                 Evaluation of Privacy for DNS Private Exchange
                        draft-am-dprive-eval-02

Abstract

   The set of DNS requests that an individual makes can provide a
   monitor with a large amount of information about that individual.
   DNS Private Exchange (DPRIVE) aims to deprive this actor of this
   information. This document describes methods for measuring the
   performance of DNS privacy mechanisms, particularly it provides
   methods for measuring effectiveness in the face of pervasive
   monitoring as defined in RFC7258. The document includes example
   evaluations for common use cases.

Status of This Memo

Copyright Notice

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. Motivation

   One of the IETF's core views is that protocols should be designed to
   enable security and privacy while online [RFC3552]. In light of the
   recent reported pervasive monitoring efforts, another goal is to
   design protocols and mechanisms to make such monitoring expensive or
   infeasible to conduct. As detailed in the DPRIVE problem statement
   [RFC7626], DNS resolution is an important arena for pervasive
   monitoring, and in some cases may be used for breaching the privacy
   of individuals. The set of DNS requests that an individual makes can
   provide a large amount of information about that individual. In some
   specific use cases, the sets of DNS requests from a DNS recursive
   resolver or other entity may also provide revealing information.
   This document describes methods for measuring the performance of DNS
   privacy mechanisms; in particular, it provides methods for measuring
   effectiveness in the face of pervasive monitoring as defined in
   [RFC7258]. The document includes example evaluations for common use
   cases.

The privacy threats associated with DNS monitoring are not new,
however they were made more obvious by the issue described in
[RFC7258]. The DPRIVE working group was formed to respond and at
this time has several DNS private exchange mechanisms in
consideration, including [dns-over-tls], [confidential-dns],
[phb-dnse], and [privatedns]. There is also related work in other
working groups, including DNSOP: [qname-minimisation] and
(potentially) DANE [ipseca]. The recently published RFC on
opportunistic security [RFC7435] also has relevance to DNS private
exchange.

Each effort related to DNS privacy mechanisms asserts some privacy
assurances and operational relevance. Metrics for these privacy
assurances are needed and are in reach based on existing techniques
from the general field of privacy engineering. Systematic evaluation
of DNS privacy mechanisms will enhance the likely operational
effectiveness of DNS private exchange.

Evaluating an individual mechanism for DNS privacy could be
accomplished on a one-off basis, presumably as Privacy Considerations
within each specification, but this will not address as much
variation of operational contexts nor will it cover using multiple
mechanisms together (in composition). Section 2 of [RFC6973]
discussed both benefits and risks of using multiple mechanisms.

Definitions required for evaluating the privacy of stand-alone and
composed design are not limited to privacy notions, but also need to
include the risk model and some information about relationships among
the entities in a given system. A mechanism for providing privacy to
withstand the power and capabilities of a passive pervasive monitor
may not withstand a more powerful actor using active monitoring by
plugging itself into the path of individuals' DNS requests as a
forwarder . Having some standard models, and understanding how
applicable they are to various designs is a part of evaluating the
privacy.

Sections 2 and 3 present privacy terminology and some assumptions.
Sections 4 and 5 cover the system model or setup and the risk models
of interest. In Section 6, we review a list of DNS privacy
mechanisms, including some which are not in scope of the DPRIVE
working group. Section 7 tackles how to evaluate privacy mechanisms,
in the form of templates and outcomes. Given a specific risk model,
the guarantees with respect to privacy of an individual or an item of
interest are quantified.

[2](#). **Privacy Evaluation Definitions**

   This section provides definitions to be used for privacy evaluation
   of DNS. The verbatim source of most of those definitions is from
   [[RFC6973](#)], which are included as an aid in reasability. In some
   definitions, text is added to apply them to the DNS case. Also, a
   new section of terms has been added to include several important
   practical or conventional terms that were not included in [[RFC6973](#)]
   such as PII. For the terms from [[RFC6973](#)], we include their
   definitions rather than simply referencing them as an aid to
   readability.

