INTERNET DRAFT
Expires: 24 Apr 2005

Ryuji Wakikawa
Masafumi Watari
Keio Univ./WIDE
24 Oct 2004

Optimized Route Cache Protocol (ORC)
draft-wakikawa-nemo-orc-01.txt

Status of This Memo

This document is an Internet-Draft and is subject to all provisions of section 3 of RFC 3667. By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she 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/1id-abstracts.txt.
The list of Internet-Draft Shadow Directories can be accessed at:
http://www.ietf.org/shadow.html.

Copyright Notice

Abstract

This draft proposes Optimized Route Cache Protocol (ORC) to provide route optimization for the NEMO Basic Support protocol. ORC provides a dynamic route optimization mechanism, similar to route optimization in Mobile IPv6.

The ORC aims to manage binding information for when routing information of each mobile network are located at special routers called ``Correspondent Router''.

Contents

1. Introduction

The NEMO Basic Support protocol [4] is currently being standardized at the IETF NEMO working group. However, the NEMO Basic Support protocol does not provide route optimization, but always use a bi-directional tunnel established between a mobile router and its home agent. Optimized route cache protocol is designed as an extension to the NEMO Basic Support protocol for providing certain route optimization. ORC was proposed earlier in the paper [8].

In the NEMO Basic Support protocol, a binding of a mobile network can be treated as routing information of the mobile network. For instance, a care-of address in the binding can be treated as a next hop address in a routing table. It is against the general standard operations of the Internet for each end-node to handle route information. End nodes should be unaware of routing since managements of a binding may bother them.

2. ORC Concept

In Mobile IPv6 [5] and the NEMO Basic Support protocol, a home agent is an original anchor router of a mobile network and maintains a binding of the mobile network persistently. All packets are first routed to the home agent and are tunneled to the mobile router by the home agent unless the mobile router starts route optimization.

On the other hand, the optimized route cache protocol introduces correspondent routers that can be configured anywhere in the Internet to be an anchor router for the mobile network, providing certain level of route optimization. Practically, the correspondent routers should be scattered near the transit AS to allow direct forwarding to the mobile network before reaching the home agent. Because it is impossible to replace all routers on the Internet with the correspondent routers support, it is effective to place a correspondent router at places where traffic is converged like the Internet Exchange Point (IXP).

The optimized routing cache protocol provides optimal route path in best effort when correspondent routers exist. However, the level of optimization can be improved depending on where the correspondent routers are placed, and the path it takes. The correspondent routers can also be dynamically discovered if necessary, to provide certain route optimization.

Since the binding is processed and maintained only by the correspondent routers scattered over the Internet, mobile router does not need to handle bindings for each correspondent nodes. It is clearly redundant operations if both the mobile router and the

remote network manage bindings for the same communicating network. This also allows the end-nodes to communicate in the optimized route without requiring any additional functions.

This concept can be applied to Mobile IPv6 as well. In Mobile IPv6, correspondent nodes are required to extend its protocols suites for route optimization. However, it is not reasonable to assume all end-nodes to support the Mobile IPv6 protocol. Therefore, a mobile host can not initiate route optimization (i.e. return routability) to all correspondent nodes. In such case, the network of which the correspondent nodes are connected to can provides route optimization on behalf of the correspondent nodes.

3. Terminology

Most of the terminology is described in [5] and [3]. This document in addition defines the following terms.

Correspondent Router

An edge router of correspondent nodes' network. A Correspondent Router is well defined in [7] and is also defined as ``ORC router'' in the paper [8]. A correspondent router is capable to manage a binding of any mobile router and setup forwarding for mobile network prefixes. A correspondent router can be statically configured or dynamically discovered by the mobile router.

Correspondent Router Anycast address

An anycast address assigned to each correspondent router. It is generated by the correspondent router's 64 bit prefix and an anycast identifier. The anycast identifier is to be defined by IANA.

Managed Prefix

Prefixes which are managed by each correspondent router. The Managed Prefix is often configured with administrative policy. For example, if a correspondent router is placed for an administrative domain (let's say 2001:a:b::/32), the Managed Prefix for the correspondent router is the 2001:a:b::/32. Mobile router can tunnel any packets meant for the managed prefixes to the correspondent router. The correspondent router has responsibility to route packets correctly to the destination which is in the managed prefixes.

Proxy Route

A proxy Route is used to intercept packets by a correspondent router at an administrative domain. A proxy route is to direct a route of a mobile network prefix to a correspondent router. Proxy route contains a mobile network prefix of a correspondent router as a destination and the correspondent router's address as a next hop. The proxy route will not be aggregated in an IGP domain, but can be distributed inside the IGP domain. Proxy Route is used to intercept packets when a correspondent router is not on the path of the traffic from the correspondent node to the Mobile Networks. (i.e. The correspondent router is neither Default Router nor Core Router. See [7]).

4. ORC Overview

4.1. Correspondent Router Discovery

Each mobile router may have the list of the correspondent routers beforehand by system administrators or users.

