mipshop
Internet-Draft
Expires: January 17, 2005

C. Williams
MCSR Labs
July 19, 2004

Localized Mobility Management Goals
draft-ietf-mipshop-lmm-requirements-03

Status of this Memo

Copyright Notice

Abstract

This document describes goals for Localized Mobility Management (LMM) for IP layer mobility, such as in Mobile IP and Mobile IPv6. These goals are intended to guide the design of a protocol specification for LMM. Localized Mobility Management, in general, introduces enhancements to IP layer mobility protocols to reduce the amount of latency in IP layer mobility management messages exchanged between a Mobile Node (MN) and its peer entities. In addition, LMM seeks to reduce the amount of signaling over the global Internet when a mobile node traverses within a defined local domain. The identified goals are essential for localized mobility management functionality. They are intended to be used as a guide for analysis on the observed benefits over the identified goals for architecting and deploying LMM schemes.

Table of Contents

1. Introduction

In order to meet the demands of real-time applications and the expectations of future wireless users for service level quality similar to the one of wireline users, IP layer mobility management is facing a number of technical challenges in terms of performance and scalability [5][6][7]. These manifest themselves as increased latencies in the control signaling between a Mobile Node and its peer entities, namely the Home Agent (HA) and its Corresponding Nodes (CNs).

In the base Mobile IP protocols [1][2], movement between two subnets requires that the Mobile Node obtain a new care-of address in the new subnet. This allows the Mobile Node to receive traffic on the new subnet. In order for the routing change to become effective, however, the Mobile Node must issue a binding update (also known in Mobile IPv4 as a Home Agent registration) to the Home Agent so that the Home Agent can change the routing from the previous subnet to the new subnet. The binding update establishes a host route on the Home Agent between the Mobile Node's Home Address and its new care-of address. In addition, if route optimization is in use [2], the Mobile Node may also issue binding updates to Correspondent Nodes to allow them to send traffic directly to the new care-of address rather than tunneling their traffic through the Home Agent.

The same approach applies also to a number of other IP layer mobility management protocols. For example, in the Host Identity Protocol (HIP) mobility management scheme [8], a Mobile Node sends an Readdress (REA) packet to its peer; this is very similar the Binding Updates send in Mobile IPv6 Route Optimization. In general, IP layer mobility protocols maintain a binding between a host identifier and either one care-of address or a set of care-of addresses. In Mobile IP, the home address acts as the host identifier, and only one care-of address is allowed. In other IP layer mobility protocols the host identifier is typically something else and more than one care-of addresses may be allowed.

After movement, traffic destined for the Mobile Node is sent to the old care-of address and is, effectively, dropped until the peer entities process mobility management messages. If the Mobile Node is at some geographical and topological distance away from a peer entity, the amount of time involved in sending the binding updates may be greater than 100 hundred milliseconds. This latency in routing update may cause some packets for the Mobile Node to be lost at the old Access Router. For instance, [10] is one such solution for extending Mobile IP to alleviate the above performance limitations. In general such proposals are identified as hierarchical/regional or more generically Localized Mobility Management (LMM).

LMM Localized mobility management schemes allow the Mobile Node to continue receiving traffic on the new subnet without any change in the bindings at the peer entities. The latency involved in updating the care-of-address bindings at far geographical and topological distances is eliminated or reduced until such time as the Mobile Node is in a position to manage the latency cost.

Having provided some motivation and brief summary of the underlying principles of LMM, it is important to enumerate goals for LMM.

Goals for LMM:

- reduce the signaling induced by changes in the point of attachment due to the movement of a host; reduction in signaling delay will minimize packet loss and possible session loss;

- reduce the usage of air-interface and network resources for mobility;

- reduce the processing overhead at the peer nodes, thereby improving protocol scalability;

- avoid or minimize the changes of, or impact to the Mobile Node, Home Agent or the Correspondent Node;

- avoid creating single points of failure;

- simplify the network design and provisioning for enabling LMM capability in a network;

- allow progressive LMM deployment capabilities.

- LMM should introduce no new security vulnerabilities.

Identifying a set of desired properties that will render the protocol internals, for some LMM scheme, robust enough to cater for the aforementioned considerations becomes essential in designing a widely accepted solution. The remainder of this document present a set of goals that encompass essential considerations for the design of an LMM scheme. It is with this foundation that we can seek to ensure that the resulting LMM solution will best preserve the fundamental philosophies and architectural principles of the Internet in practice today.

