement between the secondary and primary to use a port other than port 853 for XFR-over-TLS.

QUESTION: Should there be a requirement that the SOA is always done on a TLS connection if the XFR is? For the case when no transfer is required this could be unnecessary overhead.

### 5.3. IXoT mechanism

The figure below provides an outline of the IXoT mechanism including NOTIFYs.

Figure 4: IXoT mechanism [6]

The connection for IXFR-over-TLS SHOULD be established using port 853, as specified in [RFC7858], unless there is mutual agreement between the secondary and primary to use a port other than port 853 for XFR-over-TLS.

[RFC1995] says nothing with respect to optimizing IXFRs over TCP or re-using already open TCP connections to perform IXFRs or other queries. We provide guidance here that aligns with the guidance in [RFC5936] for AXFR and with that for performant TCP/TLS usage in [RFC7766] and [RFC7858].

An IXoT client MAY use an already opened TLS connection to send a IXFR request. Using an existing open connection is RECOMMENDED over opening a new connection. (Non-IXoT session traffic can also use an open connection.)

An IXoT client MAY use an already open TLS connection to send an SOA query. Using an existing open connection is RECOMMENDED over opening a new connection.

An IXoT server MUST be able to handle multiple IXoT requests on a single TLS connection, as well as to handle other query/response transactions over it.

An IXoT client MAY keep an existing TLS session open in the expectation it is likely to need to perform an IXFR in the near future. The client may use the frequency of recent IXFRs to calculate an average update rate and then use EDNS0 Keepalive to request an appropriate timeout from the server (if the server supports EDNS0 Keepalive). If the server does not support EDNS0 Keepalive the client MAY keep the connection open for a few seconds ([RFC7766] recommends that servers use timeouts of at least a few seconds).

An IXoT client MAY pipeline IXFR requests for different zones on a single TLS connection. AN IXoT server MAY respond to those requests out of order.

### 5.3.1. Fallback to AXFR

Fallback to AXFR can happen, for example, if the server is not able to provide an IXFR for the requested SOA. Implementations differ in how long they store zone deltas and how many may be stored at any one time.

After a failed IXFR a IXoT client SHOULD request the AXFR on the already open TLS connection.

## 6. Zone Transfer with DoT - Authentication

### 6.1. TSIG

TSIG [RFC2845] provides a mechanism for two parties to exchange secret keys which can then be used to create a message digest to protect individual DNS messages. This allows each party to authenticate that a request or response (and the data in it) came from the other party, even if it was transmitted-over-an unsecured channel or via a proxy. It provides party-to-party data authentication, but not hop-to-hop channel authentication or confidentiality.

### 6.2. TLS

#### 6.2.1. Opportunistic

Opportunistic TLS [RFC8310] provides a defence against passive surveillance, providing on-the-wire confidentiality.

#### 6.2.2. Strict

Strict TLS [RFC8310] requires that a client is configured with an authentication domain name (and/or SPKI pinset) that should be used to authenticate the TLS handshake with the server. This additionally provides a defense for the client against active surveillance, providing client-to-server authentication and end-to-end channel confidentiality.

#### 6.2.3. Mutual

This is an extension to Strict TLS [RFC8310] which requires that a client is configured with an authentication domain name (and/or SPKI pinset) and a client certificate. The client offers the certificate for authentication by the server and the client can authentic the server the same way as in Strict TLS. This provides a defense for both parties against active surveillance, providing bi-directional authentication and end-to-end channel confidentiality.



### 6.3. IP Based ACL on the Primary

Most DNS server implementations offer an option to configure an IP based Access Control List (ACL), which is often used in combination with TSIG based ACLs to restrict access to zone transfers on primary servers.

This is also possible with XoT but it must be noted that as with TCP the implementation of such an ACL cannot be enforced on the primary until a XFR request is received on an established connection.

If control were to be any more fine-grained than this then a separate port would be required for XoT such that implementations would be able to refuse connections on that port to all clients except those configured as secondaries.

### 6.4. ZONEMD

Message Digest for DNS Zones (ZONEMD)

[I-D.ietf-dnsop-dns-zone-digest] digest is a mechanism that can be used to verify the content of a standalone zone. It is designed to be independent of the transmission channel or mechanism, allowing a general consumer of a zone to do origin authentication of the entire zone contents. It is not considered suitable for highly dynamic zones. It is complementary to the above mechanisms and can be used in conjunction with XFR-over-TLS but is not considered further.

### 6.5. Comparison of Authentication Methods

The Table below compares the properties of each of the above methods in terms of what protection they provide to the secondary and primary servers during XoT in terms of:

- o 'Data Auth': Authentication that the DNS message data is signed by the party with whom credentials were shared (the signing party may or may not be party operating the far end of a TCP/TLS connection in a 'proxy' scenario). For the primary the TSIG on the XFR request confirms that the requesting party is authorized to request zone data, for the secondary it authenticates the zone data that is received.
- o 'Channel Conf': Confidentiality of the communication channel between the client and server (i.e. the two end points of a TCP/TLS connection).
- o Channel Auth: Authentication of the identity of party to whom a TCP/TLS connection is made (this might not be a direct connection between the primary and secondary in a proxy scenario).

It is noted that zone transfer scenarios can vary from a simple single primary/secondary relationship where both servers are under the control of a single operator to a complex hierarchical structure which includes proxies and multiple operators. Each deployment scenario will require specific analysis to determine which authentication methods are best suited to the deployment model in question.

Table 1: Properties of Authentication methods for XoT [7]

Based on this analysis it can be seen that:

- o A combination of Opportunistic TLS and TSIG provides both data authentication and channel confidentiality for both parties. However this does not stop a MitM attack on the channel which could be used to gather zone data.
- o Using just mutual TLS can be considered a standalone solution if the secondary has reason to place equivalent trust in channel authentication as data authentication e.g. the same operator runs both the primary and secondary.
- o Using TSIG, Strict TLS and an ACL on the primary provides all 3 properties for both parties with probably the lowest operational overhead.

## 7. Policies for Both AXFR and IXFR

We call the entire group of servers involved in XFR (all the primaries and all the secondaries) the 'transfer group'.

Within any transfer group both AXFRs and IXFRs for a zone SHOULD all use the same policy e.g. if AXFRs use AXoT all IXFRs SHOULD use IXoT.

In order to assure the confidentiality of the zone information, the entire transfer group MUST have a consistent policy of requiring confidentiality. If any do not, this is a weak link for attackers to exploit.

A XoT policy should specify

- o If TSIG is required
- o What kind of TLS is required (Opportunistic, Strict or mTLS)
- o If IP based ACLs should also be used.

Since this may require configuration of a number of servers who may be under the control of different operators the desired consistency could be hard to enforce and audit in practice.

Certain aspects of the Policies can be relatively easily tested independently e.g. by requesting zone transfers without TSIG, from unauthorized IP addresses or over cleartext DNS. Other aspects such as if a secondary will accept data without a TSIG digest or if secondaries are using Strict as opposed to Opportunistic TLS are more challenging.

NOTE: The authors request feedback on this challenge and welcome suggestions of how to practically manage this.

## 8. Multi-primary Configurations

Also known as multi-master configurations this model can provide flexibility and redundancy particularly for IXFR. A secondary will receive one or more NOTIFY messages and can send an SOA to all of the configured primaries. It can then choose to send an IXFR request to the primary with the highest SOA (or other criteria e.g. RTT).

When using persistent connections the secondary may have a TLS connection already open to one or more primaries. Should a secondary preferentially request an IXFR from a primary to which it already has an open TLS connection or the one with the highest SOA (assuming it doesn't have a connection open to it already)?

Two extremes can be envisaged here. In the first case the secondary continues to use one persistent connection to a single primary until it has reason not to. Reasons not to might include the primary repeatedly closing the connection, long RTTs on transfers or the SOA of the primary being an unacceptable lag behind the SOA of an alternative primary.

At the other extreme a primary could keep multiple persistent connections open to all available primaries and only request IXFRs from the primary with the highest serial number. Since normally the number of secondaries and primaries in direct contact in a transfer group is reasonably low this might be feasible if latency is the most significant concern.

## 9. Implementation Considerations

TBD

## 10. Implementation Status

The 1.9.2 version of Unbound [8] includes an option to perform AXFR-over-TLS (instead of TCP). This requires the client (secondary) to authenticate the server (primary) using a configured authentication domain name.

It is noted that use of a TLS proxy in front of the primary server is a simple deployment solution that can enable server side XoT.

## 11. IANA Considerations

TBD

## 12. Security Considerations

This document specifies a security measure against a DNS risk: the risk that an attacker collects entire DNS zones through eavesdropping on clear text DNS zone transfers. It presents a new Security Consideration for DNS. Some questions to discuss are:

- o Should DoT in this new case be required to use only TLS 1.3 and higher to avoid residual exposure?
- o How should padding be used in IXFR?
- o Should there be an option to 'pad' an AXFR response (i.e. a set of AXFR responses on a given connection) to hide the zone size?

## 13. Acknowledgements

The authors thank Benno Overeinder, Shumon Huque and Tim Wicinski for review and discussions.

## 14. Changelog

draft-hzpa-dprive-xfr-over-tls-01

- o Substantial re-work of the document.

draft-hzpa-dprive-xfr-over-tls-01

- o Editorial changes, updates to references.

draft-hzpa-dprive-xfr-over-tls-00

- o Initial commit

## 15. References

## 15.1. Normative References

- [I-D.bortzmeyer-dprive-rfc7626-bis]  
Bortzmeyer, S. and S. Dickinson, "DNS Privacy Considerations", draft-bortzmeyer-dprive-rfc7626-bis-02 (work in progress), January 2019.
- [I-D.vcelak-nsec5]  
Vcelak, J., Goldberg, S., Papadopoulos, D., Huque, S., and D. Lawrence, "NSEC5, DNSSEC Authenticated Denial of Existence", draft-vcelak-nsec5-08 (work in progress), December 2018.
- [RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, DOI 10.17487/RFC1995, August 1996, <<https://www.rfc-editor.org/info/rfc1995>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://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, <<https://www.rfc-editor.org/info/rfc2845>>.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, DOI 10.17487/RFC5077, January 2008, <<https://www.rfc-editor.org/info/rfc5077>>.
- [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, DOI 10.17487/RFC5155, March 2008, <<https://www.rfc-editor.org/info/rfc5155>>.
- [RFC5936] Lewis, E. and A. Hoenes, Ed., "DNS Zone Transfer Protocol (AXFR)", RFC 5936, DOI 10.17487/RFC5936, June 2010, <<https://www.rfc-editor.org/info/rfc5936>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.

- [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, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <<https://www.rfc-editor.org/info/rfc8310>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.

## 15.2. Informative References

- [I-D.ietf-dnsop-dns-zone-digest]  
Wessels, D., Barber, P., Weinberg, M., Kumari, W., and W. Hardaker, "Message Digest for DNS Zones", draft-ietf-dnsop-dns-zone-digest-00 (work in progress), June 2019.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC1982] Elz, R. and R. Bush, "Serial Number Arithmetic", RFC 1982, DOI 10.17487/RFC1982, August 1996, <<https://www.rfc-editor.org/info/rfc1982>>.
- [RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, DOI 10.17487/RFC1996, August 1996, <<https://www.rfc-editor.org/info/rfc1996>>.

- [RFC5953] Hardaker, W., "Transport Layer Security (TLS) Transport Model for the Simple Network Management Protocol (SNMP)", RFC 5953, DOI 10.17487/RFC5953, August 2010, <<https://www.rfc-editor.org/info/rfc5953>>.
- [RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and D. Wessels, "DNS Transport over TCP - Implementation Requirements", RFC 7766, DOI 10.17487/RFC7766, March 2016, <<https://www.rfc-editor.org/info/rfc7766>>.

### 15.3. URIs

- [1] [https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02\\_updates/02-draft-svg/AXFR\\_mechanism.svg](https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02_updates/02-draft-svg/AXFR_mechanism.svg)
- [2] [https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02\\_updates/02-draft-svg/IXFR%20mechanism.svg](https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02_updates/02-draft-svg/IXFR%20mechanism.svg)
- [3] <https://www.isc.org/bind/>
- [4] <https://www.nlnetlabs.nl/projects/nsd/about/>
- [5] [https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02\\_updates/02-draft-svg/AXoT\\_mechanism\\_1.svg](https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02_updates/02-draft-svg/AXoT_mechanism_1.svg)
- [6] [https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02\\_updates/02-draft-svg/IXoT\\_mechanism\\_1.svg](https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02_updates/02-draft-svg/IXoT_mechanism_1.svg)
- [7] [https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02\\_updates/02-draft-svg/Properties\\_of\\_Authentication\\_methods\\_for\\_XoT.svg](https://github.com/hanzhang0116/hzpa-dprive-xfr-over-tls/blob/02_updates/02-draft-svg/Properties_of_Authentication_methods_for_XoT.svg)
- [8] <https://github.com/NLnetLabs/unbound/blob/release-1.9.2/doc/Changelog>

### Authors' Addresses

Han Zhang  
Salesforce  
San Francisco, CA  
United States

Email: [hzhang@salesforce.com](mailto:hzhang@salesforce.com)

Pallavi Aras  
Salesforce  
Herndon, VA  
United States

Email: [paras@salesforce.com](mailto:paras@salesforce.com)

Willem Toorop  
NLnet Labs  
Science Park 400  
Amsterdam 1098 XH  
The Netherlands

Email: [willem@nlnetlabs.nl](mailto:willem@nlnetlabs.nl)

Sara Dickinson  
Sinodun IT  
Magdalen Centre  
Oxford Science Park  
Oxford OX4 4GA  
United Kingdom

Email: [sara@sinodun.com](mailto:sara@sinodun.com)

Allison Mankin  
Salesforce  
Herndon, VA  
United States

Email: [allison.mankin@gmail.com](mailto:allison.mankin@gmail.com)



dprive  
Internet-Draft  
Intended status: Best Current Practice  
Expires: January 14, 2021

S. Dickinson  
Sinodun IT  
B. Overeinder  
R. van Rijswijk-Deij  
NLnet Labs  
A. Mankin  
Salesforce  
July 13, 2020

Recommendations for DNS Privacy Service Operators  
draft-ietf-dprive-bcp-op-14

Abstract

This document presents operational, policy, and security considerations for DNS recursive resolver operators who choose to offer DNS Privacy services. With these recommendations, the operator can make deliberate decisions regarding which services to provide, and how the decisions and alternatives impact the privacy of users.

This document also presents a non-normative framework to assist writers of a Recursive operator Privacy Statement (analogous to DNS Security Extensions (DNSSEC) Policies and DNSSEC Practice Statements described in RFC6841).

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 January 14, 2021.

Copyright Notice

Copyright (c) 2020 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. Scope . . . . .                                                                           | 5  |
| 3. Privacy-related documents . . . . .                                                       | 5  |
| 4. Terminology . . . . .                                                                     | 6  |
| 5. Recommendations for DNS privacy services . . . . .                                        | 6  |
| 5.1. On the wire between client and server . . . . .                                         | 7  |
| 5.1.1. Transport recommendations . . . . .                                                   | 7  |
| 5.1.2. Authentication of DNS privacy services . . . . .                                      | 8  |
| 5.1.3. Protocol recommendations . . . . .                                                    | 9  |
| 5.1.4. DNSSEC . . . . .                                                                      | 11 |
| 5.1.5. Availability . . . . .                                                                | 12 |
| 5.1.6. Service options . . . . .                                                             | 12 |
| 5.1.7. Impact of Encryption on Monitoring by DNS Privacy<br>Service Operators . . . . .      | 13 |
| 5.1.8. Limitations of fronting a DNS privacy service with a<br>pure TLS proxy . . . . .      | 13 |
| 5.2. Data at rest on the server . . . . .                                                    | 14 |
| 5.2.1. Data handling . . . . .                                                               | 14 |
| 5.2.2. Data minimization of network traffic . . . . .                                        | 15 |
| 5.2.3. IP address pseudonymization and anonymization methods . . . . .                       | 16 |
| 5.2.4. Pseudonymization, anonymization, or discarding of<br>other correlation data . . . . . | 16 |
| 5.2.5. Cache snooping . . . . .                                                              | 17 |
| 5.3. Data sent onwards from the server . . . . .                                             | 17 |
| 5.3.1. Protocol recommendations . . . . .                                                    | 17 |
| 5.3.2. Client query obfuscation . . . . .                                                    | 18 |
| 5.3.3. Data sharing . . . . .                                                                | 19 |
| 6. Recursive operator Privacy Statement (RPS) . . . . .                                      | 20 |
| 6.1. Outline of an RPS . . . . .                                                             | 20 |
| 6.1.1. Policy . . . . .                                                                      | 20 |
| 6.1.2. Practice . . . . .                                                                    | 21 |
| 6.2. Enforcement/accountability . . . . .                                                    | 22 |
| 7. IANA considerations . . . . .                                                             | 23 |
| 8. Security considerations . . . . .                                                         | 23 |
| 9. Acknowledgements . . . . .                                                                | 23 |
| 10. Contributors . . . . .                                                                   | 23 |

|                                                                   |    |
|-------------------------------------------------------------------|----|
| 11. Changelog . . . . .                                           | 24 |
| 12. References . . . . .                                          | 27 |
| 12.1. Normative References . . . . .                              | 27 |
| 12.2. Informative References . . . . .                            | 29 |
| Appendix A. Documents . . . . .                                   | 34 |
| A.1. Potential increases in DNS privacy . . . . .                 | 34 |
| A.2. Potential decreases in DNS privacy . . . . .                 | 34 |
| A.3. Related operational documents . . . . .                      | 35 |
| Appendix B. IP address techniques . . . . .                       | 35 |
| B.1. Categorization of techniques . . . . .                       | 36 |
| B.2. Specific techniques . . . . .                                | 37 |
| B.2.1. Google Analytics non-prefix filtering . . . . .            | 37 |
| B.2.2. dnswasher . . . . .                                        | 38 |
| B.2.3. Prefix-preserving map . . . . .                            | 38 |
| B.2.4. Cryptographic Prefix-Preserving Pseudonymization . . . . . | 38 |
| B.2.5. Top-hash Subtree-replicated Anonymization . . . . .        | 39 |
| B.2.6. ipcipher . . . . .                                         | 39 |
| B.2.7. Bloom filters . . . . .                                    | 39 |
| Appendix C. Current policy and privacy statements . . . . .       | 40 |
| Appendix D. Example RPS . . . . .                                 | 40 |
| D.1. Policy . . . . .                                             | 40 |
| D.2. Practice . . . . .                                           | 43 |
| Authors' Addresses . . . . .                                      | 44 |

## 1. Introduction

The Domain Name System (DNS) is at the core of the Internet; almost every activity on the Internet starts with a DNS query (and often several). However the DNS was not originally designed with strong security or privacy mechanisms. A number of developments have taken place in recent years which aim to increase the privacy of the DNS system and these are now seeing some deployment. This latest evolution of the DNS presents new challenges to operators and this document attempts to provide an overview of considerations for privacy focused DNS services.

In recent years there has also been an increase in the availability of "public resolvers" [RFC8499] which users may prefer to use instead of the default network resolver either because they offer a specific feature (e.g., good reachability or encrypted transport) or because the network resolver lacks a specific feature (e.g., strong privacy policy or unfiltered responses). These public resolvers have tended to be at the forefront of adoption of privacy-related enhancements but it is anticipated that operators of other resolver services will follow.

Whilst protocols that encrypt DNS messages on the wire provide protection against certain attacks, the resolver operator still has

(in principle) full visibility of the query data and transport identifiers for each user. Therefore, a trust relationship (whether explicit or implicit) is assumed to exist between each user and the operator of the resolver(s) used by that user. The ability of the operator to provide a transparent, well documented, and secure privacy service will likely serve as a major differentiating factor for privacy conscious users if they make an active selection of which resolver to use.

It should also be noted that the choice of a user to configure a single resolver (or a fixed set of resolvers) and an encrypted transport to use in all network environments has both advantages and disadvantages. For example, the user has a clear expectation of which resolvers have visibility of their query data. However, this resolver/transport selection may provide an added mechanism to track them as they move across network environments. Commitments from resolver operators to minimize such tracking as users move between networks are also likely to play a role in user selection of resolvers.

More recently the global legislative landscape with regard to personal data collection, retention, and pseudonymization has seen significant activity. Providing detailed practice advice about these areas to the operator is out of scope, but Section 5.3.3 describes some mitigations of data sharing risk.

This document has two main goals:

- o To provide operational and policy guidance related to DNS over encrypted transports and to outline recommendations for data handling for operators of DNS privacy services.
- o To introduce the Recursive operator Privacy Statement (RPS) and present a framework to assist writers of an RPS. An RPS is a document that an operator should publish which outlines their operational practices and commitments with regard to privacy, thereby providing a means for clients to evaluate both the measurable and claimed privacy properties of a given DNS privacy service. The framework identifies a set of elements and specifies an outline order for them. This document does not, however, define a particular privacy statement, nor does it seek to provide legal advice as to the contents.

A desired operational impact is that all operators (both those providing resolvers within networks and those operating large public services) can demonstrate their commitment to user privacy thereby driving all DNS resolution services to a more equitable footing. Choices for users would (in this ideal world) be driven by other

factors, e.g., differing security policies or minor difference in operator policy, rather than gross disparities in privacy concerns.