In addition, when a mobile network node starts communication with a correspondent node, a mobile router may dynamically discover a correspondent router for the correspondent node. The discovery is triggered when a mobile router receives tunneled packets from its Home Agent. The mobile router first sends a correspondent router Discovery Request to the correspondent router anycast address. The anycast address is created with the prefix address of the correspondent node and the well-known anycast identifier. The prefix length for the anycast address is always 64. When one of correspondent router receives the Correspondent Router Discovery Request, it replies back a Correspondent Router Discovery Reply including all correspondent routers of the administrative domain of which the correspondent node is located.

4.2. Binding Registration to Correspondent Router

After receiving the correspondent router addresses, the mobile router can attempt Binding Registration to the correspondent router. The mobile router sends a Binding Update which is protected by IPsec to the correspondent router. The mobile router MUST set ORC flag 'O' in the Binding Update and include the Mobile Network Prefix sub-options. The Binding Update message is same as a Binding Update sent to a Home Agent except for the flag field (i.e. 'O' flag set and 'H' flag unset).

After processing the Binding Update successfully, the correspondent router MUST return a Binding Acknowledgment including the managed prefix list. The managed prefix is used when the mobile router decides which packets are sent to which correspondent router. The correspondent router setup forwarding for all the mobile network prefixes notified by the mobile router.

After getting successful Binding Acknowledgment, the mobile router set up forwarding for all managed prefixes. If the mobile router gets error status code in the Binding Acknowledgment or cannot get any Binding Acknowledgment, it SHOULD stop sending Binding Update to the correspondent router and SHOULD mark the correspondent router as invalid.

There is alternative mechanism based on Return Routability to send secured Binding Update. The detailed operation is introduced in Appendix.

4.3. Forwarding between Mobile Router and Correspondent Router

Once a mobile router registers its binding to a correspondent router, it forwards packets destined to an address which is in range of correspondent router's managed prefix. The correspondent router also intercepts packets sent to the Mobile Network and tunnels them to mobile router by IP-in-IP encapsulation.

5. Extensions to Mobile IPv6 and the Basic NEMO protocol

5.1. Forwarding Table Data Structure

Forwarding table is maintained by each mobile router. It has conceptually the following fields. How to implement forwarding table is up to implementations.

- Correspondent Router Address

- Managed Prefix Lists
  The list of Managed Prefix which is notified by the correspondent router

5.2. Mobility Header Messages

5.2.1. Binding Update

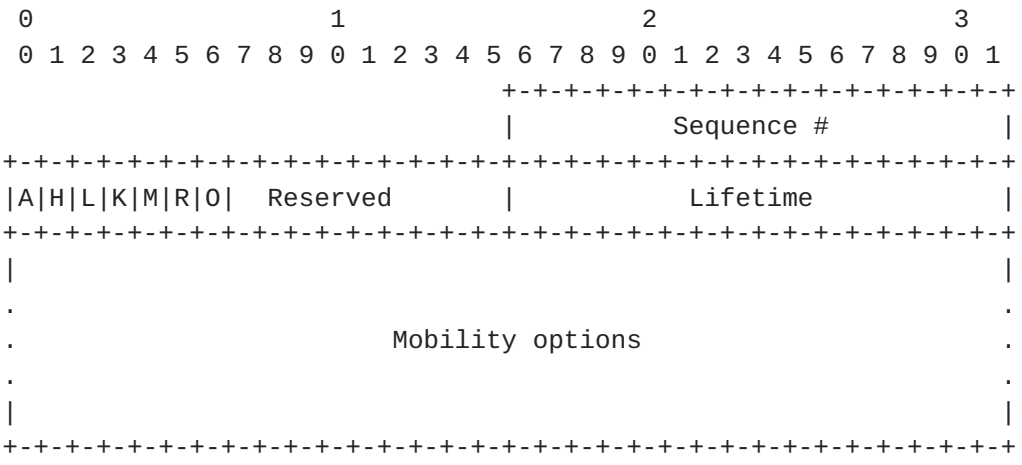

ORC Flag (O)
The flag is used to identify a Binding Update sent for a correspondent router.