This document is meant to capture the thinking and analysis of LMM that began in Spring 2001 and is not meant to bind any futures efforts that, due to gained wisdom, may wish to depart from it.

2. Terminology

See [9] for mobility terminology used in this document.

Peer entity A generic name used for nodes that communicate with a mobile node using mobility-specific signaling. In Mobile IP and IPv6, the Home Agent and Correspondent Node are peer entities.

3. Goals

This section describes the goals for a LMM solution. These desired properties are relevant to both all IP layer local mobility management schemes, independent of the actual IP layer macro mobility protocol.

3.1 Intra-domain mobility

LMM is introduced to minimize the signaling traffic to the peer entities, e.g. Home Agent and/or Correspondent Node(s), for intra-domain mobility (within a Local Coverage Area). This is the fundamental reason for introducing localized mobility management. In the LMM infrastructure a peer entity outside the administration domain must always be able to address the mobile host by the same IP address, so that from the point of view of hosts outside the administration domain, the IP address of the mobile host remains fixed regardless of any changes in the Mobile Node's subnet. The peer entities are not aware of the Mobile Node's Intra-domain movement.

3.1.1 Optimized signaling within the Local Coverage Area

By its very nature, LMM reintroduces triangle routing into the base IP layer mobility protocol in that all traffic must go through the LMM agent. There is no way to avoid this. The LMM framework should be designed in such a way as to reduce the length of the unwanted triangle leg. The LMM design should not prohibit optimal placement of LMM agents to reduce or eliminate additional triangle routing introduced by LMM.
NOTE: It is not required that a LMM scheme specify LMM agents as part of its solution.

3.2 Security

3.2.1 Security services

LMM protocol must provide security services within the respective local coverage area.

The security of exchanging LMM specific information and signaling must be ensured. In general, an LMM scheme must cater for authentication mechanisms that prevent malicious deflection of traffic destined to a legitimate MN. If applicable, replay protection must exist mutually between the LMM agents.

3.2.2 Non-interference with peer entities

LMM protocol must not interfere with the security provisioning that

exists between the peer entities and the Mobile Node.

3.2.3 No new vulnerabilities

LMM protocol must not introduce new security holes or the possibility for DOS-style attacks with the base global mobility management protocol that may be used for inter-domain mobility.

3.3 Induced LMM functional requirements

3.3.1 No additional functionality at peer entities

Any Localized Mobility Management protocol must not inject any additional functionality over the base IP layer macro mobility protocol employed. Thus, the LMM framework must not add any modifications or extensions to the peer entities, including Mobile IP or Mobile IPv6 Correspondent Node(s) and Home Agents. It is essential to minimize the involvement of the Mobile Node in routing beyond what is in the basic mobility protocol. Preferences, load balancing, and other complex schemes requiring heavy mobile node involvement in the mobility management task should BE avoided.

3.3.2 No changes to existing components

Non-LMM-aware routers, hosts, Mobile IP and IPv6 Home Agents, and Mobile Nodes must be able to interoperate with LMM agents.

3.3.3 No additional messages to peer entities

By definition a localized mobility management scheme strives to minimize excessive IP mobility management signaling toward its peer entities, caused by frequent changes in the IP network location (i.e, change in the Care-of Address (CoA)). The amount of regional signaling must not surpass the amount of global signaling that would have otherwise occurred if LMM were not present [4].

3.4 Scalability, Reliability, and Performance

3.4.1 Linear complexity

The LMM complexity must increase at most linearly with the size of the local domain and the number of Mobile Nodes.

3.4.2 Linear routing state growth

Any Localized Mobility Management protocol must assure that that LMM routing state scales at most linearly with the number of Mobile Nodes registered, and that the increase in routing state is confined to

those Access Routers/Access Network Routers (ANR) involved in implementing the LMM protocol at hand. While host routes apparently cannot be eliminated by any mobility management protocol including base IP mobility, any LMM protocol must keep the number of host routes to a minimum.

3.4.3 No additional points of failure

The LMM framework must not introduce additional points of failure in the network. The current access router would be excluded from this requirement.

3.4.4 No worse performance

The LMM framework should not degrade in any way the basic IP mobility performance of a mobile host communications with a peer entity.

3.4.5 Scalable expansion of the network

It is imperative that the LMM function must afford larger operational scales by means of incremental deployment. The LMM framework must not introduce any additional restrictions in how wireless ISPs configure their network, nor how they interconnect with other networks beyond those introduced by standard IP routing. In addition, the amount of regional signaling should not increase as the Local Domain expands in size.