Community insight [or judgment?] about operational practices can change quickly, and experience shows that a Best Current Practice (BCP) document about privacy and security is a point-in-time statement. Readers are advised to seek out any updates that apply to this document.

## 2. Scope

"DNS Privacy Considerations" [RFC7626] describes the general privacy issues and threats associated with the use of the DNS by Internet users and much of the threat analysis here is lifted from that document and from [RFC6973]. However this document is limited in scope to best practice considerations for the provision of DNS privacy services by servers (recursive resolvers) to clients (stub resolvers or forwarders). Choices that are made exclusively by the end user, or those for operators of authoritative nameservers are out of scope.

This document includes (but is not limited to) considerations in the following areas:

1. Data "on the wire" between a client and a server.
2. Data "at rest" on a server (e.g., in logs).
3. Data "sent onwards" from the server (either on the wire or shared with a third party).

Whilst the issues raised here are targeted at those operators who choose to offer a DNS privacy service, considerations for areas 2 and 3 could equally apply to operators who only offer DNS over unencrypted transports but who would otherwise like to align with privacy best practice.

## 3. Privacy-related documents

There are various documents that describe protocol changes that have the potential to either increase or decrease the privacy properties of the DNS in various ways. Note this does not imply that some documents are good or bad, better or worse, just that (for example) some features may bring functional benefits at the price of a reduction in privacy and conversely some features increase privacy with an accompanying increase in complexity. A selection of the most relevant documents are listed in Appendix A for reference.

#### 4. 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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

DNS terminology is as described in [RFC8499] with one modification: we restate the clause in the original definition of Privacy-enabling DNS server in [RFC8310] to include the requirement that a DNS over (D)TLS server should also offer at least one of the credentials described in Section 8 of [RFC8310] and implement the (D)TLS profile described in Section 9 of [RFC8310].

Other Terms:

- o RPS: Recursive operator Privacy Statement, see Section 6.
- o DNS privacy service: The service that is offered via a privacy-enabling DNS server and is documented either in an informal statement of policy and practice with regard to users privacy or a formal RPS.

#### 5. Recommendations for DNS privacy services

In the following sections we first outline the threats relevant to the specific topic and then discuss the potential actions that can be taken to mitigate them.

We describe two classes of threats:

- o Threats described in [RFC6973] 'Privacy Considerations for Internet Protocols'
  - \* Privacy terminology, threats to privacy, and mitigations as described in Sections 3, 5, and 6 of [RFC6973].
- o DNS Privacy Threats
  - \* These are threats to the users and operators of DNS privacy services that are not directly covered by [RFC6973]. These may be more operational in nature such as certificate management or service availability issues.

We describe three classes of actions that operators of DNS privacy services can take:

- o Threat mitigation for well understood and documented privacy threats to the users of the service and in some cases to the operators of the service.
- o Optimization of privacy services from an operational or management perspective.
- o Additional options that could further enhance the privacy and usability of the service.

This document does not specify policy - only best practice, however for DNS Privacy services to be considered compliant with these best practice guidelines they SHOULD implement (where appropriate) all:

- o Threat mitigations to be minimally compliant.
- o Optimizations to be moderately compliant.
- o Additional options to be maximally compliant.

The rest of this document does not use normative language but instead refers only to the three differing classes of action which correspond to the three named levels of compliance stated above. However, compliance (to the indicated level) remains a normative requirement.

#### 5.1. On the wire between client and server

In this section we consider both data on the wire and the service provided to the client.

##### 5.1.1. Transport recommendations

[RFC6973] Threats:

- o Surveillance:

- \* Passive surveillance of traffic on the wire

DNS Privacy Threats:

- o Active injection of spurious data or traffic.

Mitigations:

A DNS privacy service can mitigate these threats by providing service over one or more of the following transports

- o DNS over TLS (DoT) [RFC7858] and [RFC8310].

- o DNS over HTTPS (DoH) [RFC8484].

It is noted that a DNS privacy service can also be provided over DNS over DTLS [RFC8094], however this is an Experimental specification and there are no known implementations at the time of writing.

It is also noted that DNS privacy service might be provided over IPSec, DNSCrypt, or VPNs. However, there are no specific RFCs that cover the use of these transports for DNS and any discussion of best practice for providing such a service is out of scope for this document.

Whilst encryption of DNS traffic can protect against active injection on the paths traversed by the encrypted connection this does not diminish the need for DNSSEC, see Section 5.1.4.

#### 5.1.2. Authentication of DNS privacy services

[RFC6973] Threats:

- o Surveillance:

- \* Active attacks on client resolver configuration

Mitigations:

DNS privacy services should ensure clients can authenticate the server. Note that this, in effect, commits the DNS privacy service to a public identity users will trust.

When using DoT, clients that select a 'Strict Privacy' usage profile [RFC8310] (to mitigate the threat of active attack on the client) require the ability to authenticate the DNS server. To enable this, DNS privacy services that offer DNS over TLS need to provide credentials that will be accepted by the client's trust model, in the form of either X.509 certificates [RFC5280] or Subject Public Key Info (SPKI) pin sets [RFC8310].

When offering DoH [RFC8484], HTTPS requires authentication of the server as part of the protocol.

Server operators should also follow the best practices with regard to certificate revocation as described in [RFC7525].



#### 5.1.2.1. Certificate management

Anecdotal evidence to date highlights the management of certificates as one of the more challenging aspects for operators of traditional DNS resolvers that choose to additionally provide a DNS privacy service as management of such credentials is new to those DNS operators.

It is noted that SPKI pin set management is described in [RFC7858] but that key pinning mechanisms in general have fallen out of favor operationally for various reasons such as the logistical overhead of rolling keys.

##### DNS Privacy Threats:

- o Invalid certificates, resulting in an unavailable service which might force a user to fallback to cleartext.
- o Mis-identification of a server by a client e.g., typos in DoH URL templates [RFC8484] or authentication domain names [RFC8310] which accidentally direct clients to attacker controlled servers.

##### Mitigations:

It is recommended that operators:

- o Follow the guidance in Section 6.5 of [RFC7525] with regards to certificate revocation.
- o Automate the generation, publication, and renewal of certificates. For example, ACME [RFC8555] provides a mechanism to actively manage certificates through automation and has been implemented by a number of certificate authorities.
- o Monitor certificates to prevent accidental expiration of certificates.
- o Choose a short, memorable authentication domain name for the service.

#### 5.1.3. Protocol recommendations

##### 5.1.3.1. DoT

##### DNS Privacy Threats:

- o Known attacks on TLS such as those described in [RFC7457].

- o Traffic analysis, for example: [Pitfalls-of-DNS-Encryption].
- o Potential for client tracking via transport identifiers.
- o Blocking of well known ports (e.g., 853 for DoT).

#### Mitigations:

In the case of DoT, TLS profiles from Section 9 of [RFC8310] and the Countermeasures to DNS Traffic Analysis from section 11.1 of [RFC8310] provide strong mitigations. This includes but is not limited to:

- o Adhering to [RFC7525].
- o Implementing only (D)TLS 1.2 or later as specified in [RFC8310].
- o Implementing EDNS(0) Padding [RFC7830] using the guidelines in [RFC8467] or a successor specification.
- o Servers should not degrade in any way the query service level provided to clients that do not use any form of session resumption mechanism, such as TLS session resumption [RFC5077] with TLS 1.2, section 2.2 of [RFC8446], or Domain Name System (DNS) Cookies [RFC7873].
- o A DoT privacy service on both port 853 and 443. If the operator deploys DoH on the same IP address this requires the use of the 'dot' ALPN value [dot-ALPN].

#### Optimizations:

- o Concurrent processing of pipelined queries, returning responses as soon as available, potentially out of order as specified in [RFC7766]. This is often called 'OOOR' - out-of-order responses (providing processing performance similar to HTTP multiplexing).
- o Management of TLS connections to optimize performance for clients using [RFC7766] and EDNS(0) Keepalive [RFC7828]

#### Additional Options:

Management of TLS connections to optimize performance for clients using DNS Stateful Operations [RFC8490].

#### 5.1.3.2. DoH

##### DNS Privacy Threats:

- o Known attacks on TLS such as those described in [RFC7457].
- o Traffic analysis, for example: [DNS-Privacy-not-so-private].
- o Potential for client tracking via transport identifiers.

##### Mitigations:

- o Clients must be able to forgo the use of HTTP Cookies [RFC6265] and still use the service.
- o Use of HTTP/2 padding and/or EDNS(0) padding as described in Section 9 of [RFC8484]
- o Clients should not be required to include any headers beyond the absolute minimum to obtain service from a DoH server. (See Section 6.1 of [I-D.ietf-httpbis-bcp56bis].)

#### 5.1.4. DNSSEC

##### DNS Privacy Threats:

- o Users may be directed to bogus IP addresses which, depending on the application, protocol and authentication method, might lead users to reveal personal information to attackers. One example is a website that doesn't use TLS or its TLS authentication can somehow be subverted.

##### Mitigations:

- o All DNS privacy services must offer a DNS privacy service that performs Domain Name System Security Extensions (DNSSEC) validation. In addition they must be able to provide the DNSSEC RRs to the client so that it can perform its own validation.

The addition of encryption to DNS does not remove the need for DNSSEC [RFC4033] - they are independent and fully compatible protocols, each solving different problems. The use of one does not diminish the need nor the usefulness of the other.

While the use of an authenticated and encrypted transport protects origin authentication and data integrity between a client and a DNS privacy service it provides no proof (for a non-validating client) that the data provided by the DNS privacy service was actually DNSSEC

authenticated. As with cleartext DNS the user is still solely trusting the AD bit (if present) set by the resolver.

It should also be noted that the use of an encrypted transport for DNS actually solves many of the practical issues encountered by DNS validating clients e.g. interference by middleboxes with cleartext DNS payloads is completely avoided. In this sense a validating client that uses a DNS privacy service which supports DNSSEC has a far simpler task in terms of DNSSEC Roadblock avoidance [RFC8027].

#### 5.1.5. Availability

##### DNS Privacy Threats:

- o A failed DNS privacy service could force the user to switch providers, fallback to cleartext or accept no DNS service for the outage.

##### Mitigations:

A DNS privacy service should strive to engineer encrypted services to the same availability level as any unencrypted services they provide. Particular care should be taken to protect DNS privacy services against denial-of-service attacks, as experience has shown that unavailability of DNS resolving because of attacks is a significant motivation for users to switch services. See, for example Section IV-C of [Passive-Observations-of-a-Large-DNS].

Techniques such as those described in Section 10 of [RFC7766] can be of use to operators to defend against such attacks.

#### 5.1.6. Service options

##### DNS Privacy Threats:

- o Unfairly disadvantaging users of the privacy service with respect to the services available. This could force the user to switch providers, fallback to cleartext or accept no DNS service for the outage.

##### Mitigations:

A DNS privacy service should deliver the same level of service as offered on un-encrypted channels in terms of options such as filtering (or lack thereof), DNSSEC validation, etc.

#### 5.1.7. Impact of Encryption on Monitoring by DNS Privacy Service Operators

##### DNS Privacy Threats:

- o Increased use of encryption can impact DNS privacy service operator ability to monitor traffic and therefore manage their DNS servers [RFC8404].

Many monitoring solutions for DNS traffic rely on the plain text nature of this traffic and work by intercepting traffic on the wire, either using a separate view on the connection between clients and the resolver, or as a separate process on the resolver system that inspects network traffic. Such solutions will no longer function when traffic between clients and resolvers is encrypted. Many DNS privacy service operators still have need to inspect DNS traffic, e.g., to monitor for network security threats. Operators may therefore need to invest in alternative means of monitoring that relies on either the resolver software directly, or exporting DNS traffic from the resolver using e.g., [dnstap].

##### Optimization:

When implementing alternative means for traffic monitoring, operators of a DNS privacy service should consider using privacy conscious means to do so (see section Section 5.2 for more details on data handling and also the discussion on the use of Bloom Filters in Appendix B.

#### 5.1.8. Limitations of fronting a DNS privacy service with a pure TLS proxy

##### DNS Privacy Threats:

- o Limited ability to manage or monitor incoming connections using DNS specific techniques.
- o Misconfiguration (e.g., of the target server address in the proxy configuration) could lead to data leakage if the proxy to target server path is not encrypted.

##### Optimization:

Some operators may choose to implement DoT using a TLS proxy (e.g. [nginx], [haproxy], or [stunnel]) in front of a DNS nameserver because of proven robustness and capacity when handling large numbers of client connections, load balancing capabilities and good tooling. Currently, however, because such proxies typically have no specific

handling of DNS as a protocol over TLS or DTLS using them can restrict traffic management at the proxy layer and at the DNS server. For example, all traffic received by a nameserver behind such a proxy will appear to originate from the proxy and DNS techniques such as ACLs, RRL, or DNS64 will be hard or impossible to implement in the nameserver.

Operators may choose to use a DNS aware proxy such as [dnsdist] which offers custom options (similar to that proposed in [I-D.bellis-dnsop-xpf]) to add source information to packets to address this shortcoming. It should be noted that such options potentially significantly increase the leaked information in the event of a misconfiguration.

## 5.2. Data at rest on the server

### 5.2.1. Data handling

[RFC6973] Threats:

- o Surveillance.
- o Stored data compromise.
- o Correlation.
- o Identification.
- o Secondary use.
- o Disclosure.

Other Threats

- o Contravention of legal requirements not to process user data.

Mitigations:

The following are recommendations relating to common activities for DNS service operators and in all cases data retention should be minimized or completely avoided if possible for DNS privacy services. If data is retained it should be encrypted and either aggregated, pseudonymized, or anonymized whenever possible. In general the principle of data minimization described in [RFC6973] should be applied.

- o Transient data (e.g., that is used for real time monitoring and threat analysis which might be held only in memory) should be

retained for the shortest possible period deemed operationally feasible.

- o The retention period of DNS traffic logs should be only those required to sustain operation of the service and, to the extent that such exists, meet regulatory requirements.
- o DNS privacy services should not track users except for the particular purpose of detecting and remedying technically malicious (e.g., DoS) or anomalous use of the service.
- o Data access should be minimized to only those personnel who require access to perform operational duties. It should also be limited to anonymized or pseudonymized data where operationally feasible, with access to full logs (if any are held) only permitted when necessary.

Optimizations:

- o Consider use of full disk encryption for logs and data capture storage.

#### 5.2.2. Data minimization of network traffic

Data minimization refers to collecting, using, disclosing, and storing the minimal data necessary to perform a task, and this can be achieved by removing or obfuscating privacy-sensitive information in network traffic logs. This is typically personal data, or data that can be used to link a record to an individual, but may also include revealing other confidential information, for example on the structure of an internal corporate network.

The problem of effectively ensuring that DNS traffic logs contain no or minimal privacy-sensitive information is not one that currently has a generally agreed solution or any standards to inform this discussion. This section presents an overview of current techniques to simply provide reference on the current status of this work.

Research into data minimization techniques (and particularly IP address pseudonymization/anonymization) was sparked in the late 1990s/early 2000s, partly driven by the desire to share significant corpuses of traffic captures for research purposes. Several techniques reflecting different requirements in this area and different performance/resource tradeoffs emerged over the course of the decade. Developments over the last decade have been both a blessing and a curse; the large increase in size between an IPv4 and an IPv6 address, for example, renders some techniques impractical, but also makes available a much larger amount of input entropy, the

better to resist brute force re-identification attacks that have grown in practicality over the period.

Techniques employed may be broadly categorized as either anonymization or pseudonymization. The following discussion uses the definitions from [RFC6973] Section 3, with additional observations from [van-Dijkhuizen-et-al.]

- o Anonymization. To enable anonymity of an individual, there must exist a set of individuals that appear to have the same attribute(s) as the individual. To the attacker or the observer, these individuals must appear indistinguishable from each other.
- o Pseudonymization. The true identity is deterministically replaced with an alternate identity (a pseudonym). When the pseudonymization schema is known, the process can be reversed, so the original identity becomes known again.

In practice there is a fine line between the two; for example, how to categorize a deterministic algorithm for data minimization of IP addresses that produces a group of pseudonyms for a single given address.

#### 5.2.3. IP address pseudonymization and anonymization methods

A major privacy risk in DNS is connecting DNS queries to an individual and the major vector for this in DNS traffic is the client IP address.

There is active discussion in the space of effective pseudonymization of IP addresses in DNS traffic logs, however there seems to be no single solution that is widely recognized as suitable for all or most use cases. There are also as yet no standards for this that are unencumbered by patents.

Appendix B provides a more detailed survey of various techniques employed or under development in 2019.

#### 5.2.4. Pseudonymization, anonymization, or discarding of other correlation data

DNS Privacy Threats:

- o Fingerprinting of the client OS via various means including: IP TTL/Hoplimit, TCP parameters (e.g., window size, ECN support, SACK), OS specific DNS query patterns (e.g., for network connectivity, captive portal detection, or OS specific updates).



- o Fingerprinting of the client application or TLS library by, e.g., HTTP headers (e.g., User-Agent, Accept, Accept-Encoding), TLS version/Cipher suite combinations, or other connection parameters.
- o Correlation of queries on multiple TCP sessions originating from the same IP address.
- o Correlating of queries on multiple TLS sessions originating from the same client, including via session resumption mechanisms.
- o Resolvers might receive client identifiers, e.g., MAC addresses in EDNS(0) options - some Customer-premises equipment (CPE) devices are known to add them [MAC-address-EDNS].

Mitigations:

- o Data minimization or discarding of such correlation data.

#### 5.2.5. Cache snooping

[RFC6973] Threats:

- o Surveillance:
  - \* Profiling of client queries by malicious third parties.

Mitigations:

- o See [ISC-Knowledge-database-on-cache-snooping] for an example discussion on defending against cache snooping. Options proposed include limiting access to a server and limiting non-recursive queries.

#### 5.3. Data sent onwards from the server

In this section we consider both data sent on the wire in upstream queries and data shared with third parties.

##### 5.3.1. Protocol recommendations

[RFC6973] Threats:

- o Surveillance:
  - \* Transmission of identifying data upstream.

Mitigations:

As specified in [RFC8310] for DoT but applicable to any DNS Privacy services the server should:

- o Implement QNAME minimization [RFC7816].
- o Honor a SOURCE PREFIX-LENGTH set to 0 in a query containing the EDNS(0) Client Subnet (ECS) option ([RFC7871] Section 7.1.2).

Optimizations:

- o As per Section 2 of [RFC7871] the server should either:
  - \* not use the ECS option in upstream queries at all, or
  - \* offer alternative services, one that sends ECS and one that does not.

If operators do offer a service that sends the ECS options upstream they should use the shortest prefix that is operationally feasible and ideally use a policy of allowlisting upstream servers to send ECS to in order to reduce data leakage. Operators should make clear in any policy statement what prefix length they actually send and the specific policy used.

Allowlisting has the benefit that not only does the operator know which upstream servers can use ECS but also allows the operator to decide which upstream servers apply privacy policies that the operator is happy with. However some operators consider allowlisting to incur significant operational overhead compared to dynamic detection of ECS support on authoritative servers.

Additional options:

- o Aggressive Use of DNSSEC-Validated Cache [RFC8198] and [RFC8020] (NXDOMAIN: There Really Is Nothing Underneath) to reduce the number of queries to authoritative servers to increase privacy.
- o Run a copy of the root zone on loopback [RFC8806] to avoid making queries to the root servers that might leak information.

### 5.3.2. Client query obfuscation

Additional options:

Since queries from recursive resolvers to authoritative servers are performed using cleartext (at the time of writing), resolver services need to consider the extent to which they may be directly leaking information about their client community via these upstream queries

and what they can do to mitigate this further. Note, that even when all the relevant techniques described above are employed there may still be attacks possible, e.g. [Pitfalls-of-DNS-Encryption]. For example, a resolver with a very small community of users risks exposing data in this way and ought to obfuscate this traffic by mixing it with 'generated' traffic to make client characterization harder. The resolver could also employ aggressive pre-fetch techniques as a further measure to counter traffic analysis.

At the time of writing there are no standardized or widely recognized techniques to perform such obfuscation or bulk pre-fetches.

Another technique that particularly small operators may consider is forwarding local traffic to a larger resolver (with a privacy policy that aligns with their own practices) over an encrypted protocol so that the upstream queries are obfuscated among those of the large resolver.

#### 5.3.3. Data sharing

[RFC6973] Threats:

- o Surveillance.
- o Stored data compromise.
- o Correlation.
- o Identification.
- o Secondary use.
- o Disclosure.

DNS Privacy Threats:

- o Contravention of legal requirements not to process user data.

Mitigations:

Operators should not share identifiable data with third-parties.

If operators choose to share identifiable data with third-parties in specific circumstance they should publish the terms under which data is shared.

Operators should consider including specific guidelines for the collection of aggregated and/or anonymized data for research

purposes, within or outside of their own organization. This can benefit not only the operator (through inclusion in novel research) but also the wider Internet community. See the policy published by SURFnet [SURFnet-policy] on data sharing for research as an example.

