Internet Engineering Task Force
INTERNET-DRAFT

Dana Blair
Alex Tweedly
Michael Thomas
Jonathan Trostle
Michael Ramalho
Cisco Systems

File: draft-blair-rt-mobileipv6-seamoby-00.txt

November 2000
Expires: June 2000

Realtime Mobile IPv6 Framework

Status of this Memo

Abstract

This draft develops terminology, an architectural framework, a list of objectives, and a proposed solution for Realtime Mobile IPv6. Realtime Mobile IPv6 seeks to deliver realtime application data to IPv6 capable Mobile Nodes while minimizing the impact of handoff on realtime applications.

1. Introduction

The convergence of wireless networking and IP networking requires solutions for transporting realtime application data to mobile devices. Current IP mobility solutions tend to focus primarily on best effort data transport to and from a mobile device. To support realtime applications, security, admission control, and QoS need to be examined concurrently with mobility to facilitate smooth transitions from one access link to another while minimizing the impact on realtime applications.

This draft assumes IPv6 because the proliferation of mobile devices will require the use of IPv6 addresses to support the number of attached devices. Without the larger IPv6 address space, the end-to-end transparent Internet model will be broken up and made much more complex by the use of private IP addressing and result in the inclusion of

Network Address Translators, Application Layer Gateways, and Protocol Translators. Also, use of private IP addressing is not practical when the end device functions as an application server such as Voice over IP.

1.1 Problem Description

Mobility solutions require that packets be redirected to the current location of a mobile device. The Mobile IPv6 Binding Update [MIPv6] is sufficient to accomplish this for best effort traffic. However, realtime applications require coordination of admission control, security associations, and Quality of Service guarantees to minimize packet loss and jitter to satisfy the expectations of realtime applications.

Considered independently, admission control, Security Associations, Quality of Service signaling, and Binding Updates could significantly delay the transition from one access link to another, thus degrading the performance of realtime applications.

This draft develops terminology, an architectural framework, a list of objectives, and a proposed solution for Realtime Mobile IPv6.

2. Terminology

Subscripted Elements

> In this draft we will have occasion to reference the "present" element of a given type and the "next" element of that same type with indices "n" and "n+1", respectively. For example, we will reference the first access link (AL) as AL1, the next as AL2, and so on. We will use this indexing methodology as shorthand for any element (e.g., ALn+1, ARn, vPDPn. We do not intend that the actual "present" index "n" to be identical across element type (e.g., the present ALn could be AL9, while the present vPDPn could be vPDP2).

The sample network diagram of Figure 1 below is used to illustrate the terminology used in this draft. This figure will be used to describe the operation of the proposed handoff mechanisms by means of the specific network topology example shown in this figure.

