

Network Mobility
Internet-Draft
Expires: April 5, 2005

P. Thubert
Cisco
R. Wakikawa
Keio University
V. Devarapalli
Nokia
October 5, 2004

Global HA to HA protocol
draft-thubert-nemo-global-haha-00

Status of this Memo

By submitting this Internet-Draft, I certify that any applicable patent or other IPR claims of which I am aware have been disclosed, and any of which I become aware will be disclosed, in accordance with [RFC 3668](#).

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/lid-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 April 5, 2005.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This paper extends Mobile IPv6 [9] and the NEMO Basic Support [11] to remove the Layer 2 dependencies on the Home Link and allow the global (L3) distribution of the HAs that serve a same Home Network.

Internet-Draft

Global-HAHA

October 2004

Table of Contents

1.	Introduction	3
2.	Motivations	3
2.1	Requirements	3
2.2	Layer 3 operations	4
2.3	Path improvement	5
2.4	Single point of failure	7
3.	Rationale for the proposed solution	7
4.	Overview	8
4.1	Initial routing	9
4.1.1	External routing	9
4.1.2	Internal routing	10
4.2	Binding	11
4.2.1	Direct primary binding	11
4.2.2	local proxy binding	12
4.2.3	Foreign proxy binding	12
4.3	Path improvements	13
4.3.1	Leaking MNP routes in the HAHA network	14
4.3.2	On-demand proxy routes	15
5.	Terminology and concepts	16
6.	Distributed Home Network	18
7.	Message Formats	19
8.	Mobile Router Operation	19
8.1	Locating Home	19
8.2	Proxy MIP client	19
9.	Home Agent Operation	19
9.1	Locating the other HAs that serve the same Home	19
9.2	Locating the HA that owns the binding for a HoA	19
10.	Acknowledgements	19
11.	References	20
	Authors' Addresses	21
	Intellectual Property and Copyright Statements	22

Internet-Draft

Global-HAHA

October 2004

[1.](#) Introduction

The reader of this document is expected to be familiar with both the Mobile IPv6 [[9](#)] and NEMO Basic Support [[11](#)] documents. As such, the reader is expected to understand the concept of a Home Link and the Neighbor Discovery related operations that take place over that link.

Home Agent global distribution is useful when a Mobile Router moves geographically large area such as airplane, vehicle, etc... The overhead of the basic NEMO protocol is redundant route caused by the bi-directional tunnel between a Home Agent and a Mobile Router. If a Mobile Router moves far away from a Home Agent, the overhead can not be ignored.

Thus, it is reasonable to consider that a Mobile Router dynamically switches to the topologically closest Home Agent (Home Link). This distribution is also effective for load-balancing. The Home Agent is expected to serve thousands of Mobile Routers on its Home Link and tunnels all packets for the Mobile Routers by itself.

But with NEMO basic support and MIPv6, Home is locally anchored to the Home Link at Layer 2, so Home can not be distributed geographically. In particular for NEMO, what's needed is a route to a mobile prefix via a tunnel end point that is the CareOf address of the Mobile Router. The Home Address is but a practical artifact that is mostly needed as a correlator for the registration.

This draft proposes extensions that enable the HA to HA communication at Layer 3, allowing to get rid of the Home Link in configurations where it's not needed.

[2.](#) Motivations

[2.1](#) Requirements

This draft addresses two generic requirements expressed in the Nemo requirements [\[14\]](#):

Local Mobility and Global Mobility: Multihoming is mentioned as desirable. The global mobility type is not expected to be limited for any consideration other than administrative and security policies.

Scalability: NEMO support signaling and processing is expected to scale to a potentially large number of mobile networks. Thus draft extends the scalability of the NEMO basic protocol.

There is a requirement from airplane companies which want to be at

Home in the various airports that their planes visit. In fact, this is expressed in an abstract fashion by the case (1,n,1) of the NEMO multihoming issues [\[15\]](#) draft: "Single MR, Multiple HAs, Single NEMO-Prefix".

There is also a general direction that indicates that NEMO could be extended as a solution for VPN. To get there, we must ensure that NEMO is upscaled to the classical capabilities of VPN, including the global distribution of Points Of Presence. It is a classical feature for VPNs to allow the roaming users to connect to the closest point of presence into their company VPN. The same feature can not be provided with MIPv6 or NEMO, because the Home depends on a link that has a unique physical location.

[2.2](#) Layer 3 operations

Mobile IPv6 [\[9\]](#) standarizes an interface between a Mobile Node and its Home Agent and its correspondents, as well as an interface between Home Agents. One angle of the MIPv6 operation is that the protocols hides the MN mobility by making as if the Mobile Node was always connected to a Home Link. The connectivity is maintained by Home Agents that are permanently and physically attached to that Home Link.