## 6. Recursive operator Privacy Statement (RPS)

To be compliant with this Best Common Practices document, a DNS recursive operator SHOULD publish a Recursive operator Privacy Statement (RPS). Adopting the outline, and including the headings in the order provided, is a benefit to persons comparing RPSs from multiple operators.

Appendix C provides a comparison of some existing policy and privacy statements.

### 6.1. Outline of an RPS

The contents of Section 6.1.1 and Section 6.1.2 are non-normative, other than the order of the headings. Material under each topic is present to assist the operator developing their own RPS and:

- o Relates only to matters around to the technical operation of DNS privacy services, and not on any other matters.
- o Does not attempt to offer an exhaustive list for the contents of an RPS.
- o Is not intended to form the basis of any legal/compliance documentation.

Appendix D provides an example (also non-normative) of an RPS statement for a specific operator scenario.

#### 6.1.1. Policy

1. Treatment of IP addresses. Make an explicit statement that IP addresses are treated as personal data.
2. Data collection and sharing. Specify clearly what data (including IP addresses) is:
  - \* Collected and retained by the operator, and for what period it is retained.
  - \* Shared with partners.
  - \* Shared, sold, or rented to third-parties.

and in each case whether it is aggregated, pseudonymized, or anonymized and the conditions of data transfer. Where possible provide details of the techniques used for the above data minimizations.

3. Exceptions. Specify any exceptions to the above, for example, technically malicious or anomalous behavior.
4. Associated entities. Declare and explicitly enumerate any partners, third-party affiliations, or sources of funding.
5. Correlation. Whether user DNS data is correlated or combined with any other personal information held by the operator.
6. Result filtering. This section should explain whether the operator filters, edits or alters in any way the replies that it receives from the authoritative servers for each DNS zone, before forwarding them to the clients. For each category listed below, the operator should also specify how the filtering lists are created and managed, whether it employs any third-party sources for such lists, and which ones.
  - \* Specify if any replies are being filtered out or altered for network and computer security reasons (e.g., preventing connections to malware-spreading websites or botnet control servers).
  - \* Specify if any replies are being filtered out or altered for mandatory legal reasons, due to applicable legislation or binding orders by courts and other public authorities.
  - \* Specify if any replies are being filtered out or altered for voluntary legal reasons, due to an internal policy by the operator aiming at reducing potential legal risks.
  - \* Specify if any replies are being filtered out or altered for any other reason, including commercial ones.

#### 6.1.2. Practice

[NOTE FOR RFC EDITOR: Please update this section to use letters for the sub-bullet points instead of numbers. This was not done during review because the markdown tool used to write the document did not support it.]

Communicate the current operational practices of the service.

1. Deviations. Specify any temporary or permanent deviations from the policy for operational reasons.
  2. Client facing capabilities. With reference to each subsection of Section 5.1 provide specific details of which capabilities (transport, DNSSEC, padding, etc.) are provided on which client facing addresses/port combination or DoH URI template. For Section 5.1.2, clearly specify which specific authentication mechanisms are supported for each endpoint that offers DoT:
    1. The authentication domain name to be used (if any).
    2. The SPKI pin sets to be used (if any) and policy for rolling keys.
  3. Upstream capabilities. With reference to section Section 5.3 provide specific details of which capabilities are provided upstream for data sent to authoritative servers.
  4. Support. Provide contact/support information for the service.
  5. Data Processing. This section can optionally communicate links to and the high level contents of any separate statements the operator has published which cover applicable data processing legislation or agreements with regard to the location(s) of service provision.
- 6.2. Enforcement/accountability

Transparency reports may help with building user trust that operators adhere to their policies and practices.

Independent monitoring or analysis could be performed where possible of:

- o ECS, QNAME minimization, EDNS(0) padding, etc.
- o Filtering.
- o Uptime.

This is by analogy with several TLS or website analysis tools that are currently available e.g., [SSL-Labs] or [Internet.nl].

Additionally operators could choose to engage the services of a third party auditor to verify their compliance with their published RPS.

## 7. IANA considerations

None

## 8. Security considerations

Security considerations for DNS over TCP are given in [RFC7766], many of which are generally applicable to session based DNS. Guidance on operational requirements for DNS over TCP are also available in [I-D.dnsop-dns-tcp-requirements]. Security considerations for DoT are given in [RFC7858] and [RFC8310], those for DoH in [RFC8484].

Security considerations for DNSSEC are given in [RFC4033], [RFC4034] and [RFC4035].

## 9. Acknowledgements

Many thanks to Amelia Andersdotter for a very thorough review of the first draft of this document and Stephen Farrell for a thorough review at WGLC and for suggesting the inclusion of an example RPS. Thanks to John Todd for discussions on this topic, and to Stephane Bortzmeyer, Puneet Sood and Vittorio Bertola for review. Thanks to Daniel Kahn Gillmor, Barry Green, Paul Hoffman, Dan York, Jon Reed, Lorenzo Colitti for comments at the mic. Thanks to Loganaden Velvindron for useful updates to the text.

Sara Dickinson thanks the Open Technology Fund for a grant to support the work on this document.

## 10. Contributors

The below individuals contributed significantly to the document:

John Dickinson  
Sinodun Internet Technologies  
Magdalen Centre  
Oxford Science Park  
Oxford OX4 4GA  
United Kingdom

Jim Hague  
Sinodun Internet Technologies  
Magdalen Centre  
Oxford Science Park  
Oxford OX4 4GA  
United Kingdom

## 11. Changelog

draft-ietf-dprive-bcp-op-13

- o Minor edits

draft-ietf-dprive-bcp-op-12

- o Change DROP to RPS throughout

draft-ietf-dprive-bcp-op-11

- o Improve text around use of normative language
- o Fix section 5.1.3.2 bullets
- o Improve text in 6.1.2. item 2.
- o Rework text of 6.1.2. item 5 and update example DROP
- o Various editorial improvements

draft-ietf-dprive-bcp-op-10

- o Remove direct references to draft-ietf-dprive-rfc7626-bis, instead have one general reference RFC7626
- o Clarify that the DROP statement outline is non-normative and add some further qualifications about content
- o Update wording on data sharing to remove explicit discussion of consent
- o Move table in section 5.2.3 to an appendix
- o Move section 6.2 to an appendix
- o Corrections to references, typos and editorial updates from initial IESG comments.

draft-ietf-dprive-bcp-op-09

- o Fix references so they match the correct section numbers in draft-ietf-dprive-rfc7626-bis-05

draft-ietf-dprive-bcp-op-08

- o Address IETF Last call comments.



draft-ietf-dprive-bcp-op-07

- o Editorial changes following AD review.
- o Change all URIs to Informational References.

draft-ietf-dprive-bcp-op-06

- o Final minor changes from second WGLC.

draft-ietf-dprive-bcp-op-05

- o Remove some text on consent:
  - \* Paragraph 2 in section 5.3.3
  - \* Item 6 in the DROP Practice statement (and example)
- o Remove .onion and TLSA options
- o Include ACME as a reference for certificate management
- o Update text on session resumption usage
- o Update section 5.2.4 on client fingerprinting

draft-ietf-dprive-bcp-op-04

- o Change DPPPS to DROP (DNS Recursive Operator Privacy) statement
- o Update structure of DROP slightly
- o Add example DROP statement
- o Add text about restricting access to full logs
- o Move table in section 5.2.3 from SVG to inline table
- o Fix many editorial and reference nits

draft-ietf-dprive-bcp-op-03

- o Add paragraph about operational impact
- o Move DNSSEC requirement out of the Appendix into main text as a privacy threat that should be mitigated
- o Add TLS version/Cipher suite as tracking threat

- o Add reference to Mozilla TRR policy
  - o Remove several TODOs and QUESTIONS.
- draft-ietf-dprive-bcp-op-02
- o Change 'open resolver' for 'public resolver'
  - o Minor editorial changes
  - o Remove recommendation to run a separate TLS 1.3 service
  - o Move TLSA to purely a optimization in Section 5.2.1
  - o Update reference on minimal DoH headers.
  - o Add reference on user switching provider after service issues in Section 5.1.4
  - o Add text in Section 5.1.6 on impact on operators.
  - o Add text on additional threat to TLS proxy use (Section 5.1.7)
  - o Add reference in Section 5.3.1 on example policies.
- draft-ietf-dprive-bcp-op-01
- o Many minor editorial fixes
  - o Update DoH reference to RFC8484 and add more text on DoH
  - o Split threat descriptions into ones directly referencing RFC6973 and other DNS Privacy threats
  - o Improve threat descriptions throughout
  - o Remove reference to the DNSSEC TLS Chain Extension draft until new version submitted.
  - o Clarify use of allowlisting for ECS
  - o Re-structure the DPPPS, add Result filtering section.
  - o Remove the direct inclusion of privacy policy comparison, now just reference dnsprivacy.org and an example of such work.
  - o Add an appendix briefly discussing DNSSEC

- o Update affiliation of 1 author

draft-ietf-dprive-bcp-op-00

- o Initial commit of re-named document after adoption to replace draft-dickinson-dprive-bcp-op-01

## 12. References

### 12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7457] Sheffer, Y., Holz, R., and P. Saint-Andre, "Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS)", RFC 7457, DOI 10.17487/RFC7457, February 2015, <<https://www.rfc-editor.org/info/rfc7457>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and D. Wessels, "DNS Transport over TCP - Implementation Requirements", RFC 7766, DOI 10.17487/RFC7766, March 2016, <<https://www.rfc-editor.org/info/rfc7766>>.

- [RFC7816] Bortzmeyer, S., "DNS Query Name Minimisation to Improve Privacy", RFC 7816, DOI 10.17487/RFC7816, March 2016, <<https://www.rfc-editor.org/info/rfc7816>>.
- [RFC7828] Wouters, P., Abley, J., Dickinson, S., and R. Bellis, "The edns-tcp-keepalive EDNS0 Option", RFC 7828, DOI 10.17487/RFC7828, April 2016, <<https://www.rfc-editor.org/info/rfc7828>>.
- [RFC7830] Mayrhofer, A., "The EDNS(0) Padding Option", RFC 7830, DOI 10.17487/RFC7830, May 2016, <<https://www.rfc-editor.org/info/rfc7830>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc7871>>.
- [RFC8020] Bortzmeyer, S. and S. Huque, "NXDOMAIN: There Really Is Nothing Underneath", RFC 8020, DOI 10.17487/RFC8020, November 2016, <<https://www.rfc-editor.org/info/rfc8020>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8198] Fujiwara, K., Kato, A., and W. Kumari, "Aggressive Use of DNSSEC-Validated Cache", RFC 8198, DOI 10.17487/RFC8198, July 2017, <<https://www.rfc-editor.org/info/rfc8198>>.
- [RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles for DNS over TLS and DNS over DTLS", RFC 8310, DOI 10.17487/RFC8310, March 2018, <<https://www.rfc-editor.org/info/rfc8310>>.
- [RFC8467] Mayrhofer, A., "Padding Policies for Extension Mechanisms for DNS (EDNS(0))", RFC 8467, DOI 10.17487/RFC8467, October 2018, <<https://www.rfc-editor.org/info/rfc8467>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.

- [RFC8490] Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S., Lemon, T., and T. Pusateri, "DNS Stateful Operations", RFC 8490, DOI 10.17487/RFC8490, March 2019, <<https://www.rfc-editor.org/info/rfc8490>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.
- [RFC8806] Kumari, W. and P. Hoffman, "Running a Root Server Local to a Resolver", RFC 8806, DOI 10.17487/RFC8806, June 2020, <<https://www.rfc-editor.org/info/rfc8806>>.

## 12.2. Informative References

- [Bloom-filter]  
van Rijswijk-Deij, R., Rijnders, G., Bomhoff, M., and L. Allodi, "Privacy-Conscious Threat Intelligence Using DNSBLOOM", 2019, <<http://dl.ifip.org/db/conf/im/im2019/189282.pdf>>.
- [Brenker-and-Arnes]  
Brekne, T. and A. Arnes, "CIRCUMVENTING IP-ADDRESS PSEUDONYMIZATION", 2005, <<https://pdfs.semanticscholar.org/7b34/12c951cebe71cd2cddac5fda164fb2138a44.pdf>>.
- [Crypto-PAn]  
CESNET, "Crypto-PAn", 2015, <<https://github.com/CESNET/ipfixcol/tree/master/base/src/intermediate/anonymization/Crypto-PAn>>.
- [DNS-Privacy-not-so-private]  
Silby, S., Juarez, M., Vallina-Rodriguez, N., and C. Troncosol, "DNS Privacy not so private: the traffic analysis perspective.", 2019, <<https://petsymposium.org/2018/files/hotpets/4-siby.pdf>>.
- [dnsmist] PowerDNS, "dnsmist Overview", 2019, <<https://dnsmist.org>>.
- [dnstap] dnstap.info, "DNSTAP", 2019, <<http://dnstap.info>>.
- [DoH-resolver-policy]  
Mozilla, "Security/DOH-resolver-policy", 2019, <<https://wiki.mozilla.org/Security/DOH-resolver-policy>>.

- [dot-ALPN] IANA (iana.org), "TLS Application-Layer Protocol Negotiation (ALPN) Protocol IDs", 2020, <<https://www.iana.org/assignments/tls-extensiontype-values/tls-extensiontype-values.xhtml#alpn-protocol-ids>>.
- [Geolocation-Impact-Assesment] Conversion Works, "Anonymize IP Geolocation Accuracy Impact Assessment", 2017, <<https://support.google.com/analytics/answer/2763052?hl=en>>.
- [haproxy] haproxy.org, "HAPROXY", 2019, <<https://www.haproxy.org/>>.
- [Harvan] Harvan, M., "Prefix- and Lexicographical-order-preserving IP Address Anonymization", 2006, <[http://mharvan.net/talks/noms-ip\\_anon.pdf](http://mharvan.net/talks/noms-ip_anon.pdf)>.
- [I-D.bellis-dnsop-xpf] Bellis, R., Dijk, P., and R. Gacogne, "DNS X-Proxied-For", draft-bellis-dnsop-xpf-04 (work in progress), March 2018.
- [I-D.ietf-dnsop-dns-tcp-requirements] Kristoff, J. and D. Wessels, "DNS Transport over TCP - Operational Requirements", draft-ietf-dnsop-dns-tcp-requirements-06 (work in progress), May 2020.
- [I-D.ietf-httpbis-bcp56bis] Nottingham, M., "Building Protocols with HTTP", draft-ietf-httpbis-bcp56bis-09 (work in progress), November 2019.
- [Internet.nl] Internet.nl, "Internet.nl Is Your Internet Up To Date?", 2019, <<https://internet.nl>>.
- [IP-Anonymization-in-Analytics] Google, "IP Anonymization in Analytics", 2019, <<https://support.google.com/analytics/answer/2763052?hl=en>>.
- [ipcipher1] Hubert, B., "On IP address encryption: security analysis with respect for privacy", 2017, <<https://medium.com/@bert.hubert/on-ip-address-encryption-security-analysis-with-respect-for-privacy-dabel1201b476>>.

- [ipcipher2] PowerDNS, "ipcipher", 2017, <<https://github.com/PowerDNS/ipcipher>>.
- [ipcrypt] veorq, "ipcrypt: IP-format-preserving encryption", 2015, <<https://github.com/veorq/ipcrypt>>.
- [ipcrypt-analysis] Aumasson, J., "Analysis of ipcrypt?", 2018, <<https://www.ietf.org/mail-archive/web/cfrg/current/msg09494.html>>.
- [ISC-Knowledge-database-on-cache-snooping] ISC Knowledge Database, "DNS Cache snooping - should I be concerned?", 2018, <<https://kb.isc.org/docs/aa-00482>>.
- [MAC-address-EDNS] DNS-OARC mailing list, "Embedding MAC address in DNS requests for selective filtering IDs", 2016, <<https://lists.dns-oarc.net/pipermail/dns-operations/2016-January/014143.html>>.
- [nginx] nginx.org, "NGINX", 2019, <<https://nginx.org/>>.
- [Passive-Observations-of-a-Large-DNS] de Vries, W., van Rijswijk-Deij, R., de Boer, P., and A. Pras, "Passive Observations of a Large DNS Service: 2.5 Years in the Life of Google", 2018, <[http://tma.ifip.org/2018/wp-content/uploads/sites/3/2018/06/tma2018\\_paper30.pdf](http://tma.ifip.org/2018/wp-content/uploads/sites/3/2018/06/tma2018_paper30.pdf)>.
- [pcap] tcpdump.org, "PCAP", 2016, <<http://www.tcpdump.org/>>.
- [Pitfalls-of-DNS-Encryption] Shulman, H., "Pretty Bad Privacy: Pitfalls of DNS Encryption", 2014, <<https://dl.acm.org/citation.cfm?id=2665959>>.
- [policy-comparison] dnsprivacy.org, "Comparison of policy and privacy statements 2019", 2019, <<https://dnsprivacy.org/wiki/display/DP/Comparison+of+policy+and+privacy+statements+2019>>.
- [PowerDNS-dnswasher] PowerDNS, "dnswasher", 2019, <<https://github.com/PowerDNS/pdns/blob/master/pdns/dnswasher.cc>>.

- [Ramaswamy-and-Wolf]  
Ramaswamy, R. and T. Wolf, "High-Speed Prefix-Preserving IP Address Anonymization for Passive Measurement Systems", 2007,  
<<http://www.ecs.umass.edu/ece/wolf/pubs/ton2007.pdf>>.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005,  
<<https://www.rfc-editor.org/info/rfc4034>>.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, DOI 10.17487/RFC4035, March 2005,  
<<https://www.rfc-editor.org/info/rfc4035>>.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, DOI 10.17487/RFC5077, January 2008, <<https://www.rfc-editor.org/info/rfc5077>>.
- [RFC6235] Boschi, E. and B. Trammell, "IP Flow Anonymization Support", RFC 6235, DOI 10.17487/RFC6235, May 2011,  
<<https://www.rfc-editor.org/info/rfc6235>>.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", RFC 6265, DOI 10.17487/RFC6265, April 2011, <<https://www.rfc-editor.org/info/rfc6265>>.
- [RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <<https://www.rfc-editor.org/info/rfc7626>>.
- [RFC7873] Eastlake 3rd, D. and M. Andrews, "Domain Name System (DNS) Cookies", RFC 7873, DOI 10.17487/RFC7873, May 2016,  
<<https://www.rfc-editor.org/info/rfc7873>>.
- [RFC8027] Hardaker, W., Gudmundsson, O., and S. Krishnaswamy, "DNSSEC Roadblock Avoidance", BCP 207, RFC 8027, DOI 10.17487/RFC8027, November 2016, <<https://www.rfc-editor.org/info/rfc8027>>.
- [RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/RFC8094, February 2017, <<https://www.rfc-editor.org/info/rfc8094>>.



- [RFC8404] Moriarty, K., Ed. and A. Morton, Ed., "Effects of Pervasive Encryption on Operators", RFC 8404, DOI 10.17487/RFC8404, July 2018, <<https://www.rfc-editor.org/info/rfc8404>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <<https://www.rfc-editor.org/info/rfc8555>>.
- [RFC8618] Dickinson, J., Hague, J., Dickinson, S., Manderson, T., and J. Bond, "Compacted-DNS (C-DNS): A Format for DNS Packet Capture", RFC 8618, DOI 10.17487/RFC8618, September 2019, <<https://www.rfc-editor.org/info/rfc8618>>.
- [SSL-Labs] SSL Labs, "SSL Server Test", 2019, <<https://www.ssllabs.com/ssltest/>>.
- [stunnel] ISC Knowledge Database, "DNS-over-TLS", 2018, <<https://kb.isc.org/article/AA-01386/0/DNS-over-TLS.html>>.
- [SURFnet-policy] SURFnet, "SURFnet Data Sharing Policy", 2016, <<https://surf.nl/datasharing>>.
- [TCPdpriv] Ipsilon Networks, Inc., "TCPdpriv", 2005, <<http://ita.ee.lbl.gov/html/contrib/tcpdpriv.html>>.
- [van-Dijkhuizen-et-al.] Van Dijkhuizen, N. and J. Van Der Ham, "A Survey of Network Traffic Anonymisation Techniques and Implementations", 2018, <<https://doi.org/10.1145/3182660>>.
- [Xu-et-al.] Fan, J., Xu, J., Ammar, M., and S. Moon, "Prefix-preserving IP address anonymization: measurement-based security evaluation and a new cryptography-based scheme", 2004, <<http://an.kaist.ac.kr/~sbmoon/paper/intl-journal/2004-cn-anon.pdf>>.

## Appendix A. Documents

This section provides an overview of some DNS privacy-related documents, however, this is neither an exhaustive list nor a definitive statement on the characteristic of the document.

### A.1. Potential increases in DNS privacy

These documents are limited in scope to communications between stub clients and recursive resolvers:

- o 'Specification for DNS over Transport Layer Security (TLS)' [RFC7858].
- o 'DNS over Datagram Transport Layer Security (DTLS)' [RFC8094]. Note that this document has the Category of Experimental.
- o 'DNS Queries over HTTPS (DoH)' [RFC8484].
- o 'Usage Profiles for DNS over TLS and DNS over DTLS' [RFC8310].
- o 'The EDNS(0) Padding Option' [RFC7830] and 'Padding Policy for EDNS(0)' [RFC8467].