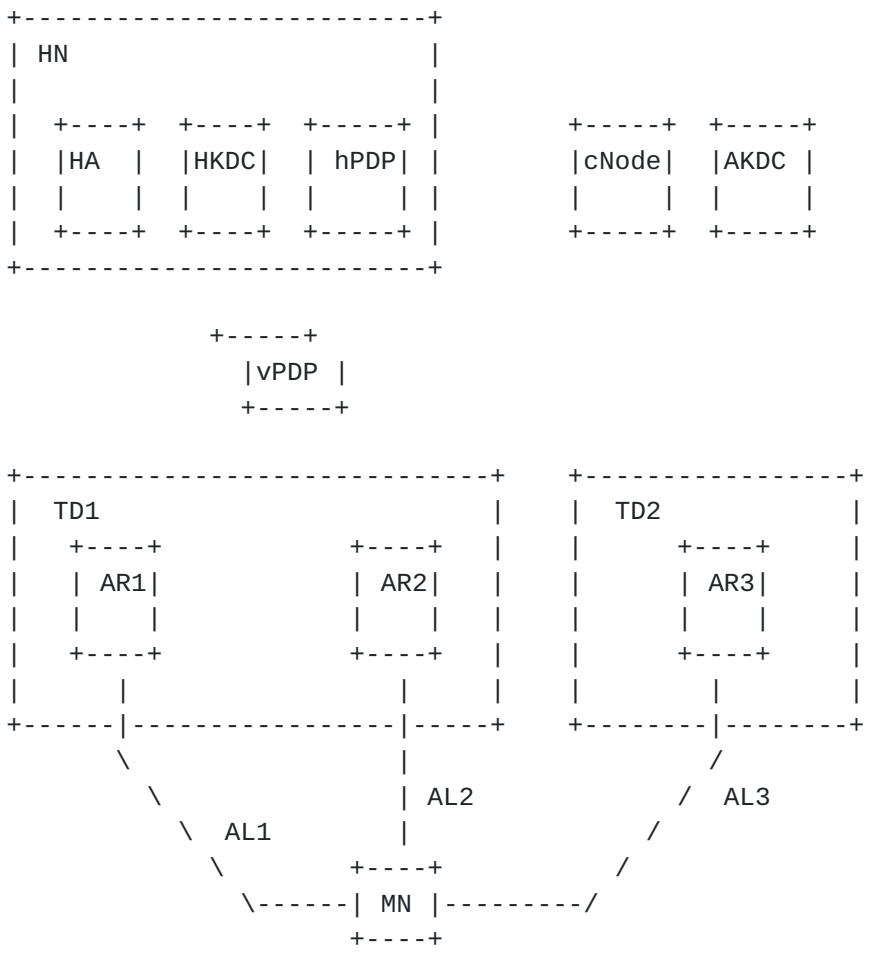

Figure 1: Sample Network Diagram showing AR1 and AR2 in a particular trust domain and AR3 in a different trust domain.

Common Terms

Mobile Node (MN)

Device which supports IPv6 mobility according to [MIPv6]. A MN may be a host or a router.

Home Agent (HA)

Device that supports the Home Agent function described in [MIPv6].

Correspondent Node

A peer node with which a mobile node is communicating. Defined in [MIPv6].

Home Address (hAddr)

An IP address assigned to a MN as defined in [MIPv6].

Care-of-Address (CoA)

An IP address assigned to a MN as defined in [MIPv6]

New CoA (NCoA) means that CoA on ALn+1 is different than CoA on ALn. Same CoA (SCoA) means that CoA on ALn+1 is the same as CoA on ALn.

Access Router (AR)

An IP router between an Access Network and one or more access links.

Access Network (AN)

An IP network which includes one or more Access Routers.

Policy Decision Point (PDP)

A network service that is responsible for handling policy decisions (e.g. access authorization, qos authorization, etc.) and also billing and settlement issues. Defined in [COPS1]

Home PDP (hPDP) is the policy decision point used by the Home Network.

Visited (vPDP) is the policy decision point used by the visited network.

Key Distribution Center (KDC)

A network service that supplies tickets and temporary session keys. Defined in [KRB1].

The Home KDC (HKDC) provides security credentials for the MN to register with the Home Agent.

Trust Domain (TD)

A set of access routers within a particular network that have direct trust relationships between all ARs in that set. In the figure above TD1 is comprised by the set {AR1, AR2} and TD2 is the set {AR3}..

Access Link (AL)

A link layer connection between MN and AR. Multiple Access Routers may share the same access link. An AL consists one downlink or one uplink or both.

Downlink communication flows from AR to MN.

Uplink communication flows from MN to AR.

Handoff Trigger

A Handoff Trigger initiates a Handoff Decision. The MN or an element within the AN may create a Handoff Trigger. The result of the handoff decision may create a Handoff Command.

Example sources of a Handoff Trigger are MN detects a stronger signal for a different AL, or AN detects that a MN is moving and a new AL would provider better transmission. Many other sources of Handoff Trigger exists.

This draft does not discuss events that result in a Handoff Trigger.

Handoff Decision

A Handoff Decision is an algorithm which decides whether or not to issue a Handoff Command based on one or more Handoff Triggers.

This draft does not propose algorithms for the Handoff Decision.

Handoff Command

The Handoff Command initiates a Handoff Sequence. The MN or an element within the AN may issue a Handoff Command. The Handoff Trigger and Handoff Command may be created by different network components or the same network component.

Handoff Sequence

A Handoff Sequence is a series of steps executed by various network components to attempt a handoff. A Handoff Sequence may be successful in establishing ALn+1. However, establishment of ALn+1 does not guarantee communication with a cNode. Since IP is connectionless, a Handoff Sequence cannot rely on IP or link layer indications to verify that the MN can communicate with a cNode over ALn+1.

The solutions section of this draft contains an abstract Handoff Sequence proposal. Detailed Handoff Sequences are in the appendix.

Initialization Sequence

An Initialization Sequence is a series of steps executed by various network components to establish communication over an AL for the first time.

