Network Working Group   S. Brim
Internet-Draft   M. Linsner
Intended status: Informational   B. McLaughlin
Expires: April 10, 2011   K. Wierenga
                                                               Cisco
                                                     October 7, 2010

                           Mobility and Privacy
                     draft-brim-mobility-and-privacy-00

Abstract

   Choices in Internet mobility architecture may have profound effects
   on privacy. This draft revisits this issue, stresses its increasing
   importance, and makes recommendations.

Requirements Language

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
   document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

   This Internet-Draft is submitted in full conformance with the
   provisions of BCP 78 and BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF). Note that other groups may also distribute
   working documents as Internet-Drafts. The list of current Internet-
   Drafts is at http://datatracker.ietf.org/drafts/current/.

   Internet-Drafts are draft documents valid for a maximum of six months
   and may be updated, replaced, or obsoleted by other documents at any
   time. It is inappropriate to use Internet-Drafts as reference
   material or to cite them other than as "work in progress."

   This Internet-Draft will expire on April 10, 2011.

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

   Significant steps are being taken right now to make the Internet's
   architecture more scalable and robust in routing, addressing,
   multihoming, mobility, and locator/identifier separation. However,
   since the Internet is rapidly becoming an essential part of daily
   life for people around the world, our architectural changes need to
   take fundamental social issues into account as a primary
   consideration. One of those is privacy, and in this case
   particularly privacy of geographic location. If we do not, we run
   the risk of colliding with established IETF principles (see for
   example [RFC3693]) as well as public policy in many countries around
   the world.

   When the Internet was designed, IP addresses were associated with
   timesharing machines and not with particular users. In the 1980s it
   began to be likely that a device and thus an IP address would be
   associated with a single user. Now a single IP address is very
   likely to be associated with a single human being. Meanwhile, at the
   top of the stack, there has been a convergence of life functions
   using single devices using single addresses. A person now uses his
   or her personal device and associated IP address for many activities:
   work, shopping, talking, exchanging mail and files, reading,
   listening to music, etc.

   It is this convergence at both the top and bottom of the stack -- to
   a single person per device and to many applications on that device --
   that makes the social issues more and more significant in IETF work.
   People use the Internet for many, more personal, activities than
   before. The Internet needs to fulfill the obligations expected of a
   communications system essential to modern human society. Our lower
   layer protocol designs have privacy implications beyond their
   intended scope.

2. The Risks of Being Traceable

   Issues with revealing geographic location are well-established
   elsewhere. For example the RAND review of the European Data
   Directive [RAND-EDPD] points out that "the interpretation of location
   data (e.g. which locations are visited, suggesting which shops are
   frequented, and which products and services are bought), may in the
   future permit the identification of the health, social, sexual or
   religious characteristics of the data subject" (section 3.3.1). The
   less well-known problem that this document focuses on is tracing the
   movement of mobile devices. Because mobile devices are used for so
   many things, any possibility of tracing them has significant,
   probably unpredictable, social implications, perhaps more so than

revealing a single location. If an association can be made between a
mobile device and a person at any location, if that device can be
traced to a different geographic location then the association with
the person can be inferred, usually correctly, even if the person
believes they are anonymous at the new location. Consider scenarios
such as:

o You are looking for a job, interviewing at other companies over
   your lunch hour, but you don't want your current management to
   know.

o You are planning a surprise gift or party for your spouse and are
   visiting specialty stores.

o You are a journalist gathering information on a corrupt politician
   from sources who wish to hide that they are dealing with you.

o You are infiltrating an organized crime ring and don't want them
   to know when you sneak in the back door of police headquarters.

o You are a very famous person trying to avoid paparazzi and
   assassins who are able to find you sporadically.

Mobility mechanisms need to take this issue into account. Obviously
a mobile node must be reachable somehow, but a mobile node must be
able to hide its actual movement from public view if it wishes.

3. Basic Mobility Requirements

A mobile node may need to be reachable by others, or it may act
purely as a client of Internet-based services. Even if it is purely
a client, it still needs at least two things:

o An authentication and authorization identifier that it can use
   with each access network it connects to. (Not required for open
   access networks.)

o A Layer 3 way for its correspondents to get packets back to it.
   This may no longer be simple due to potential innovations in
   routing architecture.