These documents apply to recursive and authoritative DNS but are relevant when considering the operation of a recursive server:

- o 'DNS Query Name minimization to Improve Privacy' [RFC7816].

### A.2. Potential decreases in DNS privacy

These documents relate to functionality that could provide increased tracking of user activity as a side effect:

- o 'Client Subnet in DNS Queries' [RFC7871].
- o 'Domain Name System (DNS) Cookies' [RFC7873]).
- o 'Transport Layer Security (TLS) Session Resumption without Server-Side State' [RFC5077] referred to here as simply TLS session resumption.
- o [RFC8446] Appendix C.4 describes Client Tracking Prevention in TLS 1.3
- o 'A DNS Packet Capture Format' [RFC8618].
- o Passive DNS [RFC8499].

- o Section 8 of [RFC8484] outlines the privacy considerations of DoH. Note that (while that document advises exposing the minimal set of data needed to achieve the desired feature set) depending on the specifics of a DoH implementation there may be increased identification and tracking compared to other DNS transports.

### A.3. Related operational documents

- o 'DNS Transport over TCP - Implementation Requirements' [RFC7766].
- o 'Operational requirements for DNS over TCP' [I-D.ietf-dnsop-dns-tcp-requirements].
- o 'The edns-tcp-keepalive EDNS0 Option' [RFC7828].
- o 'DNS Stateful Operations' [RFC8490].

### Appendix B. IP address techniques

The following table presents a high level comparison of various techniques employed or under development in 2019, and classifies them according to categorization of technique and other properties. Both the specific techniques and the categorisations are described in more detail in the following sections. The list of techniques includes the main techniques in current use, but does not claim to be comprehensive.

| Categorization/Property   | GA | d | TC | C | TS | i | B |
|---------------------------|----|---|----|---|----|---|---|
| Anonymization             | X  | X | X  |   |    |   | X |
| Pseudoanonymization       |    |   |    | X | X  | X |   |
| Format preserving         | X  | X | X  | X | X  | X |   |
| Prefix preserving         |    |   | X  | X | X  |   |   |
| Replacement               |    |   | X  |   |    |   |   |
| Filtering                 | X  |   |    |   |    |   |   |
| Generalization            |    |   |    |   |    |   | X |
| Enumeration               |    | X |    |   |    |   |   |
| Reordering/Shuffling      |    |   | X  |   |    |   |   |
| Random substitution       |    |   | X  |   |    |   |   |
| Cryptographic permutation |    |   |    | X | X  | X |   |
| IPv6 issues               |    |   |    |   | X  |   |   |
| CPU intensive             |    |   |    | X |    |   |   |
| Memory intensive          |    |   | X  |   |    |   |   |
| Security concerns         |    |   |    |   |    | X |   |

Table 1: Classification of techniques

Legend of techniques: GA = Google Analytics, d = dnswasher, TC = TCPdpriv, C = CryptoPAN, TS = TSA, i = ipcipher, B = Bloom filter

The choice of which method to use for a particular application will depend on the requirements of that application and consideration of the threat analysis of the particular situation.

For example, a common goal is that distributed packet captures must be in an existing data format such as PCAP [pcap] or C-DNS [RFC8618] that can be used as input to existing analysis tools. In that case, use of a format-preserving technique is essential. This, though, is not cost-free - several authors (e.g., [Brenker-and-Arnes] have observed that, as the entropy in an IPv4 address is limited, if an attacker can

- o ensure packets are captured by the target and
- o send forged traffic with arbitrary source and destination addresses to that target and
- o obtain a de-identified log of said traffic from that target

any format-preserving pseudonymization is vulnerable to an attack along the lines of a cryptographic chosen plaintext attack.

#### B.1. Categorization of techniques

Data minimization methods may be categorized by the processing used and the properties of their outputs. The following builds on the categorization employed in [RFC6235]:

- o Format-preserving. Normally when encrypting, the original data length and patterns in the data should be hidden from an attacker. Some applications of de-identification, such as network capture de-identification, require that the de-identified data is of the same form as the original data, to allow the data to be parsed in the same way as the original.
- o Prefix preservation. Values such as IP addresses and MAC addresses contain prefix information that can be valuable in analysis, e.g., manufacturer ID in MAC addresses, subnet in IP addresses. Prefix preservation ensures that prefixes are de-identified consistently; e.g., if two IP addresses are from the same subnet, a prefix preserving de-identification will ensure that their de-identified counterparts will also share a subnet. Prefix preservation may be fixed (i.e. based on a user selected prefix length identified in advance to be preserved ) or general.

- o Replacement. A one-to-one replacement of a field to a new value of the same type, for example, using a regular expression.
- o Filtering. Removing or replacing data in a field. Field data can be overwritten, often with zeros, either partially (truncation or reverse truncation) or completely (black-marker anonymization).
- o Generalization. Data is replaced by more general data with reduced specificity. One example would be to replace all TCP/UDP port numbers with one of two fixed values indicating whether the original port was ephemeral ( $\geq 1024$ ) or non-ephemeral ( $> 1024$ ). Another example, precision degradation, reduces the accuracy of e.g., a numeric value or a timestamp.
- o Enumeration. With data from a well-ordered set, replace the first data item data using a random initial value and then allocate ordered values for subsequent data items. When used with timestamp data, this preserves ordering but loses precision and distance.
- o Reordering/shuffling. Preserving the original data, but rearranging its order, often in a random manner.
- o Random substitution. As replacement, but using randomly generated replacement values.
- o Cryptographic permutation. Using a permutation function, such as a hash function or cryptographic block cipher, to generate a replacement de-identified value.

## B.2. Specific techniques

### B.2.1. Google Analytics non-prefix filtering

Since May 2010, Google Analytics has provided a facility [IP-Anonymization-in-Analytics] that allows website owners to request that all their users IP addresses are anonymized within Google Analytics processing. This very basic anonymization simply sets to zero the least significant 8 bits of IPv4 addresses, and the least significant 80 bits of IPv6 addresses. The level of anonymization this produces is perhaps questionable. There are some analysis results [Geolocation-Impact-Assessment] which suggest that the impact of this on reducing the accuracy of determining the user's location from their IP address is less than might be hoped; the average discrepancy in identification of the user city for UK users is no more than 17%.

Anonymization: Format-preserving, Filtering (truncation).

#### B.2.2. dnswasher

Since 2006, PowerDNS have included a de-identification tool dnswasher [PowerDNS-dnswasher] with their PowerDNS product. This is a PCAP filter that performs a one-to-one mapping of end user IP addresses with an anonymized address. A table of user IP addresses and their de-identified counterparts is kept; the first IPv4 user addresses is translated to 0.0.0.1, the second to 0.0.0.2 and so on. The de-identified address therefore depends on the order that addresses arrive in the input, and running over a large amount of data the address translation tables can grow to a significant size.

Anonymization: Format-preserving, Enumeration.

#### B.2.3. Prefix-preserving map

Used in [TCPdpriv], this algorithm stores a set of original and anonymised IP address pairs. When a new IP address arrives, it is compared with previous addresses to determine the longest prefix match. The new address is anonymized by using the same prefix, with the remainder of the address anonymized with a random value. The use of a random value means that TCPdpriv is not deterministic; different anonymized values will be generated on each run. The need to store previous addresses means that TCPdpriv has significant and unbounded memory requirements, and because of the need to allocated anonymized addresses sequentially cannot be used in parallel processing.

Anonymization: Format-preserving, prefix preservation (general).

#### B.2.4. Cryptographic Prefix-Preserving Pseudonymization

Cryptographic prefix-preserving pseudonymization was originally proposed as an improvement to the prefix-preserving map implemented in TCPdpriv, described in [Xu-et-al.] and implemented in the [Crypto-PAN] tool. Crypto-PAN is now frequently used as an acronym for the algorithm. Initially it was described for IPv4 addresses only; extension for IPv6 addresses was proposed in [Harvan]. This uses a cryptographic algorithm rather than a random value, and thus pseudonymity is determined uniquely by the encryption key, and is deterministic. It requires a separate AES encryption for each output bit, so has a non-trivial calculation overhead. This can be mitigated to some extent (for IPv4, at least) by pre-calculating results for some number of prefix bits.

Pseudonymization: Format-preserving, prefix preservation (general).

#### B.2.5. Top-hash Subtree-replicated Anonymization

Proposed in [Ramaswamy-and-Wolf], Top-hash Subtree-replicated Anonymization (TSA) originated in response to the requirement for faster processing than Crypto-PAn. It used hashing for the most significant byte of an IPv4 address, and a pre-calculated binary tree structure for the remainder of the address. To save memory space, replication is used within the tree structure, reducing the size of the pre-calculated structures to a few Mb for IPv4 addresses. Address pseudonymization is done via hash and table lookup, and so requires minimal computation. However, due to the much increased address space for IPv6, TSA is not memory efficient for IPv6.

Pseudonymization: Format-preserving, prefix preservation (general).

#### B.2.6. ipcipher

A recently-released proposal from PowerDNS, ipcipher [ipcipher1] [ipcipher2] is a simple pseudonymization technique for IPv4 and IPv6 addresses. IPv6 addresses are encrypted directly with AES-128 using a key (which may be derived from a passphrase). IPv4 addresses are similarly encrypted, but using a recently proposed encryption [ipcrypt] suitable for 32bit block lengths. However, the author of ipcrypt has since indicated [ipcrypt-analysis] that it has low security, and further analysis has revealed it is vulnerable to attack.

Pseudonymization: Format-preserving, cryptographic permutation.

#### B.2.7. Bloom filters

van Rijswijk-Deij et al. have recently described work using Bloom filters [Bloom-filter] to categorize query traffic and record the traffic as the state of multiple filters. The goal of this work is to allow operators to identify so-called Indicators of Compromise (IOCs) originating from specific subnets without storing information about, or be able to monitor the DNS queries of an individual user. By using a Bloom filter, it is possible to determine with a high probability if, for example, a particular query was made, but the set of queries made cannot be recovered from the filter. Similarly, by mixing queries from a sufficient number of users in a single filter, it becomes practically impossible to determine if a particular user performed a particular query. Large numbers of queries can be tracked in a memory-efficient way. As filter status is stored, this approach cannot be used to regenerate traffic, and so cannot be used with tools used to process live traffic.

Anonymized: Generalization.

## Appendix C. Current policy and privacy statements

A tabular comparison of policy and privacy statements from various DNS Privacy service operators based loosely on the proposed RPS structure can be found at [policy-comparison]. The analysis is based on the data available in December 2019.

We note that the existing set of policies vary widely in style, content and detail and it is not uncommon for the full text for a given operator to equate to more than 10 pages of moderate font sized A4 text. It is a non-trivial task today for a user to extract a meaningful overview of the different services on offer.

It is also noted that Mozilla have published a DoH resolver policy [DoH-resolver-policy], which describes the minimum set of policy requirements that a party must satisfy to be considered as a potential partner for Mozilla's Trusted Recursive Resolver (TRR) program.

## Appendix D. Example RPS

The following example RPS is very loosely based on some elements of published privacy statements for some public resolvers, with additional fields populated to illustrate the what the full contents of an RPS might look like. This should not be interpreted as

- o having been reviewed or approved by any operator in any way
- o having any legal standing or validity at all
- o being complete or exhaustive

This is a purely hypothetical example of an RPS to outline example contents - in this case for a public resolver operator providing a basic DNS Privacy service via one IP address and one DoH URI with security based filtering. It does aim to meet minimal compliance as specified in Section 5.

### D.1. Policy

1. Treatment of IP addresses. Many nations classify IP addresses as personal data, and we take a conservative approach in treating IP addresses as personal data in all jurisdictions in which our systems reside.
2. Data collection and sharing.



1. IP addresses. Our normal course of data management does not have any IP address information or other personal data logged to disk or transmitted out of the location in which the query was received. We may aggregate certain counters to larger network block levels for statistical collection purposes, but those counters do not maintain specific IP address data nor is the format or model of data stored capable of being reverse-engineered to ascertain what specific IP addresses made what queries.
2. Data collected in logs. We do keep some generalized location information (at the city/metropolitan area level) so that we can conduct debugging and analyze abuse phenomena. We also use the collected information for the creation and sharing of telemetry (timestamp, geolocation, number of hits, first seen, last seen) for contributors, public publishing of general statistics of system use (protections, threat types, counts, etc.) When you use our DNS Services, here is the full list of items that are included in our logs:

- + Request domain name, e.g., example.net
- + Record type of requested domain, e.g., A, AAAA, NS, MX, TXT, etc.
- + Transport protocol on which the request arrived, i.e. UDP, TCP, DoT, DoH
- + Origin IP general geolocation information: i.e. geocode, region ID, city ID, and metro code
- + IP protocol version - IPv4 or IPv6
- + Response code sent, e.g., SUCCESS, SERVFAIL, NXDOMAIN, etc.
- + Absolute arrival time using a precision in ms
- + Name of the specific instance that processed this request
- + IP address of the specific instance to which this request was addressed (no relation to the requestor's IP address)

We may keep the following data as summary information, including all the above EXCEPT for data about the DNS record requested:

- + Currently-advertised BGP-summarized IP prefix/netmask of apparent client origin
- + Autonomous system number (BGP ASN) of apparent client origin

All the above data may be kept in full or partial form in permanent archives.

3. Sharing of data. Except as described in this document, we do not intentionally share, sell, or rent individual personal information associated with the requestor (i.e. source IP address or any other information that can positively identify the client using our infrastructure) with anyone without your consent. We generate and share high level anonymized aggregate statistics including threat metrics on threat type, geolocation, and if available, sector, as well as other vertical metrics including performance metrics on our DNS Services (i.e. number of threats blocked, infrastructure uptime) when available with our threat intelligence (TI) partners, academic researchers, or the public. Our DNS Services share anonymized data on specific domains queried (records such as domain, timestamp, geolocation, number of hits, first seen, last seen) with our threat intelligence partners. Our DNS Services also builds, stores, and may share certain DNS data streams which store high level information about domain resolved, query types, result codes, and timestamp. These streams do not contain IP address information of requestor and cannot be correlated to IP address or other personal data. We do not and never will share any of its data with marketers, nor will it use this data for demographic analysis.
3. Exceptions. There are exceptions to this storage model: In the event of actions or observed behaviors which we deem malicious or anomalous, we may utilize more detailed logging to collect more specific IP address data in the process of normal network defence and mitigation. This collection and transmission off-site will be limited to IP addresses that we determine are involved in the event.
4. Associated entities. Details of our Threat Intelligence partners can be found at our website page (insert link).
5. Correlation of Data. We do not correlate or combine information from our logs with any personal information that you have provided us for other services, or with your specific IP address.

## 6. Result filtering.

1. Filtering. We utilise cyber threat intelligence about malicious domains from a variety of public and private sources and blocks access to those malicious domains when your system attempts to contact them. An NXDOMAIN is returned for blocked sites.
1. Censorship. We will not provide a censoring component and will limit our actions solely to the blocking of malicious domains around phishing, malware, and exploit kit domains.
2. Accidental blocking. We implement allowlisting algorithms to make sure legitimate domains are not blocked by accident. However, in the rare case of blocking a legitimate domain, we work with the users to quickly allowlist that domain. Please use our support form ([insert link](#)) if you believe we are blocking a domain in error.

## D.2. Practice

1. Deviations from Policy. None in place since (insert date).
2. Client facing capabilities.
  1. We offer UDP and TCP DNS on port 53 on (insert IP address)
  2. We offer DNS over TLS as specified in RFC7858 on (insert IP address). It is available on port 853 and port 443. We also implement RFC7766.
    1. The DoT authentication domain name used is (insert domain name).
    2. We do not publish SPKI pin sets.
  3. We offer DNS over HTTPS as specified in RFC8484 on (insert URI template).
  4. Both services offer TLS 1.2 and TLS 1.3.
  5. Both services pad DNS responses according to RFC8467.
  6. Both services provide DNSSEC validation.
3. Upstream capabilities.

1. Our servers implement QNAME minimization.
2. Our servers do not send ECS upstream.
4. Support. Support information for this service is available at (insert link).
5. Data Processing. We operate as the legal entity (insert entity) registered in (insert country); as such we operate under (insert country/region) law. Our separate statement regarding the specifics of our data processing policy, practice, and agreements can be found here (insert link).

#### Authors' Addresses

Sara Dickinson  
Sinodun IT  
Magdalen Centre  
Oxford Science Park  
Oxford OX4 4GA  
United Kingdom

Email: sara@sinodun.com

Benno J. Overeinder  
NLnet Labs  
Science Park 400  
Amsterdam 1098 XH  
The Netherlands

Email: benno@nlnetLabs.nl

Roland M. van Rijswijk-Deij  
NLnet Labs  
Science Park 400  
Amsterdam 1098 XH  
The Netherlands

Email: roland@nlnetLabs.nl

Allison Mankin  
Salesforce

Email: allison.mankin@gmail.com

dprive  
Internet-Draft  
Obsoletes: 7626 (if approved)  
Intended status: Informational  
Expires: September 10, 2021

T. Wicinski, Ed.  
March 9, 2021

DNS Privacy Considerations  
draft-ietf-dprive-rfc7626-bis-09

Abstract

This document describes the privacy issues associated with the use of the DNS by Internet users. It provides general observations about typical current privacy practices. It is intended to be an analysis of the present situation and does not prescribe solutions. This document obsoletes RFC 7626.

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 <https://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 10, 2021.

Copyright Notice

Copyright (c) 2021 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 (<https://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 . . . . .                                     | 2  |
| 2. Scope . . . . .                                            | 5  |
| 3. Risks . . . . .                                            | 5  |
| 4. Risks in the DNS Data . . . . .                            | 6  |
| 4.1. The Public Nature of DNS Data . . . . .                  | 6  |
| 4.2. Data in the DNS Request . . . . .                        | 6  |
| 4.2.1. Data in the DNS Payload . . . . .                      | 8  |
| 4.3. Cache Snooping . . . . .                                 | 8  |
| 5. Risks On the Wire . . . . .                                | 8  |
| 5.1. Unencrypted Transports . . . . .                         | 8  |
| 5.2. Encrypted Transports . . . . .                           | 10 |
| 6. Risks in the Servers . . . . .                             | 11 |
| 6.1. In the Recursive Resolvers . . . . .                     | 12 |
| 6.1.1. Resolver Selection . . . . .                           | 12 |
| 6.1.2. Active Attacks on Resolver Configuration . . . . .     | 14 |
| 6.1.3. Blocking of DNS Resolution Services . . . . .          | 15 |
| 6.1.4. Encrypted Transports and Recursive Resolvers . . . . . | 15 |
| 6.2. In the Authoritative Name Servers . . . . .              | 16 |
| 7. Other risks . . . . .                                      | 17 |
| 7.1. Re-identification and Other Inferences . . . . .         | 17 |
| 7.2. More Information . . . . .                               | 18 |
| 8. Actual "Attacks" . . . . .                                 | 18 |
| 9. Legalities . . . . .                                       | 19 |
| 10. Security Considerations . . . . .                         | 19 |
| 11. IANA Considerations . . . . .                             | 19 |
| 12. Contributions . . . . .                                   | 19 |
| 13. Acknowledgments . . . . .                                 | 19 |
| 14. References . . . . .                                      | 20 |
| 14.1. Normative References . . . . .                          | 20 |
| 14.2. Informative References . . . . .                        | 20 |
| 14.3. URIs . . . . .                                          | 26 |
| Appendix A. Updates since RFC7626 . . . . .                   | 27 |
| Appendix B. Changelog . . . . .                               | 27 |
| Author's Address . . . . .                                    | 30 |

## 1. Introduction

This document is an analysis of the DNS privacy issues, in the spirit of Section 8 of [RFC6973].

The Domain Name System (DNS) is specified in [RFC1034], [RFC1035], and many later RFCs, which have never been consolidated. It is one of the most important infrastructure components of the Internet and

often ignored or misunderstood by Internet users (and even by many professionals). Almost every activity on the Internet starts with a DNS query (and often several). Its use has many privacy implications and this document is an attempt at a comprehensive and accurate list.

Let us begin with a simplified reminder of how the DNS works (See also [RFC8499]). A client, the stub resolver, issues a DNS query to a server, called the recursive resolver (also called caching resolver or full resolver or recursive name server). Let's use the query "What are the AAAA records for www.example.com?" as an example. AAAA is the QTYPE (Query Type), and www.example.com is the QNAME (Query Name). (The description that follows assumes a cold cache, for instance, because the server just started.) The recursive resolver will first query the root name servers. In most cases, the root name servers will send a referral. In this example, the referral will be to the .com name servers. The resolver repeats the query to one of the .com name servers. The .com name servers, in turn, will refer to the example.com name servers. The example.com name server will then return the answer. The root name servers, the name servers of .com, and the name servers of example.com are called authoritative name servers. It is important, when analyzing the privacy issues, to remember that the question asked to all these name servers is always the original question, not a derived question. The question sent to the root name servers is "What are the AAAA records for www.example.com?", not "What are the name servers of .com?". By repeating the full question, instead of just the relevant part of the question to the next in line, the DNS provides more information than necessary to the name server. In this simplified description, recursive resolvers do not implement QNAME minimization as described in [RFC7816], which will only send the relevant part of the question to the upstream name server.