Wake-up Sequence

A Wake-up Sequence is a series of steps executed by various network components to establish an AL while a MN is in sleep mode. Sleep mode means the device is in a lower power state with no uplink, and a limited downlink capability.

Handoff Delay

Handoff delay is the time between when the handoff command is issued and when ALn+1 is established.

Handoff Types

Make-Before-Break Downlink Handoff (MBBDH)

ALn+1 downlink is established before removing ALn downlink. Duration of simultaneous ALn and ALn+1 downlink connections is not bounded.

Make-Before-Break Uplink Handoff (MBBUH)

ALn+1 uplink is established before removing ALn uplink. Duration

of simultaneous ALn and ALn+1 uplink connections is not bounded.

Break-Before-Make Downlink Handoff (BBMDH)

The MN stops receiving on ALn before being able to receive on ALn+1.

Break-Before-Make uplink Handoff (BBHDH)

The MN stops transmitting on ALn before being able to transmit on ALn+1.

Handoff Control

Handoff Control may be performed by various network elements. Any Handoff Sequence should be analyzed to determine which control methods are used. Some control methods may be preferable depending access network technology or administration preference.

Mobile Controlled

A MN may create any number of access links without a-prior knowledge of any other access link. The MN generates the Handoff

Trigger, performs the Handoff Decision, and issues the Handoff Command.

Mobile Assisted

The MN generates the Handoff Trigger, the network performs the Handoff Decision and issues a Handoff Command to the MN.

Network Assisted

The network generates a Handoff Trigger. If the network performs the Handoff Decision, the network issues a Handoff Command to the MN. If the MN performs the Handoff Decision, the MN issues a Handoff Command to itself.

Network Controlled.

The network generates the Handoff Trigger, performs the Handoff Decision, and issues the Handoff Command. The MN must execute a Handoff Sequence to maintain an AL.

3. Objectives

Items to consider when analyzing a Handoff Sequence.

3.1 Handoff Objectives

Minimize Handoff Delay

Minimize packet drops during Handoff Delay

Realtime applications have a relatively low tolerance for packet drops. Packet drops are more likely during BBMDH and BBMUH.

Minimize packet latency

Minimize misadaptations of adaptive jitter buffers due to handoff.

Link Layer Handoff Independence

A Realtime Mobile IPv6 Handoff Sequence must not depend on any particular link layer handoff mechanisms but may exploit layer 2 information.

Handoff link failure behavior

A Realtime Mobile IPv6 Handoff Sequence should consider behavior when ALn+1 cannot be established during the Handoff Sequence.

Handoff addressing, either SCoA or NCoA, is an important consideration when analyzing or creating a Handoff Sequence. For example, SCoA does

not require MN re-registration but may not be possible between Trusted Domains.

Handoff Control is an important policy consideration of some access networks. For example, A Handoff Sequence may not be acceptable for a given access network or network policy if the sequence is not Network Assisted or Network Controlled.

3.2 Security

Fast smooth handoffs must take into account access networks which require both authentication and authorization of mobile users in order to gain access to best effort traffic as well as any other services such as enhanced quality of service, etc.

User authentication may be done directly on the access network or via proxied authentication through a third party where the access network provider and the third party have a trust relationship. The latter arrangement allows for a settlement model where a mobile node does not need to have a direct arrangement with the visited service provider, but instead is authenticated with the third party, but settled between the third party and the access service provider.

The initial authentication and authorization state should be captured by the access service provider router in a way which allows the authorization state to be forwarded to other access routers within the same Trust Domain. This allows access routers the option of not needing to perform authentication or authorization back to the authoritative authorization source. The advantage of caching this authorization state at the edges is that potentially expensive round trips to the authorizing source can be avoided.

The MN should consider requiring mutual authentication. Mutual authentication prevents an intruder from pretending to be an AR and gathering credentials from the MN.

3.3 Quality of Service

The handoff mechanism should consider the re-establishment of QoS

instantiated in ALn when handing off to Aln+1. This should include either intserv or signaled diffserv techniques.

4. Handoff Behavior

The following is an abstraction of events that may be required before and during handoff. Not all steps are required for all handoff nor are they necessarily executed in the sequence presented. A Handoff Sequence will define a set of steps and protocols required to attempt a handoff.

