

Mobility for IPv6 (MIPv6) WG
Internet-Draft
Expires: January 10, 2008

W. Haddad
M. Naslund
Ericsson Research
P. Nikander
Ericsson Research Nomadic Lab
July 9, 2007

IP Tunneling Optimization in a Mobile Environment
draft-haddad-mip6-tunneling-optimization-01

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 10, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This memo introduces a simple tunneling optimization mechanism, which removes the need for inserting an additional header in the IP packet. The main goals are to minimize the packet size, provide a simpler protocol design and a better efficiency.

Table of Contents

1.	Introduction	3
2.	Conventions used in this document	4
3.	Motivation	5
4.	Tunneling Optimization for BT mode and HMIPv6	7
5.	Tunneling Optimization for R0 mode	11
6.	Tunneling Optimization for DSMIPv6	12
7.	Bit and Options Formats	13
8.	Security Considerations	14
9.	Acknowledgments	15
10.	References	16
10.1.	Normative References	16
10.2.	Informative References	16
	Authors' Addresses	17
	Intellectual Property and Copyright Statements	18

1. Introduction

IP tunneling is a mechanism widely used for different purposes. For example, it enables relying on a dedicated third party to modify the IP packet header, in order to re-route the packets to their new destination. This is mainly the case in a mobile environment where IP tunneling is used in most protocols. In fact, the mobile IPv6 bidirectional tunneling (BT) mode described in [[MIPv6](#)] uses IP tunneling to route data through the home agent (HA). The same mechanism is applied between the mobility anchor point (MAP) and the mobile node (MN) in the hierarchical mobile IPv6 (HMIPv6) protocol (described in [[HMIPv6](#)]). A modified IP tunneling version is also used in MIPv6 route optimization (RO) mode.

This memo introduces a simple tunneling optimization (TO) mechanism, which virtualizes the IP tunnel concept often used in traffic exchange. TO mechanism is based on securely exchanging a "pad translator" (PaT) between both sides of the (supposed) tunnel. The PaT main role is to translate incoming/outgoing packet header to another one before injecting it inside/outside the tunnel. TO main goals include a reduced packet size, a simpler protocol design and a better efficiency.

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. Motivation

The motivation behind this work is the widespread use of different forms of IP tunneling mechanisms in mobile environment and their negative impact on the protocol efficiency as translated in the data packet size, bandwidth usage and battery power consumption.

In the following, we consider IPv6 mobility protocols only and start with a brief description of the IP tunneling mechanism used in both MIPv6 BT mode and HMIPv6 protocol. Next, we describe the MIPv6 R0 mode tunneling mechanism, then show in the next section how the T0 mechanism completely eliminates the need for IP tunneling in these mobility protocols in a simple and elegant way.

Other mobility related protocols, e.g., Fast MIPv6 described in [FMIPv6] and the ongoing work on proxy MIPv6 ([PMIPv6]) can also benefit from using T0 mechanism, especially when the path between access routers (ARs) and/or between the HA and a PMIPv6 node(s) includes a wireless medium.

A closer look on the tunneling mechanism used in the BT mode (i.e., between the HA and the MN) and HMIPv6 protocol (i.e., between the MAP and the MN) highlights different issues that need to be addressed. A first one is the expansion of the packet size due to the addition of a new IP header. In the two protocols, such expansion is mainly equal to two IP addresses. This means that the MN has to send at least an additional 256 bits each time it transmits a data packet to the CN. The impact of such addition is very significant on the battery life. In fact, it has been shown in [EALDC] that wireless transmission of a single bit can require over 1000 times more energy than a single 32-bit computation.

Consequently, it would be more beneficial for the MN to avoid transmitting extra bits over the air interface in exchange for some additional computation only.

A second issue is the impact of packet size on the available bandwidth. In fact, as wireless bandwidths have gained the solid reputation of being always scarce, it would be of great importance to monitor carefully how they are managed, especially when real time multimedia applications are introduced on a larger scale.

Consequently, eliminating the extra bits added to each IP packet header would also be highly beneficial for the network access provider.

A third issue is related to privacy aspects. In fact, when exchanging data traffic with the HA, the MN has to include its IPv6 home address (HoA) in each data packet sent to the HA, in addition to its care-of address (CoA) (or its regional care-of address (RCoA) when sending data packets to the MAP in addition to the on-link CoA

(LCoA)). Similarly, the HA has to disclose the MN's HoA and CoA in each data packet sent to the MN (or the RCoA and LCoA in each data packet sent by the MAP to the MN). It follows that having the two MN's addresses disclosed in the same data packet enables a malicious node located on path between the MN and the HA or between the MN and the MAP to easily identify, link and trace the MN's movements. Note that the MN's privacy can be better protected if the traffic is encrypted, which may not be always possible for various reasons.