5.2.2. Binding Acknowledgment

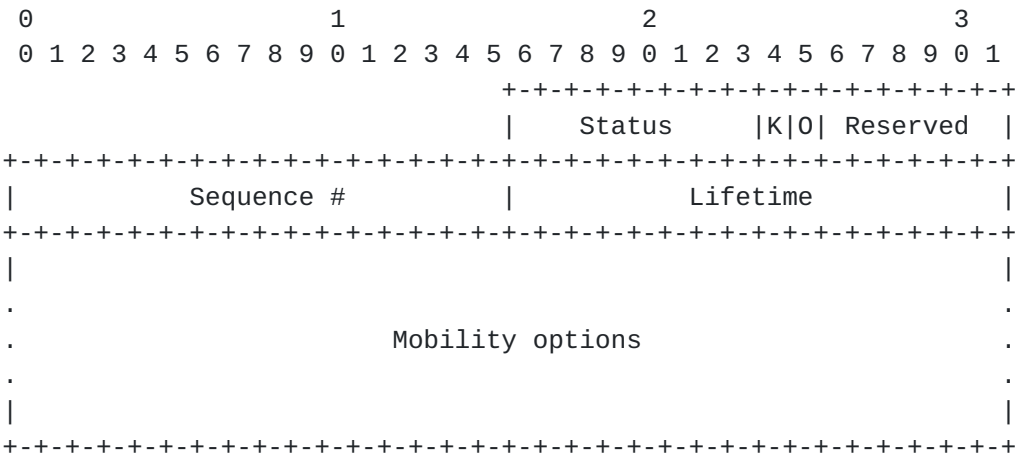

ORC Flag (O)
The flag is used to identify a Binding Acknowledgment sent from a correspondent router.

5.2.3. Managed Prefix Lists sub-option

The managed prefix lists mobility header sub-option is valid only in the Binding Acknowledgment.

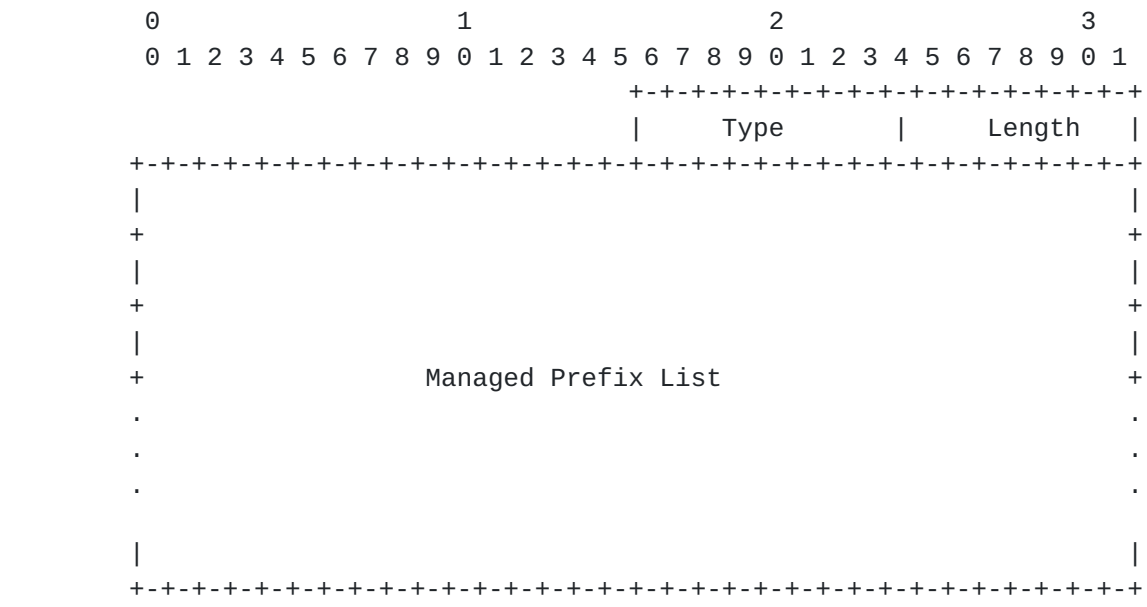

Type

To be assigned by IANA

Length

The length of sub-option.

Managed Prefix List

The list of prefixes which are managed by a correspondent router. Lower bit after prefix length should be set zero. (ex. 2001:a:b:c::/64 is expressed by 2001:a:b:c:0:0:0:0)

5.3. New ICMP Messages

Two new ICMP messages are defined for the discovery of the correspondent routers. Two messages have similar structure of home agent address discovery request and reply messages defined in Mobile IPv6 [5].

5.3.1. Correspondent Router Discovery Request

The Correspondent Router Discovery Request message is used by a mobile node to discover correspondent routers where correspondent nodes are located. The message format is as follows.

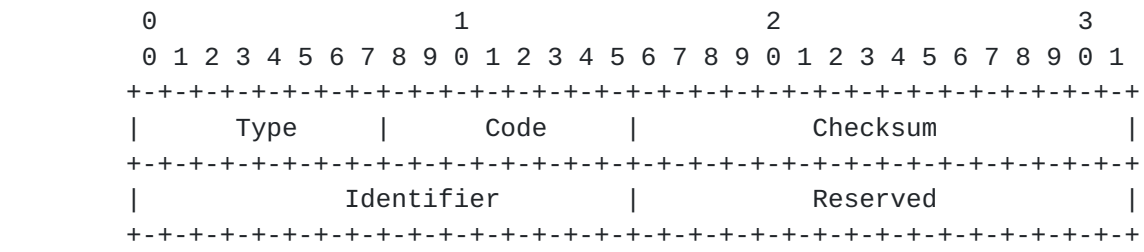

Type

To be assigned by IANA

Code

0

Checksum

The ICMP Checksum

Identifier

The 16-bit identifier to aid in matching correspondent router Discovery Reply message. The identifier should never be set to 0 and should always be more than 1.