So the model for MIPv6 is Home Link centric and it is no surprise that it extends IPv6 Neighbor Discovery [\[3\]](#) for its operations, in particular for HAs to discover each others, and to discover when one of them has a binding for a Mobile Node, and which one. An immediate

consequence of being Link centric is that Home can not be distributed at Layer 3, locally within a site or over the Internet.

the NEMO Basic Support [11] inherits the concept of Home Link and MIPv6 operations on that link, making NEMO partially a link layer operation. On the other hand, the NEMO Basic Support also operates as a routing protocol at L3, for example when it injects routes in the explicit prefix mode. So NEMO operations are somewhat half L2 and half L3.


What we are getting at with the HAHA protocol is placing NEMO fully at L3. This mostly means the replacement of all ND based exchanges by some equivalent, but at Layer 3, over the Internet Protocol. This also means the abstraction of the concept of Home Address into a globally unique router ID, as opposed to an address from a Home Link.


So even if this paper trivially applies to Mobile IPv6, we place our descriptions in the context of NEMO, and use MRs where MIPv6 MNs could fit as well.

[2.3](#) Path improvement

MIPv6 comes with a Route Optimization scheme that enables a direct MR-CN conversation, bypassing the Home Agent. With the basic support, NEMO does not have such a support yet. In any case, RO comes at an additional cost in terms of protocol, which varies with the degree of expected trust.

Without Route optimization, all the packets MR-CN flow via the Home Agent; this increases both the cost and the latency. The resulting path can be illustrated like this:





Current model with a Home Link

The routing overhead becomes costly when:

The distance $||MR, CN||$ is much smaller than the sum of the distances $||MR, HA|| + ||HA, CN||$.

AND

$||MR, HA|| + ||HA, CN||$ is costly. If the 3 points are very close, the overhead is relatively important, but small in absolute terms.


In the picture above, say that a European phone (MR) is roaming in New Jersey but Homed in Brittany. And say that the phone owner places a call in New York city to CN1. Without RO, the voice packets flow back and forth over the peering lines between Brittany and the US, and the routing overhead causes an additional latency that decreases the perceived quality of the phone call.

On the other hand, calling CN2 would result in a small, acceptable overhead, considering that the distance $||HA, CN2||$ is very small with regards to $||MR, HA||$ or $||MR, CN2||$. Now, when the MR moves back to Brittany and places a new call to CN2, going via the HA might double the distance, but the whole thing being local, it is negligible.

The geographical distribution of Home generalizes this latter situation. If we can get rid of the concept of a Home Link that anchors the HA in a single location, then we can distribute HAs geographically, and, hopefully, one is close to our MR when it's roaming.


So if a MR can locate and bind with a closeby HA, then $||MR, HA||$ is

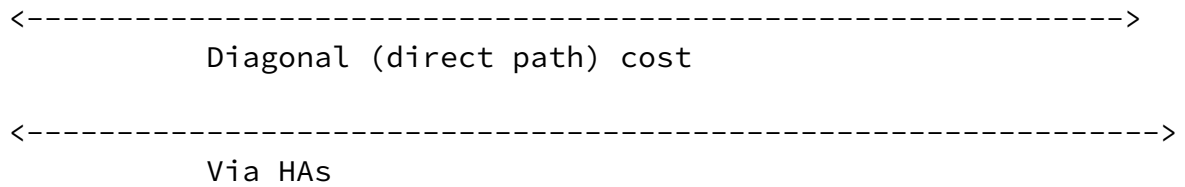
contained and the overhead is globally limited. In a same fashion, when a CN sends a packet to the MR, it finds a HA closeby and the overhead $||HA, CN||$ is contained as well.



Globally Distributed HA for World Wide service

In our previous example, we see that a HA' deployed in the East Coast saves the round trip over the Atlantic. There is a new overhead for the call to Europe, though, since an additional path is involved between MR and HA'. Then again, if both $||MR, HA'||$ and $||CN2, HA||$ are relatively small compared to $||HA, HA'||$ then the overhead is acceptable; unless all 3 points are located closeby, in which case, again, the additional cost is acceptable.





Illustrating that the overhead can be relatively small

[2.4](#) Single point of failure

The Home Link is a single point of failure for MIPv6/NEMO operations.

Should the Home Link fail, the whole set of MNs / MRs is disconnected from the rest of the world. One could decide to use a virtual link for Home, but then:

MIPv6 provides a support for multiple HAs, with the DHAAD mechanism. This mechanism helps scaling up the Home by adding HAs dynamically, and eventually load balancing the bindings between them. But this all relies on HAHA communication over the PHYSICAL Home Link; so making that link virtual implies a single Home Agent.

