

IP Address Location Privacy and Mobile IPv6: Problem Statement

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

In this document, we discuss location privacy as applicable to Mobile IPv6. We document the concerns arising from revealing a Home Address to an onlooker and from disclosing a Care-of Address to a correspondent.

Table of Contents

1. Introduction	2
2. Definitions	3
3. Problem Definition	4
3.1. Disclosing the Care-of Address to the Correspondent Node ...	4
3.2. Revealing the Home Address to Onlookers	4
3.3. Problem Scope	4
4. Problem Illustration	5
5. Conclusion	7
6. Security Considerations	7
7. Acknowledgments	8
8. References	8
8.1. Normative References	8
8.2. Informative References	8
Appendix A. Background	10

1. Introduction

The problems of location privacy, and privacy when using IP for communication, have become important. IP privacy is broadly concerned with protecting user communication from unwittingly revealing information that could be used to analyze and gather sensitive user data. Examples include gathering data at certain vantage points, collecting information related to specific traffic, and monitoring (perhaps) certain populations of users for activity during specific times of the day, etc. In this document, we refer to this as the "profiling" problem.

Location privacy is concerned with the problem of revealing roaming, which we define here as the process of a Mobile Node (MN) moving from one network to another with or without ongoing sessions. A constant identifier with global scope can reveal roaming. Examples are a device identifier such as an IP address, and a user identifier such as a SIP [RFC3261] URI [RFC3986]. Often, a binding between these two identifiers is available, e.g., through DNS [RFC1035]. Traffic analysis of such IP and Upper Layer Protocol identifiers on a single network can indicate device and user roaming. Roaming could also be inferred by means of profiling constant fields in IP communication across multiple network movements. For example, an Interface Identifier (IID) [RFC2462] in the IPv6 address that remains unchanged across networks could suggest roaming. The Security Parameter Index (SPI) in the IPsec [RFC4301] header is another field that may be subject to such profiling and inference. Inferring roaming in this way typically requires traffic analysis across multiple networks, or colluding attackers, or both. When location privacy is compromised, it could lead to more targeted profiling of user communication.

As can be seen, the location privacy problem spans multiple protocol layers. Nevertheless, we can examine problems encountered by nodes using a particular protocol layer. Roaming is particularly important to Mobile IP, which defines a global identifier (Home Address) that can reveal device roaming, and in conjunction with a corresponding user identifier (such as a SIP URI), can also reveal user roaming. Furthermore, a user may not wish to reveal roaming to correspondent(s), which translates to the use of a Care-of Address. As with a Home Address, the Care-of Address can also reveal the topological location of the Mobile Node.

This document scopes the problem of location privacy for the Mobile IP protocol. The primary goal is to prevent attackers on the path between the Mobile Node (MN) and the Correspondent Node (CN) from detecting roaming due to the disclosure of the Home Address. The attackers are assumed to be able to observe, modify, and inject traffic at one point between the MN and the CN. The attackers are

assumed not to be able to observe at multiple points and correlate observations to detect roaming. For this reason, MAC addresses, IIDs, and other fields that can be profiled to detect roaming are only in scope to the extent that they can be used by an attacker at one point. Upper layer protocol identifiers that expose roaming are out of scope except that a solution to the problem described here needs to be usable as a building block in solutions to those problems.

This document also considers the problem from the exposure of a Care-of Address to the CN.

This document is only concerned with IP address location privacy in the context of Mobile IPv6. It does not address the overall privacy problem. For instance, it does not address privacy issues related to MAC addresses or the relationship of IP and MAC addresses [HADDAD], or the Upper Layer Protocol addresses. Solutions to the problem specified here should provide protection against roaming disclosure due to using Mobile IPv6 addresses from a visited network.

This document assumes that the reader is familiar with the basic operation of Mobile IPv6 [RFC3775] and the associated terminology defined therein. For convenience, we provide some definitions below.

2. Definitions

- o Mobile Node (MN): A Mobile IPv6 Mobile Node that freely roams around networks
- o Correspondent Node (CN): A Mobile IPv6 that node corresponds with an MN
- o Home Network: The network where the MN is normally present when it is not roaming
- o Visited Network: A network that an MN uses to access the Internet when it is roaming
- o Home Agent: A router on the MN's home network that provides forwarding support when the MN is roaming
- o Home Address (HoA): The MN's unicast IP address valid on its home network
- o Care-of Address (CoA): The MN's unicast IP address valid on the visited network

- o Reverse Tunneling or Bidirectional Tunneling: A mechanism used for packet forwarding between the MN and its Home Agent
- o Route Optimization: A mechanism that allows direct routing of packets between a roaming MN and its CN, without having to traverse the home network

3. Problem Definition

3.1. Disclosing the Care-of Address to the Correspondent Node

When a Mobile IP MN roams from its home network to a visited network or from one visited network to another, use of Care-of Address in communication with a correspondent reveals that the MN has roamed. This assumes that the correspondent is able to associate the Care-of Address to the Home Address, for instance, by inspecting the Binding Cache Entry. The Home Address itself is assumed to have been obtained by whatever means (e.g., through DNS lookup).