Reserved

0

5.3.2. Correspondent Router Discovery Reply

The Correspondent Router Discovery Reply message is used by a correspondent router to send lists of correspondent routers configured at the network in reply to the Correspondent Router Discovery Request message. The message format is as follows.

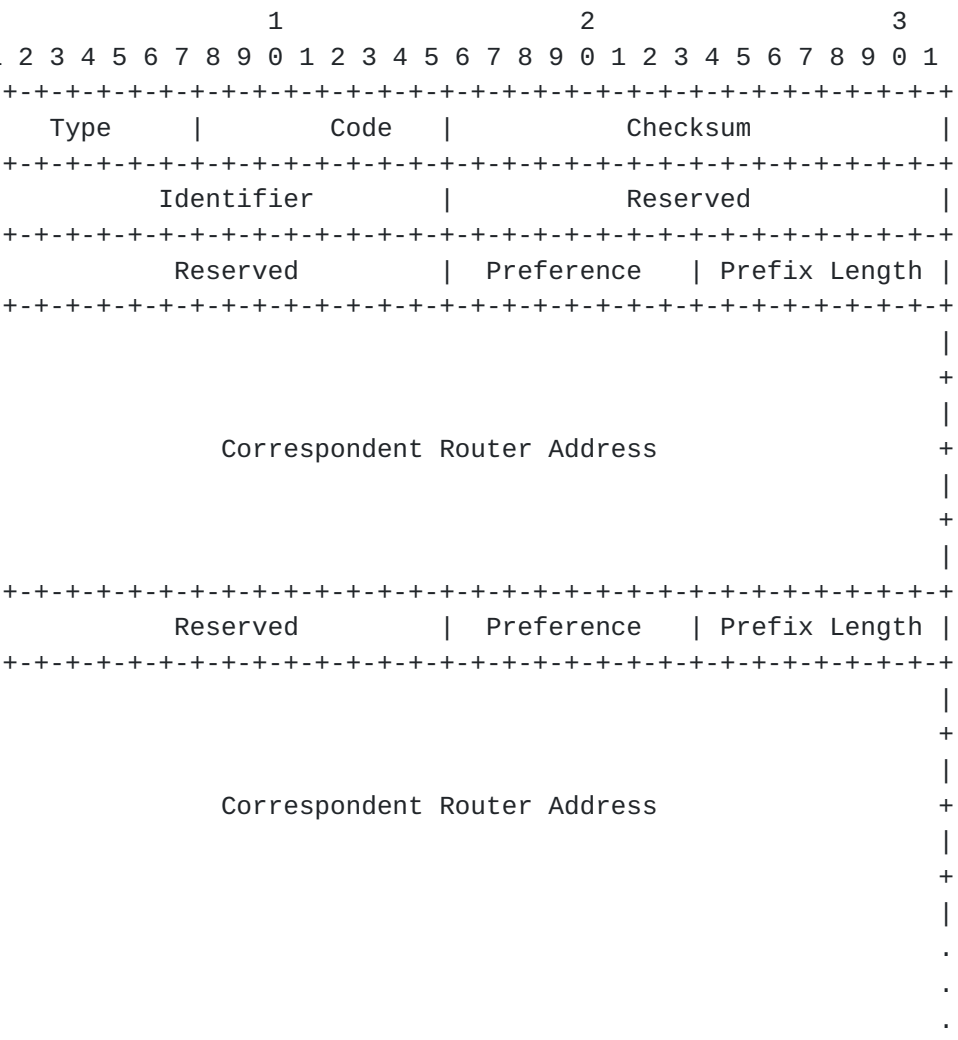

Type

To be assigned by IANA

Code

0

Checksum

The ICMP Checksum

Identifier

The 16-bit identifier to aid in matching Correspondent Router Discovery Reply message. The identifier should never be set to 0 and should always be more than 1.

Reserved

0

Preference
: The 8-bit preference value of each correspondent router. The default preference value is zero. Higher value indicate higher preference.

Prefix length
: The length of prefix with which a correspondent router is configured and responsible for.

Correspondent Router Address
: A global IPv6 address of a correspondent router.

A correspondent router replies multiple addresses of correspondent routers that are configured in same network domain by a single Correspondent Router Discovery Reply message.

6. Protocol Operations

6.1. Correspondent Router Discovery

A correspondent router is dynamically discovered with Correspondent Router Discovery and Correspondent Router Reply. The discovery mechanism is similar to the dynamic home agent discovery mechanism of Mobile IPv6 [5].

When a mobile router detects that a received packet is tunneled by its home agent, it can initiate Correspondent Router Discovery on demand by sending a Correspondent Router Discovery Request to the correspondent router anycast address. The mobile router learns the correspondent router anycast address from the correspondent node's prefix and the anycast identification. The prefix length of the correspondent node's prefix is always assumed to be 64-bit. Correspondent routers with a shorter prefix length are notified later with a Correspondent Router Discovery Reply.

