Seamoby Working Group
INTERNET DRAFT
14 March 2003

Dirk Trossen
Govind Krishnamurthi
Hemant Chaskar
Nokia Research Center

Robert C. Chalmers
UC Santa Barbara

Eunsoo Shim
NEC Labs America

A Dynamic Protocol for Candidate Access-Router Discovery

draft-trossen-seamoby-dycard-01.txt

Status of This Memo

Abstract

Many protocols currently being developed for seamless IP-level handovers, such as fast handovers and context transfers, possess an inherent requirement that the mobile node and/or it's current access router have a priori knowledge concerning the target of the handover. In particular, the target access router (TAR) should be known to have

the appropriate set of capabilities necessary to meet the requirements of the mobile node. Although the TAR selection process occurs at or

near the time of handover, applicable candidate access routers (CARs) can be identified in advance. In this draft, we propose a dynamic, distributed protocol, dyCARD, which allows geographically adjacent routers to exchange capability information, thus providing critical information to the target selection process.

Contents

1. Introduction

Mobile IP [8, 4] enables a mobile node (MN) to execute IP-level handovers between access routers (ARs) which act as points of attachment to the IP network. However, for many scenarios, the handover latency and packet loss incurred by standard Mobile IP are too high. Thus, work is underway to develop protocols intended to provide seamless handovers (low latency and low packet loss) between ARs [3, 6, 5, 7]. Many of these seamless protocols, however, make the assumption that the MN and/or the current AR have a priori knowledge of the target of the handover, the next access router. In order to provide this information to these seamless solutions, a new protocol is required to discover geographically adjacent routers, and to collect their capabilities prior to MN handover.

This document presents such a protocol, dyCARD, a dynamic protocol for candidate access-router discovery. The protocol serves three key functions:

1) To provide a reverse mapping from AP layer-2 identifiers to IP addresses of supporting ARs.

2) To identify physically neighboring access routers sufficiently in advance of MN handover such that AR capabilities may be exchanged.

3) To use these collected capabilities in addition to information provided by the MN, such as reachability and preferences, to aid the MN in selecting a target access router at or near the time of handover.

Three general approaches to CAR discovery exist, namely anticipated discovery, dynamic discovery and hybrid approaches. In the anticipated approach, the current AR identifies the CARs prior to handover, and subsequently selects the TAR on the behalf of the MN. Dynamic CAR selection is controlled by the MN who collects capabilities directly from reachable ARs and then performs target selection based on local policy. Hybrid approaches exists between anticipated and dynamic solutions, where both the AR and MN contribute to CAR discovery and TAR selection.

In this draft, we present a hybrid approach for CAR discovery. In the protocol described herein, ARs discover neighboring CARs through the

handover patterns of the mobile nodes. As a MN hands-over between two adjacent ARs, the ARs learn of one another's existence, and then proceed to exchange capability information off-line over the wired backbone. At some point prior to the next handover, the MN queries its current AR with a list of reachable access points. Using the previously collected

CAR capabilities and the MN's reachability information, the AR can then assist the MN in selecting a new target access router.

2. Terminology

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 RFC 2119[1].

Moreover, the following specific terms are used throughout the document:

Home Address (HoA)
: The IPv4 or IPv6 address of the MN that it uses when it is not moving, i.e., when it is attached to its home network.

Care-of Address (CoA)
: An IPv4 or IPv6 address configured by or assigned to the MN for use on a foreign network.

Access Router (AR)
: A layer-3 device acting as the first-hop IP router for a mobile node. In Mobile IPv4, this would be the Foreign Agent.

Access Point (AP)
: A layer-2 device to which a mobile node connects through a wireless link. A single AR may have many APs of differing technologies.

Beacon
: A layer-2 message emitted by an AP, and received by a mobile node used to announce the presence of the AP. The beacon should contain some unique identifier to allow the mobile node to distinguish between APs

Capabilities
: Information describing the services provided by a given AR.

Profile
: The requirements and preferences of a MN as they pertain to handover.

Candidate AR (CAR)
: An access router which is a possible target for a mobile node's handover. A CAR is a geographical neighbor of the mobile node's current AR. The two ARs have APs with overlapping coverage areas.

Target AR (TAR)	A particular CAR chosen as the target of the mobile node's handover.
Border Router	A globally, IP addressable router through which a privately addressed network connects to the Internet.
Disjoint Handover	A handover to a new AP from which the MN could not receive beacons while attached to its previous AP. In other words, the two APs do not have overlapping coverage areas (from the viewpoint of the MN).

3. Protocol Overview

3.1. Conceptual Data Structures

This document uses the conceptual data structures listed in this section to describe the operation of the CAR discovery protocol. A specific implementation does not necessarily need to implement these data structures as long as the external behavior is consistent with that described in this document.

Physical Neighbor Cache	A map of geographically adjacent ARs and their associated APs. Each AR and AP entry have independent lifetimes associated with them.
Mobile Node List	A list of MNs currently supported by the AR as part of the stateful mode of the CAR discovery protocol. Each MN entry has a lifetime associated with it.
Beacon Entry	An entry representing a single, uniquely identifiable AP beacon. Each Beacon Entry has an associated link-layer address or other unique identifier. A locally unique 16-bit Id is assigned to each Beacon Entry by the MN.
Beacon List	A list of AP Beacon Entries maintained by a MN or AR which represents the current reachability of a given MN. When maintained by the MN, each Beacon Entry has an associated lifetime.

3.2. Discovering Neighboring Routers

In order for an AR to be considered as a candidate for handover, the coverage area of one or more of its attached access points must overlap with the coverage of the MN's existing point of attachment. Two ARs with such overlapping coverage areas are considered to be geographically adjacent, or physical neighbors. This designation differs from logically adjacent routers with only a single IP hop separating them. Geographically adjacent routers can be separated by any number of IP hops, and could actually be in completely different domains. How then do geographically adjacent routers discover each other's existence?

One solution to the discovery problem would be to manually configure this list of physical neighbors for each access router. However, as previously described [11], such an approach has serious disadvantages, and in many cases, may be infeasible. For example, two ARs that are geographically adjacent may be under separate administrative control, and thus may not be informed of one another's presence. Even within the same administrative domain, manual configuration demands consistent network planning and maintenance.

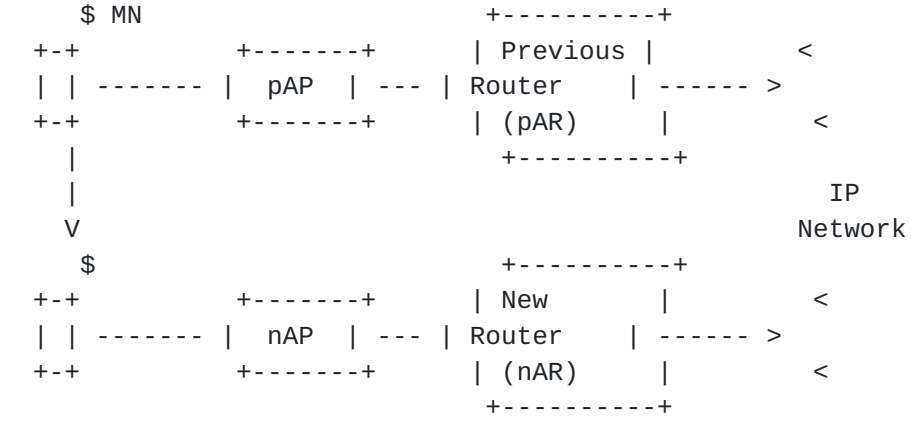

Figure 1: Reference Scenario for Handovers.

Therefore, we propose a dynamic approach where geographically adjacent routers are identified by the handover patterns of the mobile nodes. In short, if a MN can handover between two access points, then the associated ARs should be considered as candidates for future handovers. In this sense, it is crucial that a MN distinguish between handing-over between two adjacent APs, and simply attaching to a new access point after having disconnected from its previous AP.

Following a handover, a MN informs its new access router (nAR), as seen in Figure 1, about its previous point of attachment. In

particular, the MN sends to nAR a Router Identity (RI) message (see Section 4.2) containing the IP address of the previous access router (pAR), the link-layer address of the previous access point, as well as the link-layer address of the new AP. So, for the scenario depicted in Figure 1, the MN would send the tuple (pAR,pAP,nAP) to nAR.

With this information, nAR can create or refresh two entries in its Physical Neighbor Cache (PNC): 1) an entry for nAP as a locally attached access point, and 2) an entry for pAR as a physical neighbor with the associated access point pAP. Prior to trusting the MN's report, however, nAR performs a number of checks to ensure the validity of the received information.

First, nAP must be authorized as one of nAR's local access points. This can be achieved through a managed list of APs currently attached to nAR, or by a larger set of APs that could possibly be attached over a given period of time. The size and variability of this authorization list is controlled by the ARs administrator. A larger set of authorized APs provides a higher level of reconfigurability, but also increases the possibility of malicious MNs polluting the AR's PNC.

To further verify the contents of the RI message, the AR sends a Physical Neighbor Exchange (PNE) message (see Section 4.7) to pAR. The PNE includes both the new and the previous AP identifiers, pAP and nAP. pAR verifies that pAP is indeed valid via its own local PNC entries, and that the MN was recently present. pAR replies to nAR, indicating the result of the validation. If the report was valid, pAR updates its own PNC with an entry for nAR and nAP.

In this way, each handover of a MN results in a bi-directional discovery process between the two participating ARs. Further handovers simply refresh existing PNC entries. If handovers cease between two particular access routers, the associated references will eventually timeout and be removed from each AR's PNC.

The dynamic nature of the CAR discovery protocol we propose adapts well to changes in the network topology. Since the PNC is maintained as soft-state in the ARs, APs that are removed will timeout. ARs that no longer overlap, will dissociate, and new APs will be discovered once a MN performs a handover to it. The protocol design provides for a tradeoff between the stability of the PNC and the degree of reconfigurability. Longer lifetimes for cache entries provide a more stable picture of the neighborhood in the face of few handovers. On the other hand, lower lifetimes improve robustness, allowing APs to be removed or moved between ARs more frequently.

3.3. Exchanging Capabilities

In order to select a particular target from a list of CARs, it is necessary to learn which services each CAR can provide to the MN. This set of services, or capabilities, must be exchanged prior to actual handover to ensure that the information is available at the time of TAR selection.

During a validation exchange, as described in the previous section, two ARs may exchange capability information via the PNE message. Capabilities have lifetimes that are independent of, although bounded by, the PNC entries of their associated ARs. They may be refreshed during normal validation exchanges, or via explicit PNE messages. To request current capabilities, an AR simply sets the C-bit in the PNE message. The receiving AR may then return its capability set as a sub-option in a PNE reply (see Section 4.8.7).

The exact semantics of how to adequately describe capabilities are beyond the scope of this document. Actually, the CAR discovery protocol described herein works with literally any capability representation required. The Capabilities sub-option treats the actual data as a binary block. Marshalling, parsing and management of capability information is handled outside the CAR discovery protocol.

3.4. Reporting Reachability Information

As part of the TAR selection process, the current AR can use the set of APs reachable by a MN to limit the total number of CARs considered. This has two key advantages: 1) the complexity of the TAR selection algorithm is reduced when fewer CARs are considered, and 2) less information can be transmitted to the MN in the case that it controls the TAR selection process.

Reachability information can be transmitted to the AR in two ways: 1) statelessly via the Candidate List Request (CLR) message (see Section 4.5), or 2) statefully via Reachability (RM) message (see Section 4.3). The former provides a transaction-based interface where the mobile node provides the list of reachable APs with each request for resolution. The latter allows the mobile node to inform the AR about reachable APs just once. Then, subsequent resolution attempts can take advantage of the reachability information already available at the AR. We discuss the stateful mode of the protocol in more detail in Section 3.6.