3.2. Revealing the Home Address to Onlookers

When a Mobile IP MN roams from its home network to a visited network or from one visited network to another, use of a Home Address in communication reveals to an onlooker that the MN has roamed. When a binding of a Home Address to a user identifier (such as a SIP URI) is available, the Home Address can be used to also determine that the user has roamed. This problem is independent of whether the MN uses a Care-of Address to communicate directly with the correspondent (i.e., uses route optimization), or the MN communicates via the Home Agent (i.e., uses reverse tunneling). Location privacy can be compromised when an onlooker is present on the MN - CN path (when route optimization is used). It may also be compromised when the onlooker is present on the MN - HA path (when bidirectional tunneling without encryption is used; see below).

3.3. Problem Scope

With existing Mobile IPv6 solutions, there is some protection against location privacy. If a Mobile Node uses reverse tunneling with Encapsulating Security Payload (ESP) encryption, then the Home Address is not revealed on the MN - HA path. So, eavesdroppers on the MN - HA path cannot determine roaming. They could, however, still profile fields in the ESP header; however, this problem is not specific to Mobile IPv6 location privacy.

When an MN uses reverse tunneling (regardless of ESP encryption), the correspondent does not have access to the Care-of Address. Hence, it cannot determine that the MN has roamed.

Hence, the location privacy problem is particularly applicable when Mobile IPv6 route optimization is used or when reverse tunneling is used without protecting the inner IP packet containing the Home Address.

4. Problem Illustration

This section is intended to provide an illustration of the problem defined in the previous section.

Consider a Mobile Node at its home network. Whenever it is involved in IP communication, its correspondents can see an IP address valid on the home network. Elaborating further, the users involved in peer-to-peer communication are likely to see a user-friendly identifier such as a SIP URI, and the communication endpoints in the IP stack will see IP addresses. Users uninterested in or unaware of IP communication details will not see any difference when the MN acquires a new IP address. Of course, any user can "tcpdump" or "ethereal" a session, capture IP packets, and map the MN's IP address to an approximate geo-location. This mapping may reveal the home location of a user, but a correspondent cannot ascertain whether the user has actually roamed or not. Assessing the physical location based on IP addresses has some similarities to assessing the geographical location based on the area code of a telephone number. The granularity of the physical area corresponding to an IP address can vary depending on how sophisticated the available tools are, how often an ISP conducts its network re-numbering, etc. Also, an IP address cannot guarantee that a peer is at a certain geographic area due to technologies such as VPN and tunneling.

When the MN roams to another network, the location privacy problem consists of two parts: revealing information to its correspondents and to onlookers.

With its correspondents, the MN can either communicate directly or reverse-tunnel its packets through the Home Agent. Using reverse tunneling does not reveal the Care-of Address of the MN, although end-to-end delay may vary depending on the particular scenario. With those correspondents with which it can disclose its Care-of Address "on the wire", the MN has the option of using route-optimized communication. The transport protocol still sees the Home Address with route optimization. Unless the correspondent runs some packet capturing utility, the user cannot see which mode (reverse tunneling or route optimization) is being used, but knows that it is communicating with the same peer whose URI it knows. This is similar to conversing with a roaming cellphone user whose phone number, like the URI, remains unchanged.

Regardless of whether the MN uses route optimization or reverse tunneling (without ESP encryption), its Home Address is revealed in data packets. When equipped with an ability to inspect packets "on the wire", an onlooker on the MN - HA path can determine that the MN has roamed and could possibly also determine that the user has roamed. This could compromise the location privacy even if the MN took steps to hide its roaming information from a correspondent.

The above description is valid regardless of whether a Home Address is statically allocated or is dynamically allocated. In either case, the mapping of IP address to a geo-location will most likely yield results with the same level of granularity. With the freely available tools on the Internet, this granularity is the physical address of the ISP or the organization that registers ownership of a prefix chunk. Since an ISP or an organization is not, rightly, required to provide a blueprint of its subnets, the granularity remains fairly coarse for a mobile wireless network. However, sophisticated attackers might be able to conduct site mapping and obtain more fine-grained subnet information.

A compromise in location privacy could lead to more targeted profiling of user data. An eavesdropper may specifically track the traffic containing the Home Address, and monitor the movement of the Mobile Node with a changing Care-of Address. The profiling problem is not specific to Mobile IPv6, but could be triggered by a compromise in location privacy due to revealing the Home Address. A correspondent may take advantage of the knowledge that a user has roamed when the Care-of Address is revealed, and modulate actions based on such knowledge. Such information could cause concern to a mobile user, especially when the correspondent turns out to be untrustworthy. For these reasons, appropriate security measures on the management interfaces on the MN to guard against the disclosure or misuse of location information should be considered.