[2.1](#). **Entities**

   o Attacker: An entity that works against one or more privacy
      protection goals. Unlike observers, attackers' behavior is
      unauthorized, in a way similar to that of an eavesdropper.

   o Eavesdropper: A type of attacker that passively observes an
      initiator's communications without the initiator's knowledge or
      authorization. DNS example may include a passive pervasive
      monitor, defined below.

   o Enabler: A protocol entity that facilitates communication between
      an initiator and a recipient without being directly in the
      communications path. DNS examples of an enabler in this sense
      include a recursive resolver, a proxy, or a forwarder.

   o Individual: A human being (or a group of them)

   o Initiator: A protocol entity that initiates communications with a
      recipient.

   o Intermediary: A protocol entity that sits between the initiator
      (DNS example of stub resolver) and the recipient (DNS example of
      recursive resolver or authority resolver) and is necessary for the
      initiator and recipient to communicate. Unlike an eavesdropper,
      an intermediary is an entity that is part of the communication
      architecture and therefore at least tacitly authorized.

   o Observer: An entity that is able to observe and collect
      information from communications, potentially posing privacy risks,
      depending on the context. As defined in this document,
      initiators, recipients, intermediaries, and enablers can all be
      observers. Observers are distinguished from eavesdroppers by
      being at least tacitly authorized.

o We note that while the definition of an observer may include an
   initiator in the risk model, an initiator of a request is excluded
   in the context of this document, because it corresponds to the
   subject of interest being studied. Similar to the definition in
   [RFC7258], we note that an attacker is broader than an observer.
   While [RFC7258] claim that an attack does not consider the motive
   of the actor, the given context of DNS implies a motive if the
   term attacker is used to characterize the risk.

2.2. Data and Analysis

We assume the following definitions related to data and analysis from
[RFC4949]: attacker, correlation, fingerprint, fingerprinting, item
of interest (IOI), personal data, interaction, traffic analysis,
undetectability, and unlinkability. We augment some of those
definitions later in this document.

from [RFC4949], we relax the definition of IOI to exclude "the fact
that a communication interaction has taken place" as this does not
suite the evaluated context of DNS.

2.3. Identifiability

We assume the following definitions related to identifiability from
[RFC4949]: anonymity, anonymity set, anonymous, attribute, identity
provider, personal name, and relying party.

The following definitions are modified for the context of this
document from those defined in [RFC4949]

o Identifiability: The extent to which an individual is
   identifiable. [RFC6973] includes the rest of the variations on
   this (Identifiable, Identification, Identified, Identifier,
   Identity, Identity Confidentiality)

o Personal Name: A natural name for an individual. Personal names
   are often not unique and often comprise given names in combination
   with a family name. An individual may have multiple personal
   names at any time and over a lifetime, including official names.
   From a technological perspective, it cannot always be determined
   whether a given reference to an individual is, or is based upon,
   the individual's personal name(s) (see Pseudonym). NOTE: The
   reason to import this definition is that some query names that
   cause privacy leakage do so by embedding personal names as
   identifiers of host or other equipment, e.g.
   AllisonMankinMac.example.com.

o Pseudonymity: See the formal definition in section 2.4 in lieu of
   [RFC6973].

NOTE: Identifiability Definitions in [RFC6973] also include some
material not included here because the distinctions are not major for
DNS Private Exchange, such as real and official names, and variant
forms of Pseudonymity in its informal definition.

2.4. Other Central Definitions and Formalizations

Central to the presentation of this document is the definition of
personally identifiable information (PII), as well as other
definitions that supplement the definitions listed earlier or modify
them for the context of this document. In this section, we outline
such definitions we further notes on their indications.

o Personally Identifiable Information (PII): Information
   (attributes) that can be used as is, or along with other side
   information, to identify, locate, and/or contact a single
   individual or subject (c.f. item of interest).

NOTE: the definition above indicates that PII can be used on its own
or in context. In DNS privacy, the items without additional context
include IP(v4 or v6) addresses, qname, qtype, timings of queries,
etc. The additional context includes organization-level attributes,
such as a network prefix that can be associated with an organization.
The definition of PII is complementary to the definition of items of
interest.

o Subject: This term is useful as a parallel term to Individual.
   When the privacy of a group or an organization is of interest, we
   can reference the group or organization as Subject rather than
   Individual.

Often it is desirable to reference alternative identifiers known as
pseudonyms. A pseudonym is a name assumed by an individual in some
context, unrelated to the names or identifiers known by others in
that context.

o Pseudonymity/Pseudonym: a relaxation of the definition of
   anonymity for usability. In particular, pseudonymity is an
   anonymity feature obtained by using a pseudonym, an identifier
   that is used for establishing a long relationship between two
   entities.

