Network Working Group
Request for Comments: 4831
Category: Informational

J. Kempf, Ed.
DoCoMo USA Labs

April 2007

Goals for Network-Based Localized Mobility Management (NETLMM)

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, design goals for a network-based localized mobility management (NETLMM) protocol are discussed.

Table of Contents

1. Introduction

In [1], the basic problems that occur when a global mobility protocol is used for managing local mobility are described, and two currently used approaches to localized mobility management -- the host-based approach that is used by most IETF protocols, and the proprietary Wireless LAN (WLAN) switch approach used between WLAN switches in different subnets -- are examined. The conclusion from the problem statement document is that none of the approaches has a complete solution to the problem. While the WLAN switch approach is most convenient for network operators and users because it requires no software on the mobile node other than the standard drivers for WiFi, the proprietary nature limits interoperability, and the restriction to a single last-hop link type and wired backhaul link type restricts scalability. The IETF host-based protocols require host software stack changes that may not be compatible with all global mobility protocols. They also require specialized and complex security transactions with the network that may limit deployability. The conclusion is that a localized mobility management protocol that is network based and requires no software on the host for localized mobility management is desirable.

This document develops a brief functional architecture and detailed goals for a network-based localized mobility management protocol (NETLMM). Section 2 describes the functional architecture of NETLMM. In Section 3, a list of goals that is desirable in the NETLMM protocol is presented. Section 4 briefly outlines Security Considerations. More discussion of security can be found in the threat analysis document [2].

1.1. Terminology

Mobility terminology in this document follows that in RFC 3753 [10] and in [1]. In addition, the following terms are related to the functional architecture described in Section 2:

Localized Mobility Management Domain

An Access Network in the sense defined in [1] in which mobility is handled by the NETLMM protocol.

Mobile Access Gateway

A Mobile Access Gateway (MAG) is a functional network element that terminates a specific edge link and tracks mobile node IP-level mobility between edge links, through NETLMM signaling with the Localized Mobility Anchor. The MAG also terminates host routed data traffic from the Localized Mobility Anchor for mobile nodes

currently located within the edge link under the MAG's control, and forwards data traffic from mobile nodes on the edge link under its control to the Localized Mobility Anchor.

Local Mobility Anchor

A Local Mobility Anchor (LMA) is a router that maintains a collection of host routes and associated forwarding information for mobile nodes within a localized mobility management domain under its control. Together with the MAGs associated with it, the LMA uses the NETLMM protocol to manage IP node mobility within the localized mobility management domain. Routing of mobile node data traffic is anchored at the LMA as the mobile node moves around within the localized mobility management domain.

2. NETLMM Functional Architecture

The NETLMM architecture consists of the following components. Localized Mobility Anchors (LMAs) within the backbone network maintain a collection of routes for individual mobile nodes within the localized mobility management domain. The routes point to the Mobile Access Gateways (MAGs) managing the links on which the mobile nodes currently are located. Packets for a mobile node are routed to and from the mobile node through tunnels between the LMA and MAG. When a mobile node moves from one link to another, the MAG sends a route update to the LMA. While some mobile node involvement is necessary and expected for generic mobility functions such as movement detection and to inform the MAG about mobile node movement, no specific mobile-node-to-network protocol will be required for localized mobility management itself. Host stack involvement in mobility management is thereby limited to generic mobility functions at the IP layer, and no specialized localized mobility management software is required.

3. Goals for the NETLMM Protocol

Section 2 of [1] describes three problems with using a global mobility management protocol for localized mobility management. Any localized mobility management protocol must naturally address these three problems. In addition, the side effects of introducing such a solution into the network need to be limited. In this section, we address goals for NETLMM, including both solving the basic problems (Goals 1, 2, and 3) and limiting the side effects (Goals 4+).

Some basic goals of all IETF protocols are not discussed in detail here, but any solution is expected to satisfy them. These goals are fault tolerance, robustness, interoperability, scalability, and minimal specialized network equipment. A good discussion of their applicability to IETF protocols can be found in [4].

Out of scope for the initial goals discussion are Quality of Service (QoS) and dormant mode/paging. While these are important functions for mobile nodes, they are not part of the base localized mobility management problem. In addition, mobility between localized mobility management domains is not covered here. It is assumed that this is covered by the global mobility management protocols.

