Mobility Extensions   W. Haddad
Internet-Draft   G. Tsirtsis
Intended status: Informational   Qualcomm
Expires: January 15, 2009   B. Lim
                                                             Panasonic
                                                           S. Krishnan
                                                              Ericsson
                                                             F. Dupont
                                                                   ISC
                                                         July 14, 2008

                    Mobile IPv6 Residual Threats
              draft-haddad-mext-mip6-residual-threats-02

Status of this Memo

   By submitting this Internet-Draft, each author represents that any
   applicable patent or other IPR claims of which he or she is aware
   have been or will be disclosed, and any of which he or she becomes
   aware will be disclosed, in accordance with Section 6 of BCP 79.

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF), its areas, and its working groups. Note that
   other groups may also distribute working documents as Internet-
   Drafts.

   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."

   The list of current Internet-Drafts can be accessed at
   http://www.ietf.org/ietf/1id-abstracts.txt.

   The list of Internet-Draft Shadow Directories can be accessed at
   http://www.ietf.org/shadow.html.

   This Internet-Draft will expire on January 15, 2009.

Abstract

   This memo aims to highlight specific "residual" threats in Mobile
   IPv6 protocol. We call these threats "residual" simply because they
   were rightfully deemed not urgent during the design of Mobile IPv6
   protocol. However, these threats are somehow benefiting from new
   mechanisms and/or extensions built on top of Mobile IPv6 protocol to
   improve their effects and likelihood. Hence, our main goal is to
   motivate designers to re-assess their potential taking into
   consideration these new mechanisms.

Table of Contents

1. Introduction

Mobile IPv6 protocol design (described in [MIPv6]) involved extensive
security analysis in order to evaluate the potential of each threat
and suggest defensive measures when necessary or avoid adding
complexities in case a security weakness was deemed acceptable (i.e.,
does not make IPv6 Internet more secure than without IP mobility).

However, these weaknesses have been implicitly inherited in new
mobility mechanisms and/or extensions built on top of MIPv6 which may
in turn have increased their effects and thus, made them more
attractive.

This memo aims to describe these residual threats and to motivate
designers to re-assess their potential in the light of what has been
added so far to MIPv6 protocol.

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 adopted in MIPv6 protocol assumes that the mobile
   node (MN) will always refrain from misusing the relationship it
   forges with its home agent (HA). In return, the HA should treat any
   legitimate information sent by the MN as being trustable. For
   example, the HA will accept any new care-of address (CoA) claimed by
   the MN and sent in a valid binding update (BU) message(s).

   In fact, there are two interrelated factors for expecting a well
   behaving MN. First, the MN is expected to be fully aware about the
   HA tracing capabilities coupled with a strong authentication.
   Second, there is a high probability that the victim will complain to
   the operator in which case, the MN is quickly identified and punitive
   actions can be taken against it.

   However, the situation changes when the first factor is no longer
   valid. This is the case where a user gains access to the operator
   network without strong identification (e.g., prepaid phone card). In
   such scenario, a malicious behavior can go unpunished since the

operator is unable to trace the user. The malicious behavior can
consist of sending to the HA a valid BU message carrying a fake CoA,
which triggers traffic redirection towards the victim. While the
target can always complain to the attacker's operator, the latter can
do little or nothing about it (of course, the attack will eventually
stop when the credit on the prepaid card is entirely consumed or the
prepaid card itself expires). Note that while some countries have
imposed some restrictions on using prepaid card, e.g, by requiring
additional identification and denying roaming service, other
countries are still allowing them without control.

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 flooding attack against any potential target by simply
claiming a new CoA which seems to be topologically located within the
targeted network. Without testing the new CoA reachability, the HA
will simply re-route data packets to the new CoA, i.e., targeted
network, and the attacker can keep sending acknowledgment (ACK)
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 and has been analyzed in
[MROD]. In fact, when the route optimization (RO) mode is used, the
impact of such attack is mitigated by imposing a periodic return
routability procedure. Another way to protect against this attack is
to deploy ingress filtering (described in [INGRESS]).

Moreover, new added extensions to MIPv6 may enhance launching
flooding attack through the HA. This is due to the MN's ability to
register multiple CoAs with the HA. Such scenario is better
described in [Multih Sec] and is analyzed in the next section.

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 primary '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 always acknowledge any binding that the
MN requests. 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, i.e., 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).

Even with ingress filtering deployed worldwide in all networks, the
problem of the flooding attack described above could still be
achieved in the Multiple Care-of Address registration (MCoA) protocol
where the MN is able to use multiple binding identifier options in a
single binding update message to the home agent for registration
purposes. With the care-of addresses embedded inside the BU message,
it is not possible for ingress filtering to be used to verify these
CoAs. Figure 1 shows an example on how MCoA could be used to
initiate the redirection attack.