As an example, in the DNS context, a randomly generated pseudonym
might identify a set of query data with a shared context, such as
geographic origin. Such pseudonymity enables another entity

interested in breaching the privacy to link multiple queries on a long-term basis. Pseudonyms are assumed long-lived and their uniqueness may be a goal. There are many findings that indicate that pseudonymity is weaker than anonymity.

o Unlinkability: Formally, two items of interest are said to be unlinkable if the certainty of an actor concerning those items of interest is not affected by observing the system. This is, unlinkability implies that the a-posteriori probability computed a monitor that two items of interest are related is close enough to the a-priori probability computed by a monitor based on their knowledge.

Two items of interest are said to be unlinkable if there is a small chance (beta, close to 0 probability) that the monitor identifies them as associated, and they are linkable if there is a sufficiently large probability (referred to as alpha).

Informally, given two items of interest (user attributes, DNS queries, users, etc.), unlinkability is defined as the inability of the monitor to sufficiently determine whether those items are related to one another. In the context of DNS, this refers typically but not only to a monitor relating queries to the same individual.

o Undetectability: a stronger definition of privacy, where an item of interest is said to be undetectable if the monitor is not sufficiently able to know or tell whether the item exists or not.

Note that undetectability implies unlinkability. As explained below, a way of ensuring undetectability is to use encryption that is secure under known ciphertext attacks, or randomized encryption.

o Unobservability: a stronger definition of privacy that requires satisfying both undetectability and anonymity. Unobservability means that an item of interest is undetectable by any not involved individual, monitor or not.

In theory, there are many ways of ensuring unobservability by fulfilling both requirements. For example, undetectability requires that no party that is not invovled in the resolution of a DNS query shall know that query has existed or not. A mechanism to ensure this function is encryption secure under known ciphertext attacks, or randomized encryption for all other than stub resolvers, and pseudonyms for the stub resolver. An alternative mechanism to provide the anonymity property would be the use of mix networks for routing DNS queries.

3. Assumptions about Quantification of Privacy

   The quantification of privacy is connected with the privacy goals.
   Is the desired privacy property unlinkability only, or is it
   undetectability? Is pseudonymity a sufficient property? Parameters
   and entire privacy mechanism choices are affected by the choice of
   privacy goals.

   While a binary measure of privacy is sometimes possible, that is,
   being able to say that the transaction is anonymous, in this
   document, we assume that the binary is not frequently obtainable, and
   therefore we focus on methods for continuous quantification. Both
   are relevant to DNS Private Exchange. Another way to state this is
   that the quantification could be exactly the probabilities 1 and 0,
   corresponding to the binary, but the methods prefer to provide
   continuous values instead.

   Here is an example of continuous quantification, related to
   identifiability of an individual or item of interest based on
   observing queries.

   o For an individual A, and a set of observations by a monitor Y,
      where Y = [y1, y2, ... yn], we define the privacy of A as the
      uncertainty of the monitor knowing that A is itself among many
      others under the observations Y; that is, we define Privacy = 1 -
      P[A | Y]

   o For an item of interest r associated with a user A, we similarly
      define the privacy of r as Privacy = 1 - P[r | Y].

4. System Model

   A DNS client (a DNS stub resolver) may resolve a domain name or
   address into the corresponding DNS record by contacting the
   authoritative name server responsible for that domain name (or
   address) directly. However, to improve the operation of DNS
   resolution, and reduce the round trip time required for resolving an
   address, both caching and recursive resolution are implemented.
   Caching is implemented at an intermediary between the stub and the
   authoritative name server. In practice, many caching servers also
   implement the recursive logic of DNS resolution for finding the name
   server authoritative for a domain, and are thus called DNS recursive
   resolvers. Another type of entity, DNS forwarders (or proxies), acts
   as intermediaries between stub resolvers and recursive resolvers, and
   recursive resolvers and authoritative resolvers. The system model
   for DNS privacy evaluation includes the four entities quickly
   sketched here: stub resolvers, recursive resolvers, authoritative
   name servers, and forwarders.