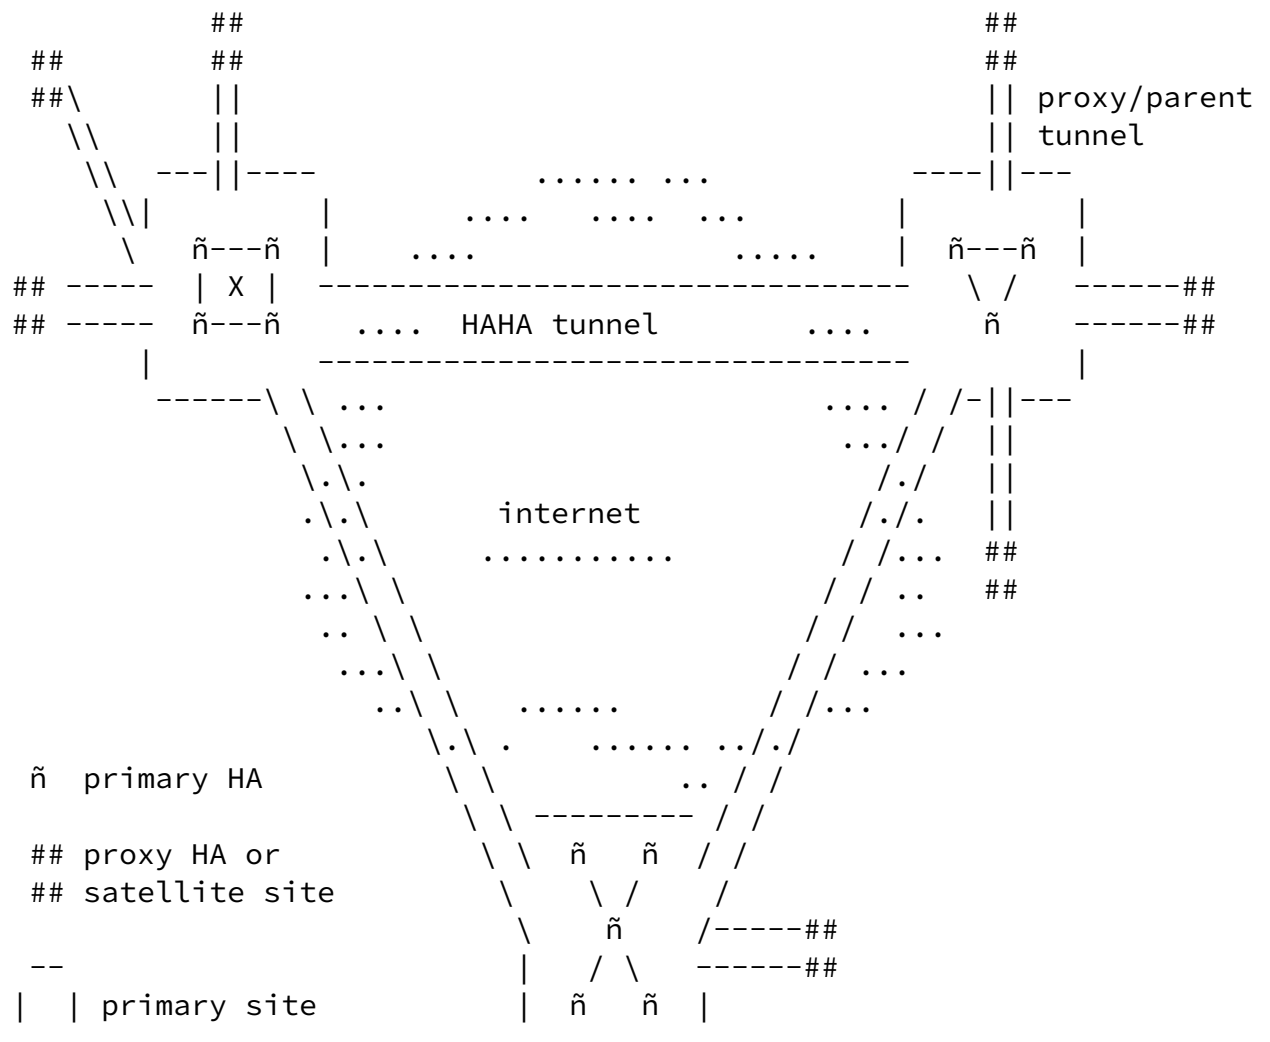
In turn this makes the HA a single point of failure, and disables the scalability that the DHAAD mechanism provides to MIPv6.

[3.](#) Rationale for the proposed solution

For the time being, the precise flows are not elaborated. One idea is that a protocol such as IS-IS or OSPFv3 could help a lot, mostly in the registration phase. Another is that HAs should be proactively preassigned to receive a given set of registration, in order to allow a certain degree of aggregation within sites and in between site. Finally, the concept of proxy is introduced to limit the number of primary sites (to 1?) and as a key element for an upcoming NEMO route optimization scheme, where routes can be exchanged in a trusted fashion between proxies.

4. Overview

This description covers the specific case of a Partitioned Home Network. Home is subnetted and the subnets are attributed to the distributed sites. As a result, in a given location, HAs will be operating both as primary HA taking the registrations for the local partition and proxy HA for registrations that belong to other sites. Additional satellite sites might be deployed around some of the main sites.