3.5. TAR Selection

At any time prior to handover, a MN may request a list of CARs from the current AR via the CLR message. The AR should use the MN's current reachability state, whether included in the current message or maintained statefully, to limit the set of CARs returned. Moreover, if the AR has the MN's profile available, it could further reduce and sort the set of CARs based on the best match between the MN's requirements and the CAR capabilities. The extent of this reduction could be a property of the profile itself.

Once a set of CARs has been chosen, the AR transmits this list via the Candidate List (CL) message (see Section 4.6) to the MN in the form of the AP identifiers known to the MN. In this way, the list truly consists of candidate access points. If requested by the MN, via the C-bit in the CLR, the AR may also add capability information relevant to the TAR selection process for each of the CARs.

TAR selection is inherently dependent upon the semantics of the MN's profile and the AR capabilities. Again, this is beyond the scope of the CAR discovery protocol, and can be handled by a companion TAR-selection algorithm. The protocol described in this document provides the mechanism to collect and transport the appropriate information to the point of TAR selection, which could be at the current AR or the MN, but plays no part in using that information to select a target AR.

3.6. Stateful Mode

In the stateful mode of the protocol, the AR maintains state on behalf of the MN that it then uses as part of the CAR/TAR selection process. This is done statefully so that the MN can provide the information well before the time of handover, and may actually request multiple CAR selections based on the state stored at the AR; i.e., the MN may periodically refresh its current TAR selection.

As a MN hears new beacons emitted by nearby APs, it reports them to its current AR via the RM message. Each beacon is reported only once upon initial reception. An acknowledgment is returned by the AR. Both the MN and the AR maintain a Beacon List which represents the MN's current reachability state. At the MN, each beacon has an associated lifetime which is dependent upon the beacon frequency. The lifetime is refreshed with each received beacon. Once a beacon expires, a RM message is sent to revoke the entry in the AR's beacon list.

The MN can set the S-bit in the RM message to initiate a

resynchronization, providing the AR with an absolute set of Beacon Entries. This is used on the initial RM issued after handover to ensure an existing MN entry at the AR will not be reused, an example being

the case where a MN moves away and then quickly back again. Moreover, if the MN detects a synchronization problem due to error conditions in the acknowledgment, it may choose to send an absolute list of current beacons to force resynchronization.

The RM message may also be used to provide extra information about the MN to the AR in the form of sub-options, such as the MN's home address or current profile. Since RM messages can be acknowledged, the MN can ensure reception, or subsequently retransmit on failure.

3.7. Security Assumptions

In the design of this protocol, we have made a few assumptions about the security model in place between the MN and the AR, and between ARs. In particular, we assume that prior to any protocol messaging, the the AR has authenticated and authorized the MN to participate in CAR discovery. Moreover, in order for two ARs to cooperate without introducing serious security concerns, they must be able to establish a security association. For intra-domain routers, this could be as simple as a shared secret key. For the inter-domain scenario, the two domains must have a previously established relationship that can be leveraged to derive an adequate session key (e.g., AAA). All messages listed herein should be protected by IPSec or TLS to provide authentication and ensure message integrity.

4. Protocol Messages

The CAR discovery protocol described in this document defines five messages and eleven sub-options. Messages are formed as options so that they can be piggy-backed on other handover messages, if appropriate. All messages exchanged between the MN and AR are presented as ICMP packets [9, 2]. Messages between ARs are presented as payload data of SCTP packets [10]. All ICMP messages share a single ICMP type and code (TBA). Likewise, all options share one ICMP option type (TBA) with distinct sub-types. All SCTP messages share a single Payload Protocol Identifier (TBA).

4.1. ICMP Container

Protocol messages between the AR and MN are sent using ICMP [9, 2]. Each ICMP packet is formed with a general message format.

IP Fields:

Source Address The current CoA of the MN (see Section 6.2).

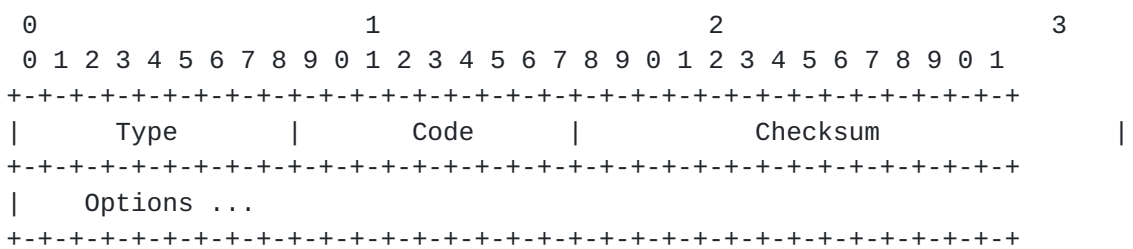

Figure 2: ICMP Message Format.

Destination Address
: The address of the access router as determined from the last router advertisement of said AR.

Hop Limit/TTL
: 1

Authentication Header
: If a security association exists between the MN and the AR, then the sender SHOULD include this header.

ICMP Fields:

Type
: TBA

Code
: TBA

Checksum
: The ICMP checksum.

Options
: Each packet SHOULD contain a single message option defined below.

4.2. Router Identity Message

The Router Identity (RI) message is sent by the MN to its current AR after handover. The message SHOULD be sent only after any necessary authentication and authorization has been performed and a new CoA has been configured, if necessary. This message has an alignment requirement of 4n.

This message MAY contain an IP sub-option with the global address of the previous AR. It MAY also include two Link-Layer sub-options containing the previous and new AP identifiers, if available. If the

handover is disjoint then the previous AR and previous AP sub-options

MUST NOT be included. If neither AP identifier is available, then the MN SHOULD NOT send this message.

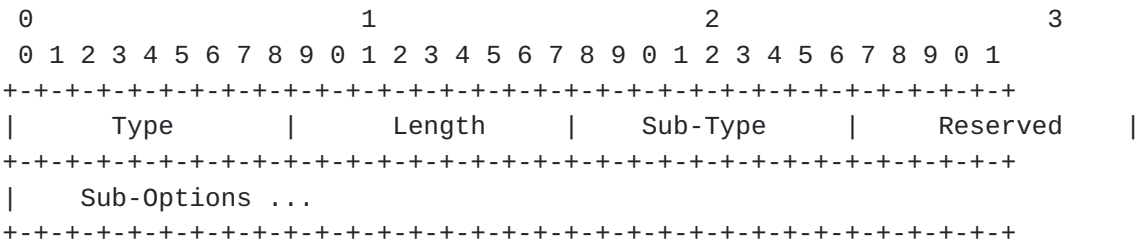

Figure 3: Router Identity Message Format.

RI Fields:

Type	TBA
Length	The size of this message including all sub-options in units of 8 octets. If the message does not completely fill the final 8 octets, then padding MUST be added to the end of the message.
Sub-Type	0
Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Sub-Options	All sub-options pertaining to this message.

4.3. Reachability Message

The Reachability (RM) message is sent by the MN in stateful mode to its current AR in order to inform the AR about which APs are reachable by the MN. This message is also used to send other information about the MN to the AR, such as the MN's HoA and its current profile, as well as to periodically refresh the lifetime associated with the MN state maintained by the AR. This message has an alignment requirement of 4n.

RM Fields:

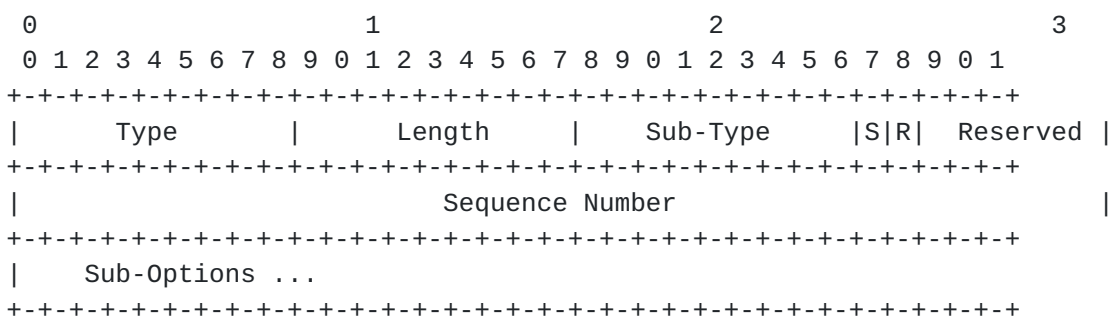

Figure 4: Reachability Message Format.

Length	The size of this message including all sub-options in units of 8 octets. If the message does not completely fill the final 8 octets, then padding MUST be added to the end of the message.
Sub-Type	1
Synchronize (S)	The Synchronize bit is set by the sending node to indicate that the receiving AR should remove any previous beacon state associated with the sending MN, treating the contents of the RM message as an absolute set of current reachability information.
Revoke (R)	The Revoke bit is set by the sender to indicate that the beacons contained within the message should be removed from the receiving AR's current Beacon List associated with the sending MN.
Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Sequence Number	The Sequence Number is assigned by the sender and used by the receiving node to sequence RM messages. Each RM message sent by a MN MUST use a sequence number greater than the Sequence Number value sent in any previous Reachability message, if any, to the same AR (modulo 2**32).

Sub-Options	All sub-options pertaining to this message.

4.4. Reachability Acknowledgement Message

The Reachability Acknowledgement (RMA) message is sent from the current AR to the MN upon reception of a RM message. This message has an alignment requirement of 4n.

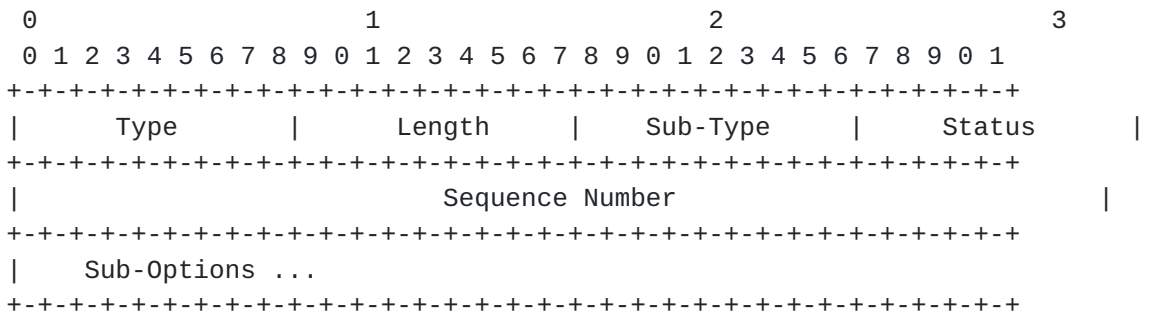

Figure 5: Reachability Acknowledgement Message Format.

RMA Fields:

Type — TBA

Length — The size of this message including all sub-options in units of 8 octets. If the message does not completely fill the final 8 octets, then padding MUST be added to the end of the message.

Sub-Type — 2

Status — An 8-bit unsigned integer indicating the status of handling the originating RM message. Values of the Status field less than 128 indicate success. The following such values are defined:

0 Success

Values of the Status field greater than or equal to 128 indicate failure to process part or all of the RM message. The following such Status values are defined:

128 Bad sequence number
129 Synchronization problem
192 MN is unauthorized

193 Insufficient resources
194 Administratively prohibited

Sequence Number — The Sequence Number in the RMA message is copied from the Sequence Number field in the RM message being acknowledged.

Sub-Options — All sub-options pertaining to this message.