4.1. DNS Resolvers (System Model)

o Stub resolver (S): a minimal resolver that does not support
   referral, and delegates recursive resolution to a recursive
   resolver. A stub resolver is a consumer of recursive resolutions.
   Per the terminology of [RFC6973], a stub resolver is an Initiator.

o Recursive resolver (R): a resolver that implements the recursive
   function of DNS resolution on behalf of a stub resolver. Per the
   terminology of [RFC6973], a recursive resolver is an Enabler.

o Authoritative resolver (A): is a server that is the origin of a
   DNS record. A recursive resolver queries the authoritative
   resolver to resolve a domain name or address. Per the terminology
   of [RFC6973], the authoritative name server is also an Enabler.

o Forwarder/proxy (P): between the stub resolver and the
   authoritative resolver there may be one or more DNS-involved
   entity. These are systems located between S and R (stub resolver
   and recursive), or between R and A (recursive and authoritative),
   which do not play a primary role in the DNS protocol. Per the
   terminology of [RFC6973], forwarders are Intermediaries.

4.2. System Setup - Putting It Together

Evaluating various privacy protection mechanisms in relation to
monitors such as the pervasive monitors defined next requires
understanding links in the system setup. We define the following
links. In relation to [RFC7258] these are the attack surface where a
monitor (eavesdropper) can collect sets of query information.

o Stub -> Recursive (S-R): a link between the stub resolver and a
   recursive resolver. At the time of writing, the scope of DPRIVE
   Working Group privacy mechanisms is supposed to be limited to S-R.

o Stub -> Proxy (S-P): a link between the stub resolver and a
   forwarder/ proxy. The intended function of this link may be
   difficult to analyze.

o Proxy -> Recursive (P-R): a link between a proxy and a recursive
   server.

o Recursive -> Authoritative (R-A): a link between a recursive and
   an authoritative name server. Although at the time of writing,
   R-A is not in the DPRIVE scope, we discuss it in the evaluation.

Rather than notating in system setup that an entity is compromised,
this is covered in the monitor model in Section 6, which has system
elements as parameters.

In the system setup, there is a possibility that S and R exist on a
single machine. The concept of the Unlucky Few relates S and R in
this case. A monitor can monitor R-A and find the query traffic of
the initiator individual. The same concept applies in the case where
a recursive is serving a relatively small number of individuals. The
query traffic of a subject group or organization (c.f. Subject in
the definitions) is obtained by the monitor who monitors this
system's R-A.

Because R-A is not in the DPRIVE scope, it is for future work to
examine the Unlucky Few circumstance fully. The general system setup
is that PII, the individual's private identifying information, is not
sent on R-A and is not seen by authoritative name server.

There could be one or more proxies between the stub resolver and a
recursive. From a functionality point of view they can all be
consolidated into a single proxy without affecting the system view.
However, the behavior of such proxies may affect the size and shape
of the attack surface. However, we believe that an additional
treatment is needed for this case and it is not included in the
discussion.

We also do not include in discussion proxies that exist along R-A,
between a recursive and an authoritative name server. We do so in
respect for the DPRIVE charter's scope at this time. According to
recent work at [openresolverproject.org], there may be multiple
intermediaries with poorly defined behavior.

The system setup here leaves out other realistic considerations for
simplicity, such as the impact of shared caches in DNS entities.

5. Risk Model

The Definitions section defines observer, attack and monitor, but not
a Risk Model, which is needed to actually evaluate privacy.

For consistency, we note that the only difference between an attacker
and an obeserver is that an attacker is an unauthorized observer with
all the capabilities it may have. We also stress that for the
context of DNS privacy, the term attacker may implicitly assume an
intent. To that end, active and passive observers are collectively
referred to as actors.

o Risk Model: a well-defined set of capabilities indicating how much
   information an observer (or eavesdropper) has, and in what
   context, in order to reach a goal of breaching the privacy of an
   individual or subject with respect to a given privacy metric.

In this document we focus on two risk models, namely a pervasive
monitor and a malicious monitor.

5.1. Risk Type-1 - Passive Pervasive Monitor

This risk corresponds to the passive pervasive monitoring model
described in [RFC7258]. This model relies on monitoring capabilities
to breach the privacy of individuals from the DNS traffic at scale
without decimation. An actor causing this risk is capable of
eavesdropping or observing traffic between two end points, including
traffic between any of the pairs of entities described in section
2.1. Per [RFC7258], this type of actor has the ability to eavesdrop
pervasively on many links at once, which is a powerful form of
attack. Type-1 monitors are passive. They do not modify traffic or
insert traffic.

5.2. Risk Type-2 - Active Monitor

This risk corresponds to an actor with the same types of capabilities
of monitoring links. Additionally, this model can select links in
order to target specific individuals. A Type-2 monitor for instance
might put into place intermediaries in order to obtain traffic on
specific links.

