

DMM Working Group
Internet-Draft
Intended status: Informational
Expires: August 18, 2013

H. Ali-Ahmad (Ed.)
France Telecom - Orange
D. Moses
H. Moustafa
Intel Corporation
P. Seite
France Telecom - Orange
February 14, 2013

Mobility Anchor Selection in DMM: Use-case Scenarios
draft-aliahmad-dmm-anchor-selection-00.txt

Abstract

This document presents and discusses different use-case scenarios of mobility anchor selection in Distributed Mobility Management (DMM).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on August 18, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as

described in the Simplified BSD License.

Table of Contents

1. Terminology	3
2. Introduction	4
3. Considered contexts	5
3.1. Mobile node context	5
3.2. Application context	6
3.3. Network context	6
4. Use-case scenarios	7
4.1. Extremely mobile nodes without any typical location	7
4.2. Mobile nodes with one or more typical locations	8
4.3. Fairly stationary nodes	9
5. Security Considerations	10
6. IANA Considerations	10
7. Acknowledgements	10
8. References	11
8.1. Normative References	11
8.2. Informative References	11
Authors' Addresses	11

1. Terminology

IP-handover:

a handover of a mobile node at the IP level resulting in an IP address change at the mobile node.

New flow:

a flow that did not undergo any IP-handover.

Handover flow:

a flow that did undergo one or more IP-handovers.

New traffic:

the data traffic of the new flows.

Handover traffic:

the data traffic of the handover flows.

Current access router:

the access router where the mobile node is currently attached at the IP level.

DMM default mode of mobility anchor selection:

new flows are always anchored at the current access router which acts as the mobility anchor for these flows after an IP-handover.

2. Introduction

Distributed Mobility Management (DMM) aims at overcoming the shortcomings of the existing IP mobility protocols, such as Mobile IPv6 [RFC6275] and Proxy Mobile IPv6 [RFC5213], that are considered centralized. It brings the mobility anchor closer to the mobile node, down at the access routers level. This is the enabler of a concept that is so-called dynamic mobility, where the mobile node changes its mobility anchor for new flows. New flows are always initiated using the mobile node's current IP address which is configured using the prefix provided by the current access router. The data traffic of these flows is then routed optimally until the mobile node undergoes an IP-handover. However, upon an IP-handover, tunneling mechanisms are needed with that access router, which is then considered the mobility anchor of those flows initiated using its prefix during the whole lifetime of those flows. In what follows, this is considered the DMM default mode of mobility anchor selection.

If most of the flows are short enough to not undergo one or more IP-handovers, it is expected that most of the data traffic is routed optimally. However, this assumption is not always valid and the mobility anchor for new flows, when initiated, could be selected in a more appropriate manner.

When a flow is initiated, it is assigned a mobility anchor that lasts during its whole lifetime. Thus, selecting the most appropriate mobility anchor for a flow when initiated can significantly enhance the mobility management performance and reduce the introduced overhead. In order to achieve this, different metrics and contexts should be taken into consideration. Distributing the mobility anchor functionalities at the access routers level allows considering several contexts such as the mobile node's mobility context, the application context, and the network context.

Hereafter, the considered contexts are presented and then the different use-case scenarios are discussed.

3.2. Application context

Based on the application, the need of IP continuity and the flow characteristics can be determined. While applications that require IP continuity cause the establishment of tunnels in the access network upon an IP-handover, applications that can tolerate an IP address change do not. The mobility anchor selection is less important in the latter case due to the possibility of changing the IP address; in fact, there is no need for a mobility anchor.

In addition, the flow characteristics are highly dependent on the application. Some applications generate in general long flows such as multimedia (e.g. video streaming), online gaming, and large files downloading; others generate in general short flows such as TCP connections for HTTP and SMTP sessions. Long flows are more likely to undergo one or more IP-handovers and therefore the mobility anchor selection can play an important role to enhance the mobility management performance. On the other hand, short flows are more likely to be initiated and terminated before an IP-handover.

3.3. Network context

When a mobility anchor is assigned to a flow (when the flow is initiated), it acts as a mobility anchor for this flow the whole flow's lifetime. It is responsible to forward the flow's data packets if the mobile node is physically attached to it. It is responsible, in addition, to encapsulate and de-capsulate the flow's data packets if the mobile node is not attached to it and tunneling mechanisms are used.

Even with distributed mobility anchors, the distribution of the active mobile nodes in the network is not necessarily even. As a result, some mobility anchors are overloaded more than others. It is then reasonable to take into consideration the current (or projected) level of load of the mobility anchors when selecting one of them for a new flow (the metrics for measuring this level are left for specific implementations).

4. Use-case scenarios

4.1. Extremely mobile nodes without any typical location

Extreme mobility could be due to either a high mobile node's speed, or a small access router's coverage area, or both.

Scenario 1: running applications generating typically short flows

Short flows are more likely to be initiated and terminated before the mobile node undergoes an IP-handover. Even if a flow experiences an IP-handover, it is expected that the flow does not last long after the IP-handover. In other words, most of the mobile node's traffic is new traffic in this scenario. As a result, the closest mobility anchor to the mobile node during the longest portion of a flow is its current access router. It is recommended then to always anchor new flows at the current access router, which is the DMM default mode of mobility anchor selection.

