Network-based Mobility Management (Netlmm)
Internet-Draft
Intended status: Standards Track
Expires: May 11, 2008

W. Haddad
S. Krishnan
Ericsson Research

November 11, 2007

On Providing Light SeND and Privacy Extensions for Proxy MIPv6 (PMIPv6)

draft-haddad-netlmm-pmipv6-privacy-00

Status of this Memo

Copyright Notice

Abstract

This document describes a light and CGA free version of the secure neighbor discovery protocol combined with a privacy extension for the Proxy Mobile IPv6 protocol.

Table of Contents

1. Introduction

Proxy Mobile IPv6 protocol (described in [PMIPv6]) is an ongoing activity, which aims essentially to provide network based mobility. The main concept is to trick the mobile node (MN) into believing that it is always attached to its home network even when in reality, it has switched to foreign network(s). Consequently, the MN can keep using it IP home address (HoA) while being located away from its home network.

This document describes a mechanism which combines a light and CGA free version of the secure neighbor discovery (described in [SeND]) combined with a privacy extension for PMIPv6 protocol. At this stage, the light SeND (LiSeND) protocol enables the MN and its access router, i.e., the mobility access gateway (MAG), to authenticate the exchange of router solicitation (RtSol) and advertisement (RtAdv) messages and removes the need for running duplicate address detection (DAD) on the MN side (for more details on RtAdv/RtSol and DAD, please refer to the Neighbor Discovery Protocol described in [NDP]). Another key feature lies in the fact that LiSeND does not require the crypto generated address technique (described in [CGA]) deployment on both the infrastructure and the MN sides.

Building on LiSeND, we then describe a simple privacy extension which enables to mask the MN's HoA in a visited network(s) and thus, prevents an eavesdropper from learning, identifying and tracing the MN. A side effect of the suggested proposal is a mechanism which removes the harmful impact on the MN's ongoing sessions in case of a possible duplicate address detection (DAD) failure.

2. Conventions used in this document

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 [TERM].

3. Terminology

Since our proposal is mainly designed for network-based mobility, we borrow the terminology used in PMIPv6 and refer to a mobility (un)aware by MN. The main privacy aspects definitions are defined in [PRITERM]. Finally, we reuse the tunneling optimization mechanism and terminology that we have introduced earlier in [TOM].

4. Background and Motivation

Being a network-base mobility, PMIPv6 achieves its goal by enabling the MN to retain the same IP address while roaming between different foreign networks and delegates the task of securely discovering and updating the MN's Local Mobility Anchor (LMA) to the MAG.

The MAG fulfills its task by sending a proxy binding update (PBU) message to request binding (and potentially assigning) the MN's home network prefix (HNP) to its own egress interface address as being the MN's care-of address (CoA). Following a successful update, the MN's LMA starts tunneling data packets sent from the correspondent node(s) (CN) (i.e., which is kept totally unaware about the MN's mobility) to the MN's CoA, i.e., MAG's egress interface address. The MAG then decapsulates each data packet sent to the MN's CoA and forwards it to the MN. It follows, as mentioned earlier, that the MN will always believe that it is still attached to its home network, especially that the MAG takes great care of nurturing the MN's belief by advertising in unicast mode its home prefix in order to convince it to (re)-configure its HoA.

Our motivation is mainly guided by a requirement in EU and some Asian countries and from a general desire in others (!!), to protect the MN's privacy while switching to and moving across foreign networks. Such protection should enable to provide anonymity and unlinkability aspects whenever possible. Consequently, privacy protection in a PMIPv6 environment means first and foremost preventing the MN from disclosing its HoA within any foreign network and removing any "linkability" when switching to a new MAG. It follows that an efficient privacy protection should enable the MN to mask its HoA and to update the mask each time it switches to a new MAG. It is noteworthy that updating the mask becomes more efficient for protecting the MN's privacy if it is applied at higher frequency than just when switching to a new MAG. While the suggested proposal enables such enhancement, we don't detail it in this version.