DNS relies heavily on caching, so the algorithm described above is actually a bit more complicated, and not all questions are sent to the authoritative name servers. If a few seconds later the stub resolver asks the recursive resolver, "What are the SRV records of \_xmpp-server.\_tcp.example.com?", the recursive resolver will remember that it knows the name servers of example.com and will just query them, bypassing the root and .com. Because there is typically no caching in the stub resolver, the recursive resolver, unlike the authoritative servers, sees all the DNS traffic. (Applications, like web browsers, may have some form of caching that does not follow DNS rules, for instance, because it may ignore the TTL. So, the recursive resolver does not see all the name resolution activity.)

It should be noted that DNS recursive resolvers sometimes forward requests to other recursive resolvers, typically bigger machines, with a larger and more shared cache (and the query hierarchy can be

even deeper, with more than two levels of recursive resolvers). From the point of view of privacy, these forwarders are like resolvers, except that they do not see all of the requests being made (due to caching in the first resolver).

At the time of writing, almost all this DNS traffic is currently sent unencrypted. However, there is increasing deployment of DNS-over-TLS (DoT) [RFC7858] and DNS-over-HTTPS (DoH) [RFC8484], particularly in mobile devices, browsers, and by providers of anycast recursive DNS resolution services. There are a few cases where there is some alternative channel encryption, for instance, in an IPsec VPN tunnel, at least between the stub resolver and the resolver. Some recent analysis on service quality of encrypted DNS traffic can be found in [dns-over-encryption].

Today, almost all DNS queries are sent over UDP [thomas-ditl-tcp]. This has practical consequences when considering encryption of the traffic as a possible privacy technique. Some encryption solutions are only designed for TCP, not UDP, although new solutions are still emerging [I-D.ietf-quic-transport] [I-D.ietf-dprive-dnsquic].

Another important point to keep in mind when analyzing the privacy issues of DNS is the fact that DNS requests received by a server are triggered by different reasons. Let's assume an eavesdropper wants to know which web page is viewed by a user. For a typical web page, there are three sorts of DNS requests being issued:

- o Primary request: this is the domain name in the URL that the user typed, selected from a bookmark, or chose by clicking on an hyperlink. Presumably, this is what is of interest for the eavesdropper.
- o Secondary requests: these are the additional requests performed by the user agent (here, the web browser) without any direct involvement or knowledge of the user. For the Web, they are triggered by embedded content, Cascading Style Sheets (CSS), JavaScript code, embedded images, etc. In some cases, there can be dozens of domain names in different contexts on a single web page.
- o Tertiary requests: these are the additional requests performed by the DNS system itself. For instance, if the answer to a query is a referral to a set of name servers, and the glue records are not returned, the resolver will have to do additional requests to turn the name servers' names into IP addresses. Similarly, even if glue records are returned, a careful recursive server will do tertiary requests to verify the IP addresses of those records.



It can also be noted that, in the case of a typical web browser, more DNS requests than strictly necessary are sent, for instance, to prefetch resources that the user may query later or when autocompleting the URL in the address bar. Both are a significant privacy concern since they may leak information even about non-explicit actions. For instance, just reading a local HTML page, even without selecting the hyperlinks, may trigger DNS requests.

For privacy-related terms, the terminology is from [RFC6973].

## 2. Scope

This document focuses mostly on the study of privacy risks for the end user (the one performing DNS requests). The risks of pervasive surveillance [RFC7258] are considered as well as risks coming from a more focused surveillance. In this document, the term 'end user' is used as defined in [RFC8890].

This document does not attempt a comparison of specific privacy protections provided by individual networks or organizations, it makes only general observations about typical current practices.

Privacy risks for the holder of a zone (the risk that someone gets the data) are discussed in [RFC5936] and [RFC5155].

Privacy risks for recursive operators (including access providers and operators in enterprise networks) such as leakage of private namespaces or blocklists are out of scope for this document.

Non-privacy risks (e.g security related considerations such as cache poisoning) are also out of scope.

The privacy risks associated with the use of other protocols that make use of DNS information are not considered here.

## 3. Risks

The following four sections outline the privacy considerations associated with different aspects of the DNS for the end user. When reading these sections it needs to be kept in mind that many of the considerations (for example, recursive resolver and transport protocol) can be specific to the network context that a device is using at a given point in time. A user may have many devices and each device might utilize many different networks (e.g. home, work, public or cellular) over a period of time or even concurrently. An exhaustive analysis of the privacy considerations for an individual user would need to take into account the set of devices used and the multiple dynamic contexts of each device. This document does not

attempt such a complex analysis, but instead it presents an overview of the various considerations that could form the basis of such an analysis.

#### 4. Risks in the DNS Data

##### 4.1. The Public Nature of DNS Data

It has been stated that "the data in the DNS is public". This sentence makes sense for an Internet-wide lookup system, and there are multiple facets to the data and metadata involved that deserve a more detailed look. First, access control lists (ACLs) and private namespaces notwithstanding, the DNS operates under the assumption that public-facing authoritative name servers will respond to "usual" DNS queries for any zone they are authoritative for without further authentication or authorization of the client (resolver). Due to the lack of search capabilities, only a given QNAME will reveal the resource records associated with that name (or that name's non-existence). In other words: one needs to know what to ask for, in order to receive a response. There are many ways in which supposedly "private" resources currently leak. A few examples are DNSSEC NSEC zone walking[RFC4470]; passive-DNS services[passive-dns]; etc. The zone transfer QTYPE [RFC5936] is often blocked or restricted to authenticated/authorized access to enforce this difference (and maybe for other reasons).

Another difference between the DNS data and a particular DNS transaction (i.e., a DNS name lookup). DNS data and the results of a DNS query are public, within the boundaries described above, and may not have any confidentiality requirements. However, the same is not true of a single transaction or a sequence of transactions; those transactions are not / should not be public. A single transaction reveals both the originator of the query and the query contents which potentially leaks sensitive information about a specific user. A typical example from outside the DNS world is: the web site of Alcoholics Anonymous is public; the fact that you visit it should not be. Furthermore, the ability to link queries reveals information about individual use patterns.

##### 4.2. Data in the DNS Request

The DNS request includes many fields, but two of them seem particularly relevant for the privacy issues: the QNAME and the source IP address. "source IP address" is used in a loose sense of "source IP address + maybe source port number", because the port number is also in the request and can be used to differentiate between several users sharing an IP address (behind a Carrier-Grade NAT (CGN), for instance [RFC6269]).

The QNAME is the full name sent by the user. It gives information about what the user does ("What are the MX records of example.net?" means he probably wants to send email to someone at example.net, which may be a domain used by only a few persons and is therefore very revealing about communication relationships). Some QNAMEs are more sensitive than others. For instance, querying the A record of a well-known web statistics domain reveals very little (everybody visits web sites that use this analytics service), but querying the A record of www.verybad.example where verybad.example is the domain of an organization that some people find offensive or objectionable may create more problems for the user. Also, sometimes, the QNAME embeds the software one uses, which could be a privacy issue. For instance, `_ldap._tcp.Default-First-Site-Name._sites.gc._msdcs.example.org`. There are also some BitTorrent clients that query an SRV record for `_bittorrent-tracker._tcp.domain.example`.

Another important thing about the privacy of the QNAME is the future usages. Today, the lack of privacy is an obstacle to putting potentially sensitive or personally identifiable data in the DNS. At the moment, your DNS traffic might reveal that you are doing email but not with whom. If your Mail User Agent (MUA) starts looking up Pretty Good Privacy (PGP) keys in the DNS [RFC7929], then privacy becomes a lot more important. And email is just an example; there would be other really interesting uses for a more privacy-friendly DNS.

For the communication between the stub resolver and the recursive resolver, the source IP address is the address of the user's machine. Therefore, all the issues and warnings about collection of IP addresses apply here. For the communication between the recursive resolver and the authoritative name servers, the source IP address has a different meaning; it does not have the same status as the source address in an HTTP connection. It can be typically the IP address of the recursive resolver that, in a way, "hides" the real user. However, hiding does not always work. Sometimes EDNS(0) Client subnet [RFC7871] is used (see one privacy analysis in [denis-edns-client-subnet]). Sometimes the end user has a personal recursive resolver on their machine. In both cases, the IP address originating queries to the authoritative server is as sensitive as it is for HTTP [sidn-entrada].

A note about IP addresses: there is currently no IETF document that describes in detail all the privacy issues around IP addressing in general, although [RFC7721] does discuss privacy considerations for IPv6 address generation mechanisms. In the meantime, the discussion here is intended to include both IPv4 and IPv6 source addresses. For a number of reasons, their assignment and utilization characteristics are different, which may have implications for details of information

leakage associated with the collection of source addresses. (For example, a specific IPv6 source address seen on the public Internet is less likely than an IPv4 address to originate behind an address sharing scheme.) However, for both IPv4 and IPv6 addresses, it is important to note that source addresses are propagated with queries via EDNS(0) Client subnet and comprise metadata about the host, user, or application that originated them.

#### 4.2.1. Data in the DNS Payload

At the time of writing there are no standardized client identifiers contained in the DNS payload itself (ECS [RFC7871] while widely used is only of Category Informational).

DNS Cookies [RFC7873] are a lightweight DNS transaction security mechanism that provides limited protection against a variety of increasingly common denial-of-service and amplification/forgery or cache poisoning attacks by off-path attackers. It is noted, however, that they are designed to just verify IP addresses (and should change once a client's IP address changes), but they are not designed to actively track users (like HTTP cookies).

There are anecdotal accounts of MAC addresses [1] and even user names being inserted in non-standard EDNS(0) options [RFC6891] for stub to resolver communications to support proprietary functionality implemented at the resolver (e.g., parental filtering).

#### 4.3. Cache Snooping

The content of recursive resolvers' caches can reveal data about the clients using it (the privacy risks depend on the number of clients). This information can sometimes be examined by sending DNS queries with RD=0 to inspect cache content, particularly looking at the DNS TTLs [grangeia.snooping]. Since this also is a reconnaissance technique for subsequent cache poisoning attacks, some counter measures have already been developed and deployed [cache-snooping-defence].

### 5. Risks On the Wire

#### 5.1. Unencrypted Transports

For unencrypted transports, DNS traffic can be seen by an eavesdropper like any other traffic. (DNSSEC, specified in [RFC4033], explicitly excludes confidentiality from its goals.) So, if an initiator starts an HTTPS communication with a recipient, while the HTTP traffic will be encrypted, the DNS exchange prior to it will not be. When other protocols will become more and more privacy-aware

and secured against surveillance (e.g., [RFC8446], [I-D.ietf-quic-transport]), the use of unencrypted transports for DNS may become "the weakest link" in privacy. It is noted that at the time of writing there is on-going work attempting to encrypt the SNI in the TLS handshake [RFC8744], which is one of the last remaining non-DNS cleartext identifiers of a connection target.

An important specificity of the DNS traffic is that it may take a different path than the communication between the initiator and the recipient. For instance, an eavesdropper may be unable to tap the wire between the initiator and the recipient but may have access to the wire going to the recursive resolver, or to the authoritative name servers.

The best place to tap, from an eavesdropper's point of view, is clearly between the stub resolvers and the recursive resolvers, because traffic is not limited by DNS caching.

The attack surface between the stub resolver and the rest of the world can vary widely depending upon how the end user's device is configured. By order of increasing attack surface:

- o The recursive resolver can be on the end user's device. In (currently) a small number of cases, individuals may choose to operate their own DNS resolver on their local machine. In this case, the attack surface for the connection between the stub resolver and the caching resolver is limited to that single machine. The recursive resolver will expose data to authoritative resolvers as discussed in Section 6.2.
- o The recursive resolver may be at the local network edge. For many/most enterprise networks and for some residential networks, the caching resolver may exist on a server at the edge of the local network. In this case, the attack surface is the local network. Note that in large enterprise networks, the DNS resolver may not be located at the edge of the local network but rather at the edge of the overall enterprise network. In this case, the enterprise network could be thought of as similar to the Internet Access Provider (IAP) network referenced below.
- o The recursive resolver can be in the IAP network. For most residential networks and potentially other networks, the typical case is for the user's device to be configured (typically automatically through DHCP or RA options) with the addresses of the DNS proxy in the Customer Premise Equipment (CPE), which in turns points to the DNS recursive resolvers at the IAP. The attack surface for on-the-wire attacks is therefore from the end

user system across the local network and across the IAP network to the IAP's recursive resolvers.

- o The recursive resolver can be a public DNS service (or a privately run DNS resolver hosted on the public internet). Some machines may be configured to use public DNS resolvers such as those operated by Google Public DNS or OpenDNS. The user may have configured their machine to use these DNS recursive resolvers themselves -- or their IAP may have chosen to use the public DNS resolvers rather than operating their own resolvers. In this case, the attack surface is the entire public Internet between the user's connection and the public DNS service. It can be noted that if the user selects a single resolver with a small client population (even when using an encrypted transport) it can actually serve to aid tracking of that user as they move across network environments.

It is also noted that typically a device connected only to a modern cellular network is

- o directly configured with only the recursive resolvers of the IAP and
- o afforded some level of protection against some types of eavesdropping for all traffic (including DNS traffic) due to the cellular network link-layer encryption.

The attack surface for this specific scenario is not considered here.

## 5.2. Encrypted Transports

The use of encrypted transports directly mitigates passive surveillance of the DNS payload, however there are still some privacy attacks possible. This section enumerates the residual privacy risks to an end user when an attacker can passively monitor encrypted DNS traffic flows on the wire.

These are cases where user identification, fingerprinting or correlations may be possible due to the use of certain transport layers or clear text/observable features. These issues are not specific to DNS, but DNS traffic is susceptible to these attacks when using specific transports.

There are some general examples, for example, certain studies have highlighted that IPv4 TTL, IPv6 Hop Limit, or TCP Window sizes or fingerprint [2] values can be used to fingerprint client OS's or that various techniques can be used to de-NAT DNS queries [dns-de-nat].

Note that even when using encrypted transports, the use of clear text transport options to decrease latency can provide correlation of a users' connections, e.g. using TCP Fast Open [RFC7413].

Implementations that support encrypted transports also commonly re-use connections for multiple DNS queries to optimize performance (e.g. via DNS pipelining or HTTPS multiplexing). Default configuration options for encrypted transports could in principle fingerprint a specific client application. For example:

- o TLS version or cipher suite selection
- o session resumption
- o the maximum number of messages to send or
- o a maximum connection time before closing a connections and re-opening.

If libraries or applications offer user configuration of such options (e.g. [getdns]) then they could in principle help to identify a specific user. Users may want to use only the defaults to avoid this issue.

Whilst there are known attacks on older versions of TLS, the most recent recommendations [RFC7525] and the development of TLS 1.3 [RFC8446] largely mitigate those.

Traffic analysis of unpadded encrypted traffic is also possible [pitfalls-of-dns-encryption] because the sizes and timing of encrypted DNS requests and responses can be correlated to unencrypted DNS requests upstream of a recursive resolver.

## 6. Risks in the Servers

Using the terminology of [RFC6973], the DNS servers (recursive resolvers and authoritative servers) are enablers: they facilitate communication between an initiator and a recipient without being directly in the communications path. As a result, they are often forgotten in risk analysis. But, to quote again [RFC6973], "Although [...] enablers may not generally be considered as attackers, they may all pose privacy threats (depending on the context) because they are able to observe, collect, process, and transfer privacy-relevant data." In [RFC6973] parlance, enablers become observers when they start collecting data.

Many programs exist to collect and analyze DNS data at the servers -- from the "query log" of some programs like BIND to tcpdump and more

sophisticated programs like PacketQ [packetq] and DNSmezzo [dnsmezzo]. The organization managing the DNS server can use this data itself, or it can be part of a surveillance program like PRISM [prism] and pass data to an outside observer.

Sometimes, this data is kept for a long time and/or distributed to third parties for research purposes [ditl] [day-at-root], security analysis, or surveillance tasks. These uses are sometimes under some sort of contract, with various limitations, for instance, on redistribution, given the sensitive nature of the data. Also, there are observation points in the network that gather DNS data and then make it accessible to third parties for research or security purposes ("passive DNS" [passive-dns]).

#### 6.1. In the Recursive Resolvers

Recursive Resolvers see all the traffic since there is typically no caching before them. To summarize: your recursive resolver knows a lot about you. The resolver of a large IAP, or a large public resolver, can collect data from many users.

##### 6.1.1. Resolver Selection

Given all the above considerations, the choice of recursive resolver has direct privacy considerations for end users. Historically, end user devices have used the DHCP-provided local network recursive resolver. The choice by a user to join a particular network (e.g. by physically plugging in a cable or selecting a network in a OS dialogue) typically updates a number of system resources - these can include IP addresses, availability of IPv4/IPv6, DHCP server, and DNS resolver. These individual changes, including the change in DNS resolver, are not normally communicated directly to the user by the OS when the network is joined. The choice of network has historically determined the default system DNS resolver selection; the two are directly coupled in this model.

The vast majority of users do not change their default system DNS settings and so implicitly accept the network settings for DNS. The network resolvers have therefore historically been the sole destination for all of the DNS queries from a device. These resolvers may have varied privacy policies depending on the network. Privacy policies for these servers may or may not be available and users need to be aware that privacy guarantees will vary with the network.

All major OS's expose the system DNS settings and allow users to manually override them if desired.



More recently, some networks and users have actively chosen to use a large public resolver, e.g., Google Public DNS [3], Cloudflare [4], or Quad9 [5]. There can be many reasons: cost considerations for network operators, better reliability or anti-censorship considerations are just a few. Such services typically do provide a privacy policy and the user can get an idea of the data collected by such operators by reading one e.g., Google Public DNS - Your Privacy [6].

In general, as with many other protocols, issues around centralization also arise with DNS. The picture is fluid with several competing factors contributing which can also vary by geographic region. These include:

- o ISP outsourcing, including to third party and public resolvers
- o regional market domination by one or only a few ISPs
- o applications directing DNS traffic by default to a limited subset of resolvers, see Section 6.1.1.2

An increased proportion of the global DNS resolution traffic being served by only a few entities means that the privacy considerations for users are highly dependent on the privacy policies and practices of those entities. Many of the issues around centralization are discussed in [centralisation-and-data-sovereignty].

#### 6.1.1.1. Dynamic Discovery of DoH and Strict DoT

Whilst support for opportunistic DoT can be determined by probing a resolver on port 853, there is currently no standardized discovery mechanism for DoH and Strict DoT servers.

This means that clients which might want to dynamically discover such encrypted services, and where users are willing to trust such services, are not able to do so. At the time of writing, efforts to provide standardized signaling mechanisms to discover the services offered by local resolvers are in progress [I-D.ietf-dnsop-resolver-information]. Note that an increasing numbers of ISPs are deploying encrypted DNS, for example see the Encrypted DNS Deployment Initiative [EDDI].

#### 6.1.1.2. Application-specific Resolver Selection

An increasing number of applications are offering application-specific encrypted DNS resolution settings, rather than defaulting to using only the system resolver. A variety of heuristics and

resolvers are available in different applications including hard-coded lists of recognized DoH/DoT servers.

Generally, users are not aware of application specific DNS settings, and may not have control over those settings. To address these limitations, users will only be aware of and have the ability to control such settings if applications provide the following functions:

- o communicate clearly to users the change when the default application resolver changes away from the system resolver
- o provide configuration options to change the default application resolver, including a choice to always use the system resolver
- o provide mechanisms for users to locally inspect, selectively forward, and filter queries (either via the application itself or use of the system resolver)

Application-specific changes to default destinations for users' DNS queries might increase or decrease user privacy - it is highly dependent on the network context and the application-specific default. This is an area of active debate and the IETF is working on a number of issues related to application-specific DNS settings.

#### 6.1.2. Active Attacks on Resolver Configuration

The previous section discussed DNS privacy, assuming that all the traffic was directed to the intended servers (i.e those that would be used in the absence of an active attack) and that the potential attacker was purely passive. But, in reality, there can be active attackers in the network.

The Internet Threat model, as described in [RFC3552], assumes that the attacker controls the network. Such an attacker can completely control any insecure DNS resolution, both passively monitoring the queries and responses and substituting their own responses. Even if encrypted DNS such as DoH or DoT is used, unless the client has been configured in a secure way with the server identity, an active attacker can impersonate the server. This implies that opportunistic modes of DoH/DoT as well as modes where the client learns of the DoH/DoT server via in-network mechanisms such as DHCP are vulnerable to attack. In addition, if the client is compromised, the attacker can replace the DNS configuration with one of its own choosing.

### 6.1.3. Blocking of DNS Resolution Services

User privacy can also be at risk if there is blocking of access to remote recursive servers that offer encrypted transports e.g., when the local resolver does not offer encryption and/or has very poor privacy policies. For example, active blocking of port 853 for DoT or of specific IP addresses could restrict the resolvers available to the user. The extent of the risk to user privacy is highly dependent on the specific network and user context; a user on a network that is known to perform surveillance would be compromised if they could not access such services, whereas a user on a trusted network might have no privacy motivation to do so.