In addition, if the mobile node wants to be reachable as a peer or to
offer services, it needs a few more things:

o An identifier (or identifiers) by which the node may be found by
   others, and a mechanism by which this identifier can be mapped to
   IP addresses/locators. Examples are domain names, SIP URIs, and

the corresponding services.

o   An IP address/locator for initially contacting the mobile node.
    This does not have to be associated with the mobile node's actual
    topological location. It can instead be associated with a
    rendezvous point or agent.

o   A mechanism for "route optimization", whereby such an agent can be
    eliminated from a data path between the mobile node and a
    correspondent.

o   An identifier or identifiers by which the mobile node can
    authenticate itself to its correspondents during initial contact,
    route optimization, and/or change of topological location. These
    identifiers can be at any layer, from 2 to 7. They can be
    associated with the mobile device's whole IP stack, individual
    transport sessions, or individual application instances.

o   Identifiers by which the mobile node can be referred to by third
    parties.

If all mobile nodes are reduced to being clients only -- if they are
willing to register with servers in order to use the Internet and
have others be able to reach them -- then there are fewer
requirements. However, over the evolution of the Internet we have
seen several times that it is not good to give up the symmetry of
Internet communication and "permission-free" networking, i.e. the
ability for anyone anywhere to communicate as a peer with other nodes
on the Internet. For the rest of this document we assume that the
IETF still wants to retain this model.

Every identifier listed above has a scope in which it needs to be
known, but it is only required to be known in that scope. For
example, an access authentication identifier only needs to be known
to the mobile node, the access network, and a trusted third party (a
mobile node's home network administration, or a bank, etc.). A
session identifier only needs to be known among the parties using it,
but not by the access network.

4. Avoiding Making a Mobile Node Traceable

As a mobile node moves, if L3 or higher layer mobility mechanisms are
used it will change its IP addresses/locators. The Internet already
has sophisticated publicly available services for determining where a
node is based on IP address alone. These mechanisms are not always
precise or accurate, but they are in very many cases and even
imprecise information is information. Protocol designers must assume

that whatever IP address or locator a node has, it is likely that
there is a service to turn that into a geographic location.

The tracing problem occurs when it is possible for a third party to
correlate IP addresses/locators and something unique about the mobile
node. Data can be gathered either through monitoring traffic or by
accessing public information. It does not have to be done
continuously -- periodic snapshots can make the mobile node just as
vulnerable. Once the data is gathered, the third party can search
for correlations.

Using identifiers for multiple purposes makes leakage of tracing
information more likely. Different entities in different scopes may
know different things about a mobile node or a person. Using
overlapping identifiers mixes scopes and may make new, perhaps
unexpected, correlations easier. For example if an access identifier
such as a mobile phone's IMEI (hard-coded and not changeable,
primarily used for access authentication) is also used for session
continuity, or is registered in an Internet database service that is
publicly accessible, changes in that device's IP addresses (and thus
geographic location) can be traced.

Long-lasting identifiers make correlation easier as a device moves.
They should not be used in scopes where they are not necessary.

The biggest concern is if information that makes a mobile node
traceable is required to be publicly available in order for the
Internet to function. If it is, it can be accessed not only without
the mobile node's consent but even without its knowledge, perhaps
without any audit trail of who is accessing the information that
could be looked at after the fact. Some architectures for mobility
and/or routing and addressing described in
[I-D.irtf-rrg-recommendation] assume the use of DNS or other public
mapping systems. In these, the mobile node is required to publish a
mapping between its identifier and its current IP addresses/locators
in order to be reachable, even if a mobile node is acting purely as a
client (because otherwise packets would not get back to it). This
architectural assumption removes all of the mobile node's freedom of
choice about how much confidentiality to preserve -- either it
exposes all of its movement to all of the world or it is simply not
reachable. Public information systems like DNS are not designed to
support confidentiality.