3.1. Goal 1: Handover Performance Improvement

Handover packet loss occurs because there is usually latency between when the link handover starts and when the IP subnet configuration and global mobility management signaling completes. During this time, the mobile node is unreachable at its former topological location on the old link where correspondents are sending packets. Such misrouted packets are dropped. This aspect of handover performance optimization has been the subject of much work, both in other Standards Development Organizations (SDOs) and in the IETF, in order to reduce the latency in IP handover. Many solutions to this problem have been proposed at the link layer and at the IP layer. One aspect of this goal for localized mobility management is that the processing delay for changing the forwarding after handover must approach as closely as possible the sum of the delay associated with link-layer handover and the delay required for active IP-layer movement detection, in order to avoid excessive packet loss. Ideally, if network-side link-layer support is available for handling movement detection prior to link handover or as part of the link handover process, the routing update should complete within the time required for link handover. This delay is difficult to quantify, but for voice traffic, the entire handover delay, including Layer 2 handover time and IP handover time should be between 40-70 ms to avoid any degradation in call quality. Of course, if the link-layer handover latency is too high, sufficient IP-layer handover performance for good real-time service cannot be matched.

A goal of the NETLMM protocol -- in networks where the link-layer handover latency allows it -- is to reduce the amount of latency in IP handover, so that the combined IP-layer and link-layer handover latency is less than 70 ms.

3.2. Goal 2: Reduction in Handover-Related Signaling Volume

Considering Mobile IPv6 [9] as the global mobility protocol (other mobility protocols require about the same or somewhat less), if a mobile node using address autoconfiguration is required to reconfigure on every move between links, the following signaling must be performed:

1) Link-layer signaling required for handover and reauthentication. For example, in 802.11 [7], this is the Reassociate message together with 802.1x [8] reauthentication using EAP.

2) Active IP-level movement detection, including router reachability. The Detecting Network Attachment (DNA) protocol [5] uses Router Solicitation/Router Advertisement for this purpose. In addition, if SEcure Neighbor Discovery (SEND) [3] is used and the mobile node does not have a certificate cached for the router, the mobile node must use Certification Path Solicitation/Certification Path Advertisement to obtain a certification path.

3) Two Multicast Listener Discovery (MLD) [14] REPORT messages, one for each of the solicited node multicast addresses corresponding to the link local address and the global address.

4) Two Neighbor Solicitation (NS) messages for duplicate address detection, one for the link local address and one for the global address. If the addresses are unique, no response will be forthcoming.

5) Two NS messages from the router for address resolution of the link local and global addresses, and two Neighbor Advertisement messages in response from the mobile node.

6) Binding Update/Binding Acknowledgement between the mobile node and home agent to update the care of address binding.

7) Return routability signaling between the correspondent node and mobile node to establish the binding key, consisting of one Home Test Init/Home Test and Care of Test Init/Care of Test.

8) Binding Update/Binding Acknowledgement between the correspondent node and mobile node for route optimization.

Note that Steps 1-2 may be necessary, even for intra-link mobility, if the last-hop link protocol doesn’t provide much help for IP handover. Steps 3-5 will be different if stateful address configuration is used, since additional messages are required to obtain the address. Steps 6-8 are only necessary when Mobile IPv6 is

used. The result is approximately 18 messages at the IP level, where the exact number depends on various specific factors, such as whether or not the mobile node has a router certificate cached before a mobile node can be ensured that it is established on a link and has full IP connectivity. In addition to handover related signaling, if the mobile node performs Mobile IPv6 route optimization, it may be required to renew its return routability key periodically (on the order of every 7 minutes), even if it is not moving, resulting in additional signaling.

The signaling required has a large impact on the performance of handover, impacting Goal 1. Perhaps more importantly, the aggregate impact from many mobile nodes of such signaling on expensive shared links (such as wireless where the capacity of the link cannot easily be expanded) can result in reduced last-hop link capacity for data traffic. Additionally, in links where the end user is charged for IP traffic, IP signaling is not without cost.