A malicious node (acting independently), has two topological locations from which, it can learn/detect the MN's HoA (or maybe the HNP only) and use it to trace the MN's subsequent movements. The first location is anywhere on-path between the MAG and the MN's LMA. In such location, the eavesdropper is able to check the inner header carried in each data packet sent by the MN's HA to its current MAG. The data packet inner header carries the MN's HoA and the CN's IP address. With such ability, the eavesdropper can rely for example, on som prior knowledge/hint to uncover the targeted MN's current whereabout and even "lock" on it.

The second location is the link to which the MN is attached. In such

location, the eavesdropper is also able to identify the MN and trace its movements. In addition, other static identifier(s), e.g., the MN's MAC address, become available and may significantly contribute in detecting and tracing the MN. However, we consider out of scope all parameters below the IP layer but we carefully note that our proposal can also be extended to cover the MAC address, i.e., by extending the scope of the mask. Privacy issues related to the MAC layer are detailed in [ANON].

The unlinkability protection can be seriously endangered each time the MN switches to a foreign network and keeps using its CGA public key from which, it has generated its HoA. In fact, the requirement behind PMIPv6 design to enable the MN to retain its HoA means also that the MN won't be able to change its CGA public key as its HoA won't remain the same. Consequently, if an eavesdropper learns the MN's public key (which is far from being a problem!), then it will always be able to trace the MN after switching to a new MAG(s). In addition, as applying the mask will generate a pseudo-IPv6 (pIPv6) address, it is of high importance to make sure that pIPv6 won't be reused when switching to a new MAG.

The picture of the two separate threats scenarios described above becomes rapidly more complicated when they are combined. This is the case where at least two malicious nodes with each following one of the two scenarios, are coordinating their "search, identify and trace" activities. In such case, an efficient anonymity and unlinkability protection can be obtained by simply merging the two solutions addressing each of the the two scenarios. In other words, the MN's HoA MUST NOT be disclosed neither on the link between the MN and its MAG nor anywhere between the MAG and the MN's LMA. Also, the MN MUST avoid (re)-using its CGA public key and the pIPv6 MUST be refreshed after the handoff.

In the following section, we address the above scenarios separately and describe two mechanisms to reduce the eavesdropper's ability to learn the MN's HoA and/or trace its movements. The combination of the two protections is highlighted in the last section.

5. Proposed Solutions

5.1. Solution for On-path attacks

Our proposed mechanism addresses the first scenario where an eavesdropper is located on-path between the MN's MAG and the LMA by completely eliminating the need for disclosing the MN's HoA in any data packet sent to the MN's PMA. This is achieved by removing it from each data packet exchanged between the MN's LMA and its current MAG, via applying the tunneling optimization (TO) mechanism. As in the Mobile IPv6 case [MIP6], the TO mechanism can be securely applied during the PBU/PBA messages exchange between the MAG and the LMA nodes. In this case, the PMIPv6 signaling exchange should lead both sides to create a PaT, which can be immediately applied on each data packet sent by the MN to the correspondent node (CN) and/or tunneled from the HA to the MN's CoA, i.e., the MAG. This results in a complete removal of the MN's HoA from the path between the MAG and the LMA. Consequently, implementing such optimization significantly complicates the eavesdropper's task of identifying the targeted MN from snooping into data packets flow(s) exchanged between the MN's MAG and its LMA. Note that if the MN is using multiple HoAs as it may be talking to different CNs, then the PMA and HA will have to generate one PaT for each HoA.

In addition to removing the MN's HoA from data packets and in order to enhance the unlinkability aspect, it is highly recommended that the MAG refreshes periodically the MN's CoA, i.e., its own IP egress interface address, and updates all associated PaT(s) accordingly. This can be done by re-sending PBU message(s) to the LMA to update its binding cache entries table.