4.5. Candidate List Request Message

The Candidate List Request (CLR) message is sent by the MN to its current AR to request a list of CARs or a single TAR. This message MAY contain Link-Layer and Profile options for immediate use in CAR selection. This message has an alignment requirement of 4n.

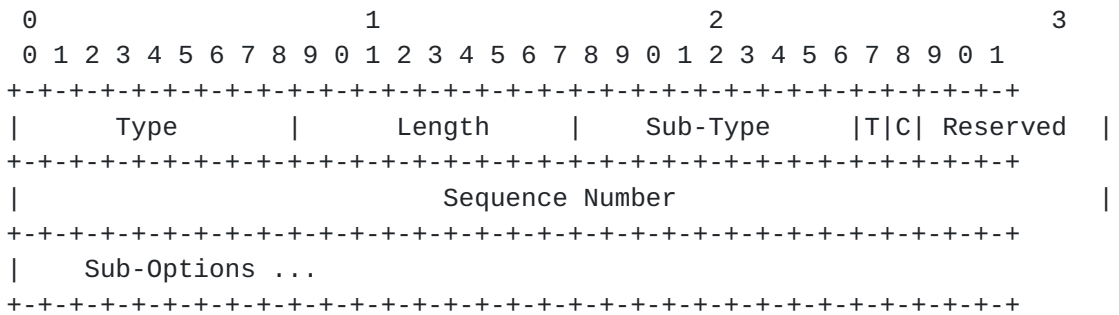

Figure 6: Candidate List Request Message Format.

CLR Fields:

Type — TBA

Length — The size of this message including all sub-options in units of 8 octets. If the message does not completely fill the final 8 octets, then padding MUST be added to the end of the message.

Sub-Type — 3

TAR Selection (T) — The TAR Selection bit is set by the sending node to request that the receiving AR attempt to perform TAR selection on behalf of the sending MN.

Capabilities (C)	The Capabilities bit is set by the sending node to request that the receiving AR include capabilities for the CARs returned in the associated CL message.
Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Sequence Number	The Sequence Number is assigned by the sender and used by the receiving node to sequence CLR messages. Each CLR message sent by a MN MUST use a sequence number greater than the Sequence Number value sent in any previous Candidate List Request message, if any, to the same AR (modulo 2**32).
Sub-Options	All sub-options pertaining to this message.

4.6. Candidate List Message

The Candidate List (CL) message is sent by the current AR to the MN in response to a Candidate List Request. Each CAR should be represented by one or more Link-Layer options indicating the link-layer addresses of the APs associated with this AR. IP, Privacy and Capabilities options may follow each Link-Layer option to provide extended information about each CAR (see Section 4.8). This message has an alignment requirement of 4n.

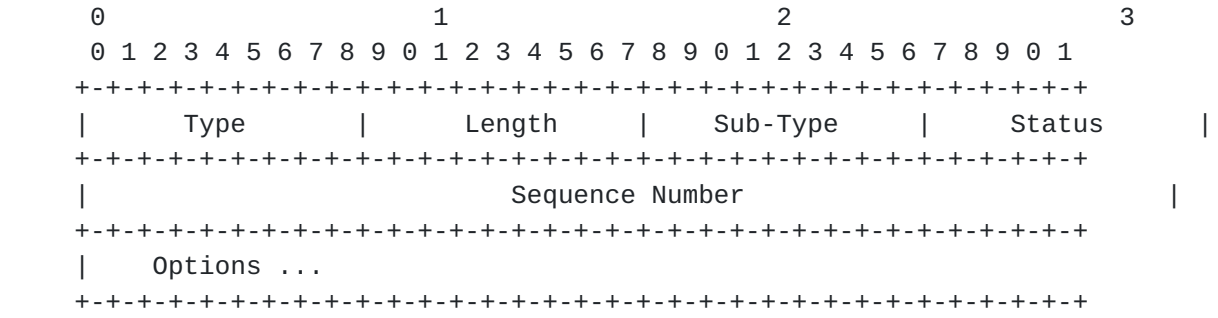

Figure 7: Candidate List Message Format.

CL Fields:

Type	TBA

Length	The size of this message including all sub-options in units of 8 octets. If the message does not completely fill the final 8 octets, then padding MUST be added to the end of the message.
Sub-Type	4
Status	An 8-bit unsigned integer indicating the status of handling the originating CLR message. Values of the Status field less than 128 indicate success. The following such values are defined: 0 Success Values of the Status field greater than or equal to 128 indicate failure to process part or all of the CLR message. The following such Status values are defined: 128 Bad sequence number 129 MN is not currently supported 130 CAR/TAR selection failed 192 MN is unauthorized 193 Insufficient resources 194 Administratively prohibited
Sequence Number	The Sequence Number in the CL message is copied from the Sequence Number field in the associated Candidate List Request message.
Sub-Options	All sub-options pertaining to this message.

4.7. Physical Neighbor Exchange Message

The Physical Neighbor Exchange (PNE) message is sent from one AR to another in order to propagate topology information generated by RI messages. PNE messages are also used to exchange capabilities between ARs, as well as to periodically refresh those capabilities.

A single message is used for both requests and replies since a single exchange may actually consist of more than two messages (a verification,

and a two-way capability exchange). To indicate the start of an exchange, the initiating node MUST set the S-bit in the first message. Subsequent messages carry the Identifier from the initial message to indicate that they are part of the same exchange. This message has an alignment requirement of 4n.

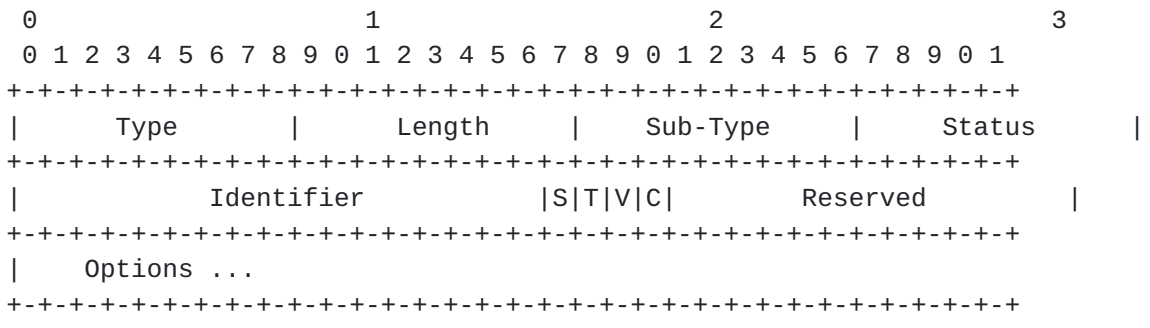

Figure 8: Physical Neighbor Exchange Message Format.

PNE Fields:

Type	TBA
Length	The size of this message including all sub-options in units of 8 octets. If the message does not completely fill the final 8 octets, then padding MUST be added to the end of the message.
Sub-Type	5
Status	An 8-bit unsigned integer indicating the status of handling the originating PNE message. Values of the Status field less than 128 indicate success. The following such values are defined: 0 Success 1 Verified Values of the Status field greater than or equal to 128 indicate failure to process part or all of the RM message. The following such Status values are defined: 128 Local AP not verified 129 MN not verified
Identifier	A unique identifier used to delineate independent message exchanges. The Identifier

field should be set by the initiator of a message exchange with a value not previously used when communicating with the neighboring AR. All subsequent messages of that exchange

MUST copy the Identifier field from the received message into any reply.

Start (S)	The Start bit is set by the sending node to indicate the start of a new message exchange.
Topology (T)	The Topology bit is set by the sending node to indicate that this message contains topology related options as part of the CAR discovery process.
Verify (V)	The Verify bit is set by the sending node to indicate that a verification of the associated topology information should be performed by the receiving AR, and a reply be returned with the status of said verification. The V-bit MUST NOT be set if the T-bit is unset.
Capabilities (C)	The Capabilities bit is set by the sending node to request that the receiving AR provide its own capabilities as an option in the reply.
Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Sub-Options	All sub-options pertaining to this message.

4.8. Message Sub-Options

The CAR discovery protocol described in this document defines eleven new sub-options used in conjunction with the messages defined in the previous subsections.

Each sub-option, other than padding, may have an Entity identifier assigned to distinguish between multiple options of the same type within the same message. The Entity associates the option with a particular participant or target of the protocol. The following entities are defined:

0 None - no entity defined

1 Source - the source of the message

2 Destination - the destination of the message

3 MN - the identity of the mobile node

4 Previous CoA - the previous CoA of the MN

5 New CoA - the new/current CoA of the MN

6 Previous AR - the previous access router

7 New AR - the new/current access router

8 Previous AP - the previous access point

9 New AP - the new/current access point

10 Link-Layer - an entity identified by its link-layer address

Moreover, the order of sub-options in the message are important. All sub-options not associated with a Link-Layer entity MUST be placed in the message prior to the start of any Link-Layer sub-options. Any sub-option following a Link-Layer sub-option is automatically associated with the immediately preceding Link-Layer sub-option; the Entity field MUST be set to 10, Link-Layer.

New sub-options may be added to the protocol in the future. If a receiving node does not recognize a sub-option, it SHOULD silently ignore the sub-option and continue processing the message.

4.8.1. Pad1 Sub-Option

The Pad1 sub-option is used to insert one octet of padding into the Sub-Options area of a message. If more than one octet of padding is required, the PadN option, described next, should be used, rather than multiple Pad1 options. This sub-option has no alignment requirements.

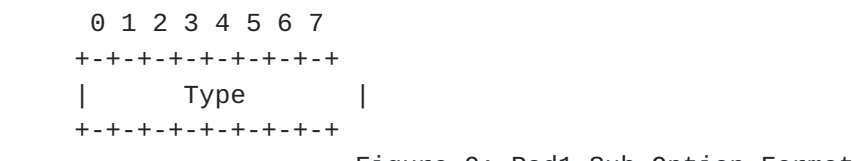

Figure 9: Pad1 Sub-Option Format

Pad1 Sub-Options Fields:

Type 0

4.8.2. PadN Sub-Option

The PadN sub-option is used to insert two or more octets of padding into the Sub-Options area of a message. For N octets of padding, the Opt Data Len field contains the value N, and the Padding consists of N-2 zero-valued octets. This sub-option has no alignment requirements.

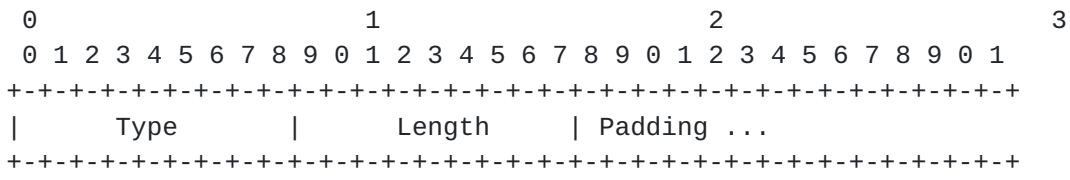

Figure 10: PadN Sub-Option Format

PadN Sub-Option Fields:

Type	1
Length	The length of the sub-option in octets including the Type and Length fields.
Padding	N-2 octets of padding for a sub-option Length of N.

4.8.3. Lifetime Sub-Option

The Lifetime sub-option is added to RM, RMA and PNE messages in order to request or grant a lifetime on the state maintained with regard to the message's peer. This sub-option has an alignment requirement of 4n.