To address the issue of signaling impact described above, the goal is that handover signaling volume from the mobile node to the network should be no more than what is needed for the mobile node to perform secure IP-level movement detection, in cases where no link-layer support exists. Furthermore, NETLMM should not introduce any additional signaling during handover beyond what is required for IP-level movement detection. If link-layer support exists for IP-level movement detection, the mobile node may not need to perform any additional IP-level signaling after link-layer handover.

3.3. Goal 3: Location Privacy

In any IP network, there is a threat that an attacker can determine the physical location of a network node from the node’s topological location. Depending on how an operator deploys their network, an operator may choose to assign subnet coverage in a way that is tightly bound to geography at some timescale, or it may choose to assign it in ways in which the threat of someone finding a node physically based on its IP address is smaller. Allowing the L2 attachment and L3 address to be less tightly bound is one tool for reducing this threat to location privacy.

Mobility introduces an additional threat. An attacker can track a mobile node’s geographical location in real-time, if the victim mobile node must change its IP address as it moves from one subnet to another through the covered geographical area. If the granularity of the mapping between the IP subnets and geographical area is small for the particular link type in use, the attacker can potentially assemble enough information to find the victim in real time.

In order to reduce the risk from location privacy compromises as a result of IP address changes, the goal for NETLMM is to remove the need to change IP address as a mobile node moves across links in an access network. Keeping the IP address fixed over a large geographical region fuzzes out the resolution of the mapping between the IP subnets and geographical area, regardless of how small the natural deployment granularity may be. This reduces the chance that the attacker can deduce the precise geographic location of the mobile node.

3.4. Goal 4: Limit Overhead in the Network

Access networks, including both the wired and wireless parts, tend to have somewhat stronger bandwidth and router processing constraints than the backbone. In the wired part of the network, these constraints are a function of the cost of laying fiber or wiring to the wireless access points in a widely dispersed geographic area. In the wireless part of the network, these constraints are due to the limitation on the number of bits per Hertz imposed by the physical layer protocol. Therefore, any solutions for localized mobility management should minimize overhead within the access network.

3.5. Goal 5: Simplify Mobile Node Mobility Management Security by Deriving from IP Network Access and/or IP Movement Detection Security

Localized mobility management protocols that have host involvement may require an additional security association between the mobile node and the mobility anchor, and establishing this security association may require additional signaling between the mobile node and the mobility anchor (see [13] for an example). The additional security association requires extra signaling and therefore extra time to negotiate. Reducing the complexity of mobile-node-to-network security for localized mobility management can therefore reduce barriers to deployment and improve responsiveness. Naturally, such simplification must not come at the expense of maintaining strong security guarantees for both the network and mobile node.

In NETLMM, the network (specifically, the MAG) derives the occurrence of a mobility event, requiring a routing update for a mobile node from link-layer handover signaling, or IP-layer movement detection signaling from the mobile node. This information is used to update routing for the mobile node at the LMA. The handover, or movement detection signaling, must provide the network with proper authentication and authorization so that the network can definitively identify the mobile node and determine its authorization. The authorization may be at the IP level -- for example, using something like SEND [3] to secure IP movement detection signaling -- or it at

the link level. Proper authentication and authorization must be implemented on link-layer handover signaling and/or IP-level movement detection signaling in order for the MAG to securely deduce mobile node movement events. Security threats to the NETLMM protocol are discussed in [2].

The goal is that security for NETLMM mobile node mobility management should derive from IP network access and/or IP movement detection security, such as SEND or network access authentication, and not require any additional security associations or additional signaling between the mobile node and the network.

3.6. Goal 6: Link Technology Agnostic

The number of wireless link technologies available is growing, and the growth seems unlikely to slow down. Since the standardization of a wireless link physical and medium access control layers is a time-consuming process, reducing the amount of work necessary to interface a particular wireless link technology to an IP network is necessary. When the last-hop link is a wireless link, a localized mobility management solution should ideally require minimal work to interface with a new wireless link technology.

In addition, an edge mobility solution should provide support for multiple wireless link technologies. It is not required that the localized mobility management solution support handover from one wireless link technology to another without a change in the IP address, but this possibility should not be precluded.

The goal is that the localized mobility management protocol should not use any wireless link specific information for basic routing management, though it may be used for other purposes, such as securely identifying a mobile node.