    [Start of packet header]

    Source Address   : CoA
    Destination Address   : HA's address

    [Mobility Options]

    Binding Unique Identifier: BID1

```
    Binding Unique Identifier: BID2
    Care-of Address         : V1's address

    Binding Unique Identifier: BID3
    Care-of Address         : V2's address

    [End of packet header]
```

                  Figure 1: Binding update message for MCoA

    CoA is a valid care-of address owned by MN. MN is attempting to bind
    addresses of two victims, V1 and V2, at HA in order to launch an
    attack towards the victims.

    When HA receives this BU message, it will accept it based on the
    following. First, the BU message is deemed authorized as the correct
    IPSec SA is used for the message. Second, the trust relationship
    that HA has with the routing infrastructure allows it to understand
    that this BU message is sent from MN. Finally, after the first two
    checks have succeeded, the trust relationship that HA has with the MN
    permits it to trust the care-of addresses that are specified in this
    BU message. Hence, the binding cache at HA will record three
    bindings for MN tied to MN's home address (HoA) as shown in Figure 2
    below.


```
    Binding 1 [HoA, CoA        , BID1]
    Binding 2 [HoA, V1's address, BID2]
    Binding 2 [HoA, V2's address, BID3]
```

                  Figure 2: Binding Cache at Home Agent

    The lack of any reachability test between the mobile node 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 new care-of addresses.

. Creating Routing Loops Among Home Agents

   In MIPv6, it is possible for a malicious MN to create a routing loop
   amongst HAs. This can be achieved when a MN binds one home address
   located on a first HA to another home address on a second HA. This
   type of binding will force HAs to route the same packet among each
   other without knowledge that a routing loop has been created. Such
   looping problem is limited to cases where a MN has multiple HAs. For
   the single case, MIPv6 has a policy at the HA to prevent the binding
   of one home address to another home address hosted by the same home
   agent. Figure 3 below shows such threat of routing loop between home
   agents.

```
                   +-----+
                   |     |
                   |  I  | cccc +-----+ Binding cache
                   |  N  |------| HA1 | [HoA1, HoA2]
       HoA1        |  T  |      +-----+
      +----+       |  E  |
      | MN |------ |  R  |
      +----+       |  N  | cccc +-----+ Binding cache
       HoA2        |  E  |------| HA2 | [HoA2, HoA1]
                   |  T  |      +-----+
                   |     |
                   +-----+
```
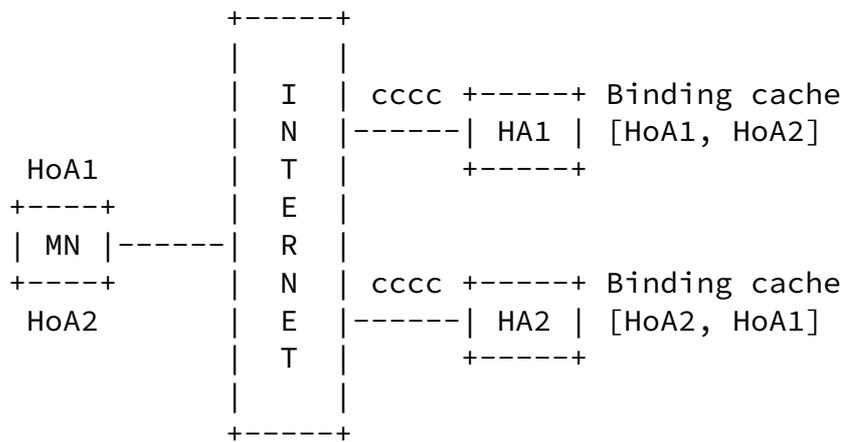
                 Figure 3: Packet flooding attack scenario

   The mobile node (MN) sends a binding update (BU) message to its first
   home agent (HA1) to bind its home address (HoA1) to a care-of address
   (HoA2). Since HoA2 is not a home address on HA1, HA1 accepts this
   binding thinking that HoA2 is indeed a CoA. Likewise, on HA2, MN
   sends a BU to bind its home address (HoA2) to a care-of address
   (HoA1). Such bindings created among the two HAs create a routing
   loop between them. For example, when HA1 wants to forward a packet
   (shown as 'c') from a CN to MN, HA1 searches the binding cache to
   find the relevant MN's CoA. In this case, HA1 tunnels the packet to
   MN via HoA2. This will cause HA2 to intercept the packet for MN.
   Now, at HA2, it sees that the packet is addressed to HoA2. Searching
   the respective binding entry in its binding cache, HA2 will tunnel
   this packet to MN via HoA1. This will cause HA1 to intercept the
   packet for MN. This repetition will continue until one of the HA
   discover that it can no longer encapsulate the packet (i.e., due to
   the tunnel encapsulation limit == 0 described in [GPT6]). In this
   case, the packet would be dropped and a flood of error message will

be sent between both HAs to indicate that the packet has failed to
reach its intended destination (the MN). Thus, it can be seen that
such attacks consumes the resources of the home agent and if launched
in full scale (e.g., multiple sets of HoAs) could 'shut down' the HA.