5.2. Solution for On-link attacks

As mentioned in the second threat scenario, when an eavesdropper is attached to the same link than the MN, it can easily detect the unicast RtAdv message sent by the MAG to the MN following a successful authentication and the receipt of a PBA message (unless the HNP is obtained via another way, e.g., from the AAA). However, as disclosing the MN's home network prefix (HNP) alone may be very sufficient for an eavesdropper to identify the MN, providing privacy protection to the MN requires a complete "blackout" on its HNP on any foreign link. Such requirement may also be raised within the MN's home network as it blocks malicious nodes from learning the MN's HoA even when it is still attached to its home network. Imposing an HNP blackout requires the MAG and/or LMA to send special parameters to the MN in order to enable it to (re)-configure its HoA and at the same time, generate a special PaT to translate its HoA to another pIPv6 address in each IP packet sent by the MN as well as in each IP

packet forwarded by the LMA/MAG to the MN. In addition, these parameters MUST be sent encrypted, which makes it tempting at this stage to turn to the CGA technique to achieve this particular goal (i.e., using the MN's CGA public key), then send them in the unicast RtAdv message. However, as mentioned earlier, using the MN's CGA public key provides the eavesdropper just another valuable tool to identify and trace the MN. In order to avoid the CGA impasse, a new secret called "privacy key (Kp)" should be computed and stored in the AAA. Computing Kp should be performed when authenticating the MN for the first time. Note that generating Kp can occur when the MN is attached to a foreign network. The mechanism(s) to be used to compute Kp is out of scope of this document.

5.3. Operation

After generating and storing Kp, the AAA may decide to share it with the MN's LMA. But, in general, Kp is used by the MN and the AAA to derive the "transient handoff key (THK)", which is then sent to the MN's current MAG. THK can be sent to the MAG directly by the AAA, following a successful authentication or by the LMA in the PBA message. This means that each time the MN has to perform an authentication, a new THK is computed and sent to the current MAG. In addition to refreshing THK, the AAA and the MN SHOULD also generate a pseudo-NAI (pNAI) and bind it to the new TKh to be used. The new pNAI is used by the MN during the next authentication and is carried by the PBU message sent by the MAG to the LMA. It follows that the pNAI MUST be sent to the LMA prior to receiving a PBU message (e.g., after a successful authentication).

Upon receiving a THK, the MAG SHOULD use it to derive a PaT and to encrypt the HNP sent by the LMA in order to be advertised to the MN. For this purpose, the MAG has to signal to the MN its capability to offer anonymity and unlinkability services. This is done by setting a new bit called "Privacy (P)" bit in the unicast RtAdv message sent to the MN. The presence of the "P" bit also indicates to the MN that it has to generate the PaT which corresponds to the THK computed by the MN. One way to handle the "P" bit is to set it in the same new option (called "Unicast RtAdv Authentication (URA)") carrying the message authentication. Moreover, the MAG MUST use THK to authenticate all unicast RtAdv messages sent to the MN. Similarly, the MN MUST use THK to authenticate any RtSol message sent to the MAG.

Upon receiving a RtAdv message carrying a URA (i.e., following a successful authentication), the MN proceeds first to check the message validity with its own THK. If the message is valid then the MN decrypts the HNP, configures its HoA and generates the corresponding PaT. Otherwise, the MN should silently discard the

message.

The PaT SHOULD be applied by the MN on each data packet sent to the CN. The MAG SHOULD apply the PaT on each IP packet sent to the MN's HoA and on each data packet sent by the MN. It follows that the MN's HNP is never disclosed on the link. It should be noted that for the purpose of enhancing the unlinkability while being attached to the same link, it is highly recommended to periodically refresh the PaT.

In addition to hiding the HNP advertised to the MN, the MAG SHOULD also run the DAD procedure on the MN's new pIPv6 address before advertising the HNP to the MN. In case of a collision, the MAG SHOULD randomly generate a unique pseudo-HNP then share it with the MN. This is achieved by XORing a random 64-bit parameter with the corresponding PaT then with the HNP and testing its uniqueness. The MAG SHOULD then send the 64-bit parameter to the MN in a new option (called "Pseudo Home Network Prefix (PHNP)") carried by the RtAdv message. Note that the PHNP MUST be encrypted with THK.