When the RO mode is used, the tunneling mechanism applied between the MN and the CN differs from the one used in the BT mode and HMIPv6 protocol in the fact that it adds only one additional IP address, i.e., the MN's HoA, in each data packet exchanged between the two endpoints. This means that each data packet sent/received by the MN MUST carry three IP addresses instead of four: the MN's CoA, the CN's address and the MN's HoA. For this purpose, the MN uses a Home Address Option (HAO) to carry its HoA in each data packet sent to the CN and the CN uses a Routing Header (RH) to carry the MN's HoA in each data packet sent to the MN. The use of HAO and RH are in fact degenerated tunnels as shown in [TUN]. This form of tunneling is mandated as long as the MN's corresponding binding lifetime has not expired.

The impact of the RO tunneling on the MN's privacy is the same as described earlier in the BT mode and HMIPv6 protocol. The fact that each data packet discloses the MN's HoA and CoA enables a malicious node located on path between the two endpoints to identify, link and trace the MN's movements across the Internet.

Note that using a PaT does not completely solve the privacy problem. However, it offers a significant advantage in the fact that it narrows the problem to the critical signaling messages, i.e., Binding Update and Acknowledgment (BU/BA) in the case of BT and RO modes and the Local BU (LBU) in the HMIPv6 case, where applying security measures is not just an option. It follows that hiding/replacing the HoA in and encrypting the PaT sent in the BU message are two measures which can provide privacy protection for the MN against eavesdroppers located on path between the MN and the CN/HA/MAP.

4. Tunneling Optimization for BT mode and HMIPv6

TO mechanism addresses tunneling issues described earlier by keeping the original packet size unmodified and removing the need for an additional IP header.

TO mechanism is based on using a PaT to easily translate IP packets headers to new ones which reflect the topologies of the new chosen origin and destination. We limit the scope of the PaT in this document to IPv6 addresses only but it is important to note that its usefulness can also be expanded to translate content(s) of other particular field(s). Another way to describe the PaT functionality is to consider it as a mechanism which virtualizes the need for explicit tunneling.

In order to better describe the suggested mechanism, we apply it to the BT mode and describes the different steps required between the MN and the HA to implement it. In case of HMIPv6 protocol, we note that the same approach can be applied but with HMIPv6 specific signaling messages and IPv6 addresses. After that, we apply the suggested mechanism in the RO mode.

The BT mode starts after the MN switches to a foreign network. In this case, the MN configures a new CoA and notifies its HA by securely exchanging a BU and BA messages. After creating the binding between the two MN's addresses, both nodes start tunneling data packets between them. From the MN side, IP tunneling is applied on each data packet sent to the HA by attaching an outer IP header which contains the MN's CoA as source address and the HA's IP address as destination address. The inner IP header carries the MN's HoA as source address and the CN's address as destination address. On the HA side, IP tunneling is applied on incoming data packets (i.e., sent by the CN to the MN's HoA) by attaching to each packet an outer which carries the MN and the HA IP addresses i.e., the source address becomes the HA's address and the destination address is the MN's CoA.

The data packet tunneling format on the MN and HA sides is as follows:

```

<-- outer IPv6 header --> <-- inner IPv6 header -->
+---+-----+-----+ +---+-----+-----+ +-----+
|   |         |         | |   |         |         | |
|oNAF|  oSRC  |  oDEST | |iNAF|  iSRC  |  iDEST | |  iPAYLOAD
|   |         |         | |   |         |         | |
+---+-----+-----+ +---+-----+-----+ +-----+

```

where:

- NAF represents the non-address fields of an IPv6 header (I.e., the first 8 octets of an IPv6 header)
- SRC is an IPv6 source address
- DEST is an IPv6 destination address
- EXT is zero or more IPv6 extension headers
- the prefix "o" means "outer"

In the HMIPv6 case, the MN attaches an outer header which contains its LCoA as source address and the selected MAP's IP address as destination address. The inner header contains the MN's RCoA as source address and the CN's IP address as destination address. Similarly, the MAP attaches an outer header to all data packets sent by the CN to the MN's RCoA. The outer header is the same as the one used by the MN but with the two addresses inverted.