4.1 Initialization Events

The Initialization Events allow a MN to establish connectivity with one or more cNodes before attempting to handoff. Although Initialization Events are not part of a Handoff Sequence some parts of initialization may be reused during a Handoff Sequence to simplify operation within the MN.

1. Obtain credentials to allow use of the visited network, if necessary.

2. Obtain access to and through a visited network. Use credentials from 1., if needed.

3. Obtain credentials to allow MN to register with the Home Agent.

4. Obtain credentials to be used when establishing Quality of Service for application flows, if necessary.

5. Establish security association with the Home Agent using credentials from 3.

6. Mobile Node registers with the Home Agent as defined in [MIPv6].

7. Establish zero or more application flows over the first access link, AL1, to one ore more cNodes. QoS credentials from 4. may be used to authorize Quality of Service via a PDP for a specific application flow.

4.2 Handoff Events

The Handoff Events allow a MN to attempt to move packet forwarding and receiving from ALn to ALn+1. Some events may not be required, or may not be executed in the order presented when applied to a Handoff Sequence.

1. Obtain credentials, if necessary, to allow use of ALn+1.

2. Obtain access to and through the visited network. If necessary, use credentials from 1.

3. QoS signaling for 0 or more application flows. If necessary, use QoS credentials obtained during initialization.

4. Establish tunneling and/or forwarding from ARn to MN using ALn+1.

5. If New CoA (NCoA), MN re-registers with Home Agent as defined in [MIPv6].

6. If NCoA and if desired, MN notifies cNodes NCoA.

7. If possible, Remove ALn when it is no longer needed.

4.3 Application of Handoff Objectives.

4.3.1 Minimize Handoff Delay

Authorization state transfer may be pulled from ARn to ARn+1 or pushed from ARn to ARn+1.

Minimize the number of required transactions before MN may send and receive packet to and from a cNode over ALn+1.

Minimize the number of home roundtrip transactions required before MN may send or receive a packet to or from a cNode over ALn+1.

Minimize the size of all packets transmitted and received before MN may send or receive a packet to or from a cNode over ALn+1.

4.3.2 Minimize Packet Drops

Limit packet drops by forwarding packets from ARn to MN over ALn+1.

4.3.3 Minimize packet latency

For NCoA handoff, re-register with HA and cNodes as soon as possible.

For SCoA handoff, minimize the delay of forwarding path convergence.

4.3.4 Security

If ALn and ALn+1 are on different routers, ARn and ARn+1, then:

A. ARn and ARn+1 may have a pre-established trust relationship, or
B. ARn and ARn+1 may be able to create a dynamic trust relationship, or
C. ARn and ARn+1 will never have a trust relationship.

4.3.5 Quality of Service

Maintain quality of service by establishing acceptable QoS over ALn+1. QoS credentials may be needed.

5. References

[MIPv6] David B. Johnson, Charles Perkins, "Mobility Support in IPv6", draft-ietf-mobileip-ipv6-12.txt, April 2000.

[KRB1] Kohl, J., Neuman, c., "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.

[COPS1] Durham, D., Boyle, J., Cohen, R., Herzog, S., Rajan, R., Sastry, A., "The COPS (Common Open Policy Service) Protocol",

RFC 2748, January 2000.

[COPS2] Herzog, S., Boyle, J., Cohen, R., Durham, D., Rajan, R., Sastry, A., "COPS usage for RSVP", RFC 2749, January 2000.

[RSVP1] Herzog, S., "RSVP Extensions for Policy Control", RFC 2750, January 2000.

[RSVP2] Baker, F., Lindell, B., Talwar, M., "RSVP Cryptographic Authentication", RFC 2747, January 2000.

Authors Information

Dana Blair - dblair@cisco.com
Alex Tweedly - agt@cisco.com
Michael Thomas - thomasm@cisco.com
Jonathon Trostle - jtrostle@cisco.com
Michael Ramalho - mramalho@cisco.com

Appendix A: Make Before Break Handoff Sequence using RSVP, Kerberos, and COPS.

This is Handoff Sequence is a work in progress. No reader should assume that this is complete or covers all required handoff scenarios.

Authorizing PDP (aPDP) is a policy decision point, which is trusted by the hPDP.

Authorizing KDC provides credentials and authorizes use of the Home KDC

Access Service Request is a message which initiates handoff.

INITIALIZATION SEQUENCE

The MN is hard-configured with Home Address and prefix length, and either shared secret with the Home Agent, Authorizing KDC, or certificate.

