Privacy Considerations for ISC Passive DNS

Introduction

Internet Systems Consortium is a US non-profit public benefit corporation dedicated to supporting the infrastructure of the Internet by developing and maintaining production quality software and systems.

The ISC Passive DNS project gathers DNS data from multiple worldwide sensors and combines it into a central database [1]. Those Passive DNS sensors watch DNS traffic, record some of what they see, and send the data back to a database in the USA. The resulting database records the way that DNS is actually being used, and can capture transient names and servers, geographically diverse patterns, and trends.

ISC is committed to respecting personal privacy when collecting this data, and also to meeting the legal requirements of the country in which the sensor is located. However, privacy laws vary from country to country. This document describes what the Passive DNS sensors do and don't do, to enable the operators of each sensor site to ensure that no privacy laws are being violated. These sensors were engineered by Internet Systems Consortium to ensure that they do not gather private information and to ensure that they cannot be compromised by hackers who might modify their behavior.

What is an ISC Passive DNS Sensor?

The ISC Passive DNS sensor platform is a standard commercial server computer running a high-security operating system. The sensing software is an application running on that server. The server is configured and operated according to recognized industry best practices for security. The sensing software can also be installed directly onto a DNS recursive server.

The sensor application itself is open source software and its source code is accessible for anyone to look at on ISC's public repository ftp://ftp.isc.org/isc/nmsg/. The global security community has found that exposing software to global review results in better security than keeping it secret, because it greatly increases the number of people who study it and make contributions. Having the sensor software source code be transparently available also enables anyone to validate our statements about its privacy aspects.

What does an ISC Passive DNS Sensor do?

The DNS Sensor application watches DNS traffic on its server's local network and records certain requests and their replies (but not the identity of the client making the request). If the server computer is connected to a network segment that carries DNS traffic, the sensor can see that traffic. No server-based measurement device can observe or measure traffic on a network segment unless it is somehow connected to that segment.

After capturing a DNS request, that request is transmitted to ISC's central database server. That transmission uses SSH link encryption with best-practice settings to ensure that the data is not monitored, altered, misdirected, or counterfeited.

How the ISC Passive DNS system works

From the point of view of the user of an online device, the DNS system is very simple. Your device asks a question and receives an answer. A device asks a question by sending that question to a "DNS resolver". That resolver does whatever it needs to determine the answer, and then sends the answer back to whatever device asked the question.

If the DNS resolver knows the answer (because it saved a copy the last time that question was asked) it answers, and the transaction is done. If it doesn't know the answer, then it must make queries of its own, asking other servers, before it can return the answer to the original device (and keep a copy). The quantity and nature of the DNS queries that the resolver will make to find the answer to the original question varies widely depending on many factors.

What information does the ISC Passive DNS sensor collect?

The ISC Passive DNS sensor collects DNS queries and replies made by DNS resolvers. The DNS data itself is, by design, public. An Internet standards document says "It is part of the design philosophy of the DNS that the data in it is public" [3].

The ISC Passive DNS sensor can only collect information from the network to which it is connected. Typically a sensor is installed by the operator of a network, who would not connect it to sensitive or secure networks. The placement of the sensor is the key factor in controlling the information that the sensor gathers.

Further, the ISC Passive DNS sensor can only collect information that it sees in DNS queries and the responses to them. If the information isn't in a DNS query or response, it cannot be collected. The information in those queries and responses is standardized and widely documented; the Wikipedia article "Domain Name System" is an excellent and accurate starting point [2].

The sensor does not collect client queries. It collects queries made by DNS resolvers in the process of answering a client query. Client queries contain the IP address of the client (so that it can receive the reply). Resolver queries contain the IP address of the resolver and do not contain the IP address of the original client. Thus by never collecting client queries, the sensor never collects the client's IP address, which might enable identification of the client.

When a Passive DNS sensor is installed, part of the installation process is to give it a list of the IP addresses of the resolvers that it is to monitor. A query issued by (or a reply sent to) a resolver whose address is not on that list will not be recorded.

The full specification of the contents of a DNS query and a DNS reply are documented in the footnotes to the above-mentioned Wikipedia article, e.g. in RFC 1035. The important items that the DNS sensor collects are these:

- The IP address of the resolver
- The IP address of the authoritative server with which the resolver is communicating
- The name being looked up
- The reply from the authoritative server

How might collected information compromise privacy?

There are two circumstances under which the information collected by the ISC Passive DNS sensor might accidentally reveal personal information. Neither is likely, and both are under the control of the person whose information would be revealed.

Private information in a name being looked up or returned

There is nothing to prevent a person from looking up the name

```
my.password.is.abcd1234.example.com
```

If in truth the person's password really is abcd1234, that fact will be recorded by the sensor, whether or not the queried name is actually defined. Variations on this are also possible; an ISP could use customer credit card numbers to identify name servers:

```
    james.jones.visa.is.4716399086646802.example.com
```

Both of these uses are preposterously unlikely, but are “possible”, so we note them.

Single-user DNS resolvers

If a DNS resolver is known to be used by only one person, then its IP address could identify that person. But the ISC Passive DNS sensor is installed by the owner of the resolver that it is monitoring. The sensor can easily be configured to remove the resolver’s IP address from data sent back to the central database.

Conclusion

The ISC Passive DNS sensor software has been designed and built to protect the personal privacy of any clients using the network that it is monitoring. The only information collected by the sensor and sent to the ISC’s Passive DNS infrastructure is data that relates to DNS resolver and authoritative server interaction. The operator of the sensor can optionally remove the identity of the DNS resolver should that be considered a risk to privacy; this will ensure that no user of the data will be able to identify any source.

References

1. “ISC Passive DNS Architecture”. Robert Edmonds, March 2012. https://kb.isc.org/article/AA-00654

2. Domain Name System. Wikipedia, http://en.wikipedia.org/wiki/Domain_Name_System

3. Domain Name System Security Extensions. D. Eastlake and C. Kaufman, January 1997. http://tools.ietf.org/html/rfc2065, §2.1.