If no replies are received, the mobile router stops further discovery for correspondent routers to the network of which the correspondent node are located. The mobile router must then communicate through its home agent.

If the mobile router receives a Correspondent Router Discovery Reply, it first verifies the message header (e.g. ICMP checksum and identifier). If all verifications are passed, it retrieves the correspondent router addresses. If more than one addresses are included, the mobile router selects one of the addresses and starts explicit binding registration described in section 6.2. The determination of which correspondent routers to select are handled

with preference values and prefix length. Some examples are shown in Appendix B.

6.2. Binding Registration to Correspondent Router

6.2.1. Sending Binding Update

A mobile router MAY maintain IPsec Security Association with correspondent routers. Alternatively, it MAY use Return Routability mechanism to protect Binding Update described in Section 6.2.2.

A mobile router creates a Binding Update as indicated in the basic NEMO protocol. It MUST set 'O' flag in the Flag field of a Binding Update. The Binding Update MUST be always protected by IPsec or Return Routability mechanism.

A mobile router also records the sent Binding Update as a Binding Update list entry for each correspondent router.

6.2.2. Return Routability

In Mobile IPv6, a mobile host provides reasonable assurance with the correspondent nodes through the return routability mechanism, and securely register its binding. A correspondent router is similar to the correspondent node in terms of security relationship with a mobile router. Although IPsec provides stubborn security for binding registration, it is expensive operations for both a mobile router and a correspondent router. The ORC follows similar approach to Mobile IPv6 so that secured binding registration is performed with the return routability mechanism.

It is necessary to extend the return routability procedure to register mobile network prefix information. To complete return routability for a mobile network, a mobile router is required to generate its home address from its mobile network prefix instead of its home network.

In Mobile IPv6, return routability procedure plays two roles when authenticating a binding update. One is to verify if the binding between the home address and the care-of address is legit, The other role is to exchange keys for authorizing binding update. In the optimized route cache protocol, following extensions are required in addition to return routability procedure. The home agent must verify the HoTI that is securely tunneled from the mobile router. The HoTI should be checked for its source address and prefix length.

HoTI will be sent with the home address as the source address, generated from the mobile network prefix. Thus, if the source address does not match the home address registered in home agent's binding, the home agent discards the HoTI. Furthermore, if the prefix length registered for the mobile router is different from the prefix sub-option sent, the home agent also discards the HoTI. On the other hand, CoTI will be sent with the care-of address as its source address. Once the mobile router receives both HoT and CoT back from the correspondent router, it is assured that the mobile router exists in topologically correct attachment point and also assures that it is the router of the network with the mobile network prefix. The mobile router can now send a binding update to the correspondent router with the keys exchanged in return routability. If the correspondent router can recompute the encryption, the binding update completes in success.

When a mobile router sends a binding update, it must set the binding acknowledge flag in order for it to receive a binding acknowledgment message from the recipient. The correspondent router must return a binding acknowledgment message containing a list of managed prefixes of its IGP domain in the managed prefix mobility option. The managed prefix mobility option is defined in section 5.2.3. If the binding update is successfully processed by the correspondent router, the mobile router establishes a bi-directional tunnel with the correspondent router as in [5]. The mobile router also records the pair of the prefixes retrieved from the managed prefix mobility option and the correspondent router's address as route entries in its routing table. These routes may be used to search a correspondent router in a routing table when the mobile router sends packet to correspondent nodes described in section 6.4.

6.3. Intercepting Packets by Correspondent Router

A correspondent router basically intercepts packets for a registered mobile router by IP level routing. However, there is different operations depending on correspondent router's topological location.

If a correspondent router is located as a gateway router of a network, it intercepts packets by parsing all packets' destination address with registered bindings.

On the other hand, if a correspondent router is located in a network, it MUST advertise a proxy route for a mobile network prefix of registered binding to its routing domain. All routers in the same routing domain forward packets meant for the mobile network prefix to the correspondent router who is advertising the prefix route.

6.4. Routing to Mobile Network

Whenever a correspondent router receives packets and query routing table as general router operations, it also searches for binding cache for a destination address in the IPv6 header just like any home agent. The correspondent router should select the prefix longest matched binding and route for the destination. When the correspondent router finds the prefix longest matched binding for the destination, it must search binding cache database recursively for the next hop address of the binding and must select the last matched binding for the destination. This recursive operation is aimed to support nested mobility.

Once the correspondent router finds a binding instead of an IGP route for outgoing packets, it tunnels the packets directly to the care-of address of the destination according to the registered binding. For the opposite direction, the mobile router may reverse tunnel packets to the correspondent router at correspondent node's IGP domain which is found with route of the correspondent router's managed prefixes in mobile router's routing table. The correspondent router then decapsulates packets and route them to a correspondent node. The mobile router does not insert the home address option as Mobile IPv6, since falsification of mobile network node's packets on intermediate nodes like the mobile router should be avoided for security considerations. The encapsulation of packets adds additional IPv6 header, and it does not change original packets.