It is out of the scope of this document to discuss how the subnetting was decided and configured. It is also out of the scope of this document to describe the operations within a site where more than one HA is deployed. It is expected that in each primary site, HAs discover each other, mesh using tunnels, and form an area that owns the partition of Home that was assigned to that site.

[illegible]

Thus, a site attracts the DHAAD requests from any MR that happens to be roaming close to the site, regardless of the MR primary site. So MRs bind to the closest site from their physical location. In a same fashion, CNs send all packets to LFNs via the closest Home site. But packets back flow directly from the site where the MR is bound.

[4.1.2](#) Internal routing

In each area, border HAs are elected to mesh with peers in other sites. Again, they discover each other, mesh using tunnels, and form a backbone. It might be preferable to form a fully meshed backbone, in order to limit the cost of routing between sites. Also, making the backbone a transit area enables the use of a specific HA for the HAHA protocol that acts as an OSPF Designated Router or a BGP Route Reflector.

```

      .....
      ..... .../
-----
| site1 | ..... ; .... | Site2 |
-----
| Home:Site2::/48 ->          <- Home:Site1::/48 |
| -----HAHA-tunnel----- |
| ñ ñ ñ | / . | ñ ñ |
| ñ ñ | <- Home::/32 ; Home::/32 -> | ñ ñ ñ |
-----
      . / ... -----
      ... ..; .....
      /.....

```

If a link state protocol such as OSPFv3 is deployed between the primary HAs, then it uses the meshed backbone as its area zero, and each site as an area. Sites exchange their partitions of Home in the form of summarized routes.

It can be expected that, in order to scale, satellite sites would be deployed to take the proxy bindings but would not participate to the HAHA protocol that happens between the primary sites - at least when a proactive version of HAHA is being used.

```

      .....
      ..... .../.
-----
| Sat1 | ..... ; .... | Site1 |
-----

```

```

|-----|
| Home::/32 ->                               <- Home:Site1:Sat1:/64 |
| -----proxyHAHA-tunnel-----|
| #### | / . | ñ ñ ñ |
| #### | <- Home::/32 ; Home::/32 -> | ñ ñ |
|----- . / ... -----|
| ... ...; .....|
| /.....|

```

In a satellite site, HAs are only proxy, never primary. Each proxy HA has at least one assigned parent HA, which belongs to a primary

site. A tunnel is established between the proxy HA and the parent HA. The parent advertises the Home Aggregation to the proxy over that tunnel, as it does over the internet. In return, the proxy advertises its own prefixes, and redistributes the Home Aggregation over the internet. Finally, the parent redistributes the route to the proxy's network into its area, via itself, as an external route.

[4.2](#) Binding

At that point, the primary sites are ready to accept bindings, either directly from Mobile Routers or via proxy HAs. This is the runtime phase for HAHA.

A MR that is located close to its primary site will register there for its primary binding. In that case, the binding is direct. Otherwise, the MR will use a proxy in order to bind locally, and the proxy will perform the primary binding on behalf of the MR. If the proxy is parented at the primary site, the binding is local; otherwise, it is called a foreign binding.

[4.2.1](#) Direct primary binding

When the primary HA accepts a direct binding from a MR, then it must let the other primaries know that it owns the binding for that Home Address, in a fashion that is discussed in [Section 9.2](#).

```

.....
...../. ... -----
... ; .. | Site1 |

```

```

..      / Home::/32 ->.|   ñ--ñ--ñ
...      ;      .. |   /
....      /      MR ==MRHA==== ñ <- Home:site1:MNP::/64
...      ;      |      .. |   \
...      /      ----- ... |   ñ--ñ
...      ;      MNP      |
... / ..      ...      -----
.....

```

Primary HA injects necessary MR routes in area

The primary HA installs all (implicit and explicit) routes to the MR MNPs over the MRHA tunnel. It must also inject any required route, such as explicit prefix routes, into the primary area, as external routes via itself. All these routes are summarized at the area border and the other areas are not affected by the routing change.

[4.2.2](#) local proxy binding

When a MR binds to a satellite site, a HA acts as a proxy and binds in turn with a primary site, on behalf of that MR, to create the primary binding. The proxy binding can only succeed if the primary binding does. If the primary accepts the binding, then it returns a positive Binding Ack, with the list of the prefixes that are routed via the Mobile Router.

```

.. .....
----- .. .
| Sat1 | .. | / . | Site1 |
|      | <- Home::/32 ; . |
|      | .. | / .. |
| ----> ===== ñ <- Home:site1
|      | . | / .. | :MNP::/64
| -- # ===== MR ; ... |
|      | . | / |
|----- . ----- ; ...
|      | MNP.../..


```

Then the proxy HA installs the routes that it got from the the

positive Binding Acknowledgement over the proxy MRHA tunnel, and Acknowledges the proxy BU. Once a primary binding has succeeded, the proxy might establish secondary bindings with other sites.

4.2.3 Foreign proxy binding


When a MR binds to a foreign site, whether the site is primary or satellite, a HA from the site acts as a proxy as if the site was a satellite from the primary.