Applying existing techniques to thwart profiling may have implications to Mobile IPv6 signaling performance. For instance, changing the Care-of Address often would cause additional Return Routability [RFC3775] and binding management signaling. And, changing the Home Address often has implications on IPsec security association management. Careful consideration should be given to the signaling cost introduced by changing either the Care-of Address or the Home Address.

When roaming, an MN may treat its home network nodes as any other correspondents. Reverse tunneling is perhaps sufficient for home network communication, since route-optimized communication will traverse the identical path. Hence, an MN can avoid revealing its Care-of Address to its home network correspondents simply by using

reverse tunneling. The Proxy Neighbor Advertisements [RFC2461] from the Home Agent could serve as hints to the home network nodes that the Mobile Node is away. However, they will not be able to know the Mobile Node's current point of attachment unless the MN uses route optimization with them.

5. Conclusion

In this document, we have discussed the location privacy problem as applicable to Mobile IPv6. The problem can be summarized as follows: disclosing the Care-of Address to a correspondent and revealing the Home Address to an onlooker can compromise the location privacy of a Mobile Node, and hence that of a user. We have seen that bidirectional tunneling allows an MN to protect its Care-of Address to the CN. And, ESP encryption of an inner IP packet allows the MN to protect its Home Address from the onlookers on the MN - HA path. However, with route optimization, the MN will reveal its Care-of Address to the CN. Moreover, route optimization causes the Home Address to be revealed to onlookers in the data packets as well as in Mobile IPv6 signaling messages. The solutions to this problem are expected to be protocol specifications that use the existing Mobile IPv6 functional entities, namely, the Mobile Node, its Home Agent, and the Correspondent Node.

6. Security Considerations

This document discusses the location privacy problem specific to Mobile IPv6. Any solution must be able to protect (e.g., using encryption) the Home Address from disclosure to onlookers in data packets when using route optimization or reverse tunneling. In addition, solutions must consider protecting the Mobile IPv6 signaling messages from disclosing the Home Address along the MN - HA and MN - CN paths.

Disclosing the Care-of Address is inevitable if an MN wishes to use route optimization. Regardless of whether the Care-of Address is an on-link address of the MN on the visited link or that of a cooperating proxy, mere presence of a Binding Cache Entry is sufficient for a CN to ascertain roaming. Hence, an MN concerned with location privacy should exercise prudence in determining whether to use route optimization with, especially previously unacquainted, correspondents.

The solutions should also consider the use of temporary addresses and their implications on Mobile IPv6 signaling as discussed in Section 4, "Problem Illustration". Use of IP addresses with privacy extensions [RFC3041] could be especially useful for Care-of Addresses

if appropriate trade-offs with Return Routability signaling are taken into account.

7. Acknowledgments

James Kempf, Qiu Ying, Sam Xia, and Lakshminath Dondeti are acknowledged for their review and feedback. Thanks to Jari Arkko and Kilian Weniger for the last-call review and for suggesting improvements and text. Thanks to Sam Hartman for providing text to improve the problem scope.

8. References

8.1. Normative References

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

8.2. Informative References

[HADDAD] Haddad, W., et al., "Privacy for Mobile and Multi-homed Nodes: Problem Statement" Work in Progress, June 2006.

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

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

- [RFC3825] Polk, J., Schnizlein, J., and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information", RFC 3825, July 2004.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

Appendix A. Background

The location privacy topic is broad and often has different connotations. It also spans multiple layers in the OSI reference model. Besides, there are attributes beyond an IP address alone that can reveal hints about location. For instance, even if a correspondent is communicating with the same endpoint it is used to, the "time of day" attribute can reveal a hint to the user. Some roaming cellphone users may have noticed that their SMS messages carry a timestamp of their "home network" time zone (for location privacy or otherwise), which can reveal that the user is in a different time zone when messages are sent during a "normal" time of the day. Furthermore, tools exist on the Internet that can map an IP address to the physical address of an ISP or the organization that owns the prefix chunk. Taking this to another step, with built-in GPS receivers on IP hosts, applications can be devised to map geo-locations to IP network information. Even without GPS receivers, geo-locations can also be obtained in environments where "Geopriv" is supported, for instance, as a DHCP option [RFC3825]. In summary, a user's physical location can be determined or guessed with some certainty and with varying levels of granularity by different means, even though IP addresses themselves do not inherently provide any geo-location information. It is perhaps useful to bear this broad scope in mind as the problem of IP address location privacy in the presence of IP Mobility is addressed.

Author's Address

Rajeev Koodli
Nokia Siemens Networks
313 Fairchild Drive
Mountain View, CA 94043

EEmail: rajeev.koodli@nokia.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.