1. Obtain sufficient access over AL1 to reach Authorizing KDC (AKDC).

   The AKDC may be separate from the Home Network.

2. Obtain ticket for Authorizing Policy Decision Point (APDP) from AKDC.

   One of the requirements of any successful scheme is the need for authentication and authorization for the large number of mobile hosts. This requires a scalable authentication mechanism. The prime candidates are Kerberos V5 [RFC 1510] and public key based mechanisms.

   The desire to achieve fast handoff, assuming relatively limited computing power on the mobile hosts, argue against any scheme involving public key computations in the handoff phase. Furthermore, the revocation checking that is required for most existing public key based mechanisms would cause additional delays. Therefore the scheme discussed proposes Kerberos for authentication. Kerberos still allows for some of the main benefits of a public key based approach since Kerberos includes extensions for public key authentication. At the same time, the speed of Kerberos secret key authentication is leveraged during the handoff.

3. Obtain sufficient access over AL1 to contact Home KDC (HKDC).

4. Obtain ticket from HKDC for HA.

5. Using Home Agent Discovery defined in [MIPv6], find the address the HA which the MN will use.

HANDOFF SEQUENCE

6. Access Service Request (ASR)

   The Access Service request is sent from MN to AR1.
   The ASR contains the following:

- description of the qos signaling mechanism in use
- some description of the type/level of service being requested (e.g. RSVP flowspec)
- flag whether tunneling is desired. A MN-AR1 key may be needed if tunneling is desired.
- address and ticket for PDP
- address and BU for the HA.

If the MN has an application which uses RSVP, then RSVP [RSVP1] is proposed as the ASR for a variety of reasons.
A layer three admission control mechanism is needed to provide both a means to access the network as well as a means of signaling that a handoff operation is desired. If enhanced quality of service is required for some flows, as may be the case with real time applications, a common admission control scheme for both basic network access and admission of quality of service flows is desirable. Modeling best effort traffic as a subset of all qualities of service leads to the possibility of using RSVP as the signaling protocol.

RSVP also has some other advantageous properties. RSVP already has the means of carrying credentials in the form of policy objects [RSVP2] which can be sent to a remote policy decision point. RSVP already has the concept of soft state. Soft state in the context of mobility is quite important as there needs to be some means of clearing stale state in interior network elements without explicit signaling.

RSVP's major advantage, however, comes when it is used with enhanced quality of service flows. Since it is likely that real time traffic in some important situations will require explicit admission control to deal with scarce bandwidth, a fast, smooth handoff cannot be considered complete until the quality of service on the new access router and beyond is installed. While RSVP messages could be sent in addition to an ASR, a single message which requests the flow at the attachment point along with the necessary credentials would be more economical. The message could also carry the request for any forward tunneling desired for a smooth handoff.

Additions to RSVP would include a handoff object and changes to support Best Effort, non-flow based admission control.

6a. Authorization

6a-1. AR1 determines whether it can make the authorization decision on its own. If AR1 cannot authorize on its own, continue with 6a-2.

6a-2. AR1 takes the level of service request, and hPDP info, and forwards those to its own PDP, the visited PDP1 (vPDP1). The vPDP1 may then decide it can authorize (or deny) service.

Otherwise, continue with 6a-3.

COPS [COPS1] is proposed as the admission control policy protocol. COPS and RSVP are well integrated via RFC 2749. As for RSVP, a single admission control mechanism for both best effort and enhanced quality of service is desirable. COPS has been

specifically designed for enhanced quality of service admission control and could be easily extended, if needed, to support admission of best effort traffic. As with RSVP, COPS has the desirable property of allowing for policy objects which can be used for authentication of a MN. Also, COPS supports cascaded policy decision points to allow a visited PDP to authorize with a home PDP.

6a-3. vPDP1 forwards the credential to hPDP for a decision, and/or settlement/billing decision between PDPs or via 3rd party broker, or other means.

The result whether local to AR1 or from vPDP1 is a boolean decision to allow access of MN through AR1.

6b. If tunneling is desired, get new key for MN/AR1

If necessary, vPDP1 sends to hPDP the ticket and address for MN, type/level of service, etc. 6a may have already done this through interaction with the hPDP. However if not, this step is required.

hPDP replies to vPDP1 with a new credential for use between MN and AR1.

6a and 6b may occur in parallel. However, 6c and 6d must wait for a successful result from 6a to reach AR1.