As a matter of policy, some recursive resolvers use their position in the query path to selectively block access to certain DNS records. This is a form of Rendezvous-Based Blocking as described in Section 4.3 of [RFC7754]. Such blocklists often include servers known to be used for malware, bots or other security risks. In order to prevent circumvention of their blocking policies, some networks also block access to resolvers with incompatible policies.

It is also noted that attacks on remote resolver services, e.g., DDoS, could force users to switch to other services that do not offer encrypted transports for DNS.

### 6.1.4. Encrypted Transports and Recursive Resolvers

#### 6.1.4.1. DoT and DoH

Use of encrypted transports does not reduce the data available in the recursive resolver and ironically can actually expose more information about users to operators. As described in Section 5.2 use of session based encrypted transports (TCP/TLS) can expose correlation data about users.

#### 6.1.4.2. DoH Specific Considerations

DoH inherits the full privacy properties of the HTTPS stack and as a consequence introduces new privacy considerations when compared with DNS over UDP, TCP or TLS [RFC7858]. Section 8.2 of [RFC8484] describes the privacy consideration in the server of the DoH protocol.

A brief summary of some of the issues includes:

- o HTTPS presents new considerations for correlation, such as explicit HTTP cookies and implicit fingerprinting of the unique set and ordering of HTTP request header fields.

- o The User-Agent and Accept-Language request header fields often convey specific information about the client version or locale.
- o Utilizing the full set of HTTP features enables DoH to be more than an HTTP tunnel, but it is at the cost of opening up implementations to the full set of privacy considerations of HTTP.
- o Implementations are advised to expose the minimal set of data needed to achieve the desired feature set.

[RFC8484] specifically makes selection of HTTPS functionality vs privacy an implementation choice. At the extremes, there may be implementations that attempt to achieve parity with DoT from a privacy perspective at the cost of using no identifiable HTTP headers, there might be others that provide feature rich data flows where the low-level origin of the DNS query is easily identifiable. Some implementations have, in fact, chosen to restrict the use of the 'User-Agent' header so that resolver operators cannot identify the specific application that is originating the DNS queries.

Privacy focused users should be aware of the potential for additional client identifiers in DoH compared to DoT and may want to only use DoH client implementations that provide clear guidance on what identifiers they add.

## 6.2. In the Authoritative Name Servers

Unlike what happens for recursive resolvers, observation capabilities of authoritative name servers are limited by caching; they see only the requests for which the answer was not in the cache. For aggregated statistics ("What is the percentage of LOC queries?"), this is sufficient, but it prevents an observer from seeing everything. Similarly the increasing deployment of QNAME minimisation [ripe-qname-measurements] reduces the data visible at the authoritative name server. Still, the authoritative name servers see a part of the traffic, and this subset may be sufficient to violate some privacy expectations.

Also, the user often has some legal/contractual link with the recursive resolver (he has chosen the IAP, or he has chosen to use a given public resolver), while having no control and perhaps no awareness of the role of the authoritative name servers and their observation abilities.

As noted before, using a local resolver or a resolver close to the machine decreases the attack surface for an on-the-wire eavesdropper. But it may decrease privacy against an observer located on an authoritative name server. This authoritative name server will see

the IP address of the end client instead of the address of a big recursive resolver shared by many users.

This "protection", when using a large resolver with many clients, is no longer present if ECS [RFC7871] is used because, in this case, the authoritative name server sees the original IP address (or prefix, depending on the setup).

As of today, all the instances of one root name server, L-root, receive together around 50,000 queries per second. While most of it is "junk" (errors on the Top-Level Domain (TLD) name), it gives an idea of the amount of big data that pours into name servers. (And even "junk" can leak information; for instance, if there is a typing error in the TLD, the user will send data to a TLD that is not the usual one.)

Many domains, including TLDs, are partially hosted by third-party servers, sometimes in a different country. The contracts between the domain manager and these servers may or may not take privacy into account. Whatever the contract, the third-party hoster may be honest or not but, in any case, it will have to follow its local laws. For example, requests to a given ccTLD may go to servers managed by organizations outside of the ccTLD's country. Users may not anticipate that, when doing a security analysis.

Also, it seems (see the survey described in [aeris-dns]) that there is a strong concentration of authoritative name servers among "popular" domains (such as the Alexa Top N list). For instance, among the Alexa Top 100K [7], one DNS provider hosts today 10% of the domains. The ten most important DNS providers host together one third of the domains. With the control (or the ability to sniff the traffic) of a few name servers, you can gather a lot of information.

## 7. Other risks

### 7.1. Re-identification and Other Inferences

An observer has access not only to the data he/she directly collects but also to the results of various inferences about this data. The term 'observer' here is used very generally, it might be one that is passively observing cleartext DNS traffic, one in the network that is actively attacking the user by re-directing DNS resolution, or it might be a local or remote resolver operator.

For instance, a user can be re-identified via DNS queries. If the adversary knows a user's identity and can watch their DNS queries for a period, then that same adversary may be able to re-identify the user solely based on their pattern of DNS queries later on regardless

of the location from which the user makes those queries. For example, one study [herrmann-reidentification] found that such re-identification is possible so that "73.1% of all day-to-day links were correctly established, i.e., user u was either re-identified unambiguously (1) or the classifier correctly reported that u was not present on day t+1 any more (2)." While that study related to web browsing behavior, equally characteristic patterns may be produced even in machine-to-machine communications or without a user taking specific actions, e.g., at reboot time if a characteristic set of services are accessed by the device.

For instance, one could imagine that an intelligence agency identifies people going to a site by putting in a very long DNS name and looking for queries of a specific length. Such traffic analysis could weaken some privacy solutions.

The IAB privacy and security program also have a document [RFC7624] that considers such inference-based attacks in a more general framework.

## 7.2. More Information

Useful background information can also be found in [tor-leak] (about the risk of privacy leak through DNS) and in a few academic papers: [yanbin-tsudik], [castillo-garcia], [fangming-hori-sakurai], and [federrath-fuchs-herrmann-piosechny].

## 8. Actual "Attacks"

A very quick examination of DNS traffic may lead to the false conclusion that extracting the needle from the haystack is difficult. "Interesting" primary DNS requests are mixed with useless (for the eavesdropper) secondary and tertiary requests (see the terminology in Section 1). But, in this time of "big data" processing, powerful techniques now exist to get from the raw data to what the eavesdropper is actually interested in.

Many research papers about malware detection use DNS traffic to detect "abnormal" behavior that can be traced back to the activity of malware on infected machines. Yes, this research was done for the good, but technically it is a privacy attack and it demonstrates the power of the observation of DNS traffic. See [dns-footprint], [dagon-malware], and [darkreading-dns].

Passive DNS systems [passive-dns] allow reconstruction of the data of sometimes an entire zone. Well-known passive DNS systems keep only the DNS responses, and not the source IP address of the client, precisely for privacy reasons. Other passive DNS systems may not be

so careful. And there is still the potential problems with revealing QNAMES.

The revelations from the Edward Snowden documents, which were leaked from the National Security Agency (NSA), provide evidence of the use of the DNS in mass surveillance operations [morecowbell]. For example the MORECOWBELL surveillance program, which uses a dedicated covert monitoring infrastructure to actively query DNS servers and perform HTTP requests to obtain meta information about services and to check their availability. Also the QUANTUMTHEORY [8] project which includes detecting lookups for certain addresses and injecting bogus replies is another good example showing that the lack of privacy protections in the DNS is actively exploited.

## 9. Legalities

To our knowledge, there are no specific privacy laws for DNS data, in any country. Interpreting general privacy laws like [data-protection-directive] or GDPR [9] applicable in the European Union in the context of DNS traffic data is not an easy task, and there is no known court precedent. See an interesting analysis in [sidn-entrada].

## 10. Security Considerations

This document is entirely about security, more precisely privacy. It just lays out the problem; it does not try to set requirements (with the choices and compromises they imply), much less define solutions. Possible solutions to the issues described here are discussed in other documents (currently too many to all be mentioned); see, for instance, 'Recommendations for DNS Privacy Operators' [I-D.ietf-dprive-bcp-op].

## 11. IANA Considerations

This document makes no requests of the IANA.

## 12. Contributions

Sara Dickinson and Stephane Bortzmeyer were the original authors on the document, and their contribution on the initial version is greatly appreciated.

## 13. Acknowledgments

Thanks to Nathalie Boulevard and to the CENTR members for the original work that led to this document. Thanks to Ondrej Sury for the interesting discussions. Thanks to Mohsen Souissi and John Heidemann

for proofreading and to Paul Hoffman, Matthijs Mekking, Marcos Sanz, Tim Wicinski, Francis Dupont, Allison Mankin, and Warren Kumari for proofreading, providing technical remarks, and making many readability improvements. Thanks to Dan York, Suzanne Woolf, Tony Finch, Stephen Farrell, Peter Koch, Simon Josefsson, and Frank Denis for good written contributions. Thanks to Vittorio Bertola and Mohamed Boucadair for a detailed review of the -bis. And thanks to the IESG members for the last remarks.

## 14. References

### 14.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.

### 14.2. Informative References

- [aeris-dns] Vinot, N., "Vie privée: et le DNS alors?", (In French), 2015, <<https://blog.imirhil.fr/vie-privee-et-le-dns-alors.html>>.
- [cache-snooping-defence] ISC, "ISC Knowledge Database: DNS Cache snooping - should I be concerned?", 2018, <<https://kb.isc.org/docs/aa-00482>>.
- [castillo-garcia] Castillo-Perez, S. and J. Garcia-Alfaro, "Anonymous Resolution of DNS Queries", 2008, <<http://deic.uab.es/~joaquin/papers/is08.pdf>>.



[centralisation-and-data-sovereignty]

De Filippi, P. and S. McCarthy, "Cloud Computing: Centralization and Data Sovereignty", October 2012, <[https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2167372](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2167372)>.

[dagon-malware]

Dagon, D., "Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority", ISC/OARC Workshop, 2007, <<https://www.dns-oarc.net/files/workshop-2007/Dagon-Resolution-corruption.pdf>>.

[darkreading-dns]

Lemos, R., "Got Malware? Three Signs Revealed In DNS Traffic", InformationWeek Dark Reading, May 2013, <<http://www.darkreading.com/analytics/security-monitoring/got-malware-three-signs-revealed-in-dns-traffic/d-d-id/1139680>>.

[data-protection-directive]

European Parliament, "Directive 95/46/EC of the European Parliament and of the council on the protection of individuals with regard to the processing of personal data and on the free movement of such data", Official Journal L 281, pp. 0031 - 0050, November 1995, <<http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML>>.

[day-at-root]

Castro, S., Wessels, D., Fomenkov, M., and K. Claffy, "A Day at the Root of the Internet", ACM SIGCOMM Computer Communication Review, Vol. 38, Number 5, DOI 10.1145/1452335.1452341, October 2008, <<http://www.sigcomm.org/sites/default/files/ccr/papers/2008/October/1452335-1452341.pdf>>.

[denis-edns-client-subnet]

Denis, F., "Security and privacy issues of edns-client-subnet", August 2013, <<https://00f.net/2013/08/07/edns-client-subnet/>>.

[ditl]

CAIDA, "A Day in the Life of the Internet (DITL)", 2002, <<http://www.caida.org/projects/ditl/>>.

[dns-footprint]

Stoner, E., "DNS Footprint of Malware", OARC Workshop, October 2010, <<https://www.dns-oarc.net/files/workshop-201010/OARC-ers-20101012.pdf>>.

[dns-over-encryption]

Lu, C., Liu, B., Li, Z., Hao, S., Duan, H., Zhang, M., Leng, C., Liu, Y., Zhang, Z., and J. Wu, "An End-to-End, Large-Scale Measurement of DNS-over-Encryption", IMC '19 Amsterdam, Netherlands, DOI 10.1145/3355369.3355580, October 2019, <<http://dl.acm.org/citation.cfm?id=3355369.3355580>>.

[dnsmezzo]

Bortzmeyer, S., "DNSmezzo", 2009, <<http://www.dnsmezzo.net/>>.

[EDDI]

EDDI, "Encrypted DNS Deployment Initiative", 2020, <<https://www.encrypted-dns.org>>.

[fangming-hori-sakurai]

Fangming, Z., Hori, Y., and K. Sakurai, "Analysis of Privacy Disclosure in DNS Query", 2007 International Conference on Multimedia and Ubiquitous Engineering (MUE 2007), Seoul, Korea, ISBN: 0-7695-2777-9, pp. 952-957, DOI 10.1109/MUE.2007.84, April 2007, <<http://dl.acm.org/citation.cfm?id=1262690.1262986>>.

[federrath-fuchs-herrmann-piosecny]

Federrath, H., Fuchs, K., Herrmann, D., and C. Piosecny, "Privacy-Preserving DNS: Analysis of Broadcast, Range Queries and Mix-based Protection Methods", Computer Security ESORICS 2011, Springer, page(s) 665-683, ISBN 978-3-642-23821-5, 2011, <[https://svs.informatik.uni-hamburg.de/publications/2011/2011-09-14\\_FFHP\\_PrivacyPreservingDNS\\_ESORICS2011.pdf](https://svs.informatik.uni-hamburg.de/publications/2011/2011-09-14_FFHP_PrivacyPreservingDNS_ESORICS2011.pdf)>.

[getdns]

getdns, "getdns - A modern asynchronous DNS API", January 2020, <<https://getdnsapi.net>>.

[grangeia.snooping]

Grangeia, L., "DNS Cache Snooping or Snooping the Cache for Fun and Profit", 2005, <<https://www.semanticscholar.org/paper/Cache-Snooping-or-Snooping-the-Cache-for-Fun-and-1-Grangeia/9b22f606e10b3609eafbdc9090b63be8778c3>>.

[herrmann-reidentification]

Herrmann, D., Gerber, C., Banse, C., and H. Federrath, "Analyzing Characteristic Host Access Patterns for Re-Identification of Web User Sessions", DOI 10.1007/978-3-642-27937-9\_10, 2012, <[http://epub.uni-regensburg.de/21103/1/Paper\\_PUL\\_nordsec\\_published.pdf](http://epub.uni-regensburg.de/21103/1/Paper_PUL_nordsec_published.pdf)>.

- [I-D.ietf-dnsop-resolver-information]  
Sood, P., Arends, R., and P. Hoffman, "DNS Resolver Information Self-publication", draft-ietf-dnsop-resolver-information-01 (work in progress), February 2020.
- [I-D.ietf-dprive-bcp-op]  
Dickinson, S., Overeinder, B., Rijswijk-Deij, R., and A. Mankin, "Recommendations for DNS Privacy Service Operators", draft-ietf-dprive-bcp-op-14 (work in progress), July 2020.
- [I-D.ietf-dprive-dnssoquic]  
Huitema, C., Mankin, A., and S. Dickinson, "Specification of DNS over Dedicated QUIC Connections", draft-ietf-dprive-dnssoquic-01 (work in progress), October 2020.
- [I-D.ietf-quic-transport]  
Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", draft-ietf-quic-transport-34 (work in progress), January 2021.
- [morecowbell]  
Grothoff, C., Wachs, M., Ermert, M., and J. Appelbaum, "NSA's MORECOWBELL: Knell for DNS", GNUnet e.V., January 2015, <<https://pdfs.semanticscholar.org/2610/2b99bdd6a258a98740af8217ba8da8a1e4fa.pdf>>.
- [packetq] DNS-OARC, "PacketQ, a simple tool to make SQL-queries against PCAP-files", 2011, <<https://github.com/DNS-OARC/PacketQ>>.
- [passive-dns]  
Weimer, F., "Passive DNS Replication", April 2005, <<https://www.first.org/conference/2005/papers/florian-weimer-slides-1.pdf>>.
- [pitfalls-of-dns-encryption]  
Shulman, H., "Pretty Bad Privacy:Pitfalls of DNS Encryption", <<https://dl.acm.org/citation.cfm?id=2665959>>.
- [prism] Wikipedia, "PRISM (surveillance program)", July 2015, <[https://en.wikipedia.org/w/index.php?title=PRISM\\_\(surveillance\\_program\)&oldid=673789455](https://en.wikipedia.org/w/index.php?title=PRISM_(surveillance_program)&oldid=673789455)>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.

- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC4470] Weiler, S. and J. Ihren, "Minimally Covering NSEC Records and DNSSEC On-line Signing", RFC 4470, DOI 10.17487/RFC4470, April 2006, <<https://www.rfc-editor.org/info/rfc4470>>.
- [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, DOI 10.17487/RFC5155, March 2008, <<https://www.rfc-editor.org/info/rfc5155>>.
- [RFC5936] Lewis, E. and A. Hoenes, Ed., "DNS Zone Transfer Protocol (AXFR)", RFC 5936, DOI 10.17487/RFC5936, June 2010, <<https://www.rfc-editor.org/info/rfc5936>>.
- [RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <<https://www.rfc-editor.org/info/rfc6269>>.
- [RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <<https://www.rfc-editor.org/info/rfc6891>>.
- [RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/info/rfc7413>>.
- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<https://www.rfc-editor.org/info/rfc7525>>.
- [RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <<https://www.rfc-editor.org/info/rfc7624>>.

- [RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <<https://www.rfc-editor.org/info/rfc7721>>.
- [RFC7754] Barnes, R., Cooper, A., Kolkman, O., Thaler, D., and E. Nordmark, "Technical Considerations for Internet Service Blocking and Filtering", RFC 7754, DOI 10.17487/RFC7754, March 2016, <<https://www.rfc-editor.org/info/rfc7754>>.
- [RFC7816] Bortzmeyer, S., "DNS Query Name Minimisation to Improve Privacy", RFC 7816, DOI 10.17487/RFC7816, March 2016, <<https://www.rfc-editor.org/info/rfc7816>>.
- [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, <<https://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, <<https://www.rfc-editor.org/info/rfc7871>>.
- [RFC7873] Eastlake 3rd, D. and M. Andrews, "Domain Name System (DNS) Cookies", RFC 7873, DOI 10.17487/RFC7873, May 2016, <<https://www.rfc-editor.org/info/rfc7873>>.
- [RFC7929] Wouters, P., "DNS-Based Authentication of Named Entities (DANE) Bindings for OpenPGP", RFC 7929, DOI 10.17487/RFC7929, August 2016, <<https://www.rfc-editor.org/info/rfc7929>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.

- [RFC8744] Huitema, C., "Issues and Requirements for Server Name Identification (SNI) Encryption in TLS", RFC 8744, DOI 10.17487/RFC8744, July 2020, <<https://www.rfc-editor.org/info/rfc8744>>.
- [RFC8890] Nottingham, M., "The Internet is for End Users", RFC 8890, DOI 10.17487/RFC8890, August 2020, <<https://www.rfc-editor.org/info/rfc8890>>.
- [ripe-qname-measurements] Vries, W., "Making the DNS More Private with QNAME Minimisation", April 2019, <[https://labs.ripe.net/Members/wouter\\_de\\_vries/make-dns-a-bit-more-private-with-qname-minimisation](https://labs.ripe.net/Members/wouter_de_vries/make-dns-a-bit-more-private-with-qname-minimisation)>.
- [sidn-entrada] Hesselman, C., Jansen, J., Wullink, M., Vink, K., and M. Simon, "A privacy framework for 'DNS big data' applications", November 2014, <[https://www.sidnlabs.nl/downloads/yBW6hBoaSZe4m6GJc\\_0b7w/2211058ab6330c7f3788141ea19d3db7/SIDN\\_Labs\\_Privacyraamwerk\\_Position\\_Paper\\_V1.4\\_ENG.pdf](https://www.sidnlabs.nl/downloads/yBW6hBoaSZe4m6GJc_0b7w/2211058ab6330c7f3788141ea19d3db7/SIDN_Labs_Privacyraamwerk_Position_Paper_V1.4_ENG.pdf)>.
- [thomas-ditl-tcp] Thomas, M. and D. Wessels, "An Analysis of TCP Traffic in Root Server DITL Data", DNS-OARC 2014 Fall Workshop, October 2014, <<https://indico.dns-oarc.net/event/20/session/2/contribution/15/material/slides/1.pdf>>.
- [tor-leak] Tor, "DNS leaks in Tor", 2013, <<https://www.torproject.org/docs/faq.html.en#WarningsAboutSOCKSAndDNSInformationLeaks>>.
- [yanbin-tsudik] Yanbin, L. and G. Tsudik, "Towards Plugging Privacy Leaks in the Domain Name System", October 2009, <<http://arxiv.org/abs/0910.2472>>.

#### 14.3. URIs

- [1] <https://lists.dns-oarc.net/pipermail/dns-operations/2016-January/014143.html>
- [2] <http://netres.ec/?b=11B99BD>
- [3] <https://developers.google.com/speed/public-dns>

- [4] <https://developers.cloudflare.com/1.1.1.1/setting-up-1.1.1.1/>
- [5] <https://www.quad9.net>
- [6] <https://developers.google.com/speed/public-dns/privacy>
- [7] <https://www.alexa.com/topsites>
- [8] <https://theintercept.com/document/2014/03/12/nsa-gchqs-quantumtheory-hacking-tactics/>
- [9] <https://www.eugdpr.org/the-regulation.html>

#### Appendix A. Updates since RFC7626

Update many references; Added discussions of encrypted transports including DoT and DoH; Added section on DNS payload; Added section on authentication of servers; Added section on blocking of services. With the publishing of RFC7816 on QNAME minimisation, text, references, and initial attempts to measure deployment were added to reflect this. The text and references on the Snowden revelations were updated.