3.4.6 Resilience to topological changes

The LMM protocols must be topology-independent. The LMM protocols must be able to adapt to topological changes within the domain. The topological changes may include the addition or removal/failure of LMM agents or that of changes in the routing of the local domain over which the LMM scheme is applied.

3.4.7 Header or Tunneling overhead

The LMM framework must not prevent header compression from being applied. It is recommended that candidate LMM designs that require additional header overhead for tunnel be reviewed by the ROHC working group to determine if the header compressor can be restarted from transferred compressor context when handover occurs without requiring any full header packet exchange on the new link.

3.5 Mobility Management Support

The following LMM requirements pertain to both inter-domain and intra-domain hand-off.

3.5.1 No increase of latency or packet loss

The LMM framework must not increase the amount of latency or amount of packet loss, compared to what exists with the core Mobile IP and Mobile IPv6 specification [1][2]. Indeed, the LMM framework should decrease the amount of latency or amount of packet loss that exists with the core mobility protocols.

3.5.2 No increase of service disruption

The LMM framework must not increase the amount of service disruption, compared to that already exists with the Mobile IP and Mobile IPv6 core mobility specifications. Again, the LMM framework should decrease the amount of service disruption that already exists with the IP layer mobility management protocols.

3.5.3 No new messages to peer entities

The LMM framework must not increase the number of messages between the mobile host and the respective peer entities. The LMM framework should decrease the number of messages between the mobile host and the respective peer entities. With respect to Mobile IP and Mobile IPv6, the current number of messages is defined in the Mobile IP core mobility specifications [1][2].

3.6 Auto-configuration capabilities for LMM constituents

It is essential that the configuration tasks of the LMM scheme can adapt to topological changes with minimal (or no) human intervention; manual configuration is usually tedious, prone to human error and for large-scale deployment, impossible. Automating the configuration task of the LMM function is elemental for addressing realistically incremental deployment of LMM agents within an expanding network domain of large numbers of MNs requiring robust IP services.

By introducing self-organizing LMM agents that caters for dynamic discovery, configuration and management while embracing resiliency with respect to state consistency or failure, an LMM scheme can address successfully the previously mentioned scalability requirements.

3.7 LMM inter-working with IP routing infrastructure

The LMM framework must not disrupt core IP routing outside the local domain.

3.8 Sparse routing element population

Any LMM protocol must be designed to be geared towards incremental deployment capabilities; the latter implies that the LMM scheme itself imposes minimum requirements on the carriers' network. Incremental deployment capabilities for an LMM protocol signifies that an initial set of sparse LMM agents can populate the administration domain of a network provider and operate sufficiently. In addition, any LMM scheme must be compatible with any additional deployment of LMM agents in future infrastructure expansions; that is to say, allow progressive LMM deployment capabilities.

It is for this reason that the LMM framework must not require that all routing elements be assumed to be LMM-aware in the signaling interactions of an LMM protocol. The LMM framework must BE supported, at the very minimum, by a sparse (proper subset) LMM agent population that is co-located within the routing topology of a single administration domain.

3.9 Support for Mobile IPv4 or Mobile IPv6 Handover

Since one of the primary goals of LMM is to minimize signaling during handover, an LMM solution must be available for the standardized Mobile IPv4 or Mobile IPv6 handover algorithms. LMM and the Mobile IP or Mobile IPv6 handover algorithms must maintain compatibility in their signaling interactions for fulfilling complementary roles with respect to each other.

This requirement should not be interpreted as ruling out useful optimizations of LMM and Mobile IP or Mobile IPv6 handoff schemes that simplify the implementation or deployment of LMM or Mobile IP or Mobile IPv6 handoff.

3.10 Simple Network design

LMM should simplify the network design and provisioning for enabling LMM capability in a network and allow progressive LMM deployment capabilities.

3.11 Stability

LMM must avoid any forwarding loops.

3.12 Quality of Service requirements

3.12.1 Co-exist with end-to-end QoS

The LMM must have the ability to coexist with QoS schemes to hide the

mobility of the MN to its peer by avoiding end-to-end QoS signaling.

3.12.2 Co-exist with link-local QoS

The LMM must have the ability to coexist with the QoS schemes to facilitate the new provisioning of both uplink and downlink QoS after a handoff.