3.7. Goal 7: Support for Unmodified Mobile Nodes

In the WLAN switching market, no modification of the software on the mobile node is required to achieve localized mobility management. Being able to accommodate unmodified mobile nodes enables a service provider to offer service to as many customers as possible, the only constraint being that the customer is authorized for network access.

Another advantage of minimizing mobile node software for localized mobility management is that multiple global mobility management protocols can be supported. There are a variety of global mobility management protocols that might also need support, including proprietary or link technology-specific protocols needing support for backward compatibility reasons. Within the Internet, both Host

Identity Protocol (HIP) [11] and IKEv2 Mobility and Multihoming (MOBIKE) [6] are likely to need support in addition to Mobile IPv6 [9], and Mobile IPv4 [12] support may also be necessary.

Note that this goal does NOT mean that the mobile node has no software at all associated with mobility. The mobile node must have some kind of global mobility protocol if it is to move from one domain of edge mobility support to another and maintain session continuity, although no global mobility protocol is required if the mobile node only moves within the coverage area of the localized mobility management protocol or no session continuity is required during global movement. Also, if the last-hop link is a wireless link, every wireless link protocol requires handover support on the mobile node in the physical and medium access control layers, typically in the wireless interface driver. Information passed from the medium access control layer to the IP layer on the mobile node may be necessary to trigger IP signaling for IP handover. Such movement detection support at the IP level may be required in order to determine whether the mobile node's default router is still reachable after the move to a new access point has occurred at the medium access control layer. Whether or not such support is required depends on whether the medium access control layer can completely hide link movement from the IP layer. For cellular type wireless link protocols, the mobile node and network undergo an extensive negotiation at the medium access control layer prior to handover, so it may be possible to trigger a routing update without any IP protocol involvement. However, for a wireless link protocol such as IEEE 802.11 [7] in which the decision for handover is entirely in the hands of the mobile node, IP-layer movement detection signaling from the mobile node may be required to trigger a routing update.

The goal is that the localized mobility management solution should be able to support any mobile node that joins the link and that has an interface that can communicate with the network, without requiring localized mobility management software on the mobile node.

3.8. Goal 8: Support for IPv4 and IPv6

While most of this document is written with IPv6 in mind, localized mobility management is a problem in IPv4 networks as well. A solution for localized mobility that works for both versions of IP is desirable, though the actual protocol may be slightly different due to the technical details of how each IP version works. From Goal 7 (Section 3.7), minimizing mobile node support for localized mobility means that ideally no IP version-specific changes should be required on the mobile node for localized mobility, and that global mobility protocols for both IPv4 and IPv6 should be supported. Any IP version-specific features should be confined to the network protocol.

3.9. Goal 9: Reuse of Existing Protocols Where Sensible

Many existing protocols are available as Internet Standards upon which the NETLMM protocol can be built. The design of the protocol should have a goal to reuse existing protocols where it makes architectural and engineering sense to do so. However, the design should not attempt to reuse existing protocols where there is no real architectural or engineering reason. For example, the suite of Internet Standards contains several good candidate protocols for the transport layer, so there is no real need to develop a new transport protocol specifically for NETLMM. Reuse is clearly a good engineering decision in this case, since backward compatibility with existing protocol stacks is important. On the other hand, the network-based, localized mobility management functionality being introduced by NETLMM is a new piece of functionality, and therefore any decision about whether to reuse an existing global mobility management protocol should carefully consider whether reusing such a protocol really meets the needs of the functional architecture for network-based localized mobility management. The case for reuse is not so clear in this case, since there is no compelling backward compatibility argument.

3.10. Goal 10: Localized Mobility Management Independent of Global Mobility Management

Localized mobility management should be implementable and deployable independently of any global mobility management protocol. This enables the choice of local and global mobility management to be made independently of particular protocols that are implemented and deployed to solve the two different sorts of mobility management problems. The operator can choose a particular localized mobility management protocol according to the specific features of their access network. It can subsequently upgrade the localized mobility management protocol on its own, without even informing the mobile nodes. Similarly, the mobile nodes can use a global mobility management protocol that best suits their requirements, or not use one at all. Also, a mobile node can move into a new access network without having to check that it understands the localized mobility management protocol being used there.

