        **Confidentiality Aspects of DNS Data, Publication, and Resolution**
              **draft-koch-perpass-dns-confidentiality-00**

Abstract

   This document describes aspects of DNS data confidentiality in the
   light of recent IETF discussions on pervasive monitoring.  It focuses
   on potential information leaks rather than prescribing methods of
   mitigation.

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 May 16, 2014.

Copyright Notice

## [1](#). Introduction

The Domain Name System (DNS) [RFC1034] [RFC1035] is the Internet's
primary name lookup system.  It consists of a publication aspect,
represented by authoritative name servers providing access to DNS
data covering parts of the DNS tree in units of zones, and a
resolution aspect.  The latter consists of applications that initite
DNS requests, DNS stub resolvers and DNS full resolvers (sometimes
also called recursive resolvers or recursive name servers).
Resolvers might be chained using a forwarding mechanism.  In today's
reality, there is a variety of intercepting DNS proxies and other
middle boxes which are currently out of scope but may be addressed in
future versions of this memo.

Threats to the DNS are described in [RFC3833] and have been addressed
by DNSSEC [RFC4033] [RFC4034] [RFC4035], both to the extent that data
origin authentication is concerned.  Confidentiality was not a DNSSEC
design goal, although in subsequent discussion that eventually led to
the specification and deployment of NSEC3 [RFC5155], confidentiality
of zone content was a major issue.

### [1.1](#). The alleged public nature of DNS data

It has long been claimed that "the data in the DNS is public".  While
this sentence makes sense for an Internet wide lookup system, there
are multiple facets to data and meta data that deserve a more
detailed look.  First, access control lists and private name spaces
nonwithstanding, 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).  A DNS
query consists of QNAME, QCLASS and QTYPE.  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 to receive a response.  The
zone transfer QTYPE [RFC5936] is often blocked or restricted to
authenticated/authorized access to enforce this difference (and maybe
for other, more dubious reasons).

Another differentiation to be applied is between the DNS data as
mentioned above and a particular transaction, most prominently but
not limited to a DNS name lookup.  The fact that the results of a DNS
query are public within the boundaries described in the previous
paragraph and therefore might have no confidentiality requirements
does not imply the same for a single or a sequence of transactions.
Any transaction has meta data associated with the query data, e.g., a
source address and a timestamp.

## 1.2.  Disclaimer

   The practices listed in this document appear only to support an
   informed discussion.  Their presence (or absence) does not imply any
   form of support, engagement, applicability, appropriateness, fitness,
   or stance on legal status.

## 2.  DNS Element walk through

   This section will address the specific confidentiality issues of
   verious elements of the DNS ecosystem.  We will start at the
   authoritative servers, leaving the provisioning side out of scope,
   cover the resolution and recursive resolvers and finally address DNS
   queries at large and packet capturing.

## 2.1.  Authoritative Name Servers

   DNS zone data is published by authoritative name servers.  Starting
   at the primary master, zone data is transfered in full (AXFR) or
   increments (IXFR) to secondary servers along the XFR dependency
   graph.  The zone data thereby is inevitably revealed to any of the
   authoritative servers.  Some zones, including the DNS root zone, are
   deliberatly published by methods other than DNS AXFR.

   While client as well as server authentication and data integrity are
   usually achieved by TSIG [RFC2845], there is no DNS protocol feature
   that provides zone transfer confidentiality.  However, VPNs or other
   private arrangements are occasionally used.  [RFC2182] is the most
   recent IETF document potentially dealing with this issue.

## 2.2.  DNS Name Resolution

   Since the communication between an application and the local resolver
   or between the local (stub) resolver and a full recursive resolver is
   rarely authenticated, DNS queries can and hve been redirected.  This
   has mostly been done with the malicius intent to inject forged
   responses, but could also be used as a man-in-the-middle (MITM)
   attack to learn a particular system's DNS queries and the response
   content.

   The same queries (and responses) could be captured on the wire, even
   on the way to (and from) the correct, intended full resolver.
   Usually it has been assumed that the DNS resolution would not add
   additional intelligence given that subsequent communication would
   most likely reveal more than the DNS lookup.  However, with recent
   suggestions to encrypt, say, web (HTTP) and mail (SMTP) connections,
   the DNS information could be of increased interest, disclosing
   otherwise unavailable information.

Operators of recursive resolvers could collect and examine queries
directed to their systems.

The content of resolvers can reveal data about the clients using it.
This information can sometimes be examined by sending DNS queries
with RD=0 to inspect cache content, particularly looking at the DNS
TTLs.  Since this also is a reconnaissance technique for subsequent
cache poisoning attacks, some counter measures have already been
developed and deployed.

## 2.3.  DNS Queries

DNS queries are initiated by an application handed over to a stub
resolver, sometimes involving a host dependent name caching mechanism
that is out of scope of this document.  They consist of a QNAME,
QCLASS and QTYPE, a DNS query ID and other parameters at the IP or
transport layer.  Among those are an IP source address, an IP ID and
a source port number [RFC5452].  While some of these parameters have
received increased attention due to their significance for DNS
response spoofing mitigation, they do not contribute to
confidentiality and may in fact deliver additional intelligence by
supportig correlation of multiple queries from one system or even a
single process or application at the same source.  This is sometimes
used in resolver software fingerprinting or behavioural analysis.

The source address in a DNS query is necessary to direct the
response, but it may help to identify the requesting entity, be that
a system, a process or an end user.  For recursive resolvers it is
sometimes argued that the size of the population 'behind' that
resolver contributes to the noise.  However, a private extension
[I-D.vandergaast-edns-client-subnet] exists that will disclose the
source address, or some prefix of the source address to the receiver,
usually an authoritative name server.