The "Risks overview" section was changed to "Scope" to help clarify the risks being considered. Text was adding on cellular network DNS, blocking and security. Considerations for recursive resolvers were collected and placed together. Added a discussion on resolver selection.

#### Appendix B. Changelog

draft-ietf-dprive-rfc7626-bis-08

- o Second batch of Editorial updates from IESG last call

draft-ietf-dprive-rfc7626-bis-07

- o First batch of Editorial updates from IESG last call

draft-ietf-dprive-rfc7626-bis-06

- o Removed Sara and Stephane as editors, made chairs as Editor.
- o Replaced the text in 6.1.1.2 with the text from the -04 version.
- o Clarified text about resolver selection in 6.1.1.

draft-ietf-dprive-rfc7626-bis-05

- o Editorial updates from second IESG last call
  - o Section renumbering as suggested by Vittorio Bertola
- draft-ietf-dprive-rfc7626-bis-04
- o Tsvart review: Add reference to DNS-over-QUIC, fix typo.
  - o Secdir review: Add text in Section 3 on devices using many networks.
  - o Update bullet in 3.4.1 on cellular encryption.
  - o Section 3.5.1.1 - re-work the section to try to address multiple comments.
  - o Section 3.5.1.4 - remove this section as now covered by 3.5.1.1.
  - o Section 3.5.1.5.2 - Remove several paragraphs and more directly reference RFC8484 by including bullet points quoting text from Section 8.2 of RFC8484. Retain the last 2 paragraphs as they are information for users, not implementors.
  - o Section 3.4.2 - some minor updates made based on specific comments.

draft-ietf-dprive-rfc7626-bis-03

- o Address 2 minor nits (typo in section 3.4.1 and adding an IANA section)
- o Minor updates from AD review

draft-ietf-dprive-rfc7626-bis-02

- o Numerous editorial corrections thanks to Mohamed Boucadair and
  - \* Minor additions to Scope section
  - \* New text on cellular network DNS
- o Additional text from Vittorio Bertola on blocking and security

draft-ietf-dprive-rfc7626-bis-01

- o Re-structure section 3.5 (was 2.5)
  - \* Collect considerations for recursive resolvers together



- \* Re-work several sections here to clarify their context (e.g., 'Rogue servers' becomes 'Active attacks on resolver configuration')
  - \* Add discussion of resolver selection
  - o Update text and old reference on Snowden revelations.
  - o Add text on and references to QNAME minimisation RFC and deployment measurements
  - o Correct outdated references
  - o Clarify scope by adding a Scope section (was Risks overview)
  - o Clarify what risks are considered in section 3.4.2
- draft-ietf-dprive-rfc7626-bis-00
- o Rename after WG adoption
  - o Use DoT acronym throughout
  - o Minor updates to status of deployment and other drafts
- draft-bortzmeyer-dprive-rfc7626-bis-02
- o Update various references and fix some nits.
- draft-bortzmeyer-dprive-rfc7626-bis-01
- o Update reference for dickinson-bcp-op to draft-dickinson-dprive-bcp-op
- draft-borztmeyer-dprive-rfc7626-bis-00:
- Initial commit. Differences to RFC7626:
- o Update many references
  - o Add discussions of encrypted transports including DoT and DoH
  - o Add section on DNS payload
  - o Add section on authentication of servers
  - o Add section on blocking of services

Author's Address

Tim Wicinski (editor)  
Elkins, WV 26241  
USA

Email: [tjw.ietf@gmail.com](mailto:tjw.ietf@gmail.com)

dprive  
Internet-Draft  
Intended status: Standards Track  
Expires: January 9, 2020

H. Zhang  
P. Aras  
Salesforce  
W. Toorop  
NLnet Labs  
S. Dickinson  
Sinodun IT  
A. Mankin  
Salesforce  
July 8, 2019

DNS Zone Transfer using DNS Stateful Operations  
draft-zatda-dprive-xfr-using-dso-00

Abstract

DNS zone transfers are transmitted in clear text, which gives attackers the opportunity to collect the content of a zone by eavesdropping on network connections. This document specifies use of DNS Stateful Operations to enable a subscribe/publish mechanism for zone transfers reducing the over head introduced by NOTIFY/SOA interactions prior to zone transfer request. This additionally prevents zone contents collection via passive monitoring of zone transfers by restricting XFR using DSO to require TLS.

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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 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 . . . . .                                  | 2  |
| 2. Terminology . . . . .                                   | 3  |
| 3. Use Cases for XFR-using-DSO . . . . .                   | 4  |
| 4. Overview . . . . .                                      | 4  |
| 5. Transport . . . . .                                     | 5  |
| 6. State Considerations . . . . .                          | 5  |
| 7. Protocol Operation . . . . .                            | 6  |
| 7.1. XuD SUBSCRIBE-XFR . . . . .                           | 6  |
| 7.1.1. SUBSCRIBE-XFR Request . . . . .                     | 7  |
| 7.1.2. SUBSCRIBE-XFR Response . . . . .                    | 9  |
| 7.2. XuD Notifications . . . . .                           | 12 |
| 7.2.1. DSO-IXFR Message . . . . .                          | 13 |
| 7.2.2. Fallback to AXFR . . . . .                          | 14 |
| 7.3. XuD UNSUBSCRIBE-XFR . . . . .                         | 15 |
| 7.3.1. UNSUBSCRIBE-XFR Message . . . . .                   | 15 |
| 7.4. Authentication . . . . .                              | 17 |
| 7.5. Multi-primary configurations . . . . .                | 17 |
| 7.6. DNS Stateful Operations TLV Context Summary . . . . . | 17 |
| 8. IANA Considerations . . . . .                           | 18 |
| 9. Implementation Considerations . . . . .                 | 18 |
| 10. Implementation Status . . . . .                        | 18 |
| 11. Security Considerations . . . . .                      | 19 |
| 12. Acknowledgements . . . . .                             | 19 |
| 13. Changelog . . . . .                                    | 19 |
| 14. References . . . . .                                   | 19 |
| 14.1. Normative References . . . . .                       | 19 |
| 14.2. Informative References . . . . .                     | 20 |
| 14.3. URIs . . . . .                                       | 20 |
| Authors' Addresses . . . . .                               | 20 |

## 1. Introduction

[I-D.hzpa-dprive-xfr-over-tls] enumerates the existing issues with clear text XFR mechanisms, outlines some use cases for using encrypted channels for zone transfer and also describes using TLS for zone transfers. It additionally discusses the various authentication

mechanisms that can be used to provide data and channel authentication, and channel confidentiality.

This draft describes the use of a DSO [RFC8490] based protocol to perform zone transfers. This mechanism is heavily based on an existing use of DSO where DNS clients can subscribe to receive asynchronous notifications of changes to RRSets of interest: DNS PUSH Notifications [I-D.ietf-dnssd-push]. That specification was developed with DNS Service Discovery in mind, this document describes an analogous protocol (XFR-using-DSO) where DNS clients can subscribe to receive asynchronous notifications of changes to zones of interest, it is developed with efficient and confidential zone transfers between primaries and secondaries in mind.

In the XFR-using-DSO model, a DSO connection is first opened between the client and server, the client can then subscribe to one or more zones to be notified of changes and the server can publish changes to the zone over the connection. Clients can choose to unsubscribe from zone updates at any time.

Servers could also use the DSO session to send command-style messages to the client, for example, to instruct a client to stop serving a zone or delete a zone. No such commands are defined in this version of the specification, but will likely be added in a future version.

## 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 BCP 14 [RFC2119] and [RFC8174] when, and only when, they appear in all capitals, as shown here.

Privacy terminology is as described in Section 3 of [RFC6973].

DNS terminology is as described in [RFC8499].

Note that in this document we choose to use the terms 'primary' and 'secondary' for two servers engaged in zone transfers.

DoT: DNS-over-TLS as specified in [RFC7858]

XuD: XFR-using-DOS mechanisms as specified in this document

### 3. Use Cases for XFR-using-DSO

This section includes additional use cases in addition to those specified in [I-D.hzpa-dprive-xfr-over-tls] that XuD can offer.

- o Confidentiality. Since this mechanism could, in principle, eliminate the need for NOTIFY and SOA queries it can provide complete confidentiality for the entire zone transfer mechanism.
- o Security. For some network configurations it is not desirable to have port 53 on the secondary open to an untrusted network for the sole purpose of receiving NOTIFYs. NOTIFYs can also be trivially spoofed unless secured with TSIG. For the DSO case, secondaries could initiate DSO connections to the primary and following that server-initiated DSO NOTIFY messages could be sent on that connection which could simultaneously be used for SOA and IXFR requests. This would allow a firewall to be restricted to just allowing outgoing connections from secondary to primary. Note that a similar but more constrained mechanism exists for IXFR whereby a short refresh period can be configured which triggers periodic SOA/IXFR requests from the secondary. TODO: Look at the details of the NSD implementation.
- o Performance. For the DSO case, a new subscribe/publish mechanism could be envisaged that greatly reducing the number of messages required to perform one transfer.
- o Improved error handling and retries. In the DSO case new explicit error codes could be defined that allow a server to indicate the reason for a failed or aborted XFR request. Also a new client initiated message could be used to gracefully cancel AXFRs.
- o New command channel. For the DSO case it would be possible to include new server-initiated 'control' commands e.g. 'stop serving this zone', 'delete this zone'.

QUESTION: Is there any case where the primary might want to initiate the DSO connection to the secondary?

### 4. Overview

The figure below provides an outline of the XuD protocol.

Figure 1: XuD protocol [1]

A DNS XuD client subscribes for zone notifications for a particular zone by connecting to the appropriate authoritative server for that zone, and sending DSO message(s) indicating the zone(s) of interest.

When the client loses interest in receiving further updates to these zones, it unsubscribes.

The authoritative server for a DNS zone is any server capable of generating the correct change notifications for a zone. It may be a primary, secondary, or stealth name server [RFC7719].

Standard DNS Queries MAY be sent over a XuD (i.e., DSO) session. For any zone for which the server is authoritative, it MUST respond authoritatively for queries on names falling within that zone both for normal DNS queries and for XuD subscriptions. For names for which the server is acting as a recursive resolver, e.g. when the server is the local recursive resolver, for any query for which it supports XuD subscriptions, it MUST also support standard queries.

XuD imposes less load on the responding server than rapid polling would, but XuD notifications do still have a cost, so XuD clients MUST only create XuD subscriptions for zones they are authorised to transfer.

Generally, as described in the DNS Stateful Operations specification [RFC8490], a client must not keep a session to a server open indefinitely if it has no subscriptions (or other operations) active on that session. A client MAY close a session as soon as it becomes idle, and then if needed in the future, open a new session when required. Alternatively, a client MAY speculatively keep an idle session open for some time, subject to the constraint that it MUST NOT keep a session open that has been idle for more than the session's idle timeout (15 seconds by default) [RFC8490].

## 5. Transport

XuD clients MUST use DNS Stateful Operations [RFC8490] running over TLS over TCP [RFC7858].

The connection for XuD SHOULD be established using port 853, as specified in [RFC7858], unless there is mutual agreement between the secondary and primary to use a port other than port 853 for XuD.

QUESTION: Is there a use case to allow XuD over TCP where confidentiality is not an issue e.g when the zone contents are already publicly available?

## 6. State Considerations

Each XuD server is capable of handling some finite number of XuD subscriptions. This number will vary from server to server and is based on physical machine characteristics, network bandwidth, and

operating system resource allocation. After a client establishes a session to a DNS server, each subscription is individually accepted or rejected. Servers may employ various techniques to limit subscriptions to a manageable level. Correspondingly, the client is free to establish simultaneous sessions to alternate DNS servers that support XuDs for the zone and distribute subscriptions at the client's discretion. In this way, both clients and servers can react to resource constraints.

## 7. Protocol Operation

The XuD protocol is a session-oriented protocol, and makes use of DNS Stateful Operations (DSO) [RFC8490].

For details of the DSO message format refer to the DNS Stateful Operations specification [RFC8490]. Those details are not repeated here.

XuD clients and servers MUST support DSO. A single server can support DNS Queries, DNS Updates, and XuD (using DSO) on the same TCP port.

A XuD exchange begins with the client making a TLS/TCP connection to the appropriate server.

A typical XuD client will immediately issue a DSO Keepalive operation to request a session timeout and/or keepalive interval longer than the the 15-second default values, but this is not required. A XuD client MAY issue other requests on the session first, and only issue a DSO Keepalive operation later if it determines that to be necessary. Sending either a DSO Keepalive operation or a XuD subscription over the TLS/TCP connection to the server signals the client's support of DSO and serves to establish a DSO session.

In accordance with the current set of active subscriptions, the server sends relevant asynchronous XuD notifications to the client. Note that a client MUST be prepared to receive (and silently ignore) XuD notifications for subscriptions it has previously removed, since there is no way to prevent the situation where a XuD notification is in flight from server to client while the client's unsubscribe message cancelling that subscription is simultaneously in flight from client to server.

### 7.1. XuD SUBSCRIBE-XFR

After connecting, and requesting a longer idle timeout and/or keepalive interval if necessary, a XuD client then indicates its desire to receive XuD notifications for a given zone by sending a



SUBSCRIBE-XFR request to the server. A SUBSCRIBE-XFR request is encoded in a DSO message [RFC8490]. This specification defines a primary DSO TLV for XuD SUBSCRIBE-XFR Requests (tentatively DSO Type Code 0x50).

DSO messages with the SUBSCRIBE-XFR TLV as the Primary TLV are not permitted in early data.

The entity that initiates a SUBSCRIBE-XFR request is by definition the client. A server MUST NOT send a SUBSCRIBE-XFR request over an existing session from a client. If a server does send a SUBSCRIBE-XFR request over a DSO session initiated by a client, this is a fatal error and the client should immediately abort the connection with a TLS close\_notify alert. See Section 6.1 of [RFC8446].

TODO: Need to define a DSO version of TSIG to cover the SUBSCRIBE-XFR and DSO-XFR responses, since the Additional section count in DSO message MUST be zero. Note the client only needs to use TSIG in the SUBSCRIBE-XFR message to prove it is authorised to request zone transfers, but all DSO-XFR messages should be signed if primary TSIG is required for the authentication model in use.

#### 7.1.1. SUBSCRIBE-XFR Request

A SUBSCRIBE-XFR request begins with the standard DSO 12-byte header [RFC8490], followed by the SUBSCRIBE-XFR primary TLV. A SUBSCRIBE-XFR request message is illustrated in Figure 2.

The MESSAGE ID field MUST be set to a unique value, that the client is not using for any other active operation on this DSO session. For the purposes here, a MESSAGE ID is in use on this session if the client has used it in a request for which it has not yet received a response, or if the client has used it for a subscription which it has not yet cancelled using UNSUBSCRIBE-XFR. In the SUBSCRIBE-XFR response the server MUST echo back the MESSAGE ID value unchanged.

The other header fields MUST be set as described in the DSO specification [RFC8490]. The DNS OPCODE field contains the OPCODE value for DNS Stateful Operations (6). The four count fields MUST be zero, and the corresponding four sections MUST be empty (i.e., absent).

The DSO-TYPE is SUBSCRIBE-XFR (tentatively 0x50).

The DSO-LENGTH is the length of the DSO-DATA that follows, which specifies the name and class of the zone and optionally the SOA value of the client's version of the zone.

If the client has no copy of the zone it MUST omit the SOA value to indicate to the server that a DSO-AXFR is required in response (see the next section).

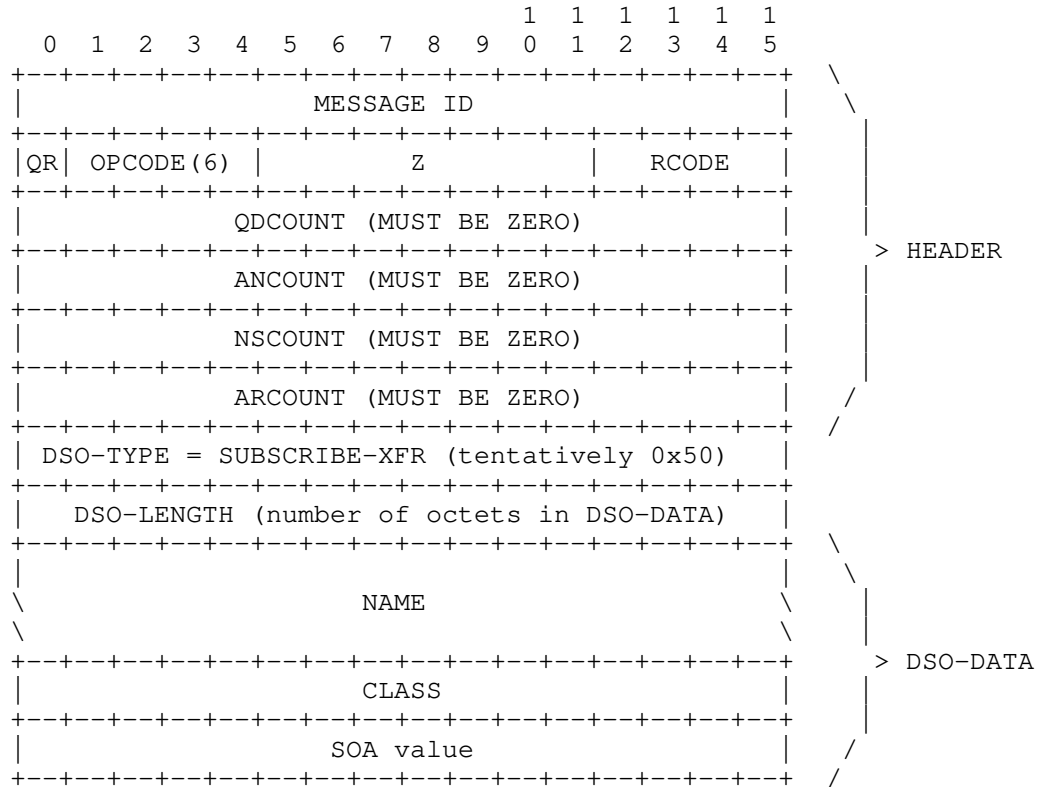


Figure 2: SUBSCRIBE-XFR Request

The DSO-DATA for a SUBSCRIBE-XFR request MUST contain exactly one NAME, CLASS and SOA value. Since SUBSCRIBE-XFR requests are sent over TCP, multiple SUBSCRIBE-XFR DSO request messages can be concatenated in a single TCP stream and packed efficiently into TCP segments.

If accepted, the subscription will stay in effect until the client cancels the subscription using UNSUBSCRIBE-XFR or until the DSO session between the client and the server is closed.

SUBSCRIBE-XFR requests on a given session MUST be unique. A client MUST NOT send a SUBSCRIBE-XFR message that duplicates the NAME, CLASS and SOA value of an existing active subscription on that DSO session. For the purpose of this matching, the established DNS case-

insensitivity for US-ASCII letters applies (e.g., "example.com" and "Example.com" are the same). If a server receives such a duplicate SUBSCRIBE-XFR message this is an error and the server MUST immediately terminate the connection with a TLS close\_notify alert.

QUESTION: Is there a use case where a client may want to signal that the version of the zone it holds has been updated via another mechanism and the zone transfer should restart from a different SOA than that currently exchanged between client and server?

DNS wildcarding is not supported. SUBSCRIBE-XFR requests received for zones containing wildcards are considered an error (see below).

A CLASS of 'ANY' (255) is not supported.

#### 7.1.2. SUBSCRIBE-XFR Response

Each SUBSCRIBE-XFR request generates exactly one SUBSCRIBE-XFR response from the server. A SUBSCRIBE-XFR request message is illustrated in Figure 3.

A SUBSCRIBE-XFR response begins with the standard DSO 12-byte header [RFC8490]. The QR bit in the header is set indicating it is a response. The header MAY be followed by one or more optional TLVs, such as a Retry Delay TLV.

The MESSAGE ID field MUST echo the value given in the Message ID field of the SUBSCRIBE-XFR request. This is how the client knows which request is being responded to.

A SUBSCRIBE-XFR response message MUST NOT include a SUBSCRIBE-XFR TLV. If a client receives a SUBSCRIBE-XFR response message containing a SUBSCRIBE-XFR TLV then the response message is processed but the SUBSCRIBE-XFR TLV MUST be silently ignored.