Note that we exclude malicious monitoring from this document since,
by definition, a malicious actor has an intent associated with his
actions.

5.3. Risks in the System Setup

To evaluate the privacy provided by a given mechanism or mechanisms
in a particular system model, we characterize the risk using a
template where parameters from the system model in which the risk
actor (eavesdropper or observer as monitors) is located are used.
The general template is: Risk(Type, [Entities], [Links]). For
example, the template Risk(Type-2, R, S-R) passed as a parameter in
the evaluation of a privacy mechanism indicates a Type-2 monitor that
controls a recursive resolver and has the capability of eavesdropping
on the link between the stub and recursive resolvers (S-R). Other
risk templates include the appropriate parameterizations based on the
above description of those monitors, including monitors that have the
capabilities of monitoring multiple links and controlling multiple
pieces of infrastructure.

6. Privacy Mechanisms

   Various mechanisms for enhancing privacy in networks are applicable
   to DNS private exchange. Some mechanisms common to privacy research
   include mix networks, dummy traffic, and private information
   retrieval techniques. Applicable protocol mechanisms include
   encryption-based techniques - encrypting the channel carrying the
   queries using IPSEC [ipseca], TLS [dns-over-tls] or special-purpose
   encryption [confidential-dns]. [privatedns] includes special-purpose
   encryption but also depends on a trusted service broker.

   o Mix Networks: in this type of mechanism, the initiator uses a mix
      network such as Tor to route the DNS queries to the intended DNS
      server entity. A monitor observing part of the system finds it
      difficult to determine which individual sends which queries, and
      will not be able to tell which individual has sent them (ideally,
      though it is known that attacks exist that allow correlation and
      privacy breaches against mix networks). The privacy property is
      unlinkability of the queries; the probability that two queries
      coming from one exit node in the mix network belong to the same
      individual is uniform among all the individuals using the network.

   o Dummy Traffic: a simple mechanism in which the initiator of a DNS
      request will also generate k dummy queries and send the intended
      query along with those queries. As such, the adversary will not
      be able to tell which query is of interest to the initiator. For
      a given k, the probability that the adversary will be able to
      detect which query is of interest to the initiator is equal to
      $1-1/(k+1)$. In that sense, and for the proper parameterization of
      the protocol, the monitor is bounded to the undetectability of the
      queries.

   o Private Information Retrieval: a mechanism that allows a user s to
      retrieve a record r from a database DB on a server without
      allowing the server to learn the recordr. A trivial solution to
      the problem requires that s downloads the entire database DB and
      then perform the queries locally. While this provides privacy to
      the queries of the user, the solution is communication inefficient
      at the scale of the DNS. More sophisticated cryptographic
      solutions are multi-round, and thus reduce the communication
      overhead, but are still inefficient for the DNS.

   o Query Minimization: a mechanism that allows the resolver to
      minimize the amount of information it sends on behalf of a stub
      resolver. A method of query minimization is specified in
      [qname-minimisation]. Qname minimization deprives a Type-1 risk
      on R-A of information from correlating queries, unless the
      individuals have an Unfortunate Few problem.

o NOTE: queries on R-A generally do not include an identifier of the
   individual making the query, because the source address is that of
   R. With respect R or A themselves, they may have well established
   policies for respecting the sensitivity of queries they process,
   while still using summary analysis of those queries to improve
   security, stability of their business operation.

o Encrypted Channel Mechanisms: Using these mechanisms, an initiator
   has an encrypted channel with a corresponding enabler, so that the
   queries are not available to eavesdropping Pervasive Monitor risk.
   Examples include [dns-over-tls], [ipseca], and [confidential-dns].
   Depending on the characteristics of the channel, various privacy
   properties are ensured. For instance, undetectability of queries
   is ensured for encryption-based mechanisms once the encrypted
   channel is established. Unlinkability of the queries may depend
   on the type of crypto-suite; it is provided as long as randomized
   encryption is used.

o Composed (Multiple) Mechanisms: the use of multiple mechanisms is
   a likely scenario and results in varied privacy guarantees.
   Consider a hypothetical system in which mix networks (for
   unlinkability) and randomized encryption (for undetectability) can
   both be applied, thus providing for unobservability, a stronger
   property than either of the two along. On the other hand,
   consider another hypothetical system in which mix networks are
   used to reach a third party broker requiring sign-in and
   authorization. Depending on the risk type, this could mean that
   the mix network unlinkability was cancelled out by the linkability
   due to entrusting the third party with identifying information in
   order to be authorized.