6c. QoS signaling

If using RSVP - send the RSVP packet towards HA.
If no QoS signaling, then AR1 simply imposes DiffServ policing if needed.

6d. Send Binding Update (BU) [MIPv6] to HA.

The BU may be attached as a routing header on the ASR or imbedded in an object within the ASR. If the BU is embedded in an object within the ASR, the BU must be in the form of a complete IP packet including IPSEC encapsulation.

If the BU is in an embedded object, AR1 extracts the BU and sends it on to HA.

End of 6: The MN may now communicate with the HA including proper

level of QoS assuming RSVP was successful.

As long as the neither MN nor the AN requires a handoff, no more steps are required.

7. MN attempts to handoff to AR2.

The Handoff Trigger may occur because MN heard a beacon from AR2 and decided AL2 was a better connection, or because the network containing AR1 determined that the MN should moved to AL1. There are many other possibilities as well.

7a. MN acquires suitable local address by either stateless or stateful address configuration. This may be either SCoA or NCoA.

8. MN sends an ASR to AR2.

The ASR is sent from MN to AR2, contains:
- QoS : description of the qos signaling mechanism in use
- QoS : description of the type/level of service being requested (e.g. RSVP flowspec)
- Security : flag whether tunneling back to AR1 is desired (i.e. is a MN-AR2 key needed)
- Security : one or more tuples of
  < vPDP, ticket vPDP, address/name of hPDP, ticket hPDP >
- Mobility : [if New-CoA] address and embedded (authenticated) BU for the HA.
- Mobility : [optional] address and BU for the previous AR

Note that this allows multiple tuples of PDP address + ticket. These will be considered in turn, with AR2 deciding whether or not to trust the vPDPn-1. Thus, if MN has recently been accessing the network via AR1, and has successfully established a key for use with AR1, then MN may (will?, must?) make the first tuple be <AR1, hPDP, ticket from 1a>, and the second tuple be <*, hPDP, ticket from 1a>.

Thus in the authorization step AR2 will consider each of these tuples for a possible authorization decision in turn.

Otherwise, this should be just same as the sub-steps of 2 above, with the addition of (optionally) AR2 sending a BU to AR1. Also, if using Same-Coa, AR2 will kick off the routing update / convergence.

Note - if RSVP is the protocol to be used here, it most likely is addressed to HA, but will be intercepted by AR2, and the following steps taken before the packet is allowed to continue on its way to HA. But if it is some new "Access Service Request", then it can simply be destined to AR2.

Therefore :

8a. Authorization

8a-1. AR2 determines whether it can make the authorization decision on its own. If AR2 does not authorize on its own, continue with 8a-2.

8a-2. AR2 tries first PDP address + PDP ticket tuple where the PDP address is the address of AR1. If AR2 trusts AR1, then AR2 will treat AR1 as a vPDP by asking AR1 to authorize the Access Service Request.

AR1 may then decide it can authorize (or deny) service, based on having previously done so. Otherwise, continue with ..

8a-3. AR2 tries second PDP address + PDP ticket tuple where PDP address may be *.

The vPDP specified is null (blank), so AR2 uses his own, normal vPDP.

AR2 forwards hPDP info to vPDP2, and vPDP2 forwards the credential to hPDP for a decision, and/or settlement/billing decision between PDPs or via 3rd part broker, or other means.

8b. If tunneling is desired, get new key for MN/AR2

vPDP sends to hPDP the ticket and address for MN, type/level of service, etc. In general, step 8a above would include doing this, but in the case where it did not, they need to be send on to hPDP anyway.

hPDP replies to vPDP with a new credential for use between MN and AR2. This used the pre-existing trust between vPDP2-hPDP and between MN-hPDP to generate this new credential. vPDP2 sends this on to AR2 and to MN.

8a and 8b may occur in parallel. 8c and 8d must wait for a successful result from 8a to reach AR2.

8c. QoS Signaling

If using RSVP - send the RSVP packet towards HA.
If no QoS signaling, then AR2 simply imposes DiffServ policing, if required.

8d. If the optional BU for AR1 is included, AR2 sends the BU, which was IPSEC encapsulated using the key derived in 6b, to AR1. MN is responsible for the construction of this BU, and therefore MN determines whether SCoA or NCoA is used.

For SCoA, the BU will cause AR1 to tunnel to AR2 where AR2 will forward the packet to MN.