4. Security Considerations

The usual threats against mobility mechanisms are [11]

unauthorized redirection of traffic, and

flooding.

An LMM scheme must cater for suitable authorization or other security mechanisms that prevent malicious flooding and other deflection of traffic. When LMM mechanisms are applied only within a single administrative domain, solving these issues may be easier than in the case of generic end-to-end mobility. It must be remembered, though, that the MNs are not necessarily trusted. In the general setting, it is possible that there are authenticated but malicious MNs that attempt to disrupt the service. The situation is complicated by the co-existence of multiple mobility mechanisms, such as Mobile IP and LMM. Since security mechanisms are, in general, non-composable, each protocol combination should be analyzed separately.

Due to administrative constraints, any LMM function should allow for any security provisioning to be negotiable or at least pre-configurable. In certain administrative domains, reduced security requirements may be allowed, so as to minimize incurred overhead.

Involvement with the LMM function into the security semantics of the end-to-end IP mobility signaling between the MN and its peers is beyond the functional scope of any LMM protocol. It is important to implement the effect of mobility localization without interfering with security mechanisms between visiting MNs and their peers. Any possibly existing security associations between the MN and its peers must be considered transparent for the LMM function.

4.1 Trust model

By necessity, the MN must trust the LMM agents to provide LMM services. In most settings the MN must probably authenticate the LMM agents before it can trust them. Whenever several LMM agents are co-operating (as may be the case in fast mobility, for example), the agents must trust each other, at least to an extent. Typically, this trust is manifested in the amount of space and resources that the LMM agent that is receiving packets from another LMM agent is ready do reserve and use.

The LMM agents should not trust the MNs. Basically, they must assume that the MNs may try to launch various kinds of attacks against them or other MNs. On the other hand, the LMM agents are considered to be

obliged to provide services to the MNs. That is, even though the LMM agents do not trust the MNs, they must still be willing to provide the LMM services to the MNs.

An LMM protocol may assume that the involved nodes do not trust any one else in the network but what has been defined above. On the other hand, it is also possible to assume one or more trusted third parties.

4.2 Potential new vulnerabilities

Beyond the possibility of failure, any LMM agents can also exhibit new security vulnerabilities, including the following.

Denial of service. It is possible that an LMM agent may receive LMM messages that incur redundant processing or resource reservation at the LMM agent, and as a result, deprive other MNs from LMM services. The LMM function should ensure that malicious nodes are excluded from further communications with the LMM agents, in the event that their mobility signals are discarded. Thus, the LMM function must cater for denial of service attacks at the LMM agent nodes.

Message replay. Signals that are sent by the MN to a LMM agent can also be captured and replayed by malicious nodes towards the LMM agents; the LMM agents must ensure that such signaling is either authenticated or have a restricted lifetime. Hence, the LMM function must ensure protection from replay attacks at the LMM agents.

Unauthorized creation of LMM state. If an attacker is able to create an unauthorized LMM state at an LMM agent, it can effectively forward packets to itself, to a black hole, or to a flood victim. The LMM agents should make sure that any LMM state is strongly linked to a known MN.

Creation of LMM state before a MN arrives. If an attacker is able to anticipate the care-of-address a MN is likely to use on a link, and if it can attach to the link using that particular address, it can use the address for a while, move away, and request LMM services on that address. Such a request would be authorized since the attacker was legally using the very address. When the victim later comes to the link, it will not get any packets since its address is forwarded away. If the care-of-address assignment is completed as part of the LMM services, then the LMM system should make sure that a new MN cannot acquire a care-of-address that has previously been assigned to another MN. In the case where by care-of-address assignment is completed in a manner

unrelated to the LMM services, then authentication must be completed prior to beginning LMM services.

Unauthorized tearing down of LMM state. If an attacker is able to cause an LMM agent to discard its LMM state before requested by the MN, the MN may experience loss of service. Since LMM is basically an optimization, this threat may not be so severe and may be ignored by an LMM mechanism.

5. Acknowledgments

Thank you to all who participated in the LMM requirement discussion on the Mobile IP working group alias. First, the editor wishes to recognize Theo Pagtzis's work on LMM requirement analysis. Theo has contributed significantly to the LMM discussion on the mailing list and at IETF working group meetings and has provided text for various requirements. Special thanks also to Gabriel Montenegro, John Loughney, Alper Yegin, Alberto Lopez Toledo, and Madjid Nakhjiri for providing input to the draft in its preliminary stage and many other comments they had. Thanks to the LMM requirement analysis design team: Hesham Soliman, Erik Nordmark, Theo Pagtzis, James Kempf, and Jari Malinen. Thanks to Pekka Nikander for changes to make the document non-Mobile IP specific and the security considerations. He also produced the final xml2rfc format that was used for submitting the draft to the RFC Editor.

