No Specific Working Group
Internet-Draft
Expires: December 30, 2004

C. Ng
Panasonic Singapore Labs
J. Hirano
Panasonic
July 2004

Host/Edge Multihoming Problem Statement
draft-ng-edge-multihoming-problem-statement-00

Status of this Memo

Copyright Notice

Abstract

This document analyzes multihoming from the perspective of the Internet edge: i.e. hosts and other edge networks. We built on top of the previous work that describes goals and benefits of multihoming, and identify problems for the provisioning of multihoming at the edge level. In this memo, we first look at the problem of multihoming for a generic IPv6 node, followed by narrowing the analysis down to mobile hosts and networks.

Table of Contents

1. Introduction

Multihoming has attracted wide interest recently. Within the IETF, Site Multihoming for IPv6 (multi6) Working Group is looking at the multihoming problem from the perspective of an IPv6 site, and provisioning of multihoming in IPv6 is being solved more or less using the core network. This document aims to look at multihoming from the perspective of the Internet edge: i.e. hosts and other edge networks.

Benefits of multihoming has been identified in RFC 3582 [1] and "draft-multihoming-generic-goals-and-benefits-00" [2]. However, there is no clear problem statement on the provisioning of multihoming at the edge level. It is an objective of this draft to identify such problems.

By "edge", we refer to hosts or stub networks attached at the edge of the Internet that does not have any transit traffic. The host and network may themselves be mobile. This document will attempt to look at multihoming problem at two different levels of details: (a) no assumption on the mobility of the edge elements; and (b) when the edge elements are mobile.

It is assumed that the readers are familiar with the goals, benefits and deployments scenarios as spelt out in [2].

We begin by first looking at the problem of multihoming for a generic edge element from three different perspective:

1. From the perspective of upper layer protocols (i.e transport and above), we explore the problem multihoming brings with regards to continuity of upper layer sessions. This is done in Section 2.

2. From the perspective of multihomed edge elements, we look into the problem of sending packets out given the source is

multihomed. This is done in Section 3.

3. From the perspective of peers of multihomed edge elements, we describe the problem of sending packets to a destination that is multihomed. This is done in Section 4.

Following these, problem specific to multihomed edge elements that are mobile is investigated in Section 5.

2. Upper Layer Transparency

By definition, a multihomed node has multiple IPv6 addresses. Here, we assumed that it has multiple global-scoped IPv6 addresses (i.e. excluding node-local, link-local, and site-local addresses).

One immediate problem faced by a multihomed node is the question of which address to use. It may be trivially selected, or chosen based on certain policy or preferences settings. No matter how chosen, there is now an issue of upper layer transparency.

The more common and widely used transport layer protocols by far are TCP and UDP. These transport protocols associate end-point addresses with each session. Thus, to these protocols, a change in either the source address or destination address would mean a different session. In order to enjoy benefits of multihoming, it is often necessary to change the address used. Thus, there is a problem of how to achieve upper layer transparency when employing multihoming mechanisms. In other words, the problem is to achieve multihoming without breaking legacy transport protocols such as TCP and UDP.

There are two parts to consider in upper layer transparency: (a) the receiver is multihomed, and (b) the sender is multihomed.

2.1 Receiver is Multihomed

When the receiver is multihomed, the sender has a choice of using one of the multiple addresses as the destination (we ignore for now the

question of how the sender learns of these multiple addresses). The problem in this case is how to use these different addresses in the destination address field and yet have the transport protocol at the receiver associates these packets to the same transport session.

As an illustration (refer to Figure 1), suppose a node, TX, sends to another node, RX, two packets belonging to the same transport session. In the first packet, node TX decided to use RX.Addr1 as the destination address, and in the second packet, node TX decided to use RX.Addr2 as the destination address. The problem is how does the transport protocol at node RX knows that the two packets belong to the same session.