MIPv6's "home agent" [I-D.ietf-mext-rfc3775bis] is an example of how
to avoid this problem: Contact with a mobile node is initially
through a home agent, a rendezvous point for both data and control
traffic. The home agent acts on behalf of the mobile node and
encapsulates traffic to it. After an exchange of packets, the mobile

node may decide, on its own, if it wants to reveal its topological
location, and thus probably its geographic location, to the
correspondent node. It controls its own location information. The
decision to reveal it can be based on anything, including local
policy.

The principle of hiding information that can expose geographic
location in both data and control planes, and deferring revealing
more until the mobile node or its agent decides what it wants to do,
is essential. This can be included in any mobility architecture that
is designed to allow it and does not insist on exposing location to a
wide audience in order to gain efficiency. The obvious way to do it
is an indirection mechanism such as a home agent, but this is just
one way to do it. Any way will do.

Monitoring is a more subtle issue than exposure in public services,
but still real, even if the mobile node is client-only. If packets
contain an identifier that uniquely identifies the mobile node for
some period of time, someone able to gather data on packet traffic
can easily trace the mobile node's movements as the IP address/
locator changes. It is not necessary for the watcher to be able to
gather this information in real time if it can access logs gathered
by others. Here, approaches to the problem are more difficult to
define because there is a conflict between three goals: to avoid
overhead, to preserve session continuity with low delay, and to keep
control over location information. Some designs such already try to
find their balance. All protocol work should consider the tradeoffs
with privacy and explicitly find a balance point.

5. Recommendations

Members of the Internet community who are creating or reviewing
proposed architectural changes, particularly in mobility but also in
other areas that impinge on mobility such as routing and addressing,
should consider the following points:

o Architectural changes MUST avoid requiring exposing a mapping
   between any of a node's identifiers and IP addresses/locators to
   unknown observers. If they require exposure, they will experience
   a head-on collision with basic principles of the IETF and with
   privacy policies around the world. It will simply not be
   acceptable to require the loss of this much individual privacy.

o An architectural proposal MAY make it possible to use public
   information systems to optimize traffic flow, but ideally it
   should do so without sacrificing privacy. If it cannot do so
   without sacrificing privacy, the default case built into the

architecture SHOULD be to preserve privacy instead of optimizing. The reason is that most users will not change defaults, and the default be one of privacy, only moving away from it by customer choice.

o If possible, information about who is gathering data about a user SHOULD be available to that user. Everyone deserves to know who is watching them.

o Proposals SHOULD address the issue of loss of geographic location privacy due to monitoring of packets crossing the Internet, and find an explicit balance between conflicting goals.

o Protocols SHOULD avoid using identifiers for multiple purposes. Different identifier scopes do not need to overlap. Confidentiality boundaries can be established by clearly defining limited interfaces.

o Protocols SHOULD avoid using long-lasting identifiers in scopes where they are not necessary.

6. IANA Considerations

   This document makes no request of IANA.

   Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

   In a sense this entire document is about security.

8. Acknowledgements

   Thanks to many with whom we have discussed this issue in recent months.

9. Normative References

   [I-D.ietf-mext-rfc3775bis]
             Perkins, C., Johnson, D., and J. Arkko, "Mobility Support
             in IPv6", draft-ietf-mext-rfc3775bis-08 (work in
             progress), October 2010.

   [I-D.irtf-rrg-recommendation]
             Li, T., "Recommendation for a Routing Architecture",
             draft-irtf-rrg-recommendation-14 (work in progress),
             September 2010.

   [RAND-EDPD]
             Robinson, N., Graux, H., Botterman, M., and L. Valeri,
             "Review of the European Data Protection Directive",
             May 2009.

   [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
             Requirement Levels", BCP 14, RFC 2119, March 1997.

   [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and
             J. Polk, "Geopriv Requirements", RFC 3693, February 2004.

Authors' Addresses

   Scott Brim
   Cisco

   Email: scott.brim@gmail.com

   Marc Linsner
   Cisco

   Email: mlinsner@cisco.com

   Bryan McLaughlin
   Cisco

   Email: brmclaug@cisco.com

   Klaas Wierenga
   Cisco

   Email: kwiereng@cisco.com