Additional comments on the LMM requirements were received from: Charlie Perkins, Muhammad Jaseemuddin, Tom Weckstr, Jim Bound, Gopal Dommety, Glenn Morrow, Arthur Ross, Samita Chakrabarti, Karim El-Malki, Phil Neumiller, Behcet Sarikaya, Karann Chew, Michael Thomas, Pat Calhoun, Bill Gage, Vinod Choyi, John Loughney, Wolfgang Schoenfeld, David Martin, Daichi Funato, Ichiro Okajima, Jari Malinen, Kacheong Poon, Koshimi Takashi, and Cedric Westphal.

An LMM requirement analysis of this body of work was completed by a number of members of the Mobile IP working group and published as [4].

In addition special thanks to the Mobile IP working group chairs, Phil Roberts, Gabriel Montenegro, and Basavaraj Patil, for their input as well as capturing/organizing the initial set of requirements from the discussions. Thomas Narten provided important AD feedback and review.

Normative references

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

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

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Informative References

[4] Pagtzis, T., Williams, C., Kirstein, P., Perkins, C. and A. Yegin, "Requirements for Localized IP Mobility Management", in Proceedings of IEEE Wireless Communications and Networking Conference (WCNC2003), Louisiana, New Orleans, March 2003.

[5] Karlsson, G., "Quality Requirements for Multimedia Network Services", in Proceedings of Radiovetenskap och kommunikation, pages 96-100, June 1996.

[6] Kurita, T., Iai, S. and N. Kitawaki, "Effects of transmission delay in audiovisual communications", Electronics and Communications in Japan, Vol 77, No 3, pages 63-74, 1995.

[7] Wang, Y., Claypool, M. and Z. Zuo, "An Empirical Study of RealVideo Performance Across the Internet", in Proceedings of ACM SIGCOMM Internet Measurement Workshop, Nov 2001.

[8] Moskowitz, R., "Host Identity Protocol", draft-ietf-hip-base-00 (work in progress), June 2004.

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

[10] Soliman, H., Castelluccia, C., Malki, K. and L. Bellier, "Hierarchical Mobile IPv6 mobility management (HMIPv6)", draft-ietf-mipshop-hmipv6-02 (work in progress), June 2004.

[11] Nikander, P., "Mobile IP version 6 Route Optimization Security Design Background", draft-nikander-mobileip-v6-ro-sec-02 (work in progress), December 2003.

Author's Address

Carl Williams
MCSR Labs
3790 El Camino Real
Palo Alto, CA 94306
USA

Phone: +1 650 279 5903
EMail: carlw@mcsr-labs.org

Appendix A. LMM requirements and HMIPv6

HMIPv6 was evaluated as a localized mobility management protocol, and that it was mostly found to satisfy the requirements put forth in this document. This section details one exception with some explanation.

Exception:

One LMM requirement that needs further clarification with respect to HMIPv6 is the requirement that states that LMM should not introduce additional single points of failure. The HMIPv6 Mobility Anchor Point (MAP) is a new single point of failure. Proposals for HMIPv6 MAP replication can be optionally incorporated in order to avoid this new single point of failure. Such proposals can also be applied to the base Mobile IPv6 specification to also allow for Home Agent fail-over as well.

Appendix B. Changes from last revision

Changes since last revision:

- o Updated all references
- o Small editorial fixes throughout the document
- o Added reference to HIP in Introduction
- o Identified HMIPv6 as a possible solution for LMM per feedback
- o Change requirements to goals and/or desired properties
- o Minor change to Peer entity definition - states using mobility-specific signaling
- o Captured more fully the definition of intra-domain movement in section 3.1
- o LMM security provisioning was updated - section 3.2.1
- o Clarification in section 3.2.3 and 3.3.3
- o must changed to should in section 3.4.4
- o must not changed to should not in section 3.5.1
- o scalability and auto-configuration sections presented more as goals
- o section 3.4.8 is now a subsection of 3.1
- o all uppercase directives changed to lowercase in an effort to present more along the lines of a goals document

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.