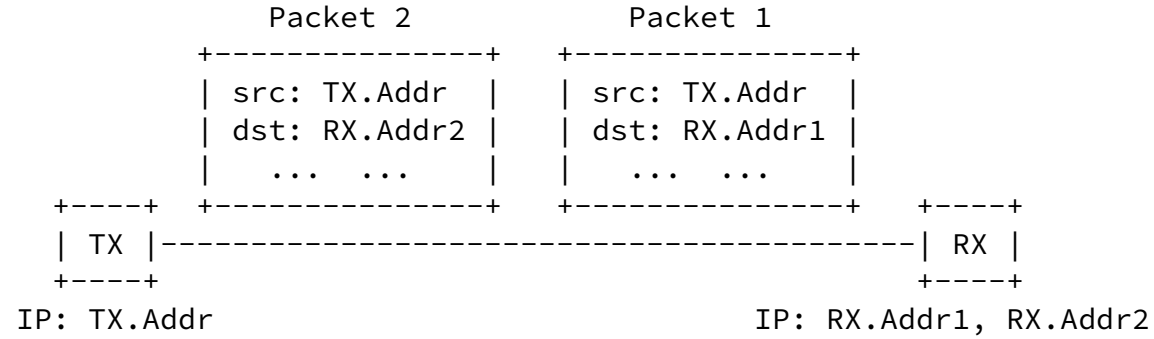

Figure 1: Packets from the same transport session with different destination addresses

2.2 Sender is Multihomed

When the sender is multihomed, the sender may have to switch among multiple source addresses for reasons of ingress filtering [3]. The problem in this case is how to use these different addresses in the source address field and yet have the transport protocol at the receiver associates these packets to the same transport session.

As an illustration (refer to Figure 2), suppose a node, TX, sends to another node, RX, two packets belonging to the same transport session. In the first packet, node TX uses TX.Addr1 as the source address, and in the second packet, node TX uses TX.Addr2 as the source address. The problem is how does the transport protocol at node RX knows that the two packets belong to the same session.

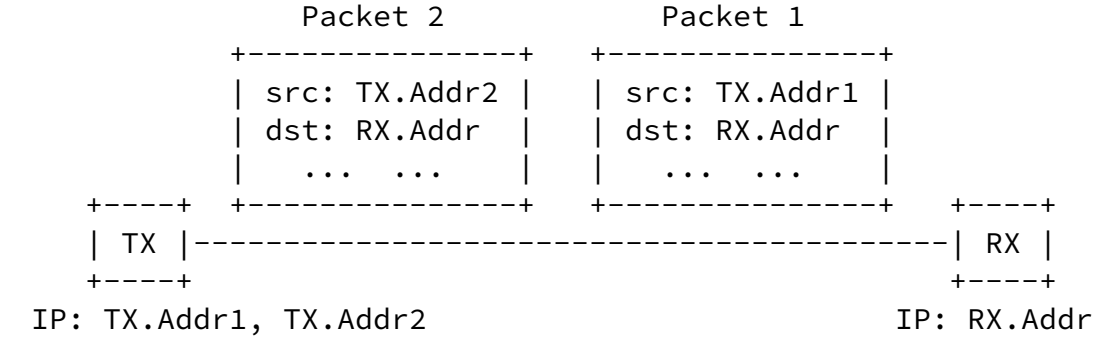

Figure 2: Packets from the same transport session with different source addresses

2.3 Discussion

It can be perceived that a possible approach is to separate the currently overloaded IP address into individual functions: an "identifying" address used to identify the multihomed node, and one

or more "locating" addresses used to route packets to/from the multihomed node. In this case, the protocol layers above IP will use the "identifying" address, and the protocol layers below IP will use the "locating" address.

3. Outgoing Path

A multihomed node would usually has multiple, independent, routes to the Internet. Given these routes, how do a multihomed node choose which route to send a packet? Such decision can be made arbitrarily, or based on certain preferences or policy.

In addition, it is sometimes necessary for the multihomed node to select the route based on the actual source address used on a packet. This is mainly due to ingress filtering. Ingress filtering occurs when an intermediate router discards packet because the source-address of the packet is not of a valid prefix [3].

As an illustration, consider a node N with two independent routes to the Internet through two Internet Service Providers, ISP-A and ISP-B, as shown in Figure 3. ISP-A assigns node N an address PA.N from the prefix PA, and ISP-B assigns node N an address PB.N from the prefix PB. Both ISPs implements ingress filtering to prevent malicious subscribers from performing IP spoofing.