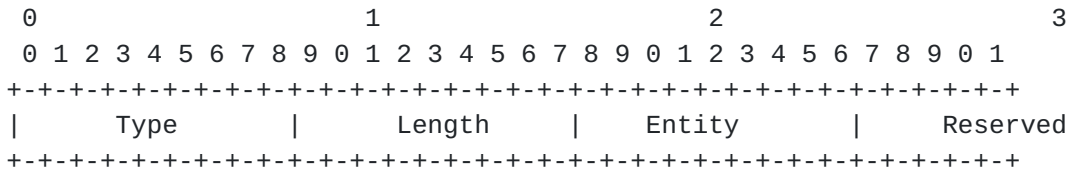

Figure 11: Lifetime Sub-Option Format

Lifetime Sub-Option Fields:

Type	2
Length	1
Entity	An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.
Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Lifetime	The lifetime in seconds associated with Entity.

4.8.4. IP Sub-Option

The IP sub-option is added to a message to indicate the global IP address of a given Entity. For example, the IP address of the associated AR can accompany the Link-Layer sub-option representing a candidate target AR in a CL message returned by the AR. An AR may also add an IP sub-option to a PNE exchange to provide a single, global address to be used by a peer (see Section 7.7). The sub-option has an alignment requirement of 4n for IPv4 addresses and 8n for IPv6 addresses.

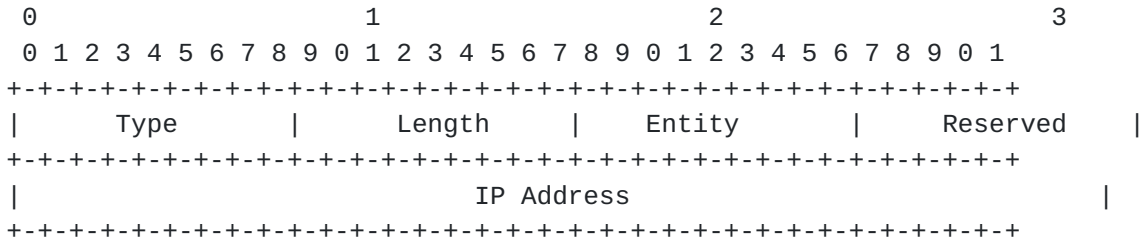

Figure 12: IPv4 Sub-Option Format

IP Sub-Option Fields:

Type	3 or 4. For an IPv4 sub-option the code is 3. The code is 4 for an IPv6 sub-option.
Length	The size of this sub-option in units of 8 octets.

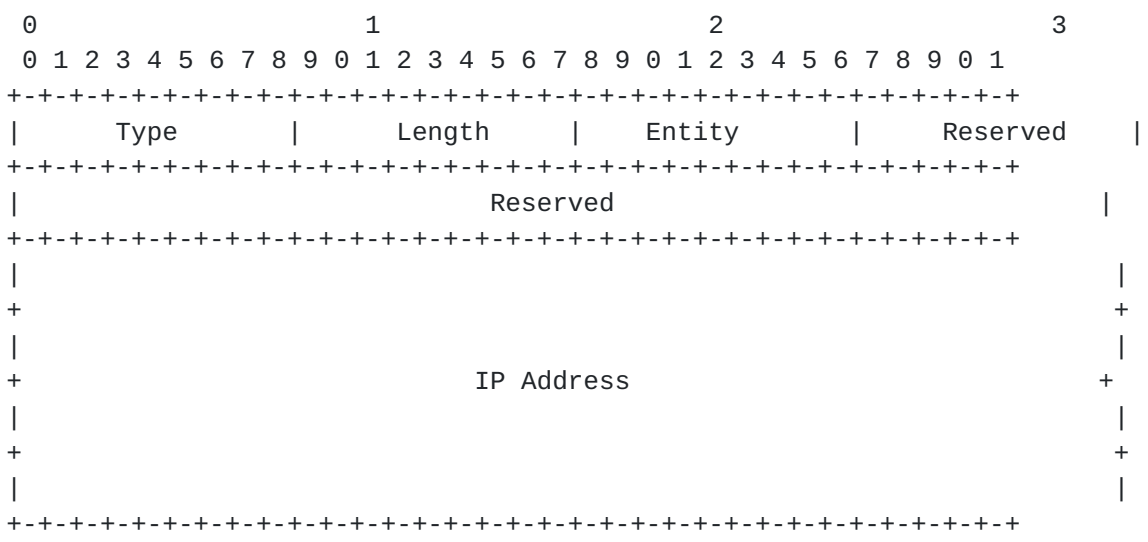

Figure 13: IPv6 Sub-Option Format

Entity — An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.

Reserved — This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

IP Address — The IP address of the target Entity.

4.8.5. Privacy Sub-Option

The Privacy sub-option is added to a message in order to associate a locally unique identifier with an Entity when the global IP address provided for that entity is not unique to that device; e.g. an AR with a private address which uses a global address from a Border Router. This sub-option has an alignment requirement of 4n.

Privacy Sub-Option Fields:

Type — 5

Length — 1

Entity — An Entity value defined in Section 4.8 indicating to which entity this sub-option

applies.

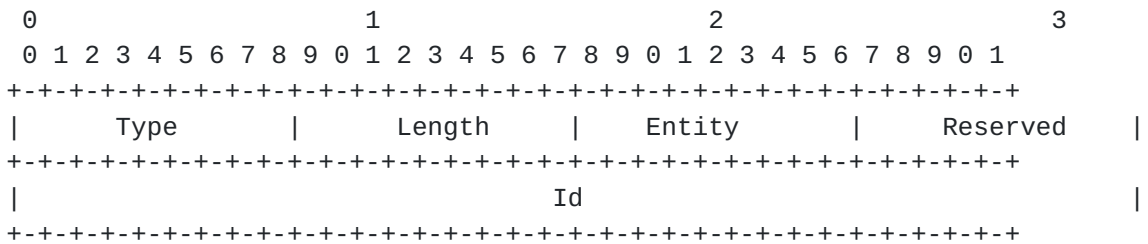

Figure 14: Privacy Sub-Option Format

Reserved — This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Id — A 32-bit unsigned identifier used to identify the target Entity when multiple nodes may share a single global IP address.

4.8.6. Profile Sub-Option

A Profile sub-option is added to RM messages in order to transmit a MN's current profile to an AR. The profile is treated as a binary block of data. It is not assumed that the entire profile must be sent within each sub-option. Rather portions of or updates to the profile could be sent independently. How partial information is managed by the receiver, however, is outside the scope of this document. This sub-option has an alignment requirement of 4n.

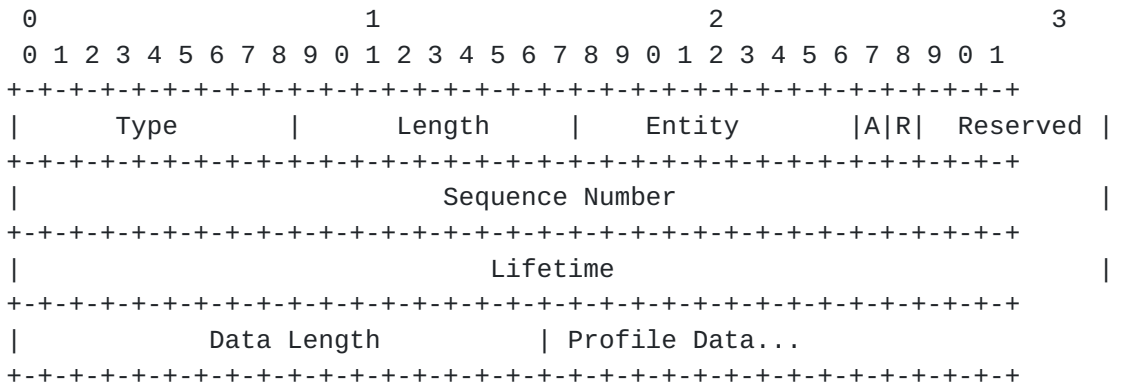

Figure 15: Profile Sub-Option Format

Profile Sub-Option Fields:

Type	6
Length	The size of this sub-option in units of 8 octets.
Entity	An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.
Absolute (A)	The Absolute bit is set by the sending node to indicate that the profile data included in the sub-option MUST override any existing profile currently cached at the receiving node with regards to the sending MN.
Revoke (R)	The Revoke bit is set by the sending node to indicate that the profile data included in the sub-option MUST be removed from the receiving node's cache for the sending AR.
Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Sequence Number	The Sequence Number is assigned by the sender and used by the receiving node to sequence profile updates. Each Profile sub-option sent by a MN MUST use a sequence number greater than the Sequence Number value sent in any previous Profile sub-option, if any, to the same receiving node (modulo 2**32).
Lifetime	The lifetime in seconds for the profile included in the sub-option.
Data Length	The actual length in bytes of the Profile Data contained within the sub-option. This field may be set to zero if no data is present, e.g., to revoke the complete profile.

Profile Data	The Profile Data associated with this entity. The field MUST be zero-padded to an integral number of 8-octet units. The actual length of the data should be placed in the Data Length field.

4.8.7. Capabilities Sub-Option

A Profile sub-option is added to CL and PNE messages in order to transmit an AR's current capabilities to a MN or another AR. Capabilities are treated as a binary block of data. It is not assumed that the entire set of capabilities must be sent within each sub-option. Rather portions of or updates to the capabilities could be sent independently. How partial information is managed by the receiver, however, is outside the scope of this document. This sub-option has an alignment requirement of 4n.

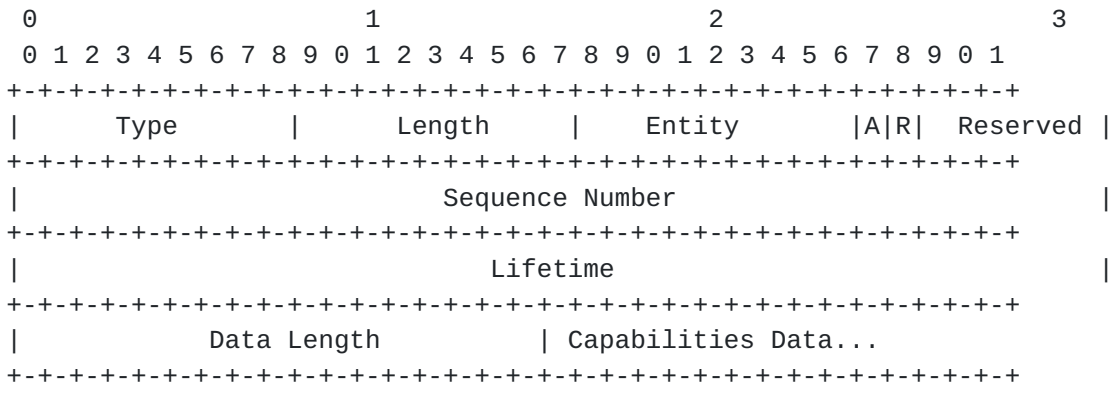

Figure 16: Capabilities Sub-Option Format

Capabilities Sub-Option Fields:

Type	7
Length	The size of this sub-option in units of 8 octets.
Entity	An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.
Absolute (A)	The Absolute bit is set by the sending node to indicate that the capabilities included in the sub-option MUST override any existing capabilities currently cached at the receiving node with regards to the sending AR.
Revoke (R)	The Revoke bit is set by the sending node to indicate that the capabilities included in the

sub-option MUST be removed from the receiving
node's cache for the sending AR.

Reserved	This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.
Sequence Number	The Sequence Number is assigned by the sender and used by the receiving node to sequence capability updates. Each Capabilities sub-option sent by an AR MUST use a sequence number greater than the Sequence Number value sent in any previous Capabilities sub-option, if any, to the same receiving node (modulo 2**32).
Lifetime	The lifetime in seconds for the capabilities included in the sub-option.
Data Length	The actual length in bytes of the Capabilities Data contained within the sub-option. This field may be set to zero if no data is present, i.e., to revoke the complete capabilities set.
Capabilities Data	The Capabilities Data associated with this entity. The field MUST be zero-padded to an integral number of 8-octet units. The actual length of the data should be placed in the Data Length field.