In the unlikely event leading to inserting a PHNP option in the RtAdv message, the MN MUST use the 64-bit pseudo-HNP to update the PaT generated from THK. This is done by XORing the 64-bit pseudo-HNP with the PaT in order to enable generating a new pIPv6 when the data packet header is XORed with the updated PaT.

It follows from the above, that in order to avoid breaking the MN's anonymity during an NDP exchange between two MNs, the MAG SHOULD also act as the "reference" node for any NDP queries. Such enhancement will enable LiSeND to provide most of features provided by SeND protocol.

In order to address merging privacy threats scenarios described earlier, the two mechanisms have to be combined in order to protect the path between the MAG and the LMA and at the same time, the one between the MN and the MAG against malicious eavesdroppings. Doing so provides anonymity and unlinkability features to the MN on the path between the MN to its LMA. This means that the MN's MAG SHOULD apply a PaT when dealing with the corresponding LMA and another one when dealing with the MN in order to mask the MN's HoA or its HNP from each data packet exchanged between the MN and the CN and/or from signaling messages exchanged between the MAG and the LMA.

5.4. Future work

Future versions of this work will probe further enhancements for the LiSeND protocol and possible stretching of anonymity and unlinkability extensions down to the MAC layer.

6. New Option Format

TBD

7. Security Considerations

This document introduces first LiSeND protocol which is a light and CGA free version of SeND protocol. LiSeND is shown to be adapted to the concept of network based mobility where PMIPv6 protocol is a leading candidate. We describe next how key privacy aspects can be build on top of LiSeND in a seamless way which does not affect the MN's unawareness about its mobility.

Our proposal relies on sharing a different "privacy key" between the MN and each MAG visited by the mobility unaware MN. For this purpose, and considering the main clients for deploying PMIPv6, we adopt a realistic approach centered around the existence of a AAA infrastructure which will enable the MN to be authenticated upon attachment to foreign network(s). We also consider that the MN is able to derive the corresponding transient handoff key (THK) after each successful authentication.

However, in order to avoid sharing the same THK between three different nodes (i.e., MN, LMA and MAG), and in order to enable LiSeND to protect the MN against compromised MAG, we suggest sending a hash of the current THK (H_THK) to the MN's LMA. It follows that each MAG MUST include H_THK in the PBU message sent to the LMA following a successful authentication of the MN.

8. References

8.1. Normative References

[CGA] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3792, March 2005.

[MIP6] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support for IPv6", RFC 3775, June 2004.

[NDP] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

[SeND] Arkko, J., Kempf, J., Sommerfield, B., Zill, B., and P. Nikander, "Secure Neighbor Discovery (SeND)", RFC 3971, March 2005.

[TERM] Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels", RFC 2119, BCP , March 1997.

8.2. Informative References

[ANON] Haddad, W., Nordmark, E., Dupont, F., Bagnulo, M., Patil, B., and H. Tschofenig, "Anonymous Identifiers (ALIEN): Privacy Threat Model for Mobile and Multi-Homed Nodes", Internet Draft, draft-haddad-alien-threat-model-00.txt, January 2007.

[PMIPv6] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", Internet Draft, draft-ietf-netlmm-proxymip6-06.txt, september 2007.

[PRITERM] Haddad, W. and E. Nordmark, "Privacy Terminology", Internet Draft, draft-haddad-alien-privacy-terminology-03.txt, October 2007.

[TOM] Haddad, W., Naslund, M., and P. Nikander, "IP Tunneling Optimization in a Mobile Environment", Internet-Draft, draft-haddad-mip6-tunneling-optimization-01.txt, July 2007.

Authors' Addresses

Wassim Haddad
Ericsson Research
Torshamnsgatan 23
SE-164 80 Stockholm
Sweden

Phone: +46 8 4044079
Email: Wassim.Haddad@ericsson.com

Suresh Krishnan
Ericsson Research
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Phone: +1 514 345 7900
Email: Suresh.Krishnan@ericsson.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).