Mobility Extensions
Internet-Draft
Intended status: Informational
Expires: August 28, 2008

W. Haddad

G. Tsirtsis
Qualcomm

B. Lim
Panasonic

S. Krishnan
Ericsson Research

February 25, 2008

Mobile IPv6 Residual Threats
draft-haddad-mext-mip6-residual-threats-01

Status of this Memo

Copyright Notice

Abstract

This memo aims to highlight specific residual threats in the Mobile IPv6 design. These threats are inherited in the design of new mechanisms built on top of the mobility protocol, and are raising concerns regarding the amplitude of their potential impacts.

Table of Contents

1. Introduction

The design of Mobile IPv6 protocol (described in [MIPv6]) did not address a set of specific threats for various reasons. In fact, these residual threats were (rightly or not) not considered of equal importance than others which required immediate action. However, as these threats are implicitely inherited in the design of new mechanisms built on top of MIPv6, their potential impact is raising some concerns.

This memo aims to describe these residual threats and to motivate designers to take a fresh look at the reasoning behind leaving them in an unconvincing state and to address them in case it is needed.

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. Residual Threats Associated with a Malicious Mobile Node

3.1. Violating Trust between the Mobile Node and its Home Agent

The trust model that guided MIPv6 protocol design was based on two main assumptions. The first one considers that the mobile node (MN) will always refrain from misusing the relationship with its home agent (HA). It follows that the HA can always blindly accept any information sent by the MN. Thus, the second assumption requests the HA to accept any new care-of address(es) (CoA(s)) claimed by the MN and sent in a valid binding update (BU) message(s).

The justification behind the first assumption is the tracing ability that the HA is supposed to always acquire over the MN. In other words, the MN is always supposed to fear being traced -as the target is also supposed to complain- thus, refrain from misusing a mutual trust which can appear now as being "imposed" on both nodes. We argue that such assumption is too naive when applied in the real world as it is simply impossible to back it with a solid proof that confirm the existence of enough anxiety to deter all potential attackers from launching malicious act in such particular context.

In fact, the lack of any reachability test between the MN and its HA, prior to or after sending a BU message, enables a malicious MN to launch a network flooding attack against any potential target by simply claiming a new CoA which is topologically located within the targeted network. This is especially possible when only the bidirectional tunneling (BT) mode is used and/or when the enhanced route optimization mode (described in [ERO]) is in use. Without testing the new CoA reachability, the HA will simply re-route data packets to the new CoA, i.e., targeted network, and the MN can keep sending acknowledgment messages to all its CN(s) in order to maintain the attack as long as needed. Note that this type of attack is not new as it has been well analyzed in [MROD] and its effects on the targeted network are mitigated by imposing a periodic return routability procedure.

Moreover, as MIPv6 is acquiring many extensions, such attack on the HA side may get amplified with the MN's ability to register multiple CoAs with the HA. Such scenario is better described in [Multih_Sec].

3.2. Violating Trust between a Multihomed Mobile Node and its Home Agents

Multiple Care-of Address registration (MCoA) protocol (described in [MCoA]) extends the MIPv6 protocol to enable a multi-interface MN to register multiple CoAs at its HA. The fundamental difference between MIPv6 and MCoA is that for a given home address (HoA) in MIPv6, the

MN is only able to bind a single 'fake' CoA. Hence, this implies that once a malicious MN binds the 'fake' CoA at HA, that MN loses its ability to use that HoA for communication. However, in MCoA, with the ability to bind several CoAs to a single HoA, a malicious MN could bind a mixture of 'real' and 'fake' CoAs. The MN can still use the HoA for communication by directing control traffic towards its 'real' CoA.

Likewise in the trust model described in [MROD] between the MN and its HA, it permits the HA to 'blindly' accept any binding that the MN makes. This trust relationship is further strengthened when one assume that ingress filtering is being used such that when the HA receives a BU message from the MN stating its CoA as the source address, the HA trusts that the incoming packets do indeed originate from the specified source address. In addition, the HA also trusts the routing infrastructure that packets forwarded by the HA would be sent to the intended destination. This assumption makes it possible for the HA to somewhat trust the MN if the MN sends the binding of each CoA individually (e.g. one BU message per CoA).

However, such a trust is no longer valid when the MN utilizes a single BU message to register its multiple CoAs at its HA. This technique is explained in [MCoA] with the aim of introducing some optimization when registering multiple CoAs for a MN. Such optimization technique is useful in scenarios when resources (e.g. bandwidth) are scarce on some of the MN's interfaces, since it allows the MN to send a BU message containing multiple CoAs to its HA from an interface that does not have such resource constraint. Moreover, in MIPv6, the use of the alternate CoA option permits the MN to achieve the same effect of registering a CoA for another interface via a specific interface. This introduces the risk of having 'fake' CoAs registered at the HA and compromise the security of the network. If these 'fake' CoAs are IP addresses of other MNs, it is possible for the malicious MN to instruct the HA to re-direct traffic towards these 'fake' CoAs, thereby flooding these MNs with useless traffic.

With such a threat towards the network, some mechanism might be need in order for the HA to verify the CoA for the MN prior to binding or using them for packet routing.

4. Security Considerations

This document is a security analysis of some specific parts in the MIPv6 protocol. It describes residual threats in the protocol which may not appear necessarily new for the reader. It is worth noting in this context that most of the above threats can be mitigated by ingress filtering. However, the particular case where a malicious mobile node provides its HA with fake CoA(s) configured with the same subnet prefix as the one it is using cannot be countered by ingress filtering and requires additional mechanisms to block it.

5. References

5.1. Normative References

[ERO] Vogt, C., Arkko, J., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6", RFC 4866, June 2006.

[MIPv6] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support for IPv6", RFC 3775, June 2004.

[MROD] Nikander, P., Arkko, J., Aura, T., and E. Nordmark, "Mobile IPv6 version 6 Route Optimization Security Design Background", RFC 4225, December 2005.

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

5.2. Informative References

[MCoA] Wakikawa, R., Ernst, T., Nagami, K., and V. Devarapalli, "Multiple Care-of Addresses Registration", Internet Draft, draft-ietf-monami6-multiplecoa-05.txt, January 2008.

[Multih_Sec] Lim, B., Ng, C., and K. Aso, "Verification of Care-of Addresses in Multiple Bindings Registration", Internet Draft, draft-lim-mext-multiple-coa-verify-00.txt, November 2007.

Authors' Addresses

Wassim Haddad

Email: wmhaddad@gmail.com

Georges Tsirtsis
Qualcomm

Phone: +908 443 8174
Email: tsirtsis@qualcomm.com

Benjamin Lim
Panasonic Singapore Laboratories Pte Ltd
Blk 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415

Phone: +65 65505478
Email: benjamin.limck@sg.panasonic.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).