The first step towards implementing the TO mechanism consists on building the PaT at the MN's side. Since we limit the PaT functionality in this document to the IP addresses only, the PaT will consist on two different 128-bit translator parameters (TPs):
 $\text{PaT} = \{\text{TP_Source (TPS)}, \text{TP_Destination (TPD)}\}$

In the BT mode, TPS and TPD are computed in the following way:

- $\text{TPS} = \text{oSrc_Addr} \text{ XOR } \text{iSrc_Addr}$

Where:

- oSrc_Addr is the Outer header Source Address (CoA)
- iSrc_Addr is the Inner header Source Address (HoA)
- $\text{TPD} = \text{oDst_Addr} \text{ XOR } \text{iDst_Addr}$

Where:

- oDst_Addr is the Outer header Destination Address (HA's IP address)
- iDst_Addr is the Inner header Destination Address (CN's IP address)

In HMIPv6 protocol, the PaT is computed in the same way as in the BT mode but with using the appropriate IP addresses defined in HMIPv6:

- CoA = LCoA
- HoA = RCoA
- HA's IP address = MAP's IP address

The next step after building the PaT is to securely share it with the HA or the MAP (i.e., HMIPv6 case). This is done by inserting the PaT

components in two new options (called PaT Source (PaTS) and PaT Destination (PaTD)) carried by the BU message. The two options MUST be sent encrypted.

Another way to share the PaT would consist on using the BU message (e.g., by setting a new bit) to request the HA to build the PaT. In such case, the MN has to send the CN's IP address(es) in the BU message. For privacy purpose, the CN's address SHOULD be sent encrypted.

It follows from the above description, that each time the MN configures a new CoA, it has to build a new PaT and send it in a BU/LBU (i.e., with the new CoA/LCoA) to the HA/MAP.

If the MN is communicating with multiple CNs, then it SHOULD configure one CoA per CN and build the corresponding PaT before sending it to the HA. Such step is required in order to avoid any confusion over which PaT to apply at the HA side on data packets sent by the MN. The same requirement arises in the HMIPv6 case which means that the MN has to configure one LCoA per CN, build the corresponding PaT then send it to the MAP.

Consequently, the HA/MAP SHOULD create one entry in its BCE for each MN's CoA/LCoA and SHOULD add the corresponding CN's IP address together with the corresponding PaT.

Upon receiving a BU message carrying a PaT, the HA creates first a binding between the two MN's addresses and stores them together with the PaT in its binding cache entries (BCE) table. Then, the HA sends back a BA message to the MN.

It follows that, when the MN has multiple CoAs, then the HA SHOULD create one entry in its BCE for each MN's CoA and SHOULD add the corresponding CN's IP address to the MN's addresses and the corresponding PaT.

Similarly, in an HMIPv6 domain, the MN sends the PaT encrypted in the LBU message to its MAP and waits for the BA message before applying the PaT on data packets.

After creating the binding and sharing the PaT with the HA/MAP, the MN applies it on each data packet sent to the CN (i.e., via the HA/MAP) and on each data packet received from the HA/MAP. The HA/MAP applies the PaT on each data packet received from the MN before sending it to the CN and on each data packet sent by the CN to the MN's HoA/RCoA before sending it to the MN.

When the MN needs to send a data packet to the CN, it applies the PaT on the IP header in the following way (Eq. 1):

- Src_Addr = iSrc_Addr XOR TPS
- Dst_Addr = iDst_Addr XOR TPD

Where Src_Addr is the source address disclosed in the IP header and Dst_Addr is the destination address used in the IP header. It becomes obvious at this stage that the Src_Addr is the MN's CoA (or LCoA in HMIPv6) and the Dst_Addr is the HA's IP address (or MAP's IP address in HMIPv6).

When the HA/MAP receives a data packet from the MN, it retrieves the corresponding PaT on the IP header and translates the IP addresses to new ones. Since the PaT used by the HA is the same as the one used by the MN and the required operation is only a XOR, then the resulting IP addresses are the ones used as iSrc_Addr and iDst_Addr in (Eq. 1). This means:

- nSrc_Addr = Src_Addr XOR TPS (= iSrc_Addr)
- nDst_Addr = Dst_Addr XOR TPD (= iDst_Addr)

Where the nSrc_Addr is the new IP source address and nDst_Addr is the new IP destination address used in the data packet sent by the HA to the CN.

Similarly, when the CN sends a data packet to the MN's HoA, the HA/MAP applies the MN's corresponding PaT to the IP packet header and translates the two IP addresses to new ones before sending the data packet to the MN.

Finally, when the MN gets a data packet from its HA/MAP, it checks first if the data packet is encapsulated or not. In the latter case, the MN applies the PaT to the IP packet header. Otherwise, the MN follows the BT mode to process the packet.