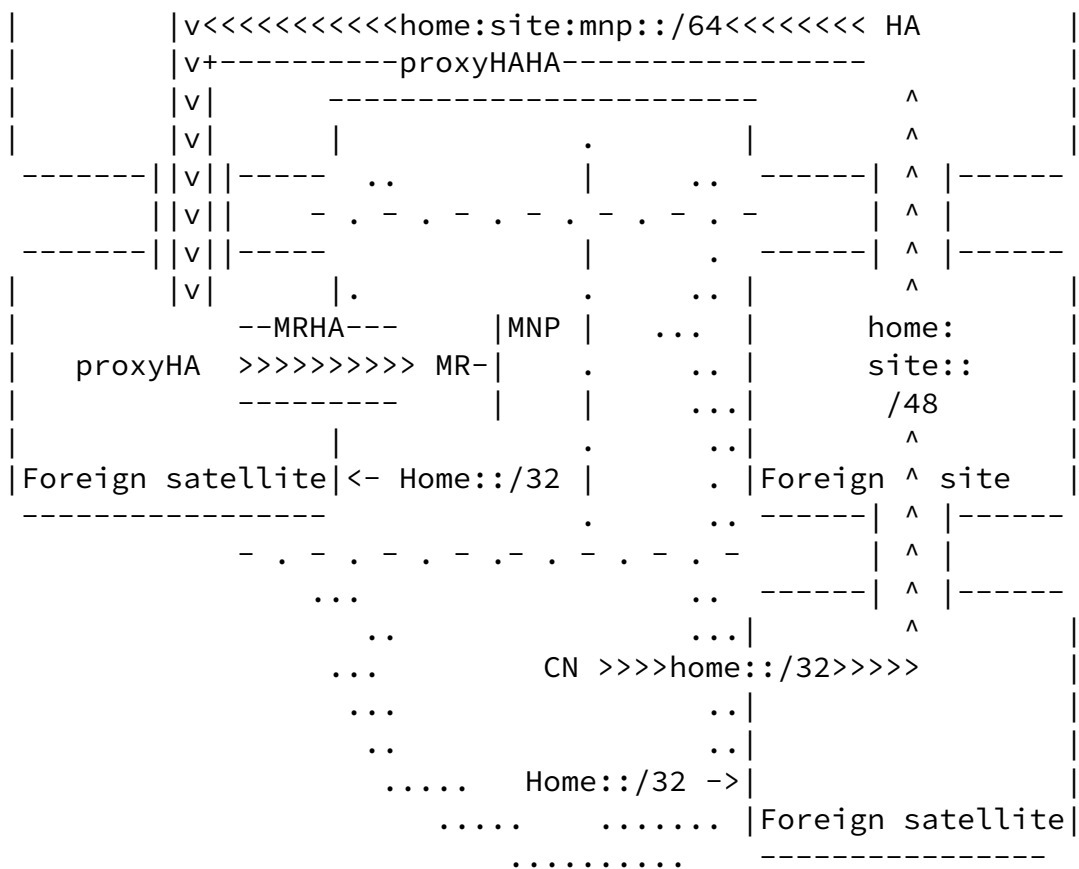


4.3 Path improvements

When the MR binds in a foreign location, the transport between an arbitrary correspondent and the MR within the HAHA network might be far from optimized.

As a result of the primary binding, a proxyHAHA tunnel is established between the proxy and the primary HA. That tunnel is itself encapsulated in the HAHA tunnels when packets flow over the internet.