7. Privacy Evaluation

Now we turn our attention to the evaluation of privacy mechanisms in
a standard form, given the risk models and system definitions, for
some of the example mechanisms.

An evaluation takes multiple parameters as input. The output of the
evaluation template is based on the analysis of the individual
algorithms, settings, and parameters passed to this evaluation
mechanism.

Here is the top level interface of the evaluation template:

```
Eval(Privacy_Mechanism(params),
     System_Setting(params),
     Risk_Model(params)
     )
```

The output of the function is a privacy guarantee for the given
settings, expressed through defined properties such as unlinkability
and unobservability, for the specified system and risk model.

7.1 Dummy Traffic Example

```
Eval(Dummy_Traffic (k=10, distribution=uniform),
    System_Setting([S, P, R, A], [S-P, P-R, R-A]),
    Risk_Model(Type-1A, S-R))
```

The dummy traffic mechanism is not presented as a practical
mechanism, though there's no way to know if there are deployments of
this type of mechanism. This example evaluation uses k=10 to
indicate that for every one query initiated by an individual, ten
queries that disguise the query of interest are selected randomly
from a pool of queries. In the parameters passed in the evaluation
function, we indicate that the privacy assurances of interest concern
the S-R link, with a Passive Pervasive Monitor (Type-1A) risk.

Here is a template format for the example:

```
Eval(Dummy_Traffic (k=10, distribution=uniform),
    System_Setting([S, P, R, A], [S-P, P-R, R-A]),
      Risk_Model(Type-1A, S-R)) {
         Privacy_Mechanism{
             Mechanism_name = Dummy_Traffic
             Parameters{
                 Queries = 10
                 Query_distribution = uniform
         }
         System_settings{
             Entities = S, P, R and A;
             Links = S-P, P-R, R-A
         }
         Risk_Model{
             Type = Type-1A
             Links = S-R
         }
         Privacy_guarantee = undetectability
         Privacy_measure = 1-(1/(queries+1)).

         Return Privacy_guarantee, Privacy_measure


}
```

Undetectability is provided with 0.91 probability (though we know
there are other weaknesses for dummy traffic) If the threat model is

replaced with Type-2, so that responses to arbitrary requests can be injected, and tracked, the undetectability probability is decreased.

7.2 Mix Network Example

Here is an input for a mix network privacy mechanism:

Eval(Mix (u=10, distribution=uniform), System_Setting(link=S-R), threat_Model(Type-1A))

This indicates that the monitor resides between the stub and resolver. While queries are not undetectable, two queries are not linkable to the same individual; the provided guarantee is unlinkability. For a given number of individuals in the mix network, indicated by the parameter u, assuming that at any time, traffic from these individuals is uniformly random, the probability that one query is comes from a given individual is (1/10 = 0.1). The probability that two queries are issued by the same initiator is 0.1^2 = 0.01, which represents the linkability probability. The unlinkability probability is given as 1-0.01 = 0.99. Thus,

(unlinkability, 0.99) < Eval(Mix (u=10, distribution=uniform), System_Setting(link=S-R), Risk_Model(type-1)).

We note that even if there is a Type-2 Risk in recursive resolver R, the same results hold.

To sum up, the above example is represented in the following template:

```
Eval(Mix (u=10, distribution=uniform),
    System_Setting([S, P, R, A], [S-P, P-R, R-A]),
      Risk_Model(Type-1A, S-R)) {
         Privacy_Mechanism{
             Mechanism_name = mix    //mix network
             Parameters{
                 Users = 10
                 Query_distribution = uniform
         }
         System_settings{
             Entities = S, P, R and A;
             Links = S-P, P-R, R-A
         }
         Risk_Model{
             Type = Type-1A
             Links = P-R
         }

         Privacy_guarantee = unlinkability
         Privacy_measure = 1-(1/users)^2.

         Return privacy_guarantee, privacy_measure
}
```