4.8.8. Link-Layer Sub-Option

The Link-Layer sub-option is added to messages to describe particular Entities identified by their link-layer address, namely APs. Each Link-Layer sub-option may contain an identifier that should be sufficient to uniquely determine the Entity with respect to the context of the message. In the case that this identifier is available, the actual link-layer address may be elided to save bandwidth. This sub-option has an alignment requirement of 4n.

Other sub-options may be associated with a Link-Layer sub-option to further describe properties of the Entity. Any sub-options following a Link-Layer sub-option, not including another Link-Layer sub-option, are automatically associated with that Link-Layer sub-option.

Link-Layer Sub-Option Fields:

Type	8
Length	The size of this sub-option in units of 8 octets.

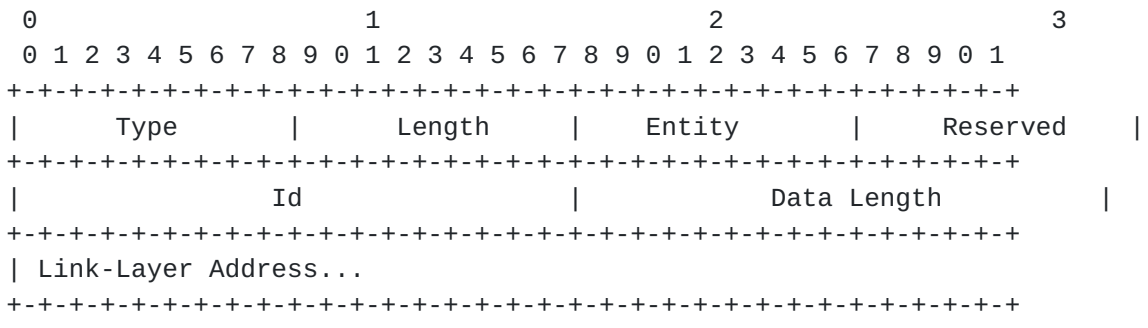

Figure 17: Link-Layer Sub-Option Format

Entity — An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.

Reserved — This field is reserved for future use. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Id — A locally unique identifier assigned by the MN to designate this link-layer address.

Data Length — The actual length in bytes of the Link-Layer Address contained within the sub-option. This field may be set to zero if no address is present; e.g., the Id is sufficient.

Link-Layer Address — The Link-Layer Address or BSSID associated with this entity. The field MUST be zero-padded to an integral number of 8-octet units. The actual length of the address should be placed in the Data Length field.

4.8.9. MN Identifier Sub-Option

The MN Identifier sub-option is used by ARs to identify a given mobile node. The option is included in PNE messages to associate the topology information with the MN that issued the original Router Identity message. Section 7.1 describes the method by which ARs can identify a mobile node. This sub-option has an alignment requirement of 4n.

MN Identifier Sub-Option Fields:

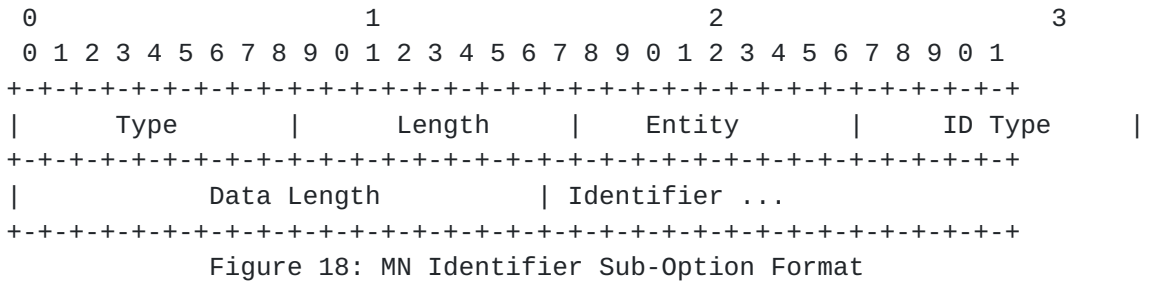

Figure 18: MN Identifier Sub-Option Format

Type	9
Length	The size of this sub-option in units of 8 octets.
Entity	An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.
ID Type	The type of identifier used (TBD).
Data Length	The actual length in bytes of the Identifier contained within the sub-option.
Identifier	The value used to uniquely identify a mobile node. The field MUST be zero-padded to an integral number of 8-octet units. The actual length of the address should be placed in the Data Length field.

4.8.10. MN Token Sub-Option

The MN Token sub-option is used to carry a token generated by an AR that is used to verify the presence of a given mobile node after handover, as described in Section 7.10. The original token is passed to the MN via either the CL message. After handover, the token is included with the RI message and the resulting PNE message. This sub-option has an alignment requirement of 4n.

MN Token Sub-Option Fields:

Type	10

Length	The size of this sub-option in units of 8 octets.

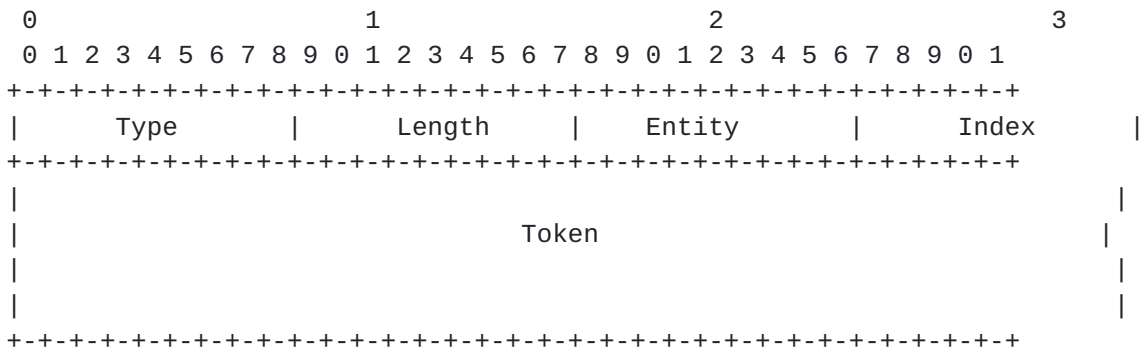

Figure 19: MN Token Sub-Option Format

Entity — An Entity value defined in Section 4.8 indicating to which entity this sub-option applies.

Index — The index of the original key used to generate the token.

Token — A 128-bit token generated by an AR, as described in Section 7.10.

5. Protocol Requirements

The protocol described in this document makes no new requirements on general IPv4 or IPv6 nodes. Only those nodes, MNs and ARs, participating in the protocol must support the features described herein.

5.1. Requirements for Mobile Nodes

A mobile node participating in this CAR discovery protocol MUST fulfill the following requirements:

- Every participating MN MUST support sending Router Identity messages as described in Section 6.3.

- Every participating MN MAY support sending Reachability messages as described in Section 6.5.

- Every participating MN MAY support receiving Reachability Acknowledgement messages as described in Section 6.6.

- Every participating MN SHOULD support sending Candidate List Request messages as described in Section 6.7.

- Every participating MN SHOULD support receiving Candidate List messages as described in Section 6.8.

- Every participating MN MUST support detecting movement at both the link-layer and the IP-layer as described in Section 6.1.

- Every participating MN SHOULD support receiving beacons from nearby APs.

- Every participating MN SHOULD support a TAR resolution protocol able to incorporate profile and capability information.

5.2. Requirements for Access Routers

An access router participating in this CAR discovery protocol MUST fulfill the following requirements:

- Every participating AR MUST support receiving Router Identity messages as described in Section 7.2.

- Every participating AR MUST support receiving Reachability messages as described in Section 7.3.

- Every participating AR MUST support sending Reachability Acknowledgement messages as described in Section 7.4.

- Every participating AR MUST support receiving Candidate List Request messages as described in Section 7.5.

- Every participating AR MUST support sending Candidate List messages as described in Section 7.6.

- Every participating AR MUST support sending and receiving Physical Neighbor Exchange messages as described in Sections 7.7 and 7.8, respectively.

- Every participating AR SHOULD be capable of uniquely identifying

attached MNs as described in Section 7.1.

- Every participating AR SHOULD be capable of verifying that a given MN was attached in the recent past, as described in Section 7.10.

- Every participating AR SHOULD have a single, global IP address used to identify itself with neighboring ARs, as described in Section 7.7.

- Every participating AR SHOULD support a CAR filtering protocol able to incorporate profile and capability information.

6. Mobile Node Operation

6.1. Movement Detection

The actual algorithm used to detect movement is necessarily outside the scope of this document, and dependent upon the mobility protocol employed by the MN. However, the MN MUST be aware of changes in its current attachment point, both at the link and network levels. The MN MUST be able to identify both its current AP, and its current AR. The AR SHOULD be identifiable by an address not of link-local scope. The AP can be identified by any unique identifier which we will refer to as the AP's link-layer address.

Beyond detecting direct handovers between two APs, a MN SHOULD also be capable of detecting when it disconnects from it current AP/AR without handing over. This is to ensure that a MN that moves between two APs with non-overlapping coverage does not incorrectly perceive the two APs as being physically adjacent. Such a ``disjoint'' change in access points could possibly be detected by the link-layer itself, or through more course grained means such as the default router entry timing out.

If the MN has more than one active interface, then the previous and current AP SHOULD be tracked independently for each interface.

6.2. Selecting a Source Address

A MN SHOULD maintain a consistent source address for all messages sent to a given AR. That address MUST uniquely identify the MN at the AR. In the case that the MN employs the address of the AR (e.g., a MIPv4 Foreign Agent) as its own CoA, then the MN SHOULD use its home address as the source address of its messages to the AR.

6.3. Sending Router Identity Messages

Upon detecting the completion of a handover, the MN SHOULD send a RI message to its new AR according to the following rules:

- In the case that the MN cannot identify either its current or previous AP, then the RI SHOULD NOT be sent.

- If the MN can identify its current AP, a Link-Layer sub-option SHOULD be added to the RI with the Entity field set to New AR. The AP's link-layer address and its length MUST be assigned to the Link-Layer Address and Data Length fields, respectively. The Id field SHOULD be set to zero.

- If the handover is not disjoint, the previous and new APs have overlapping coverage areas (see Section 6.1), then the RI MUST be sent according to the additional following rules:

  * An IP sub-option MUST be included for the global IP address of the previous AR. This address can be the one used directly by the MN in communication with said AR, or one supplied by the AR as an IP sub-option in a previous message. If a global IP address is not available, then the handover should be treated as disjoint. The Entity field MUST be set to Previous AR.

  * If the previous AR provided the MN with a Privacy sub-option in a previous RM or CL message, a new Privacy sub-option MUST be added to the RI with the Id field copied from the original sub-option sent by the AR. The Entity field MUST be set to Previous AR.

  * If the previous AR provided the MN with a MN Token sub-option in a previous RM or CL message, a new MN Token sub-option MUST be added to the RI with the Index and Token fields copied from the original sub-option sent by the AR. The Entity field MUST be set to MN.

  * If the MN can identify its previous AP, a Link-Layer sub-option SHOULD be added to the RI with the Entity field set to Previous AP. The AP identifier and its length MUST be assigned to the Link-Layer Address and Data Length fields, respectively. The Id field SHOULD be set to zero.

  * An IP sub-option MAY be included for the previous CoA of the mobile node, the same CoA used by the MN to send RM messages to the previous AR. The Entity field MUST be set to Previous CoA.