Scenario 2: running applications generating typically long flows

For extremely mobile nodes, it is more likely that a flow experiences an IP-handover soon after being initiated. And since the flows are long-lived, it is expected that a flow lasts for a long duration after the IP-handover(s). As a result, it could be said that most of the traffic is handover traffic in this scenario. Whatever is the mobility anchor selection criterion, most of (almost all) the mobile node's data traffic needs tunneling mechanisms. Thus, the mobility anchor selection cannot play a significant role regarding the route optimization or the tunneling overhead reduction.

However, there are number of consequences regarding the control plane e.g. number of simultaneous anchors/tunnels for a mobile node and the related contexts and signaling loads. First, let us consider the DMM default mode of mobility anchor selection. Since new flows are always anchored at the current access router, each flow initiated between two consecutive IP-handovers is anchored at a different mobility anchor. With extremely mobile node, long flows are expected to experience several IP-handovers and their mobility anchors are expected to be maintained for a long duration. As a result, the number of simultaneous anchors/tunnels for a mobile node may increase as well as the related contexts and signaling loads. This affects the control plane negatively.

As the DMM default mode does not achieve data plane optimization in the scenario described above, it is reasonable to consider a more centralized approach for mobility anchor selection in order to reduce the negative effects on the control plane. If data packets are going to be tunneled in both cases, managing a single tunnel to a single mobility anchor would be better than managing several tunnels to several mobility anchors at the same time.

Scenario 3: running applications generating both long and short flows

In this case, short and long flows can be distinguished when selecting a mobility anchor for a flow, based on scenario 1 and scenario 2. Short flows are always anchored at the current access router; long flows are anchored based on a more centralized approach. In this way, data packets of short flows are generally routed optimally and long flows do not introduce a large number of simultaneous anchors/tunnels.

4.2. Mobile nodes with one or more typical locations

Scenario 4: running applications generating typically short flows

As the flows are short, there is no expected benefit from having a typical location. If initiated when the mobile node is not at its typical location, such flows are more likely to end quickly before the mobile node goes back to its typical location. Otherwise, they would be initiated and terminated when the mobile node is at its typical location. As a result, the current access router is always the best mobility anchor for new flows and hence the DMM default mode of mobility anchor selection fits well this scenario.

Scenario 5: running applications generating typically long flows

In this scenario, having a typical location is expected to be beneficial for the mobile node's mobility anchor selection. As mentioned before, the best mobility anchor for a flow is the closest one to the mobile node during the longer portion of this flow. Then, the best mobility anchor for a flow could be in some cases that of the typical location even if the flow is not initiated there. For example, if the mobile node initiates a long flow and then comes back (undergoing an IP-handover) quickly to its typical location, the longer portion of the flow would be after the IP-handover. Thus, it is reasonable to select the typical location's mobility anchor for such flow when initiated. This results in tunneling part of the flow's data traffic when initiated but in routing optimally most of it afterwards.

The analysis described above would be still valid if the mobile node has more than one typical location. However, the benefits may not be in some cases as great as those of the one typical location scenario, depending on the mobile node's movements. If there is no clear benefit from selecting one out of the mobility anchors, the network context (i.e. level of load on each mobility anchor) comes into play leaning towards selecting the mobility anchor that is less loaded. Another refinement is to add the time of day to the statistics collection in the mobile node's attachments history. If it is noticed that one of the typical locations is more popular than the others, this helps in selecting a mobility anchor according to the time of attachment.

Scenario 6: running applications generating both long and short flows

If it is possible, the short and long flows should be distinguished as follows. While short flows are assigned the closest mobility anchor which is the current access router, long flows are assigned the typical location's mobility anchor. The mobile node needs a source address selection mechanism in order to distinguish between the different IP addresses when initiating a flow.

4.3. Fairly stationary nodes

Scenario 7: running similar or different applications

In fact, a fairly stationary node has one typical location for almost all the time. The current access router is then the typical location's mobility anchor and, obviously, should be selected to anchor new flows.

5. Security Considerations

TBD.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

The authors would like to express their gratitude to Tiago Condeixa, Philippe Bertin, and Wu-Chi Feng for having discussions, sharing thoughts, or providing reviews on anchor selection in DMM.

8. References

8.1. Normative References

- [RFC5213] Gundavelli, S., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", RFC 5213, August 2008.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

8.2. Informative References

- [I-D.ietf-dmm-requirements]
Chan, A., Liu, D., Seite, P., Yokota, H., and J. Korhonen,
"Requirements for Distributed Mobility Management",
draft-ietf-dmm-requirements-03 (work in progress),
December 2012.

Authors' Addresses

Hassan Ali-Ahmad (Editor)
France Telecom - Orange
Email: hassan.aliahmad@orange.com

Danny Moses
Intel Corporation
Email: danny.moses@intel.com

Hassnaa Moustafa
Intel Corporation
Email: hassnaa.moustafa@intel.com

Pierrick Seite
France Telecom - Orange
Email: pierrick.seite@orange.com