7. Security Consideration

The optimized route cache protocol enables to manage routing information across AS boundaries. In other words, it is possible for a mobile router to alter routing table of opposite routers. Wrong binding registrations will cause opposite ASs to fall into confusion or to have black-hole of routing. The ORC employees Mobile IPv6 security mechanism [5] for protecting binding updates which are the IPsec authentication header [1] and the return routability scheme. Furthermore, recipient routers can apply their IGP domain or AS routing policies to handle each binding.

8. Acknowledgements

The authors would like to thank Keisuke Uehara, Susumu Koshiba, Jun Murai, and WIDE Project for their contributions.

References

[1] R. Atkinson. IP Authentication Header. Request for Comments (Proposed Standard) 1826, Internet Engineering Task Force, August 1995.

[2] A. Conta and S. Deering. Generic Packet Tunneling in IPv6 Specification. Request for Comments (Proposed Standard) 2473, Internet Engineering Task Force, December 1998.

[3] Thierry E. et al. Network Mobility Support Terminology. Internet Draft, Internet Engineering Task Force, February 2002.

[4] Vijay Devarapalli. et al. Network Mobility (NEMO) Basic Support Protocol. Internet Draft, Internet Engineering Task Force, June 2004.

[5] David B. Johnson, C. Perkins, and Jari Arkko. Mobility Support in IPv6. Request For Comments 3775, Internet Engineering Task Force, June 2004.

[6] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). Request for Comments (Draft Standard) 1771, Internet Engineering Task Force, March 1995.

[7] P. Thubert, M. Molteni, C. Ng, and E. Ohnishi, H. Paik. Taxonomy of Route Optimization models in the Nemo context (work in progress). Internet Draft, Internet Engineering Task Force, February 2004.

[8] R. Wakikawa, S. Koshiba, K. Uehara, and J. Murai. ORC: Optimized Route Cache Management Protocol for Network Mobility. In The 10th International Conference on Telecommunication (ICT) 2003, pages 119--126, February 2003.

Authors' Addresses

Ryuji Wakikawa
Graduate School of Media and
Governance, KEIO University
5322 Endo Fujisawa
Kanagawa, 252-8520
JAPAN
Phone: +81-466-49-1100
EMail: ryuji@sfc.wide.ad.jp
Fax: +81-466-49-1395

Masafumi Watari
Graduate School of Media and
Governance, KEIO University
5322 Endo Fujisawa
Kanagawa, 252-8520
JAPAN
Phone: +81-466-49-1100
EMail: watari@sfc.wide.ad.jp
Fax: +81-466-49-1395

A. Example Scenario

Figure 1 shows the configuration of the optimized route cache protocol. In the figure, there are five ASs connected to each other by Border Gateway Protocol (BGP) [6]. This can be assumed to be typical Internet BGP routing topology.

In Mobile IPv6 and the NEMO Basic Support protocol, a home agent is an original anchor router of a mobile network and maintains a binding of the mobile network persistently. All packets are first routed to the home agent and are tunneled to the mobile router by the home agent unless the mobile router starts route optimization. Therefore, in the case when correspondent nodes in AS3 communicates with the mobile network nodes in AS5, packets must first be routed to the HA in AS2 via AS1, before being tunneled to the destination nodes.

On the other hand, the optimized route cache protocol introduces correspondent routers that can be configured anywhere in the Internet to be an anchor router, providing route optimization. Practically, the correspondent routers should be placed on expected networks where there exist correspondent nodes for a mobile router and a