7.3 Encrypted Channel (DNS-over-TLS) Example

For one of the encryption-based mechanisms, DNS-over-TLS
[dns-over-tls], we have the following template (TLS parameters are
from [RFC5246]):

```
Eval(TLS_enc (SHA256, ECDSA, port 53, uniform),
     System_Setting([S, P, R, A], [S-P, P-R, RA]),
         Risk_Model(Type-1B, S-R)) {
            Privacy_Mechanism{
              Mechanism_name = TLS-upgrade-based
                Parameters{
                  Query_distribution = uniform
                  Hash_algorithm = SHA256
                  Sig_Algorithm = ECDSA
                  Port 53
                }
                System_settings{
                  Entities = S, P, R and A;
                  Links = S-P, P-R, R-A
                }
                Risk_Model{
                  Type = Type-1B
                  Links = S-R
                }
                Privacy_guarantee = unlinkability, undetectability
                Privacy_measure (unlinkability) = 1
                Privacy_measure (undetectability) = 0 // port 53 indicates DNS
used

                Return privacy_guarantee, privacy_measure
}
```

   This template features an Active Monitor risk model (Type-2) in order
   to show how that the monitor might apply extra resources to an
   encrypted channel. Undetectability is an issue whether using
   upgrade-based TLS on port 53, or a port-based TLS on a dedicated port
   - both ports indicate the use of DNS. The source address of the
   individual is exposed in all cases. If this were a suitably
   parameterized use of [ipseca], the monitor would not be certain that
   all the traffic from S-R was DNS, and undetectability would be
   higher.

   7.4 Encrypted Channel (IPSec) Example

   In the following, we use the same template above to characterize the
   encryption capabilities provided by IPSec, as a potential mechanisms
   for enabling privacy in DNS exchanges.

```
Eval(IPSEc_enc([...]),
    System_Setting([S, P, R, A], [S-P, P-R, RA]),
      Risk_Model(Type-1B, S-R)) {
         Privacy_Mechanism{
           Mechanism_name = IPSec
           Parameters{
             Query_distribution = uniform
           }
           System_settings{
             Entities = S, P, R and A;
             Links = S-P, P-R, R-A
           }
           Risk_Model{
             Type = 2
             Links = S-R
           }
           Privacy_guarantee = unlinkability, undetectability
           Privacy_measure (unlinkability) = 1
           Privacy_measure (undetectability) = 1

           Return privacy_guarantee, privacy_measure
}
```

We note that IPSec can provide better guarantees with respect to studied privacy notions. However, whether the technique itself is widely deployable or not is worth further investigation.

7.5 QName Minimization Example (R-A) Example

Analyzing the privacy assurances of QName minimization is a non-trivial problem, given that the notions introduced in this document are techniques that do not alter items of interest. This is, the notions of privacy as outlined above are concerned with a certain IOI that is modified by this technique. To this end, we modify the aforementioned notions in ordered to be able to analyze the technique. For example, we define linkability as the ability of an adversary to link two labels of (minimized) queries to each other, and relate them to original source of query. Assuming a reasonable use of a recursive resolver that minimizes queries on behalf of users, this task is non-trivial, although quantifying the probability would depend on the number of labels in queries, the number of queries issued, and the number of users using the studied recursive resolvers. The following template captures QName minimization as a template

```
Eval(Qname_minimisation ([...],
    System_Settings([S, P, R, A], [R-A]),
      Risk_Model(Type=2){
         Privacy_Mechanism{
           Mechanism_name = Qname_minimisation
           Parameters{
             Qtype_used = NS
        }
         System_settings{
           Entities = S, P, R and A;
           Links = R-A
        }
         Risk_model{
           Type = 2
           Links = R-A
        }
         Privacy_guarantee =  unlinkability
         Privacy_measure = analytical

         Return privacy_guarantee, privacy_measure
   }
```

Note that QName minimization does not solve the problem of the
privacy for a monitoring risk between the stub and recursive.
Encrypting the channel between the recursive and the stub, utilizing
other techniques such as TDNS or IPSec, can marginalize such risk.
Furthermore, note that the risk on the link between the recursive and
authority name servers is always mitigated by the fact that recursive
name servers act as a mixer of queries, even when they are sent in
full to the authority name servers.

7.7 Private-DNS (S-R) Example

The template for [privatedns] takes note of deployments in which in
addition to S, R and A, there is another entity in the system, the
function that authenticates the individual using S prior to
permitting an encrypted channel to be formed to R or A. If the
Private-DNS connection is with R, then identifiability of S as an
individual may be similar to the identifiability of S from source
address, or it may be stronger, depending on the nature of the
account information required. If the Private-DNS connection is with
A, source address PII is provided to A, and linkability of the
queries from S has probability 1.