The QNAME itself will be an existing or a non existing domain name.
With reference to the earlier discussion of the public (or not)
nature of DNS data, the response may reveal information.  More
importantly, due to the use of search paths [RFC1535] the QNAME may
also disclose information relative to the querying entity:

   _ldap._tcp.Default-First-Site-Name._sites.gc._msdcs.example.org.

For parts of the domain name tree that more deeply enjoy the
hierarchic nature of the DNS, like the IPv6 reverse delegation
[RFC3596] or ENUM [RFC6116], the query name itself, asked for at a
particular time, may disclose related, either ongoing or subsequent
communication.  This is partly due to the fact that the DNS treats
the QNAME in full all the time.

Attempts have been made to encrypt the resource record RDATA
[I-D.timms-encrypt-naptr].

## 2.4.  DNS Packet Capturing

Both ephemeral and long term DNS captures have become DNS operational
practice [DITL1] [DITL2].  Taking these packet traces usually occurs
close to the authoritative servers, packets being captuered on the
wire, but under the control of the endpoint operator.

Initially designed to reconstruct DNS zone content from query
response data, passive DNS [FW2005] has evolved into a widely used
tool.  These traces are usually sourced by on the wire traffic
between recursive resolver and authoritative server.

## 3.  Security Considerations

This document does not define a new protocol.  It deals with
confidentiality issues of the current DNS protocol and operations.

## 4.  IANA Considerations

This document does not propose any new IANA registry nor does it ask
for any allocation from an existing IANA registry.

## 5.  Acknowledgements

This document was inspired by discussion with Wouter Wijngaards and
Alexander Mayrhofer.  Stephane Bortzmeyer and Nathalie Boulvard
raised the issue of packet captures at a CENTR workshop.  Jonathan
Spring triggered some thoughts on the same topic.

## 6.  References

## 6.1.  Normative References

[RFC1034]  Mockapetris, P., "Domain names - concepts and facilities",
           STD 13, RFC 1034, November 1987.

[RFC1035]  Mockapetris, P., "Domain names - implementation and
           specification", STD 13, RFC 1035, November 1987.

   [RFC2182]  Elz, R., Bush, R., Bradner, S., and M. Patton, "Selection
              and Operation of Secondary DNS Servers", BCP 16, RFC 2182,
              July 1997.

   [RFC2845]  Vixie, P., Gudmundsson, O., Eastlake, D., and B.
              Wellington, "Secret Key Transaction Authentication for DNS
              (TSIG)", RFC 2845, May 2000.

   [RFC3833]  Atkins, D. and R. Austein, "Threat Analysis of the Domain
              Name System (DNS)", RFC 3833, August 2004.

   [RFC5155]  Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS
              Security (DNSSEC) Hashed Authenticated Denial of
              Existence", RFC 5155, March 2008.

   [RFC5936]  Lewis, E. and A. Hoenes, "DNS Zone Transfer Protocol
              (AXFR)", RFC 5936, June 2010.

   [RFC6973]  Cooper, A., Tschofenig, H., Aboba, B., Peterson, J.,
              Morris, J., Hansen, M., and R. Smith, "Privacy
              Considerations for Internet Protocols", RFC 6973, July
              2013.

## 6.2.  Informative References

   [DITL1]    CAIDA, "A Day in the Life of the Internet (DITL)", 2011.

   [DITL2]    DNS-OARC, "DITL Traces and Analysis", 2013.

   [FW2005]   Weimer, F., "Passive DNS Replication", FIRST 17, April
              2005.

   [I-D.timms-encrypt-naptr]
              Timms, B., Reid, J., and J. Schlyter, "IANA Registration
              for Encrypted ENUM", draft-timms-encrypt-naptr-01 (work in
              progress), July 2008.

   [I-D.vandergaast-edns-client-subnet]
              Contavalli, C., Gaast, W., Leach, S., and E. Lewis,
              "Client Subnet in DNS Requests", draft-vandergaast-edns-
              client-subnet-02 (work in progress), July 2013.

   [RFC1535]  Gavron, E., "A Security Problem and Proposed Correction
              With Widely Deployed DNS Software", RFC 1535, October
              1993.

   [RFC3596]  Thomson, S., Huitema, C., Ksinant, V., and M. Souissi,
              "DNS Extensions to Support IP Version 6", RFC 3596,
              October 2003.

   [RFC4033]  Arends, R., Austein, R., Larson, M., Massey, D., and S.
              Rose, "DNS Security Introduction and Requirements", RFC
              4033, March 2005.

   [RFC4034]  Arends, R., Austein, R., Larson, M., Massey, D., and S.
              Rose, "Resource Records for the DNS Security Extensions",
              RFC 4034, March 2005.

   [RFC4035]  Arends, R., Austein, R., Larson, M., Massey, D., and S.
              Rose, "Protocol Modifications for the DNS Security
              Extensions", RFC 4035, March 2005.

   [RFC5452]  Hubert, A. and R. van Mook, "Measures for Making DNS More
              Resilient against Forged Answers", RFC 5452, January 2009.

   [RFC6116]  Bradner, S., Conroy, L., and K. Fujiwara, "The E.164 to
              Uniform Resource Identifiers (URI) Dynamic Delegation
              Discovery System (DDDS) Application (ENUM)", RFC 6116,
              March 2011.

## Appendix A.  Document Revision History

   This section is to be removed should the draft be published.

   $Id: draft-koch-perpass-dns-confidentiality.xml,v 1.1 2013/11/12
   09:29:48 pk Exp $

### A.1.  Initial Document

   First draft

Author's Address

   Peter Koch
   DENIC eG
   Kaiserstrasse 75-77
   Frankfurt  60329
   DE

   Phone: +49 69 27235 0
   Email: pk@DENIC.DE