Note that in the current design, the HA SHOULD always tunnel any new data packet sent by a new CN according to the BT mode rules and SHOULD NOT apply any PaT until it receives one from the MN. However, further optimizations are possible and will be introduced in future version of this document.

5. Tunneling Optimization for R0 mode

As mentioned earlier, MIPv6 R0 mode uses a different form of IP tunnel between the two endpoints (i.e., MN and CN) than the one used in the BT mode (i.e., between the MN and its HA). The modified IP tunnel discloses three IP addresses to the receiver and is used by both sides when exchanging data packets. Consequently, the TO mechanism requires a slight modification in order to adapt it to the degenerated tunnel used in the R0 mode.

For this purpose, when the R0 mode is used, the PaT components SHOULD be computed by both endpoints in the following way:

- TPS = CoA XOR HoA
- TPD = 0000:0000:0000:0000:0000:0000:0000:0000

Where the CoA and HoA are the MN's IP addresses. Note that when the R0 mode is used, the MN does not need to send the PaT in the BU message since the CN can build it. However, the MN SHOULD send an explicit request to the CN to check if both endpoints can use the PaT when exchanging data packets. For this purpose, the request consists on setting a new bit in the BU message sent by the MN to the CN and getting an explicit acknowledgment to the request from the CN in the BA message. If the CN is able to use the PaT, then it SHOULD set a new bit in the BA message. Otherwise, it can only send the BA message without any further action.

It follows from the above, that each time the MN needs to send a data packet to the CN following the R0 mode rules, it SHOULD apply the PaT to the IP packet header in order to translate it to the right IP address, i.e., the MN's CoA. The same rule applies when receiving a data packet from the CN (i.e., sent to the MN's CoA). On the CN side, the PaT is applied each time a data packet needs to be sent to the MN or each time a data packet is received from the MN.

It should be noted that when the R0 mode is in use, the PaT can be applied only as long as the MN's CoA binding lifetime has not expired. In addition, the MN SHOULD set the PaT bit in each BU message sent to the CN as long as it prefers to use the TO mechanism during the ongoing session. This also means that a new PaT needs to be built after each BU message carrying a new CoA.

6. Tunneling Optimization for DSMIPv6

TBD

[7.](#) Bit and Options Formats

TBD

8. Security Considerations

This memo describes a TO mechanism which is mainly used to avoid explicit IP tunneling in mobility protocols.

The proposed mechanism enhances the MN's privacy by removing the need to disclose the MN's two IP addresses in the same data packet. Consequently, it simplifies and narrows the privacy problem in a mobile environment.

In the current memo, the PaT is only applied on the IPv6 addresses and as such it does not create nor amplify any new or existing threats.

9. Acknowledgments

The authors would like to thank Laurent Marchand, Conny Larsson, Shinta Sugimoto, Suresh Krishnan, Hesham Soliman and George Tsirtsis for their valuable comments.

10. References

10.1. Normative References

- [MIPv6] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support for IPv6", [RFC 3775](#), June 2004.
- [TERM] Bradner, S., "Key Words for Use in RFCs to Indicate Requirement Levels", [RFC 2119](#), BCP , March 1997.

10.2. Informative References

- [EALDC] Barr, K. and K. Asanovic, "Energy-Aware Lossless Data Compression", ACM Transactions Computer Systems, August 2006.
- [FMIPv6] Koodli, R., "Fast Handovers for Mobile IPv6", Internet Draft, [draft-koodli-mipshop-rfc4068bis-00.txt](#), March 2007.
- [HMIPv6] Soliman, H., Castelluccia, C., ElMalki, K., and L. Bellier, "Hierarchical Mobile IPv6", Internet Draft, [draft-ietf-mipshop-4140bis-00.txt](#), March 2007.
- [PMIPv6] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", Internet Draft, [draft-ietf-netlmm-proxymip6-01.txt](#), June 2007.
- [TUN] Deering, S. and B. Zill, "Redundant Address Deletion when Encapsulating IPv6 in IPv6", Internet Draft, [draft-deering-ipv6-encap-addr-deletion-00.txt](#), November 2001.

Authors' Addresses

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

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

Mats Naslund
Ericsson Research
Torshamnsgatan 23
SE-164 80 Stockholm
Sweden

Phone: +46 8 58533739
Email: Mats.Naslund@ericsson.com

Pekka Nikander
Ericsson Research Nomadic Lab
Jorvas FI-02420
Finland

Phone: +358 9 299 1
Email: Pekka.Nikander@nomadiclab.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.

Acknowledgment

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