6.4. Tracking AP Beacons

Upon handover, a MN MUST clear its current Beacon List, or otherwise mark the existing beacons in the list such that the next beacon

from the same AP will be treated as newly heard. When a new beacon is heard, an entry SHOULD be made in the MN's Beacon List with an appropriate lifetime. The lifetime could be a property of the beacon itself or based on the expected beacon frequency. Upon receiving subsequent beacons from the same AP, the MN SHOULD update the lifetime of associated beacon in its Beacon List. When a beacon entry's lifetime expires, it MUST be removed from the Beacon List, or otherwise marked as no longer active such that the next beacon from that AP will be treated as newly heard. Only active beacons SHOULD be considered when reporting reachability information.

6.5. Sending Reachability Messages

Reachability messages SHOULD NOT be sent unless the MN is currently engaged in the stateful mode of the protocol. If in stateful mode, the MN SHOULD send an RM in any of the following cases:

- When a new beacon is received, one not already represented in the MN's Beacon List. In this case, the link-layer address of the AP (or other unique AP identifier) MUST be included via a Link-Layer sub-option. The locally-unique Id assigned to this beacon by the MN SHOULD be assigned to the Id field of the sub-option.

- When the lifetime of a beacon entry in the Beacon List expires. In this case, the R-bit of the RM message MUST be set. A Link-Layer sub-option MUST be included in the message indicating which beacon. If a locally-unique Id was previously assigned to this beacon and transmitted to the AR, then only the Id SHOULD be included.

- When the lifetime associated with the MN's state at the AR is close to expiration. In this case, a Lifetime sub-option SHOULD be included in the message with the new requested lifetime.

- After receiving a previous RMA message with a Status value indicating that the AR and MN are unsynchronized. In this case, the S-bit of the RM message MUST be set, and a set of Link-Layer sub-options SHOULD be included with the current contents of the MN's Beacon List. The MN MAY choose to limit the number of beacons sent in one message. In this case, beacon entries not sent MUST be deleted or marked in such a way that the next associated beacon received will trigger a new RM message as described in this subsection.

Each RM message sent by the MN MUST abide by the following rules:

- The first RM message sent after handover MUST have the S-bit set to indicate synchronization in case the AR has pre-existing, stale state for the MN.

- If any beacons included in the message are to be revoked, the R-bit is set, then all beacons listed in the message MUST also be marked for revocation. In other words, a single RM message may contain beacons to add to the AR's Beacon List, or beacons to remove from the list, but not both.

- Each RM message MUST have a Sequence Number assigned larger than any previous RM message sent to the same AR, if any.

Any RM sent may carry additional sub-options such as the Profile sub-option. Also, the MN MAY choose to delay a RM message a short time in order to collect a number of beacons in one message.

The MN MAY retransmit a RM message after waiting some period of time for the associated RMA. The retransmitted message MUST have a new Sequence Number assigned larger than that of the original message. Retransmissions MUST be rate limited, and SHOULD NOT be sent more frequently than one per MIN_RM_PERIOD, and should not be repeated more than MIN_RM_RETRY times. The MN SHOULD exponentially backoff the time between each retransmission. If an acknowledgement has not been received after exhausting the total number of allowed retries, a MN SHOULD consider itself unsupported by the current AR, and cease sending further Reachability messages until after handing-over to another AR.

6.6. Receiving Reachability Acknowledgement Messages

Upon receiving a Reachability Acknowledgment message, the MN MUST validate the message according to the following rules:

- The message MUST originate from the MN's current AR.

- The Sequence Number MUST match an outstanding RM message awaiting acknowledgement.

An invalid RMA MUST be silently ignored. The MN MUST act on a valid RMA according to the value of the Status field:

- A Success value indicates that the originating RM was accepted.

The MN SHOULD update its state accordingly and MUST refrain from retransmitting the originating RM message.

- For a Status value indicating that the sequence number was incorrect, the MN SHOULD increment the Sequence Number value used

in the next RM by some reasonable amount. The originating RM MAY be retransmitted with this new Sequence Number.

- For a Status value indicating a synchronization problem, the MN SHOULD transmit a new RM message with the S-bit set as described in Section 6.5.

- For Status values indicating that the MN is not currently supported by the AR due to security, administrative or resource constraints, the MN SHOULD cease sending Reachability messages until after handing-over to another AR.

6.7. Sending Candidate List Request Messages

The MN MAY send a Candidate List Request message to its current AR at any time prior to handover. The message SHOULD be sent as close to the time of handover as possible to ensure that the information provided in the associated CL message is current. The CLR message MUST be sent according to the following rules:

- Each CLR message MUST have a Sequence Number assigned larger than any previous CLR message sent to the same AR, if any.

- The MN MAY set the T-bit of the CLR message to request that the AR return a single TAR rather than a set of CARs. This, however, may not be supported by the current AR.

- The MN MAY set the C-bit of the CLR message to request that the AR include capability information with the set of CARs returned in the CL message.

- The MN MAY include Link-Layer sub-options representing reachable AP's that should be used to select CARs. This list of APs will be used by the AR only to respond to this CLR, superseding any current beacon state maintained by the AR on behalf of the MN. The previously maintained Beacon List, if any, will not be affected.

- The MN MAY include a Profile sub-option to be used by the AR in selecting appropriate CARs. This profile will be used by the AR only to respond to this message, superseding any current profile maintained by the AR on behalf of the MN. The previously maintained profile, if any, will not be affected.

The MN MAY retransmit a CLR message after waiting some period of time for the associated CL message. The retransmitted message MUST have a new Sequence Number assigned larger than that of the original message. Retransmissions MUST be rate limited, and SHOULD NOT be sent

more frequently than one per MIN_CLR_PERIOD, and should not be repeated more than MIN_CLR_RETRY times. If an acknowledgement has not been received after exhausting the total number of allowed retries, a MN SHOULD consider itself unsupported by the current AR and cease sending further protocol messages until after handing-over to another AR.

6.8. Receiving Candidate List Messages

Upon receiving a Candidate List message, the MN MUST validate the message according to the following rules:

- The message MUST originate from the MN's current AR.
- The Sequence Number MUST match an outstanding CLR message awaiting acknowledgement.

An invalid CL message MUST be silently ignored. The MN MUST act on a valid CL according to the value of the Status field:

- A Success value indicates that the originating CLR was accepted, and that at least one CAR or TAR has been selected. The MN SHOULD pass the selected CAR(s) along to the TAR resolution algorithm, if one is available, and MUST refrain from retransmitting the originating CLR message.
- For a Status value indicating that the sequence number was incorrect, the MN SHOULD increment the Sequence Number value used in the next CLR by some reasonable amount. The originating CLR MAY be retransmitted with this new Sequence Number.
- For a Status value indicating that the CAR/TAR selection process failed, MN MAY send new CLR messages in the future, but SHOULD NOT retransmit the originating CLR.
- For Status values indicating that the MN is not currently supported by the AR due to security, administrative or resource constraints, the MN SHOULD cease sending protocol message until after handing-over to another AR.

7. Access Router Operation

7.1. Identifying Mobile Nodes

The first step in verifying the presence of a mobile node is to be able to properly identify it. One option is to use the mobile node's IP address, either its care-of or home address. The problem with this approach is that IPv6 nodes can generate any number of

care-of addresses, and there exists no means for an AR to confirm that the home address provided actually belongs to the mobile node. Instead, the access router identifies the MN via the same credentials originally provided by the mobile node while authenticating with the AR. In cellular systems, this might be the International Mobile Subscriber Identifier (IMSI) from the phone's SIM card. For AAA-based authentication, the user's Network Access Identifier (NAI) would be used. In either case, this identifier will have been validated by the access router as part of the process of authentication, and thus provides a certain level of accountability against malicious activities.

7.2. Receiving Router Identity Messages

A MN sends a Router Identity message to its current AR after handover. The RI message carries the link-layer addresses of the previous and/or new APs, as well as the global address of the previous router. If an IP sub-option for the previous AR (as determined by the Entity field) is not included or set to an address local to the new AR (with the addition of a possible privacy sub-option), then the handover is considered local to the AR.

The AR MUST validate the RI message according to the following rules:

- All RI messages received MUST be rate limited as described in Section 7.9.

- The message MUST contain either a new or previous AP address, provided via Link-Layer sub-options with the Entity field set to New AR or Previous AR, respectively.

- If the new AP address is included in the message, the AR MUST verify that the link-layer address provided in the sub-option represents a viable local AP. This check could be made against an actual list of known attached APs, or a list of APs authorized to be considered local.

- If the handover is local to the AR, the AR MUST further validate the message according to the following rules:

  * If a previous AP address is included, the AR MUST verify that the link-layer address provided in the sub-option represents a viable local AP. This check could be made against an actual list of known attached APs, or a list of APs authorized to be considered local.

An invalid RI message MUST be silently ignored. Otherwise, the AR SHOULD add any local APs referenced in the message (see above) to its

Physical Neighbor Cache as locally attached, if not already. The AR SHOULD then update the lifetime of these AP entries to AP_LIFETIME.

If the handover is not local, the AR SHOULD send a PNE to the previous AR indicated in the RI message according to the rules outlined in Section 7.7.

7.3. Receiving Reachability Messages

Upon receiving a Reachability message from a MN, the AR SHOULD create a new MN entry in its list of supported MNs, if not already. The AR MAY choose not to create a new entry based on administrative or resource constraints. If the AR chooses not to support the MN, it MUST return a RMA with the appropriate error value in the Status field. All RM messages received MUST be rate limited as described in Section 7.9. The rest of this subsection assumes that the AR has created the MN entry and will support the MN for stateful CAR discovery.

Upon receiving a RM message, the AR MUST verify that the Sequence Number is greater than any received from the MN in prior RM messages. This requirement only covers the current ``session'' with the MN. Once a MN hands-over to another AR, the MN state MAY be deleted, and future sequence numbers do not have to be checked against those prior to deletion. If the sequence number check fails, the AR MUST return a RMA message with the appropriate error value in the Status field.

For valid RM messages, the AR MUST return a Reachability Acknowledgment message, and SHOULD update the MN's state according to the following rules:

- If the S-bit is set, the AR MUST delete all existing Beacon Entries from the MN's Beacon List.

- If the R-bit is set, the AR MUST delete all Beacon Entries matching those listed as Link-Layer sub-options. A Beacon Entry MAY be identified by either the Id field or the actual Link-Layer Address if provided. If any Link-Layer sub-option is not matched to an existing Beacon Entry, the AR MUST return a RMA with the Status field set to indicate that the AR has lost synchronization with the MN.

- If the R-bit is unset, the AR SHOULD create or update a Beacon Entry for each Link-Layer sub-option included in the message. If a non-zero Id is provided, the AR MUST save this Id with the

Beacon Entry. If the AR cannot handle (create or update) all of the Link-Layer sub-options included in the message, it MUST return a RMA with the Status field set to indicate that the AR has lost synchronization with the MN.

- If an IP sub-option is included with the Entity field set to MN, the AR MAY save the MN's HoA as part of MN entry.

- If a Lifetime sub-option is included in the message, the AR SHOULD update the remaining lifetime of the associated MN entry by a value no greater than that specified in the Lifetime sub-option. If the specified lifetime is zero, the AR MUST delete the MN entry immediately. The granted lifetime SHOULD be returned in the RMA as a separate Lifetime sub-option.

- If a Profile sub-option is included in the message, the AR SHOULD save the profile data with the associated MN entry for use in CAR/TAR selection.

7.4. Sending Reachability Acknowledgement Messages

The AR MUST send a Reachability Acknowledgment Message to a MN in response to a RM message. A RMA sent by the AR MUST abide by the following rules:

- The Sequence Number MUST be copied from the Sequence Number field of the originating RM.

- The Status field MUST be set to an error value if part or all of the RM message could not be processed. A non-error value SHOULD be used otherwise.