4. Exploiting Multihoming to Defeat Ingress Filtering

   A malicious multi-homed node can also use its multiple interfaces to
   emulate a home network and defeat ingress filtering. This is the
   case when an attacker with two interfaces (A) and (B) starts its
   attack by establishing sessions with a set of correspondent nodes
   (CNs) using (A)'s IPv6 address then at some point, attaches (B) to
   the targeted network and triggers a return routability (RR) procedure
   with the CN. As the RR procedure involves exchanging HoTI/HoT
   messages, the MN will use (A)'s IPv6 address for that purpose and
   receives a home keygen token. Then the MN exchanges a CoTI/CoT
   messages using (B)'s IPv6 address as a CoA and obtain a care-of
   keygen token. A BU message is then sent to the CN and the data
   traffic is redirected to (B).

   At some point, the MN detaches itself from the targeted network and
   start sending legitimate ACK messages via a legitimate address to
   each CN causing a flooding attack. Such scenario is mitigated by the
   fact that the MN MUST periodically repeat the RR procedure which
   means that it has to exchange CoTI/CoT messages with each CN.
   However, if the MN manages to position itself on-path with at least
   one CN without detaching (A) then it will be able to keep the attack
   as long as needed. Note that this attack becomes easier if the MN
   does not have to periodically repeat the RR procedure as a result of
   establishing a long lifetime security association with the CN, e.g.,
   when the enhanced RO mode ([ERO]) is used.

5. Exploiting Neighbor Discovery in a MIPv6 Environment

   Note: It may be asserted that this attack is related to Neighbor
   Discovery Protocol (described in [NDP]). However, our main goal is
   to convey a description about its potential which may go well beyond
   the local link when applied in a MIPv6 context.

   This threat offers a malicious node two edges. It requires first
   that the attacker be attached to the same foreign link as the MN, and
   the discovery of the MN's home agent IP address as well as the MN's
   IP home address (which may not pose a serious problem). After
   learning these two information, the attacker advertises the MN's home
   prefix on the link thus leading the MN to believe that it has
   returned to its home network. Such information will prompt the MN to
   send a BU message to its HA to request de-registration. However,
   such early de-registration may not be possible as the foreign network
   may have activated ingress filtering. But the main goal for the
   attacker is to get a valid copy of the MN's BU message and such goal
   is achieved. If the malicious node concludes that the MN is still
   receiving data packets tunneled by the HA to its current CoA, then it
   will get involved in the MN de-registeration procedure by forwarding
   the BU message to the MN's HA on another interface where ingress
   filtering is not activated (i.e., under the assumption that the
   attacker is multihomed). Upon receiving the BU message, the HA will
   de-register the MN and stops tunneling data packets to the MN's CoA.
   In addition, the HA sends back a BA message which will never reach
   the MN. From that moment, the data traffic sent by the CN(s) stops
   at the MN's home network. However, the lack of ACK messages sent by
   the MN will prompt the CN(s) at some point to halt sending data
   traffic and eventually tear down the session(s).

   However, the situation gets worse if the malicious node decides to

push further in his attack by sending fake ACK messages to the CN(s),
i.e., using the MN's home address. In such situation, the CN(s) will
keep sending data traffic to the MN's HA (where they eventually get
discarded) and thus, may cause severe disruption within the home
access network, possibly leading to a network flooding attack in some
specific topologies.

Note that as they may be more than one MN attached to the same
foreign link and using the same home prefix, such attack may lead to
collective de-registration.

6. Security Considerations

   This document is about security analysis of some specific parts in
   the MIPv6 protocol.

7. References

7.1. Normative References

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

   [GPT6]   Conta, A. and S. Deering, "Generic Packet Tunneling in
               IPv6", RFC 2473, December 1998.

   [INGRESS]   Ferguson, P. and D. Senie, "Network Ingress Filtering:
               Defeating Denial of Service Attacks which employ IP
               Source Address Spoofing", RFC 2827, May 2000.

   [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.

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

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

7.2. Informative References

    [MCoA]   Wakikawa, R., Ernst, T., Nagami, K., and V. Devarapalli,
                    "Multiple Care-of Addresses Registration", Internet
                    Draft, draft-ietf-monami6-multiplecoa-08.txt, May 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-01.txt,
                    February 2008.

Authors' Addresses

    Wassim Haddad
    Qualcomm
    500 Somerset Corporate Blvd
    Bridgewater, New Jersey 08807
    US

    Phone: +908 938 5027
    Email: whaddad@qualcomm.com

       Georges Tsirtsis
       Qualcomm
       London
       UK

       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
       8400 Decarie Blvd.
       Town of Mount Royal, QC
       Canada

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

       Francis Dupont
       ISC
       Rennes
       France

       Email: Francis.Dupont@fdupont.fr

Full Copyright Statement

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.