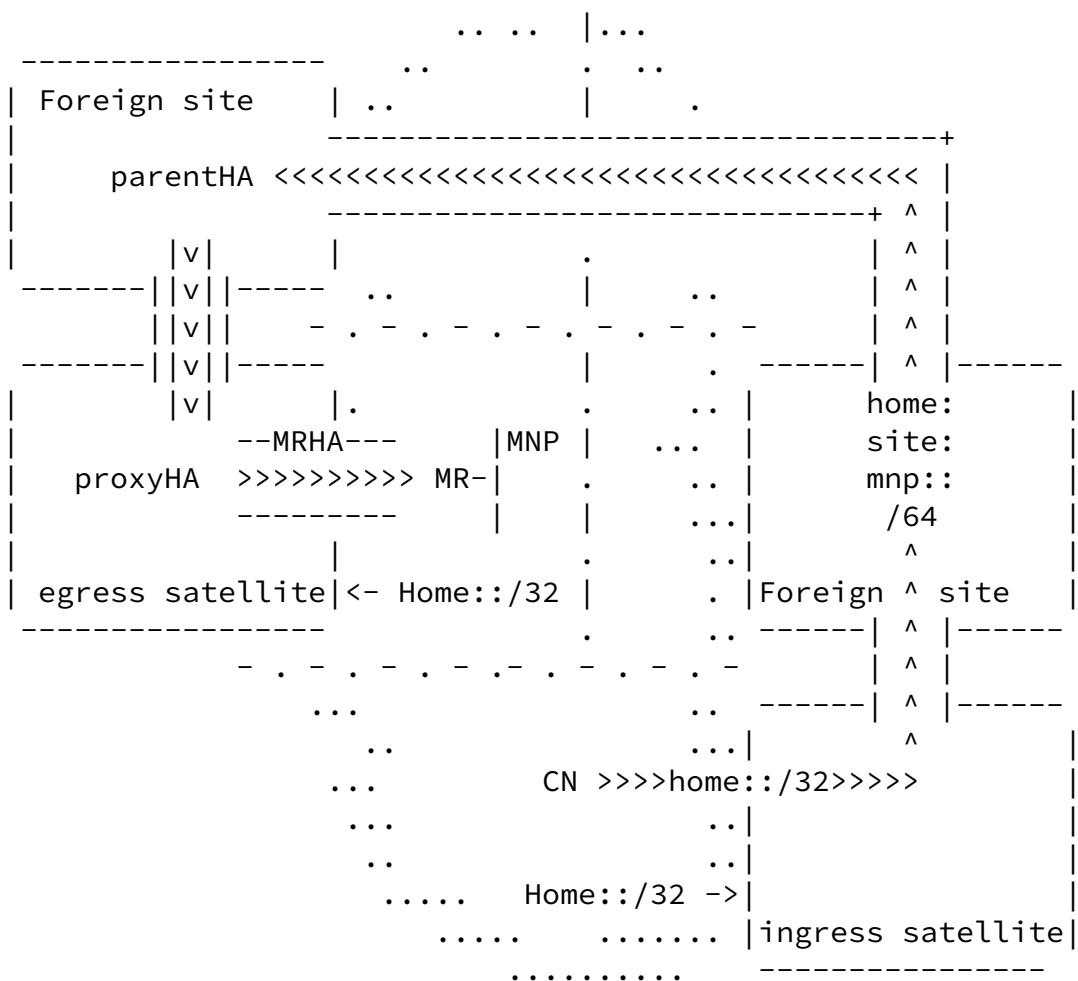
The path from CN to MR is not optimized

Also, packets from an arbitrary correspondent reach the site that is closest to the correspondent, then forwarded to the primary site for the destination. Within the primary site, they are encapsulated towards the proxy and sent across the HAHA network again. Finally they reach the proxy that decapsulates the packets and encapsulates them back.

In order to improve this, various possibilities are offered:

[4.3.1](#) Leaking MNP routes in the HAHA network

The proxy can establish a secondary binding with its parent. In return, the parent redistributes an external route to the MNP via itself, and leaks that route inside the whole HAHA network.

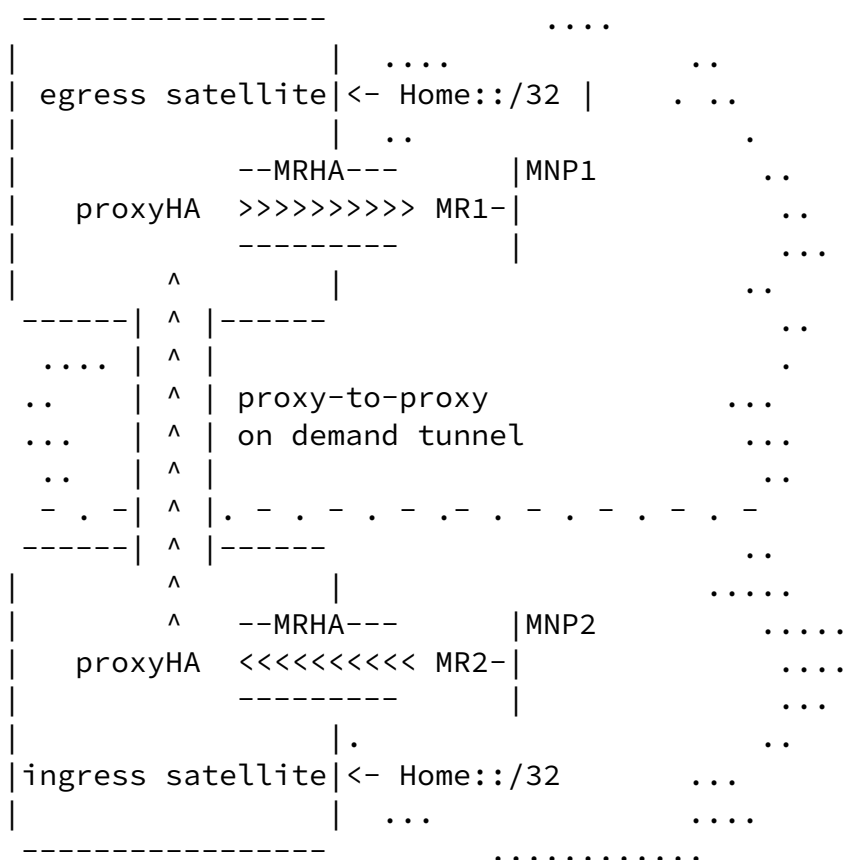


The path from CN to MR bypasses the primary HA

This bypasses the primary home agent for packet forwarding. Note that the packets still flow within the HAHA network between the ingress site close to the correspondent and the egress (satellite) site.