The goal is that the implementation and deployment of the localized mobility management protocol should not restrict, or be restricted by, the choice of global mobility management protocol.

3.11. Goal 11: Configurable Data Plane Forwarding between Local Mobility Anchor and Mobile Access Gateway

Different network operators may require different types of forwarding options between the LMA and the MAGs for mobile node data plane traffic. An obvious forwarding option that has been used in past IETF localized mobility management protocols is IP-IP encapsulation for bidirectional tunneling. The tunnel endpoints are the LMA and the MAGs. But other options are possible. Some network deployments may prefer routing-based solutions. Others may require security tunnels using IPsec Encapsulating Security Payload (ESP) encapsulation if part of the localized mobility management domain runs over a public access network and the network operator wants to protect the traffic.

A goal of the NETLMM protocol is to allow the forwarding between the LMA and MAGs to be configurable depending on the particulars of the network deployment. Configurability is not expected to be dynamic, as in controlled by the arrival of a mobile node; but rather, configuration is expected to be similar in timescale to configuration for routing. The NETLMM protocol may designate a default forwarding mechanism. It is also possible that additional work may be required to specify the interaction between a particular forwarding mechanism and the NETLMM protocol, but this work is not in scope of the NETLMM base protocol.

4. Security Considerations

There are two kinds of security issues involved in network-based localized mobility management: security between the mobile node and the network, and security between network elements that participate in the NETLMM protocol. The security-related goals in this document, described in Section 3.3 and 3.5, focus on the former, because those are unique to network-based mobility management. The threat analysis document [2] contains a more detailed discussion of both kinds of threats, which the protocol design must address.

5. Acknowledgements

The authors would like to acknowledge the following people for particularly diligent reviewing: Vijay Devarapalli, Peter McCann, Gabriel Montenegro, Vidya Narayanan, Pekka Savola, and Fred Templin.

6. Normative References

[1] Kempf, J., Ed., "Problem Statement for Network-Based Localized Mobility Management (NETLMM)", RFC 4830, April 2007.

[2] Vogt, C., and Kempf, J., "Security Threats to Network-Based Localized Mobility Management (NETLMM)", RFC 4832, April 2007.

7. Informative References

[3] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[4] Carpenter, B., "Architectural Principles of the Internet", RFC 1958, June 1996.

[5] Choi, JH. and G. Daley, "Goals of Detecting Network Attachment in IPv6", RFC 4135, August 2005.

[6] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", RFC 4555, June 2006.

[7] IEEE, "Wireless LAN Medium Access Control (MAC)and Physical Layer (PHY) specifications", IEEE Std. 802.11, 1999.

[8] IEEE, "Port-based Access Control", IEEE LAN/MAN Standard 802.1x, June, 2001.

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

[10] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.

[11] Moskowitz, R. and P. Nikander, "Host Identity Protocol (HIP) Architecture", RFC 4423, May 2006.

[12] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.

[13] Soliman, H., Castelluccia, C., El Malki, K., and L. Bellier, "Hierarchical Mobile IPv6 Mobility Management (HMIPv6)", RFC 4140, August 2005.

[14] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

8. Contributors

Kent Leung
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA
EMail: kleung@cisco.com

Phil Roberts
Motorola Labs
Schaumberg, IL
USA
EMail: phil.roberts@motorola.com

Katsutoshi Nishida
NTT DoCoMo Inc.
3-5 Hikarino-oka, Yokosuka-shi
Kanagawa,
Japan
Phone: +81 46 840 3545
EMail: nishidak@nttdocomo.co.jp

Gerardo Giaretta
Telecom Italia Lab
via G. Reiss Romoli, 274
10148 Torino
Italy
Phone: +39 011 2286904
EMail: gerardo.giaretta@tilab.com

Marco Liebsch
NEC Network Laboratories
Kurfuersten-Anlage 36
69115 Heidelberg
Germany
Phone: +49 6221-90511-46
EMail: marco.liebsch@ccrle.nec.de

Editor’s Address

James Kempf
DoCoMo USA Labs
181 Metro Drive, Suite 300
San Jose, CA 95110
USA
Phone: +1 408 451 4711
EMail: kempf@docomolabs-usa.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.