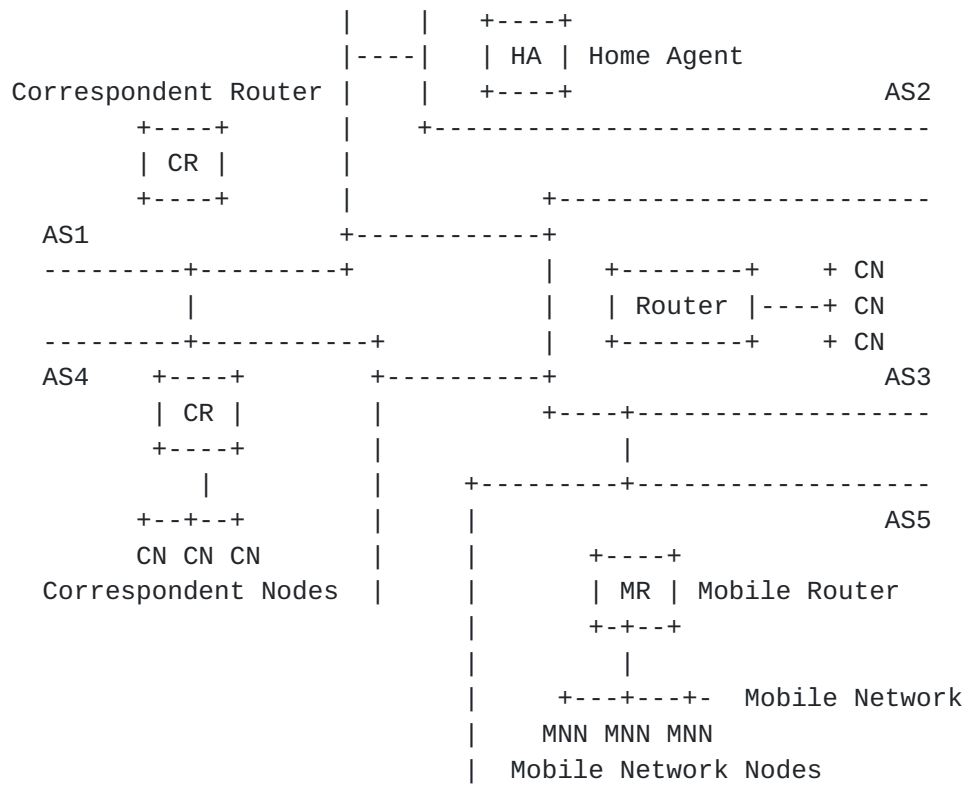

Figure 1: Optimized Route Cache Protocol Overview

mobile network because it is impossible to replace all routers on the Internet with the correspondent routers support. It is effective to place a correspondent router where traffic is converged like Internet Exchange Point (IXP).

Whenever a mobile router moves, correspondent routers may receive a binding update notification on-demand from the mobile router and cache them. The correspondent router must authorize the mobile network to receive the binding as described in section 6.2. After creation of the binding, the correspondent router intercepts packets destined to the mobile network, and tunnels them to the care-of address which is registered in the binding.

For example, as soon as one of the nodes inside AS4 communicates with the mobile network the mobile router registers its binding to the correspondent router in AS4. After the registration, any packets meant for the mobile network from AS4 are always intercepted by the correspondent router and tunneled to the mobile router. On the return path, the mobile router could tunnel packets that are sent to AS4 to the correspondent router by IP-in-IP encapsulation [2].

All correspondent routers advertise a proxy route of the mobile prefix to capture packets destined to the mobile network by routing protocols regardless of IGP or EGP. The proxy route may not be inter-exchanged by correspondent routers with any Exterior Gateway Protocol (EGP) such as BGP. The correspondent route advertises the proxy route only while the received binding is valid. After the binding expiration, the correspondent router removes the proxy route from the routing table. Thus, it may lead to frequent changes on BGP routing tables that is not desired on the Internet.

Correspondent routers can intercept packets that are from transit AS. For instance, if a correspondent node in AS3 send packets to the mobile network, the packets are routed towards the home agent since there are no correspondent routers in AS3. However, on the way to the home agent, a correspondent router in AS1 which is the transit AS of AS3, can intercept the packets and tunnels them directly to the mobile router.

The proxy route is not a binding, but it contains the mobile prefix as a destination and the correspondent router's address as the next hop. The proxy route will not be aggregated in correspondent router's IGP domain. The correspondent router can reject receiving a binding of any mobile network according to administrative policies, because the advertisement of unaggregatable routes may swell routing entries on routers. According to routing management policies of each AS, correspondent routers should be approved to provide services for mobile router from their affiliated IGP domain.

B. Correspondent Router Hierarchy

Figure 2 shows the case where the mobile router selects the correspondent router that has shorter prefix. In this case, the correspondent router advertises the proxy route(PR) for the mobile router to one router nearby. Since two routers are configured as border routers, packets sent to the mobile router are routed to one of routers according to the default route of other routers. Once the router having the proxy route intercepts the packets, it re-routes the packets to the correspondent router. Finally the correspondent router tunnels the packets to the mobile router.

Figure 3 shows the case where the mobile router selects the correspondent router configured at the leaf network. The correspondent router advertises the proxy route for the mobile router to other routers in the same network domain. Otherwise, packets are silently routed to the Internet without interception of the correspondent router. Packets sent to the mobile router are routed to one of routers according to the proxy route and then routed to the