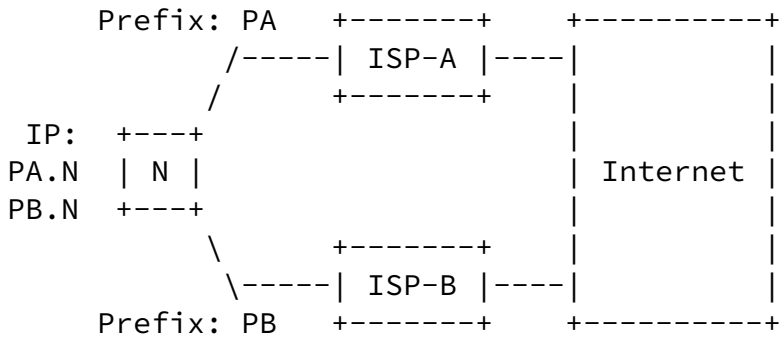

Figure 3: Multihomed node obtaining different addresses from different ISPs

In such cases, to avoid ingress filtering, node N is forced to send packets with source address PA.N through ISP-A, and send packets with source address PB.N through ISP-B.

3.1 Discussion

With relation to "Upper Layer Transparency" (Section 2), it is interesting to note that should the upper layer transparency problem be solved, the problem of ingress filtering for outgoing path may then be trivial. For instance, we assume that the problem of upper layer transparency is solved with an "identifying" address and a set of "locating" addresses. Then, the multihomed node can choose the locating address as the packet source based on the route selected such that the packet would not be discarded by ingress filtering.

Using the same scenario depicted in Figure 3, suppose node N now acquired an identifying address of ID.N. Then, the pair of addresses PA.N and PB.N will be the locating addresses. Whenever node N wants to send packet through ISP-A, it changes the source address of that packet to PA.N. Similarly, whenever node N wants to send packet through ISP-B, it changes the source address of that packet to PB.N.

4. Incoming Path

When a node is multihomed, there are multiple paths to send a packet to the node. The question is then how does a particular path get selected to be the delivery route of any given packet. If the address of the multihomed node is identical across all paths, then it is up the routing infrastructure to select the path. It will be subjected to the routing policy of the core routers to pick arbitrarily or based on certain parameters (such as congestion condition on each path, etc) a delivery route. The sender or the multihomed node has little control over which route to take.

However, if the addresses of the multihomed node is different on each path, the route taken will be decided by the destination address on each packet. The question is then how does the sender knows the different addresses belonging to the multihomed node. This is the problem of "Peer Knowledge". A slight alternative is instead of having every peer node that communicates with the multihomed node to know the set of addresses, only a smaller set of intermediate elements in the routing infrastructure know. This is the problem of "Infrastructure Knowledge".

4.1 Peer Knowledge

This is the problem of how a multihomed node notifies the peer node it is communicating with the set of addresses that it owns. A solution that solves this problem must also address the following issues:

- What is the form of signaling?
- How are the list of addresses communicated to the peer node?
- How can the peer node ascertain the specified list of addresses are indeed owned by the multihomed node?

4.2 Infrastructure Knowledge

Alternatively, perhaps it is not necessary for the peer node to learn all of the addresses of the multihomed node. It might suffice for a selected pool of nodes to know the addresses. In this case, the peer node needs only know the "identifying" address of the multihomed

node, and will only use this "identifying" address in the destination field of the packet. A set of intermediate routers will capture these packets, and translate the "identifying" address to one of the known "locating" addresses of the multihomed node.

A solution that solves the problem of infrastructure knowledge should also addresses the following issues (most are similar to those listed in Section 4.1):

- What is the form of signaling?
- How are the list of addresses communicated to the intermediate router(s)?
- How can the intermediate router(s) ascertain the specified list of addresses are indeed owned by the multihomed node?
- How can the intermediate router(s) change the "identifying" address in the destination field of the packet to one of the "locating" addresses?
- What is the impact of changing addresses by intermediate routers on the end-to-end integrity of the packet?

5. Mobility Considerations

In this section, we focused our attention to multihomed nodes that are mobile. There might be problems related to multihoming that are specific to mobile nodes. There might also be problems that are exemplified (and perhaps intensified) when multihomed nodes are mobile. We first consider the case when the mobile multihomed node is a host. Then, we consider the case when a mobile network is multihomed.