- If the originating RM message contained a Lifetime sub-option, the AR SHOULD include a Lifetime sub-option with the granted lifetime for the MN entry.

- The AR MAY include a MN Token sub-option generated according to Section 7.10.

- The AR MAY also include an IP sub-option with the Entity field set to Source and the IP Address field set to a global IP address that the MN should use when referring to this AR in subsequent RI messages. Additionally, the AR MAY include a Privacy sub-option with a 32-bit identifier to be used by a neighboring AR in subsequent communications.

7.5. Receiving Candidate List Request Messages

If an AR receives a Candidate List Request message without Link-Layer sub-options for a MN for which it has no corresponding MN state, it SHOULD return a CL message indicating that the MN is not currently supported. The AR MAY choose not to fulfill the request due to

administrative or resource constraints. If the AR chooses not to fulfill the CLR, it MUST return a CL message with the appropriate error value in the Status field. All CLR messages received MUST be rate limited as described in Section 7.9. The rest of this subsection assumes that the AR will accept the CLR message.

Upon receiving a CLR message, the AR MUST verify that the Sequence Number is greater than any received from the MN in prior CLR messages. This requirement only covers the current ``session'' with the MN. Once a MN hands-over to another AR, the MN state MAY be deleted, and future sequence numbers do not have to be checked against those prior to deletion. If the sequence number check fails, the AR MUST return a CLR message with the appropriate error value in the Status field.

For a valid CLR message, the AR MUST process the message according to the following rules:

- If Link-Layer sub-options are included in the message, the AR MUST use the associated link-layer addresses as the list of reachable APs for use in CAR/TAR selection. Otherwise, the AR MUST use the Beacon Entries from the Beacon List associated with the MN, if any.

- If a Profile sub-option is included in the message, the AR MUST use the associated profile data for use in CAR/TAR selection. Otherwise, the AR MUST use the profile stored in association with the MN, if any.

- The AR SHOULD select a limited number of CARs based on a number of factors. The exact selection algorithm is beyond the scope of this document.

  * Only those CARs SHOULD be considered that have associated AP entries in the PNC which match the list of reachable APs reported by the MN.

  * The MN's profile, if available, SHOULD be used in conjunction with the neighboring AR capabilities to select the most appropriate CARs. The total number of CARs returned to the MN could be a parameter of the MN's profile.

  * If the T-bit is set in the CLR message, the AR MAY select a

single CAR as a target AR.

- If the C-bit is set in the CLR message, the AR SHOULD attempt to return capability information for each CAR in the selected list as Capabilities sub-options associated with each Link-Layer sub-option. The capabilities sent MAY be a subset of the complete AR capabilities, dependent upon the MN's profile.

Again, how to reduce the capability sets is dependent upon the format and content of the capabilities and profile, and are beyond the scope of this document.

7.6. Sending Candidate List Messages

An AR SHOULD send a Candidate List message in response to a CLR message from an authorized MN. If the CLR message does not contain an explicit list of reachable APs and the AR does not have reachability state for the MN, the AR SHOULD return a CL message with an appropriate Status value indicating the lack of MN state at the AR.

A CL sent by the AR MUST abide by the following rules:

- The Sequence Number MUST be copied from the Sequence Number field of the originating CLR.

- The Status field MUST be set to an error value if part or all of the CLR message could not be processed. A non-error value SHOULD be used otherwise.

- Each selected CAR SHOULD be represented by one or more Link-Layer sub-options indicating the associated APs selected as targets for the MN handover. Extra information associated with each AP, such as related IP addresses or capabilities, MUST follow the Link-Layer sub-option as separate sub-options prior to the next Link-Layer sub-option in the message.

- The AR MAY include a MN Token sub-option generated according to Section 7.10.

7.7. Sending Physical Neighbor Exchange Messages

Each AR SHOULD have a global ``identifying'' IP address that it uses in communication with neighboring routers. This is necessary since a single router may have multiple assigned addresses, each interface with a different address. As the MN moves between routers, it will see only the interface-specific address of the router. Thus, a single AR would look like multiple neighbors to another AR that receives RIs from MNs handing-over from each of the different interfaces. The identifying address is used to unify the PNC entries for a neighboring AR. All Physical Neighbor Exchange messages sent to a neighboring AR SHOULD use this identifying address as the Source Address. One exception is noted below.

All PNE messages sent MUST be rate limited as described in Section 7.9. An AR MAY send a PNE message to a neighboring router in any of the following situations, however:

- The AR receives a valid RI message with a non-local previous AP address specified as a Link-Layer sub-option with the Entity field set to Previous AP. The resulting PNE MUST follow the following rules:
  * The S and T-bits MUST be set.
  * The V-bit SHOULD be set unless the AR does not require verification of the previous AP for some unspecified reason.
  * If present, the new AP and/or previous AP Link-Layer sub-options included in the RI SHOULD be copied to the PNE.
  * The AR MUST include a MN Identifier sub-option with the identity of the MN that originated the RI messages, as described in Section 7.1.
  * If present, the MN Token sub-option included in the RI message SHOULD be copied to the PNE.
- The AR attempts to verify the previous AP address of a PNE sent by a neighboring router, and the V-bit is set in the originating PNE. The reply PNE MUST follow the following rules:
  * The T-bit MUST be set and the V-bit MUST be unset.
  * The Status field MUST be set to Verified if verification succeeds or an appropriate error value otherwise.
  * If verified, the previous AP identifier from the originating PNE MUST be included as a Link-Layer sub-option with the Entity field set to Previous AP.
  * If verified, the AR MUST copy the MN Identifier sub-option from the original PNE message.

* If the Destination Address of the originating PNE is not the identifying address of the AR, the Source Address of the reply PNE MUST be set to the Destination Address of the initial PNE. Furthermore, an IP sub-option SHOULD be added to the message with the Entity field set to Source and the IP Address field assigned the identifying address.

- The AR receives a PNE from a neighboring router with the C-bit set. The AR MUST verify that the sending AR is a valid

neighboring AR by checking for a entry in the PNC. The return PNE SHOULD contain a Capabilities sub-option with the AR's capability information.

- The AR requires the capabilities of a neighboring router, or the existing set of capabilities are close to expiration. In this case, the AR MUST set the C-bit of the PNE.

For initial packets of a message stream, those with the S-bit set, the AR MUST select an Identifier with a significant chance of remaining unique between the two communicating ARs for the duration of the exchange. For subsequent packets, the Identifier field MUST be copied from the originating PNE message. The S bit MUST NOT be set in non-initial packets.

7.8. Receiving Physical Neighbor Exchange Messages

Upon receiving a Physical Neighbor Exchange message, the AR MUST validate the message according to the following rules:

- Received PNE messages MUST be rate limited as described in Section 7.9.

- If the T-bit is not set, the V-bit MUST NOT be set.

- If the T-bit is set, the previous AP address MUST be included as a Link-Layer sub-option with the Entity field set to Previous AP.

- If a MN Token sub-option is present then a MN Identifier MUST also be present in the PNE.

The AR MUST silently discard invalid PNE messages. If the S and T-bits are set, the AR MUST further verify the included topology information according to the following rules:

- The AR MUST verify that the previous AP's link-layer address provided in the sub-option represents a viable local AP. This check could be made against an actual list of known attached APs, such as the local entries in the PNC, or a list of APs authorized to be considered local.

- If a MN Token sub-option is present, the AR SHOULD verify that

the MN that sent the originating RI message, as identified by the MN Identifier sub-option, was indeed present at the AR within some reasonable amount of time, as described in Section 7.10.

If the topology verification fails and the V-bit of the original PNE is set, the AR MUST reply with another PNE as described in Section 7.7.

For verified PNE messages, processing is dependent upon the value of the S-bit. If the S-bit is set, the following rules MUST be followed in processing the PNE:

- If the T-bit is set, the AR SHOULD create a PNC entry for the previous AP as locally attached, if not already. The AR SHOULD update the lifetime of the previous AP entry to AP_LIFETIME.

- If the T-bit is set, the AR SHOULD create a PNC entry for the neighboring AR in its PNC, if not already. If the new AP is included in the PNE message, the AR SHOULD create an associated AP entry related to the neighboring AR in the PNC, if not already. The AR SHOULD update the lifetime of the new AP entry to AP_LIFETIME.

- If a Lifetime sub-option is included in the PNE, the AR SHOULD update the lifetime of the corresponding neighbor entry in the PNC, if it exists, with a value no greater than that specified in the sub-option. The granted lifetime SHOULD be returned in a PNE reply message.

If the S-bit is not set, the following rules MUST be followed in processing the PNE:

- If the T-bit is set, the AR SHOULD create PNC entries for the previous AR and previous AP, if not already. The AR SHOULD update the lifetime of the AR and AP entries to AR_LIFETIME and AP_LIFETIME respectively.

- If an IP sub-option is included with the Entity field set to Source, the address indicated by the sub-option SHOULD be used by the AR as the identifying address of the neighboring router. An existing PNC entry for the Source Address of the message SHOULD be converted to or merged with a PNC entry for the identifying address.

7.9. Rate Limiting

All messages exchanged between participants MUST be rate limited. Once the configured rate limit has been reached, subsequent packets MUST be silently discarded. For messages received from mobile nodes, rate limiting could be applied on a per-MN basis, as well as to the total number of RI, CLR and RM messages received. PNE messages SHOULD be rate limited separately for each neighboring AR, as well as for the total

number of PNE messages sent. Moreover, upon receiving an RI message, an AR MAY choose not to send a PNE message an existing cache entry exists for the previous <AR,AP> tuple with a lifetime that is above some configurable threshold.

7.10. Validating Previously Attached MNs

In order to verify that the mobile node was recently present, the previous router seemingly must maintain some short-lived state for each attached mobile node. For mobile nodes using the stateful mode of the protocol, this state is available. To support a large number of mobiles running in stateless mode, however, a more scalable solution is necessary.

Rather than track the identity of each attached mobile for some period of time, we propose that the access router generate a token that it appends to each Candidate List message that it sends to a mobile node. The mobile node then submits this token with its Router Identity message, and the token is passed back to the previous AR for verification along with the mobile's identification as part of the PNE message. To generate a token, the AR maintains a small list of random numbers used as keys to hash the identity of the mobile node. Each random number is associated with an index that is passed along with the token. Upon receiving a token for verification, the access router uses the index to lookup the associated key, hashes the MN identity passed in the PNE message and compares the hash to the token. As time progresses, new keys are generated and added to the head of the list while old keys are expired and removed. The length of the list and the frequency of generated keys are configurable and determine the total amount of time a mobile will be considered as having been recently attached.

7.11. Limiting Cache Entries

In order to mitigate the effect of erroneous reports made by possibly malicious MNs, the AR can limit the number of cache entries it creates in response to any given mobile node. When creating a cache entry, the AR SHOULD tag the entry with the identity of the mobile node (see Section 7.1). Further RI messages from the same mobile will not result in new cache entries being created once the limit is reached. We recommend a limit of two or three cache entries per mobile node.

8. Protocol Constants

The description of the CAR discovery protocol introduces a number of constants that we further define in this section. Each of these values MUST be configurable.

MIN_RM_PERIOD	2 sec.
MAX_RM_RETRY	5

MIN_CLR_PERIOD	0.5 sec.

MAX_CLR_RETRY	3
AP_LIFETIME	180 sec
AR_LIFETIME	600 sec.

9. IANA Considerations

The message containers described herein require a single ICMP type to be assigned from the available type space, as well as a single Payload Protocol Identifier for SCTP. The messages themselves require a single ICMP option type to be assigned. Message sub-types should be standardized, but do not require assignment from an existing, restricted numbering space.