Note also that when the proxy HA binds to either its parent or the primary HA, it uses an address from within the HAHA network (its HAHA Address), as CareOf.

The proxy can establish a secondary binding with the correspondent's proxy provided there's such a node. It might be envisioned to adapt NHRP [5] for IPv6 in order to discover the remote tunnel end point.



The path is now direct between the proxies

An example of application is when two proxies from a same Home establish a cross binding. In fact, the Mobile Routers are unaware of the path improvement that takes place. This feature might be desirable when the privacy of the location is an issue for the service provider.

As part of the secondary binding to the ingress proxy, the egress proxy passes all the MNPs for the MR as explicit prefix routes. It is expected that the proxies belong to a chain of trust that links the primary and the satellite sites together. This, the ingress proxy trusts the egress proxy both for the binding and for the explicit prefixes.

Note that in that case, the binding uses the proxy's external address as careof. The packets are thus routed straight between the proxies, outside of the HAHA network.

5. Terminology and concepts

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 [RFC2119](#) [1].

Most of the mobility related terms used in this document are defined in the Mobility Related Terminology document [8] and in the Mobile IPv6 specification [9].

Additionally, some terms were created or extended for NEMO. These specific terms are defined in the NEMO Terminology document [13].

This draft introduces the following definitions:

Distributed Home Network: In distributed home network, a global Home is advertised by several sites that are geographically distributed and meshed using tunnels in a VPN fashion. Mobile Nodes locate the closest site using DHAAD and bind there. More in [Section 6](#)...

Partitioned Home Network: A Partitioned Home is a specific deployment of a Distributed Home Network where each location owns a subnet of Home. The local Home Agents accept registration for the local partition. The local HAs also act as NEMO proxy HAs for the rest of Home.

Primary Home Agent: A Home Agent that can serve a Binding Update from a Mobile Router. The Mobile Router is always associated with a (set of) primary Home Agent (s) to register its binding.

Proxy Home Agent: This is a form of proxy, for the NEMO protocol. A proxy HA acts as a HA for MRs to register, but needs to register to a primary HA in order to accept the binding.

Primary site: A site is primary for a MR if at least one local HA on that site can accept a registration for that MR. When Home is not partitioned and sites overlap, primary HAs for a same subnet have to be aware of each other in order to find if a binding already exists in one of the sites and in which Home Agent.

satellite site: A site that is not primary for any binding. It is dependent on a parent primary site for HAHA operations. satellite sites are deployed around central primary sites, and one final goal for HAHA is to dynamically draw routes between satellite sites in order to shortcut the backbone of primary HAs.

Internet-Draft

Global-HAHA

October 2004

Secondary site: A site is secondary for a MR if it is primary for other MRs but not that one. HAs in a secondary site can act as proxies for that MR, and the site is its own parent.

Primary Binding: A Binding is primary if it happens with a primary Home Agent, whether the client is a MR or a proxy HA.

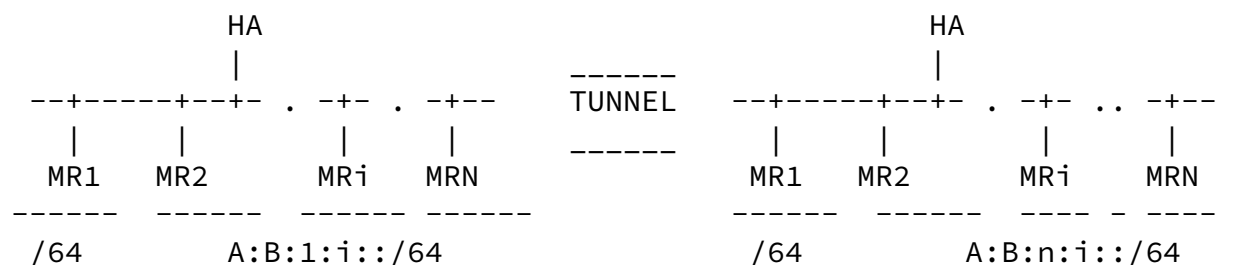
Secondary Binding: A Binding is secondary if it happens between a proxy and a non primary Home Agent. It is used to improve the path between sites towards the HA where a MR is registered.

Proxy Binding: A Binding is proxy if it happens between a MR and a proxy HA, whether the proxy is a pure proxy HA or a secondary HA acting as proxy for that MR. The proxy HA relays the proxy binding to a primary HA in a primary binding. It may maintain a set of secondary bindings, depending on the deployment.


Direct Binding: A Binding that does not pass via a proxy, straight between the MR and its Home Agent.