5.1 Multihomed Mobile IPv6 Node

When we refer to a mobile node, it is implied that the node employs Mobile IPv6 [4] to gain mobility support while roaming across different access networks. It is assumed that the readers are familiar with temrs used in Mobile IPv6. There is a draft on multihoming problem with Mobile IPv6 [5] which raise issues multihoming with a mobile node that are similar to those described here. This section however looks at the problem from another angle.

5.1.1 Upper Layer Transparency

The concept of home-address and care-of-addresses associated with a mobile node may be an effective mechanism for achieving upper layer transparency. However, a multihomed mobile node may have multiple home-addresses. Thus, there is still a need to identify the "identifying" address for use with transport (and upper) layer protocols.

5.1.2 Outgoing Path

A multihomed mobile node may have multiple care-of-addresses. In order to use more than one egress link, it might be necessary for the mobile node to use these multiple care-of-addresses simultaneously (for example, to overcome ingress filtering at the access network). Hence there is a need for the ability to bind multiple care-of-addresses to one home-address. This is currently addressed by [6].

The problem of ingress filtering, however, is two-fold. It can occur at the access network, as well as the home network. Figure 4 illustrates this case. In the access network, when mobile node MN sends a packet through AR-A, it must use the care-of-address of PA.MN; and when MN sends a packet through AR-B, it must use the care-of-address of PB.MN. In the home network, when MN tunnels the packet to home-agent HA-1, it must use the home-address P1.MN; and when MN tunnels the packet to home-agent HA-2, it must use the home-address P2.MN. This greatly limits the way MN can enjoy the

benefits of multihoming.

It must be noted that should the mobile node MN have a way of binding both care-of-addresses PA.MN and PB.MN simultaneously to both home-addresses P1.MN and P2.MN, then it can choose the care-of-address to use base on the access network it wishes to use for outgoing packets. This solves the first part of the problem: ingress filtering at the access network. It is, nonetheless, still limited to using only a specific home agent for the home-address it uses (i.e. The second problem of ingress filtering at the home network remains unsolved).

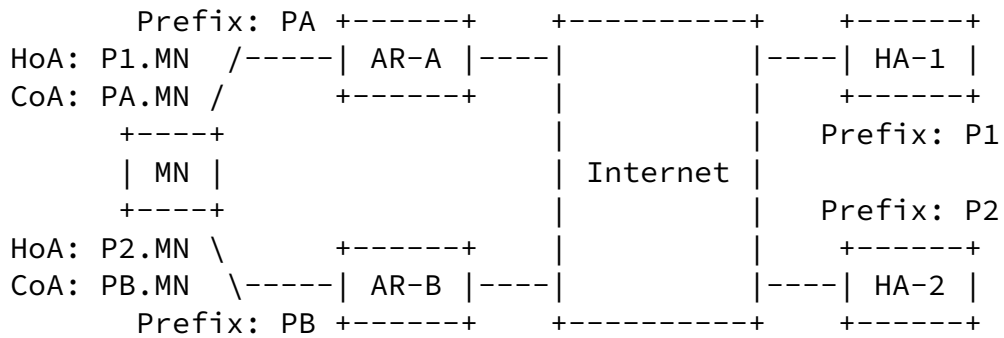

Figure 4: Multihomed mobile node with multiple home agents

5.1.3 Incoming Path

The use of binding updates in MIPv6 [4] lends itself very well to solving the problem of peer knowledge and infrastructure knowledge. The most important prerequisite is then to solve the problem of binding multiple care-of-addresses.

5.2 Multihomed Mobile Networks

The problem of mobile network is discussed in [7] and [8]. There is an extensive analysis of the multihoming issues with mobile networks [9]. This section merely highlights any problems that are specific or more relevant to mobile networks. Interested readers should refer to [9] for details.

Most of the problems relating to upper layer transparency, ingress filtering at the access network, ingress filtering at the home network, peer knowledge and infrastructure knowledge for mobile networks share similar concerns as with a mobile host (see Section 5.1). One particularly interesting problem of ingress filtering at the home network is shown in Figure 5 below.

In Figure 5, the mobile network has two mobile routers MR-A and MR-B,

with home agents HA-A and HA-B respectively. Each mobile router advertises a different mobile network prefix (PA and PB). Thus, the mobile network node MNN configures two IPv6 addresses: PA.MNN and PB.MNN. Hence, MNN is multihomed.