10. Security Considerations

In order to secure both the discovery process, as well as the capability and profile exchanges, all messages should be secured with IPSec or TLS to provide authentication and ensure message integrity. The primary question is whether the security associations between the MN and AR and between neighboring ARs explicitly provide any type of authorization to engage in the CAR discovery protocol. If not, then a separate facility is necessary to provide this type of authorization, e.g., ACLs.

Even with authorized security associations, the discovery process is not completely secure since the ARs depend upon the MNs to properly report their handovers. Although a number of verification steps, outlined in Sections 7.2, 7.8, 7.10 and 7.11, provide strong evidence that a given RI message is valid, it is possible for a set of colluding MNs to pollute the PNC of ARs. However, since all PNC entries are maintained as soft-state, collection of MNs would need to continually provide false information to sustain the polluted entries. Moreover, the polluted entries would not affect non-malicious MNs since CAR selection is based on the reachability information provided by the non-malicious MN. If the attack was persistent and widespread, though, it could possibly result in resource exhaustion at the AR. The AR SHOULD limit the size of the PNC to eliminate this possibility, and implement an appropriate cache replacement policy that favors information collected from local reports, RIs and CLRs.

11. Intellectual Property Rights Notice

Nokia Corporation and/or its affiliates hereby declare that they are in conformity with Section 10 of RFC 2026. Nokia's contributions

may contain one or more patents or patent applications. To the extent Nokia's contribution is adopted to the specification, Nokia undertakes to license patents technically necessary to implement the specification on fair, reasonable and nondiscriminatory terms based on reciprocity.

Acknowledgements

The authors would like to acknowledge Richard D. Gitlin for his efforts in the design of this protocol.

References

[1] S. Bradner. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119, Internet Engineering Task Force, March 1997.

[2] A. Conta and S. Deering. Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. RFC 2463, Internet Engineering Task Force, December 1998.

[3] K. El Malki (Editor). Low Latency Handoffs in Mobile IPv4. Technical Report draft-ietf-mobileip-lowlatency-handoffs-v4-*.txt, Internet Engineering Task Force (IETF), June 2002.

[4] D. Johnson, C. Perkins, and J. Arkko. Mobility Support in IPv6. Technical Report draft-ietf-mobileip-ipv6-*.txt, Internet Engineering Task Force (IETF), January 2003.

[5] J. Kempf (Editor). Problem Description: Reasons for Performing Context Transfers between Nodes in an IP Access Network. Request for Comments (Informational) 3374, Internet Engineering Task Force (IETF), September 2002.

[6] R. Koodli (Editor). Fast Handovers for Mobile IPv6. Technical Report draft-ietf-mobileip-fast-mipv6-*.txt, Internet Engineering Task Force (IETF), September 2002.

[7] G. Krishnamurthi, Chalmers R., and C. Perkins. Buffer Management for Smooth Handovers in Mobile IPv6. Technical Report draft-krishnamuthi-mobileip-buffer6-*.txt, Internet Engineering Task Force (IETF), March 2001.

[8] C. Perkins, editor. IP Mobility Support. RFC 2002, Internet

Engineering Task Force, October 1996.

[9] J. Postel. Internet Control Message Protocol. RFC 792, Internet Engineering Task Force, September 1981.

[10] R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson. Stream Control Transmission Protocol. RFC 2960, Internet Engineering Task Force, October 2000.

[11] D. Trossen, G. Krishnamurthi, H. Chaskar, and J. Kempf. Issues in Candidate Access Router Discovery. Technical Report draft-ietf-seamoby-cardiscovery-*.txt, Internet Engineering Task Force (IETF), November 2001.

A. Conformance to Requirements

Identifying the IP Address of a CAR

The mapping between the link-layer address of an AP, as known to the MN, and the IP address of the associated AR is achieved as part of the dynamic topology discovery process. When the MN sends an RI message to its new AR, it specifies both the IP address of the previous AR, as well as the link-layer identifier of the previous AP. This allows ARs to build the AP to AR mapping. This mapping can then be provided, if necessary, to the MN as an IP sub-option associated with each Link-Layer sub-option in a CL message returned to the MN.

Support for Inter-Technology Handoffs

Since the Link-Layer sub-option format does not dictate any particular length or form for the AP identifiers, dyCARD is independent of the actual technologies used at the link-level. The only assumption made by the protocol is that two APs of differing technology with overlapping coverage areas would not have identifiers that are of the same length and the same value. In this case, they would be viewed as a single AP by the ARs.

Identifying CARs having Site-local and Private Addresses

Site-local and private addresses are accommodated by the protocol through the use of IP and privacy sub-options. If an AR does have a private address on the interface accessed by the MN, the AR informs the MN of which global IP address and private Id to use when sending an RI to the next AR. This global address/private Id pair is then used by the new AR to contact the previous AR given that the device owning the global address, such as a Border Router, participates in the exchange. Exactly how this should be accomplished is beyond the scope of this document.

Capability Discovery

ARs are able to exchange and refresh capabilities once a neighboring relationship has been established. The PNE message can be used to request new capabilities. The actual capabilities are carried in the Capabilities sub-option. Capabilities can also be forwarded to the MN through the same sub-option, as part of the CL message.

Utilization of Network Resources

The CAR discovery protocol we have proposed makes every effort to limit the bandwidth used, especially over the wireless link between the MN and the AR.

Specifically, the semi-stateful design of the reachability message limits the number of messages to at most two per distinct beacon received, one upon first receiving the beacon and possibly another if the beacon disappears. Moreover, each RM message can contain multiple beacon reports. The use of beacon Ids also works towards reducing bandwidth. Upon reporting a new beacon, the MN assigns it a locally unique Id which it sends in the Link-Layer sub-option. Later, when the AR sends a CL message with a list of candidate APs, only the Ids need be present, the actual addresses may be elided. The same applies to revoking a previously reported beacon.

Another design feature intended to save bandwidth on the wireless link is the use of the Beacon List by the AR to reduce the set of possible CARs sent to the MN. Capabilities sent along with each CAR could also be reduced to include only those items relevant to the MN's requirements. Finally, the CLR message includes the T-bit which when set allows the AR to perform the TAR selection procedure locally and return only the selected TAR.

Format of Capabilities

The correct operation of the CAR discovery protocol is not dependent upon the format of capabilities or the MN profile. All messages exchange capabilities and profiles as binary blocks. The management and use of these blocks of data are delegated to companion specifications related to the TAR selection algorithm, and are outside the scope of this document.

Scope of CAR Discovery

The design of the CAR discovery protocol is in no way restricted by the location of the two neighboring ARs. They could be within the same domain, or in different domains administered separately. The only requirement made by the protocol is that the two ARs are capable of establishing a security association that provides the explicit authorization to engage in the CAR discovery protocol.

Introduction of Dedicated Network Elements for CAR Discovery

The CAR discovery protocol described in this document does not introduce any new network elements. The only exception is the necessary presence of some proxy device that would allow the use of a global address/private Id pair to establish communication with a privately addressed AR. This, however, is not a new requirement for supporting private address spaces.

Involvement of Non-CARs in CAR Discovery

Only the MN and the two ARs involved in a particular handover participate in the CAR discovery protocol.

Dependence on a Mobility Management Protocol

The CAR discovery protocol does not depend on any particular mobility protocol in order to operate correctly. Within this document, we use terminology, such as Care-of Address, that is normally related specifically to Mobile IP. This, however, is done only to simplify the discussion.

In the presence of other mobility-related protocols, there are opportunities for optimizing the performance of our protocol. See Appendix B for details on optimizing performance when used in conjunction with Fast Mobile IPv6.

Effect of Changes in Network Topology

The dynamic nature of the CAR discovery protocol we propose adapts well to changes in the network topology. Since the PNC is maintained as soft-state in the ARs, APs that are removed will timeout. ARs that no longer overlap, will disassociate, and new APs will be discovered once a MN handsover to it. The protocol design provides for a tradeoff between the stability of the PNC and the degree of reconfigurability. Longer lifetimes applied to AP discovery via RI messages provide a more stable picture of the neighborhood in the face of few MNs. On the other hand, lower lifetimes allow APs to be moved between ARs more frequently.

Providing the MN's Requirements to the CAR Discovery Solution

The MN may provide its current AR with a list of its requirements for CAR selection in the form of a Profile sub-option included in any RM message sent by the MN. The AR can then take advantage of the MN's

profile to better select a set of viable CARs, and possibly reduce the amount of capability information required to send back to the MN.

Secure Capability Transfer

All protocol messages should be protected by IPSec or TLS to provide authentication and ensure message integrity. The capability data may also be encrypted and/or digitally signed to further ensure secure transportation.

Verification of Router Authenticity

Verification of a neighboring AR's authenticity is approached in two ways. First, the two ARs must establish a security association via some means outside the scope of the CAR discovery protocol. This security association should provide an explicit authorization to engage in exchanging topological information.

Secondly, neighboring ARs are initially discovered through the handover pattern of MNs. In order for two ARs to have discovered one another, a MN must have moved between their respective APs. Moreover, the AP identifiers provided by the MN are verified at both the new and previous ARs in order to filter out faulty information provided by a malicious MN.

However, in the case that an incorrect PNC entry is instantiated, the result is simply a temporary allocation of memory in the AR. MN handovers will not be affected since only ARs matching the reachability information provided by the MN are selected as CARs.

Secure Inter-Operability with IETF Protocols

The CAR discovery protocol described in this document employs IPSec and/or TLS for securing communication between all parties, between two ARs and between an AR and a MN. The packet formats described in Section 4 use ICMP and SCTP, although this is not a necessary constraint. The protocol does not create any new threats to the use of existing protocols.

Secure Expression of MN's Requirements to the CAR Discovery Solution

All protocol messages should be protected by IPSec or TLS to provide authentication and ensure message integrity. The profile data may also be encrypted and/or digitally signed to further ensure secure transportation.

B. Optimization with Fast Mobile IPv6

In the presence of other mobility-enabling protocols, such as Fast Mobile IPv6 [6], it is possible to optimize the message flow of the CAR discovery protocol described by this document. In particular, the CLR and CL messages could be overlaid onto the Proxy Router Solicitation and Proxy Router Advertisement messages, respectively. This would require a slight change to Fast MIPv6. The options and sub-options defined herein could be reused to piggyback CARD information on Fast MIPv6 messaging.

Employing this type of optimization would eliminate one round-trip between the MN and the AR near the time of handover. This is quite desirable since link quality may fade significantly as a MN reaches the edge of a given APs coverage area. Moreover, the MN would receive the most current CAR list available since the transaction would occur immediately before handover.

Author Addresses

Questions about this memo can be directed to the authors:

Dirk Trossen
Communications Systems Laboratory
Nokia Research Center
5 Bayside Road
Burlington, MA 01803
USA
+1 781 993 3605
+1 781 993 1907 (fax)
dirk.trossen@nokia.com
Govind Krishnamurthi
Communications Systems Laboratory
Nokia Research Center
5 Bayside Road
Burlington, MA 01803
USA
+1 781 993 3627
+1 781 993 1907 (fax)
govind.krishnamurthi@nokia.com
Hemant Chaskar
Communications Systems Laboratory
Nokia Research Center
5 Bayside Road
Burlington, MA 01803

USA
+1 781 993 3785
+1 781 993 1907 (fax)
hemant.chaskar@nokia.com
Robert C. Chalmers
Department of Computer Science
University of California, Santa Barbara
Santa Barbara, CA 93106
USA
+1 805 893 7520
+1 805 893 8553 (fax)
robertc@cs.ucsb.edu
Eunsoo Shim
NEC Labs America
4 Independence Way
Princeton, NJ, 08540 USA
eunsoo@nec-lab.com