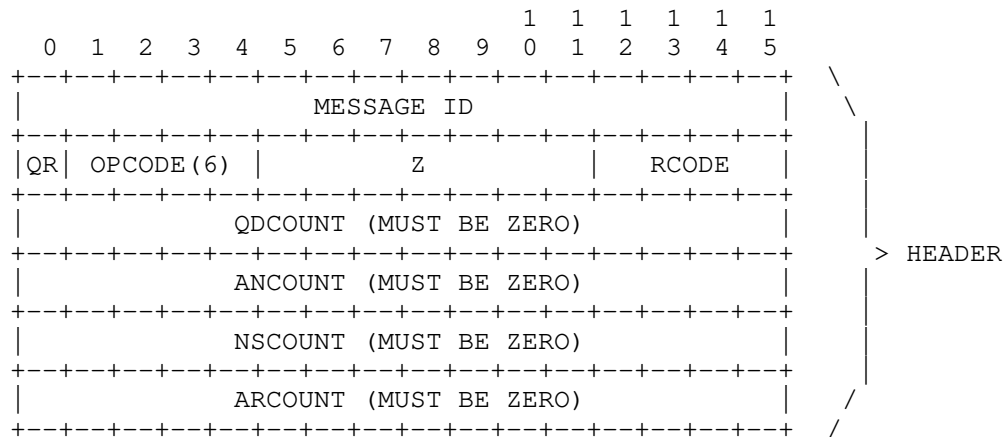


Figure 3: SUBSCRIBE-XFR Response Message

In the SUBSCRIBE-XFR response the RCODE indicates whether or not the subscription was accepted. Supported RCODEs are as follows:

| Mnemonic  | Value | Description                                                        |
|-----------|-------|--------------------------------------------------------------------|
| NOERROR   | 0     | SUBSCRIBE-XFR successful.                                          |
| FORMERR   | 1     | Server failed to process request due to a malformed request.       |
| SERVFAIL  | 2     | Server failed to process request due to a problem with the server. |
| NOTIMP    | 4     | Server does not implement DSO.                                     |
| REFUSED   | 5     | Server refuses to process request for policy or security reasons.  |
| NOTAUTH   | 9     | Server is not authoritative for the requested name.                |
| DSOTYPENI | 11    | SUBSCRIBE-XFR operation not supported.                             |

Table 1: SUBSCRIBE-XFR Response codes

This document specifies only these RCODE values for SUBSCRIBE-XFR Responses. Servers sending SUBSCRIBE-XFR Responses SHOULD use one of these values. Note that NXDOMAIN is not a valid RCODE in response to a SUBSCRIBE-XFR Request. However, future circumstances may create situations where other RCODE values are appropriate in SUBSCRIBE-XFR Responses, so clients MUST be prepared to accept SUBSCRIBE-XFR Responses with any other RCODE value.

If the server sends a nonzero RCODE in the SUBSCRIBE-XFR response, that means:

- a the client is (at least partially) misconfigured,
- b the server resources are exhausted, or
- c there is some other unknown failure on the server.

In any case, the client shouldn't retry the subscription to this server right away. If a client has other authoritative servers configured for a given zone an alternative server can be tried immediately.

If the client has other successful subscriptions to this server, these subscriptions remain even though additional subscriptions may be refused. Neither the client nor the server are required to close the connection, although, either end may choose to do so.

If the server sends a nonzero RCODE then it SHOULD append a Retry Delay TLV [RFC8490] to the response specifying a delay before the client attempts this operation again. Recommended values for the delay for different RCODE values are given below. These recommended values apply both to the default values a server should place in the Retry Delay TLV, and the default values a client should assume if the server provides no Retry Delay TLV.

For RCODE = 1 (FORMERR) the delay may be any value selected by the implementer. A value of five minutes is RECOMMENDED, to reduce the risk of high load from defective clients.

For RCODE = 2 (SERVFAIL) the delay should be chosen according to the level of server overload and the anticipated duration of that overload. By default, a value of one minute is RECOMMENDED. If a more serious server failure occurs, the delay may be longer in accordance with the specific problem encountered.

For RCODE = 4 (NOTIMP), which occurs on a server that doesn't implement DNS Stateful Operations [RFC8490], it is unlikely that the server will begin supporting DSO in the next few minutes, so the retry delay SHOULD be one hour. Note that in such a case, a server that doesn't implement DSO is unlikely to place a Retry Delay TLV in its response, so this recommended value in particular applies to what a client should assume by default.

For RCODE = 5 (REFUSED), which occurs on a server that implements XuDs, but is currently configured to disallow XuDs, the retry delay may be any value selected by the implementer and/or configured by the

operator. Since it is possible that the misconfiguration may be repaired at any time, the retry delay should not be set too high. By default, a value of 5 minutes is RECOMMENDED.

For RCODE = 9 (NOTAUTH), which occurs on a server that implements XuDs, but is not configured to be authoritative for the requested name, the retry delay may be any value selected by the implementer and/or configured by the operator. Since it is possible that the misconfiguration may be repaired at any time, the retry delay should not be set too high. By default, a value of 5 minutes is RECOMMENDED.

For RCODE = 11 (DSOTYPENI), which occurs on a server that implements DSO but doesn't implement XuD, it is unlikely that the server will begin supporting XuD in the next few minutes, so the retry delay SHOULD be one hour.

For other RCODE values, the retry delay should be set by the server as appropriate for that error condition. By default, a value of 5 minutes is RECOMMENDED.

For RCODE = 9 (NOTAUTH), the time delay applies to requests for other names falling within the same zone. Requests for names falling within other zones are not subject to the delay. For all other RCODEs the time delay applies to all subsequent requests to this server.

After sending an error response the server MAY allow the session to remain open, or MAY send a Retry Delay Operation TLV instructing the client to close the session, as described in the DSO specification [RFC8490]. Clients MUST correctly handle both cases.

## 7.2. XuD Notifications

Once a subscription has been successfully established, the server generates DSO-IXFR messages to send to the client as appropriate. In the case that the server could not provide a DSO-IXFR message based on the SOA received from the client an initial DSO-AXFR message will be sent immediately following the SUBSCRIBE-XFR Response. Subsequent changes to the zone are then communicated to the client in subsequent DSO-IXFR messages.

Until an UNSUBSCRIBE-XFR message is received the server MUST assume that the client is updating the client's version of the zone with the notifications sent and can therefore hold state on the SOA version the client holds. It MUST use this to generate the DSO-IXFR messages sent on a XuD session.

### 7.2.1. DSO-IXFR Message

A DSO-IXFR unidirectional message begins with the standard DSO 12-byte header [RFC8490], followed by the DSO-IXFR primary TLV. A DSO-IXFR message is illustrated in Figure 4.

In accordance with the definition of DSO unidirectional messages, the MESSAGE ID field MUST be zero. There is no client response to a DSO-IXFR message.

The other header fields MUST be set as described in the DSO specification [RFC8490]. The DNS OPCODE field contains the OPCODE value for DNS Stateful Operations (6). The four count fields MUST be zero, and the corresponding four sections MUST be empty (i.e., absent).

The DSO-TYPE is DSO-IXFR (tentatively 0x51).

The DSO-LENGTH is the length of the DSO-DATA that follows, which specifies the changes being communicated.

The DSO-DATA contains one or more change notifications. A DSO-IXFR Message MUST contain at least one change notification. If a DSO-IXFR Message is received that contains no change notifications, this is a fatal error, and the receiver MUST immediately terminate the connection with a TLS close\_notify alert.

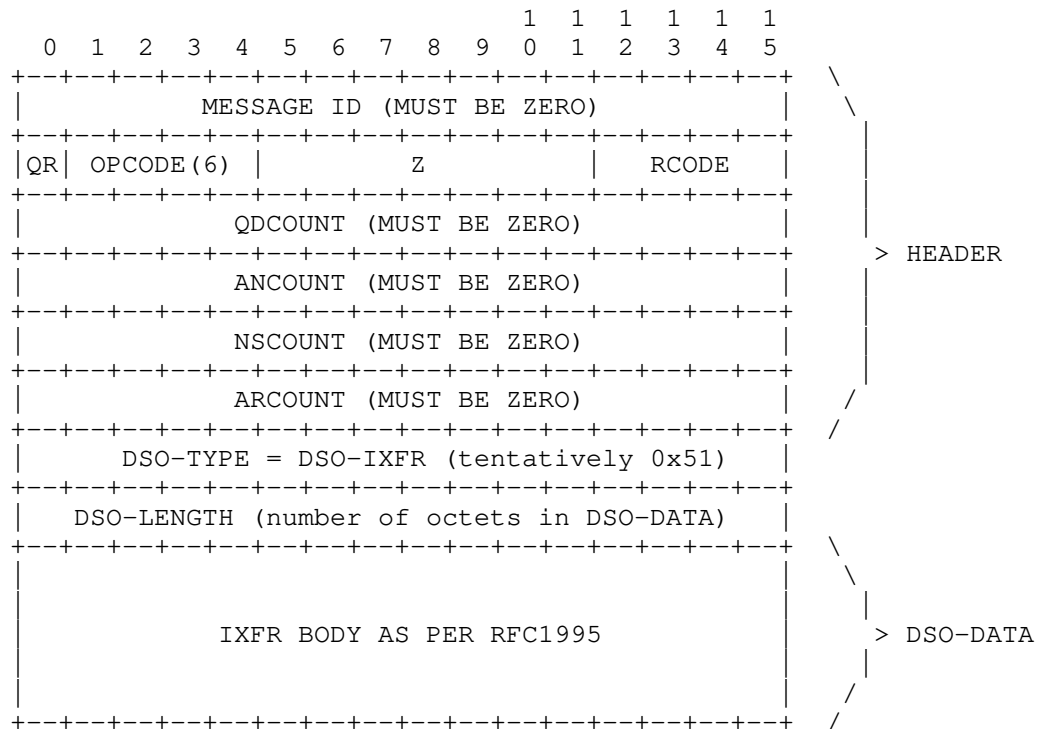


Figure 4: DSO-IXFR Message

The DSO-DATA in a DSO-IXFR message is identical to the contents of a [RFC1995] IXFR message that would be sent to communicate the same zone incremental zone transfer over UDP or TCP i.e. the set of one or more difference sequences that follow the DNS Header in an IXFR message.

When processing the records received in a DSO-IXFR Message, the receiving client MUST validate that the zone being updated correspond with at least one currently active subscription on that session. Specifically, the SOA name and CLASS MUST match the SOA name and CLASS given in a SUBSCRIBE-IXFR request, subject to the usual established DNS case-insensitivity for US-ASCII letters.

### 7.2.2. Fallback to AXFR

The format of the DSO-AXFR message is a standard DSO header with DSO-TYPE of DSO-AXFR (tentatively DSO Type Code 0x52) and the body is identical to a [RFC5936] AXFR response body.

TODO: More detail here.



If the SUBSCRIBE-XFR message contained no SOA value, the server MUST send a DSO-AXFR message as its first message on the connection.

Alternatively if incremental zone transfer is not available, the entire zone MAY be returned in a DSO-AXFR message.

QUESTION: Should we bother with a separate DSO-AXFR message or just allow full zone transfer inside the DSO-IXFR message as with [RFC1995] IXFR? A separate message type makes is more explicit and IXFR was constrained by having to respond to a IXFR request.

### 7.3. XuD UNSUBSCRIBE-XFR

To cancel an individual subscription without closing the entire DSO session, the client sends an UNSUBSCRIBE-XFR message over the established DSO session to the server. The UNSUBSCRIBE-XFR message is encoded as a DSO unidirectional message [RFC8490]. This specification defines a primary unidirectional DSO TLV for XuD UNSUBSCRIBE-XFR Messages (tentatively DSO Type Code 0x53).

A server MUST NOT initiate an UNSUBSCRIBE-XFR message. If a server does send an UNSUBSCRIBE-XFR message over a DSO session initiated by a client, this is a fatal error and the client should immediately abort the connection with a TLS close\_notify alert.

#### 7.3.1. UNSUBSCRIBE-XFR Message

An UNSUBSCRIBE-XFR unidirectional message begins with the standard DSO 12-byte header [RFC8490], followed by the UNSUBSCRIBE-XFR primary TLV. An UNSUBSCRIBE-XFR message is illustrated in Figure 5.

In accordance with the definition of DSO unidirectional messages, the MESSAGE ID field MUST be zero. There is no server response to an UNSUBSCRIBE-XFR message.

The other header fields MUST be set as described in the DSO specification [RFC8490]. The DNS OPCODE field contains the OPCODE value for DNS Stateful Operations (6). The four count fields MUST be zero, and the corresponding four sections MUST be empty (i.e., absent).

The DSO-TYPE is UNSUBSCRIBE-XFR (tentatively 0x53).

The DSO-LENGTH field contains the value 2, the length of the 2-octet MESSAGE ID contained in the DSO-DATA.

The DSO-DATA contains the value given in the MESSAGE ID field of an active SUBSCRIBE-XFR request. This is how the server knows which

SUBSCRIBE-XFR request is being cancelled. After receipt of the UNSUBSCRIBE-XFR message, the SUBSCRIBE-XFR request is no longer active.

It is allowable for the client to issue an UNSUBSCRIBE-XFR message for a previous SUBSCRIBE-XFR request for which the client has not yet received a SUBSCRIBE-XFR response. This is to allow for the case where a client starts and stops a subscription in less than the round-trip time to the server. The client is NOT required to wait for the SUBSCRIBE-XFR response before issuing the UNSUBSCRIBE-XFR message.

Consequently, it is possible for a server to receive an UNSUBSCRIBE-XFR message that does not match any currently active subscription. This can occur when a client sends a SUBSCRIBE-XFR request, which subsequently fails and returns an error code, but the client sent an UNSUBSCRIBE-XFR message before it became aware that the SUBSCRIBE-XFR request had failed. Because of this, servers MUST silently ignore UNSUBSCRIBE-XFR messages that do not match any currently active subscription.

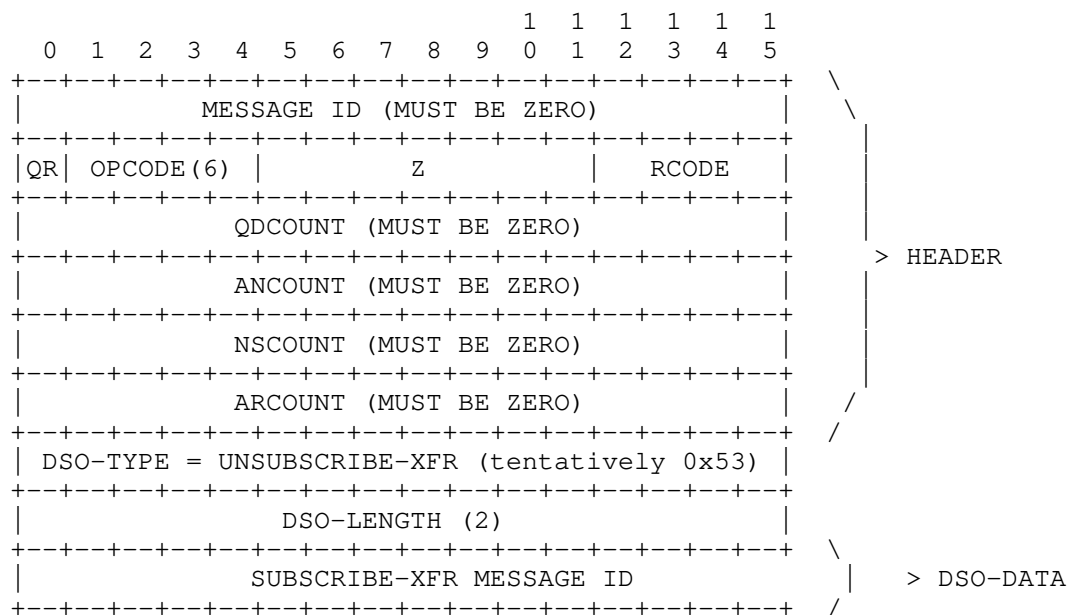


Figure 5: UNSUBSCRIBE-XFR Message

QUESTION: Do we need the equivalent of a RECONFIRM message from DNS PUSH Notifications [I-D.ietf-dnssd-push]?

#### 7.4. Authentication

The authentication considerations are largely the same as those presented in [I-D.hzpa-dprive-xfr-over-tls].

#### 7.5. Multi-primary configurations

The multi-primary considerations share some of the same issues as those presented in [I-D.hzpa-dprive-xfr-over-tls] but are different because the client is not performing SOA queries.

TODO: More detail required here.

#### 7.6. DNS Stateful Operations TLV Context Summary

This document defines four new DSO TLVs. As suggested in Section 8.2 of the DNS Stateful Operations specification [RFC8490], the valid contexts of these new TLV types are summarized below.

The client TLV contexts are:

C-P: Client request message, primary TLV

C-U: Client unidirectional message, primary TLV

C-A: Client request or unidirectional message, additional TLV

CRP: Response back to client, primary TLV

CRA: Response back to client, additional TLV

| TLV Type        | C-P | C-U | C-A | CRP | CRA |
|-----------------|-----|-----|-----|-----|-----|
| SUBSCRIBE-XFR   | X   |     |     |     |     |
| DSO-IXFR        |     |     |     |     |     |
| DSO-AXFR        |     |     |     |     |     |
| UNSUBSCRIBE-XFR |     | X   |     |     |     |

Table 2: DSO TLV Client Context Summary

The server TLV contexts are:

S-P: Server request message, primary TLV

S-U: Server unidirectional message, primary TLV

S-A: Server request or unidirectional message, additional TLV

SRP: Response back to server, primary TLV

SRA: Response back to server, additional TLV

| TLV Type        | S-P | S-U | S-A | SRP | SRA |
|-----------------|-----|-----|-----|-----|-----|
| SUBSCRIBE-XFR   |     |     |     |     |     |
| DSO-IXFR        |     | X   |     |     |     |
| DSO-AXFR        |     | X   |     |     |     |
| UNSUBSCRIBE-XFR |     |     |     |     |     |

Table 3: DSO TLV Server Context Summary

## 8. IANA Considerations

This document also defines four new DNS Stateful Operation TLV types to be recorded in the IANA DSO Type Code Registry.

| Name            | Value         | Early Data | Status             | Definition     |
|-----------------|---------------|------------|--------------------|----------------|
| SUBSCRIBE-XFR   | TBA<br>(0x50) | NO         | Standards<br>Track | Section<br>7.1 |
| DSO-IXFR        | TBA<br>(0x51) | NA         | Standards<br>Track | Section<br>7.1 |
| DSO-AXFR        | TBA<br>(0x51) | NA         | Standards<br>Track | Section<br>7.2 |
| UNSUBSCRIBE-XFR | TBA<br>(0x52) | NA         | Standards<br>Track | Section<br>7.2 |

Table 5: IANA DSO TLV Type Code Assignment

## 9. Implementation Considerations

TBD

## 10. Implementation Status

TBD

## 11. Security Considerations

This document specifies a security measure against a DNS risk: the risk that an attacker collects entire DNS zones through eavesdropping on clear text DNS zone transfers. It presents a new Security Consideration for DNS. Some questions to discuss are:

- o Should DoT in this new case be required to use only TLS 1.3 and higher to avoid residual exposure?
- o How should padding be used in IXFR?
- o Should there be an option to 'pad' an AXFR response (i.e. a set of AXFR responses on a given connection) to hide the zone size?

## 12. Acknowledgements

## 13. Changelog

draft-zatda-dprive-xfr-using-dso-00

- o Initial commit

## 14. References

### 14.1. Normative References

- [I-D.hzpa-dprive-xfr-over-tls]  
Zhang, H., Aras, P., Toorop, W., Dickinson, S., and A. Mankin, "DNS Zone Transfer over TLS", draft-hzpa-dprive-xfr-over-tls-01 (work in progress), March 2019.
- [I-D.ietf-dnssd-push]  
Pusateri, T. and S. Cheshire, "DNS Push Notifications", draft-ietf-dnssd-push-21 (work in progress), July 2019.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.

- [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, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.

#### 14.2. Informative References

- [RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, DOI 10.17487/RFC1995, August 1996, <<https://www.rfc-editor.org/info/rfc1995>>.
- [RFC5936] Lewis, E. and A. Hoenes, Ed., "DNS Zone Transfer Protocol (AXFR)", RFC 5936, DOI 10.17487/RFC5936, June 2010, <<https://www.rfc-editor.org/info/rfc5936>>.
- [RFC8490] Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S., Lemon, T., and T. Pusateri, "DNS Stateful Operations", RFC 8490, DOI 10.17487/RFC8490, March 2019, <<https://www.rfc-editor.org/info/rfc8490>>.

#### 14.3. URIs

- [1] [https://github.com/Sinodun/draft-xfr-using-dso/blob/master/draft-01-svg/XuD\\_Protocol.svg](https://github.com/Sinodun/draft-xfr-using-dso/blob/master/draft-01-svg/XuD_Protocol.svg)

#### Authors' Addresses

Han Zhang  
Salesforce  
San Francisco, CA  
United States

Email: [h Zhang@salesforce.com](mailto:h Zhang@salesforce.com)

Pallavi Aras  
Salesforce  
Herndon, VA  
United States

Email: [paras@salesforce.com](mailto:paras@salesforce.com)

Willem Toorop  
NLnet Labs  
Science Park 400  
Amsterdam 1098 XH  
The Netherlands

Email: [willem@nlnetlabs.nl](mailto:willem@nlnetlabs.nl)

Sara Dickinson  
Sinodun IT  
Magdalen Centre  
Oxford Science Park  
Oxford OX4 4GA  
United Kingdom

Email: [sara@sinodun.com](mailto:sara@sinodun.com)

Allison Mankin  
Salesforce  
Herndon, VA  
United States

Email: [allison.mankin@gmail.com](mailto:allison.mankin@gmail.com)