8. Other evaluation

   This document does not address a lot of the evaluation aspects not
   associated with privacy. For example, some of the mechanisms
   discussed in the working group are built of well-understood and
   standardized technologies, whereas others use other non-standard and
   less widely deployed techniques. A comprehensive evaluation of such
   mechanisms should take into account such facts.

9. Security Considerations

   The purpose of this document is to provide methods for those
   deploying or using DNS private exchange to assess the effectiveness
   of privacy mechanisms in depriving monitors of access to private
   information. Protecting privacy is one of the dimensions of an
   overall security strategy.

   It is possible for privacy-enhancing mechanisms to be deployed in
   ways that are vulnerable to security risks, with the result of not
   achieving security gains. For the purposes of privacy evaluation, it
   is important for the person making an evaluation to also ensure close
   attention to the content of the Security Considerations section of
   each mechanism being evaluated, for instance, to ensure if TLS is
   used for encryption of a link against surveillance, that TLS best
   security practices [RFC7525] are in use.

10. IANA Considerations

   No requests are made to IANA.

11. Acknowledgements

   We wish to thank Scott Hollenbeck, Burt Kaliski, Minsuk Kang, Paul
   Livesay and Eric Osterweil for reviewing early versions. We wish to
   thank Stephane Bortzmeyer and Tim Wicinski for their detailed review
   and feedback on the previous versions of this document. We also wish
   to thank those who commented on presentations of this work ahead of
   publication, including Simson Garfinkel, Cathy Meadows, Paul
   Syverson, and Christine Task.

12. Informative References

   [confidential-dns]
              Wijngaards, W. and G. Wiley, "Confidential DNS", draft-
              wijngaards-dnsop-confidentialdns-03 (work in progress),
              March 2015.

[dns-over-tls]
          Zhu, L., Hu, Z., Heidemann, J., Wessels, D., Mankin, A.,
          and P. Hoffman, "TLS for DNS: Initiation and Performance
          Considerations", draft-ietf-dprive-dns-over-tls-01.txt
          (work in progress), February 2015.

[ipseca]   Osterweil, E., Wiley, G., Okubo, T., Lavu, R., and A.
          Mohaisen, "Opportunistic Encryption with DANE Semantics
          and IPsec: IPSECA", draft-osterweil-dane-ipsec-03 (work in
          progress), March 2015.

[openresolverproject.org]
          Mauch, J., "The Open Resolver Project", April 2015.

[phb-dnse]
          Hallam-Baker, P., "DNS Privacy and Censorship: Use Cases
          and Requirements", draft-hallambaker-dnse-02 (work in
          progress), November 2014.

[privatedns]
          Hallam-Baker, P., "Private-DNS", draft-hallambaker-
          privatedns-01 (work in progress), November 2014.

[qname-minimisation]
          Bortzmeyer, S., "DNS query name minimisation to improve
          privacy", draft-ietf-dnsop-qname-minimisation-07 (work in
          progress), March 2015.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC
          Text on Security Considerations", BCP 72, RFC 3552, DOI
          10.17487/RFC3552, July 2003,
          <http://www.rfc-editor.org/info/rfc3552>.

[RFC4949] Shirey, R., "Internet Security Glossary, Version 2", FYI
          36, RFC 4949, DOI 10.17487/RFC4949, August 2007,
          <http://www.rfc-editor.org/info/rfc4949>.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security
          (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/
          RFC5246, August 2008,
          <http://www.rfc-editor.org/info/rfc5246>.

[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,
          <http://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, <http://www.rfc-editor.org/info/rfc7258>.

   [RFC7435]   Dukhovni, V., "Opportunistic Security: Some Protection
               Most of the Time", RFC 7435, DOI 10.17487/RFC7435,
               December 2014, <http://www.rfc-editor.org/info/rfc7435>.

   [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, <http://www.rfc-editor.org/info/rfc7525>.

   [RFC7626]   Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626,
               DOI 10.17487/RFC7626, August 2015,
               <http://www.rfc-editor.org/info/rfc7626>.

Authors' Addresses

   Aziz Mohaisen
   SUNY Buffalo
   323 Davis Hall
   Buffalo, NY 14260
   US

   Phone: +1 716 645-1592
   Email: mohaisen@buffalo.edu

   Allison Mankin
   Verisign Labs
   12061 Bluemont Way
   Reston, VA 20190
   US

   Phone: +1 703 948-3200
   Email: amankin@verisign.com