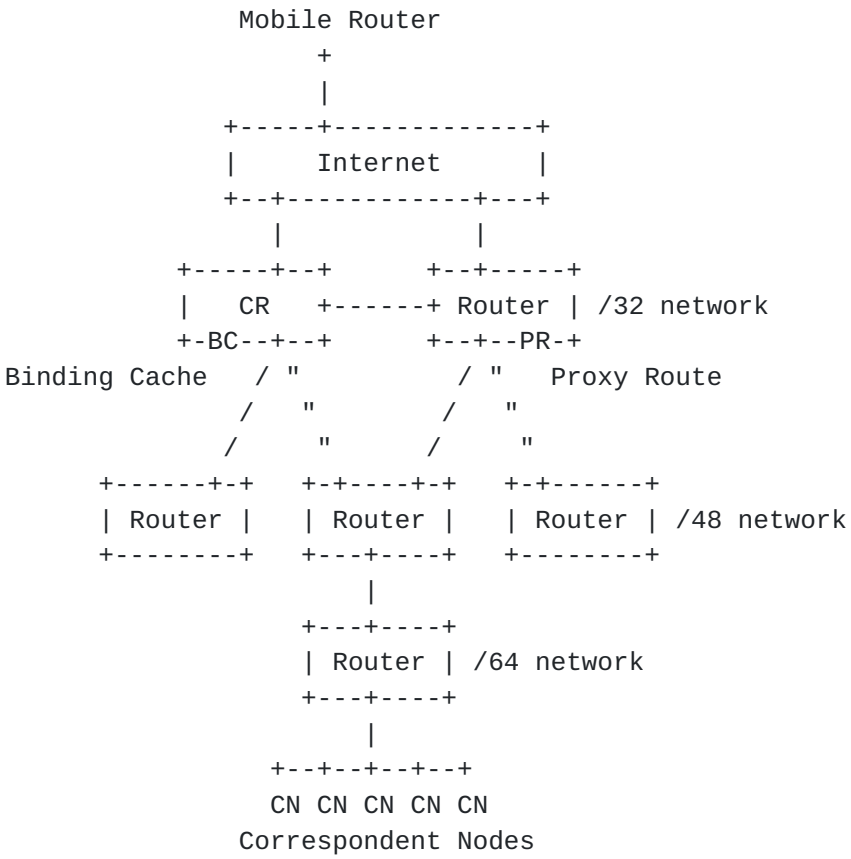

Figure 2: Registration to Higher Router

correspondent router. The correspondent router tunnels these packets to the mobile router. As a result, the route may become longer than the route between the higher route and the mobile router according to the tunnel end point.

It is better to activate a correspondent router located higher in the network hierarchy in terms of proxy route advertisements and shorter bi-directional tunnel between a correspondent router and a mobile node. However, corruption of higher router causes the network separation from the Internet. Thus, higher routers administratively prohibits the correspondent router support.

Increasing the number of correspondent routers caring the mobile network is an important factor to optimize routes between a mobile node and correspondent nodes as much as possible. By contrast, the optimized route cache protocol does not always force to have a number of correspondent routers. Binding registrations to all the correspondent routers bring considerable overheads to a mobile router and prevents scalability and quickness of movement processing.

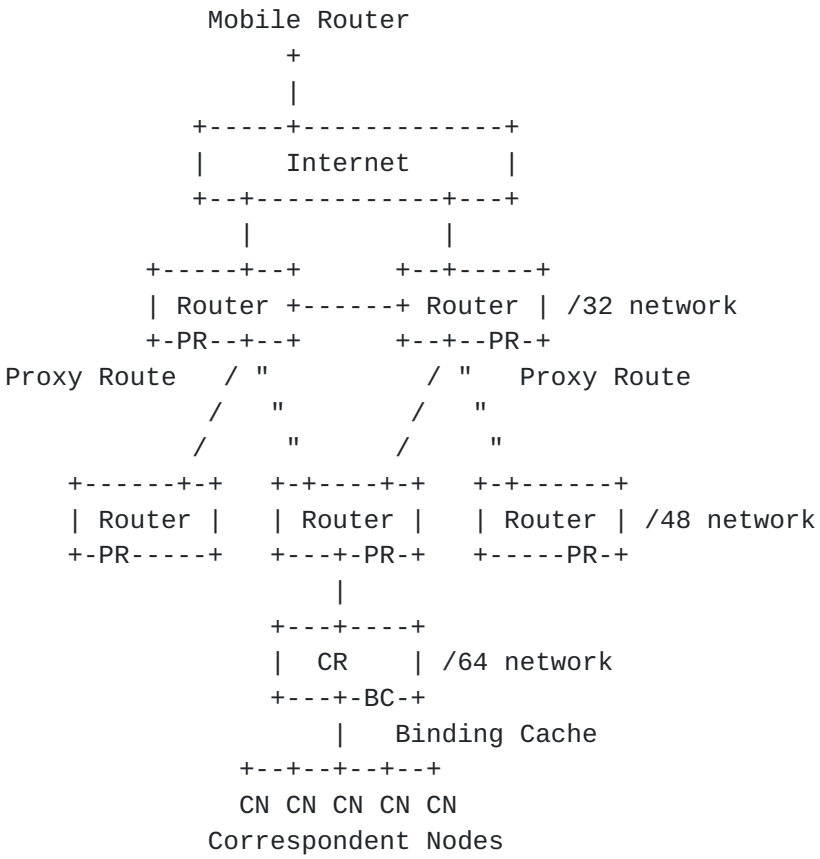

Figure 3: Registration to Lower Router

C. Modifications from the last version

- remove nested mobile networks support
- fix a few typo and packet formats

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgement

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