6. Distributed Home Network

This section describes a detailed example how multiple Home Agents are configured in different routing domains. You are encouraged to read the nemo basic Home Network Models [[12](#)] draft before going through this section.



Distributed Home Network /48



In distributed home network, a global Home is advertised by several sites that are geographically distributed and meshed using tunnels in a VPN fashion. Mobile Nodes locate the closest site using DHAAD against the global Home Network and bind there. Some form of

inter-site synchronization (e.g. a routing protocol), which Mobile IPv6 and Nemo Basic Support do not provide, must take place in order to allow packets to be routed between the incoming site to the Mobile Node. The HAHA (Home Agent to Home Agent) protocol is being designed for that purpose.

In one model, called the Partitioned Home Network each site is responsible for a subnet of Home. When a Mobile Node roams far from its natural (primary) site, it registers to a Home Agent on a remote site, that takes the registration and notifies at least the natural site of the foreign registration.

One specific advantage of not relying on a Home Link for HAHA communication is that for a large configuration, the Home Agents can be organized hierarchically and distributed geographically, as a set of local clusters linked together to form a global Home Network.

[7.](#) Message Formats

[8.](#) Mobile Router Operation

[8.1](#) Locating Home

[8.2](#) Proxy MIP client

[9.](#) Home Agent Operation

[9.1](#) Locating the other HAs that serve the same Home

[9.2](#) Locating the HA that owns the binding for a HoA

At the time of processing a binding update, a Home Agent (be it primary or simply proxy for the binding Home Address) needs to discover if the binding already exists with a primary Home Agent. There are at least 3 approaches that might be deployed for that purpose:

Reactive: This method is also referred to as 'on-demand'. In case of a binding cache miss, a primary Home Agent floods a Binding Information Request message to all the other primary Home Agents for the home address that is sought for. The reactive approach can be used between a satellite site and its parent site even when

the primary HAs use an other method in the backbone.

Proactive: The binding information is shared proactively between the primary Home Agents for the existing bindings. All primary HAs know at any point of time which Home Addresses are bound and with which primary Home Agent. This approach is preferred for stable configurations, for instance if NEMO is used as a tool to simplify the configuration and reconfiguration of mostly stable networks.

Predictive: Ranges of Home Addresses and prefixes are preassigned to the Home Agents, following a rule that is shared or commonly computed by all HAs. A partitioned Home Network is an example of that, but this is mostly useful within a site between local Home Agents.

10. Acknowledgements

The authors wish to thank:

11 References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [3] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", [RFC 2461](#), December 1998.
- [4] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", [RFC 2462](#), December 1998.
- [5] Fox, B. and B. Petri, "NHRP Support for Virtual Private Networks", [RFC 2735](#), December 1999.
- [6] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6)

Addressing Architecture", [RFC 3513](#), April 2003.

- [7] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.
- [8] Manner, J. and M. Kojo, "Mobility Related Terminology", [RFC 3753](#), June 2004.
- [9] Johnson, D., Perkins, C. and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [10] Patel, A., "Problem Statement for bootstrapping Mobile IPv6", [draft-ietf-mip6-bootstrap-ps-00](#) (work in progress), July 2004.
- [11] Devarapalli, V., "Network Mobility (NEMO) Basic Support Protocol", [draft-ietf-nemo-basic-support-03](#) (work in progress), June 2004.
- [12] Thubert, P., Wakikawa, R. and V. Devarapalli, "NEMO Home Network models", [draft-ietf-nemo-home-network-models-00](#) (work in progress), April 2004.
- [13] Ernst, T. and H. Lach, "Network Mobility Support Terminology", [draft-ietf-nemo-terminology-01](#) (work in progress), February 2004.
- [14] Ernst, T., "Network Mobility Support Goals and Requirements", [draft-ietf-nemo-requirements-02](#) (work in progress), February 2004.

Thubert, et al.

Expires April 5, 2005

[Page 20]

Internet-Draft

Global-HAHA

October 2004

- [15] Ernst, T., "Analysis of Multihoming in Network Mobility Support", [draft-ietf-nemo-multihoming-issues-00](#) (work in progress), July 2004.

Authors' Addresses

Pascal Thubert
Cisco Systems
Village d'Entreprises Green Side
400, Avenue de Roumanille

Batiment T3
Biot - Sophia Antipolis 06410
FRANCE

Phone: +33 4 97 23 26 34
EMail: pthubert@cisco.com

Ryuji Wakikawa
Keio University and WIDE
5322 Endo Fujisawa Kanagawa
252-8520
JAPAN

EMail: ryuji@sfc.wide.ad.jp

Vijay Devarapalli
Nokia Research Center
313 Fairchild Drive
Mountain View, CA 94043
USA

EMail: vijay.devarapalli@nokia.com

Intellectual Property Statement

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.

Disclaimer of Validity

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 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.

Copyright Statement

Copyright (C) The Internet Society (2004). 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.

Acknowledgment

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