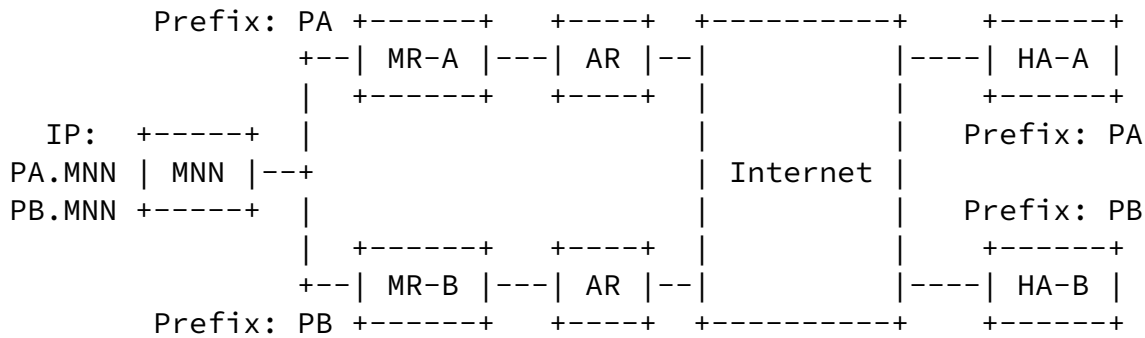

Figure 5: Multihomed mobile network with multiple mobile routers and home agents

However, MNN cannot forward packet with source address equals PA.MNN to MR-B. This will cause ingress filtering at HA-B to occur (or even at MR-B). This is contrary to normal Neighbor Discovery [10] practice that an IPv6 node is free to choose any router as its default router regardless of the prefix it choose to use. A simple solution is to impose a requirement that all mobile network nodes must set their default router to the router that advertises the prefix the mobile network nodes configured their address from. If no such requirements are to be imposed on mobile network nodes, then a multihoming solution for mobile networks must address this problem.

6. Conclusion

This document analyzed multihoming from the perspective of the Internet edge: i.e. hosts and other edge networks. We built on top of the previous work done in [2] and [1], which describe goals and benefits of multihoming, and identify problems for the provisioning of multihoming at the edge level.

We have looked at the problem of multihoming for a generic edge element from three different perspectives:

1. From the perspective of upper layer protocols (i.e transport and above), we explored the problem multihoming brings with regards to continuity of upper layer sessions.

2. From the perspective of multihomed edge elements, we looked into the problem of sending packets out given the source is multihomed.

3. From the perspective of peers of multihomed edge elements, we described the problem of sending packets to a destination that is multihomed.

Following these, problem specific to multihomed edge elements that are mobile (i.e. mobile hosts and mobile networks that are multihomed) were analyzed.

7 References

[1] Abley, J., Black, B. and V. Gill, "Goals for IPv6 Site-Multihoming Architectures", RFC 3582, August 2003.

[2] Ernst, T., "Goals and Benefits of Multihoming", draft-multihoming-generic-goals-and-benefits-00 (work in progress), February 2004.

[3] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

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

[5] Montavont, N., "Problem Statement for multihomed Mobile Nodes", draft-montavont-mobileip-multihoming-pb-statement-00 (work in progress), October 2003.

[6] Wakikawa, R., "Multiple Care-of Addresses Registration", draft-wakikawa-mobileip-multiplecoa-02 (work in progress), September 2003.

[7] Ernst, T., "Network Mobility Support Goals and Requirements", draft-ietf-nemo-requirements-02 (work in progress), February 2004.

[8] Devarapalli, V., "Network Mobility (NEMO) Basic Support Protocol", draft-ietf-nemo-basic-support-03 (work in progress), June 2004.

[9] Ng, C. and J. Charbon, "Multi-Homing Issues in Bi-Directional Tunneling", draft-ng-nemo-multihoming-issues-03 (work in progress), February 2004.

[10] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

Authors' Addresses

Chan-Wah Ng
Panasonic Singapore Laboratories Pte Ltd
Blk 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415
SG

Phone: +65 65505420
EMail: cwng@psl.com.sg

Jun Hirano
Matsushita Electric Industrial Co., Ltd. (Panasonic)
5-3 Hikarino-oka
Yokosuka, Kanagawa 239-0847
JP

Phone: +81 46 840 5123
EMail: hirano.jun@jp.panasonic.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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