For NCoA, the BU will cause AR1 to tunnel directly to CoA2

Note that AR2 or AR1 omit or reject this step if either AR2 or AR1 determines that this step violates network policy.

8e. MN sends BU to HA.

AR2 extracts the embedded BU and sends this on to HA.

8f. If SCoA, then AR2 (and also AR1 based on step 4d) will initiate routing convergence to reflect the new attachment point for the CoA.

All subsequent handoffs (e.g. ALn+1 to ALn+2) occur in like manner beginning with step 7.

Appendix B: Handoff Sequence using Kerberos and Radius/Diameter

This Handoff Sequence is a work in progress. No reader should assume that this is complete or covers all required handoff scenarios.

1. INITIALIZATION SEQUENCE

Initially, the MN obtains a ticket granting ticket (TGT) using either a password or X.509 certificate in an AS exchange with a KDC in its home realm using IAKERB [IAKERB]. Subsequently, the MN mutually authenticates with the AR in order to obtain network access using IAKERB.
An IPSEC SA is also established with the user's home agent. The IPSEC SA key management protocol could be IKE or any of the IPSEC key management protocols.

2. HANDOFF SEQUENCE

The proposal in this appendix specifies the use of Kerberos [3] for providing authentication, key establishment, and authorization for Mobile IPv6 [2,5] fast handoffs.

ARn detects the network prefix of the ARn+1. ARn sends a Neighbor Discovery Redirect Message (ND Redirect) over ALn with a list of network prefixes for ARn+1. Upon receipt of this message, the MN configures a new CoA.

We define new suboptions for the ND Redirect Message to include a negotiation flag, and the NR principal name and realm. The negotiation

flag indicates whether to pass security state from the PR, which we call local authentication, or local authentication followed by global authentication (which may include an exchange with a Kerberos KDC server), or global authentication. The MN then sends a Previous Router Notification Message to
the ARn which uses the IPSEC AH header and is keyed using the existing security association between the MN and the PR; this message informs the PR of the MN's new CoA. We define a new suboption for this message that includes a negotiation flag (the MN will choose a flag that is the same or more strict than the negotiation flag value sent to it from the PR in the Notification Message), and a suboption for including a Kerberos message.

The value of the negotiation flag that was sent from the MN to the ARn will help determine what steps occur next. If local authentication was selected, then the ARn will pass secret cryptographic key information as well as other state information to ARn+1, using an encrypted channel between ARn and the ARn+1. This data will be sent

in an unsolicited message (message TBD) from ARn to ARn+1. The MN established ALn+1 to ARn+1.
The MN and ARn+1 will exchange ND solicitation and advertisement messages over ALn+1 prior to exchange of application data. The MN and the ARn will use the shared secret keys to establish IPSEC SA's between themselves after the initial handoff operations.

If the value of the negotiation flag that was sent from the MN to ARn+1 was local authentication followed by global authentication, then the same steps as in the local case occur, except subsequently (after application data is flowing), Kerberos mutual authentication with additional key establishment occurs. We have defined new suboptions for the Neighbor Discovery Solicitation
and Advertisement [8] for inclusion of Kerberos messages to achieve global authentication. Two ND solicitation and advertisement exchanges will be needed to encapsulate the three Kerberos messages in this case.

If the negotiation flag value that was sent by the MN is global authentication, then the Previous Router Notification message will contain a ticket granting ticket (TGT) as described in the minimal messages subprotocol from IAKERB [4]. A typical exchange in this case would be:

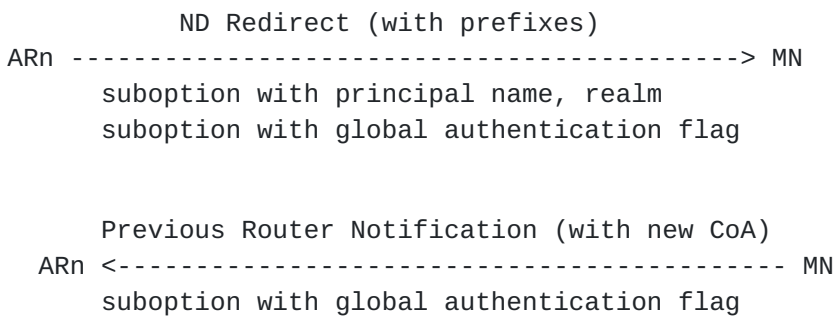

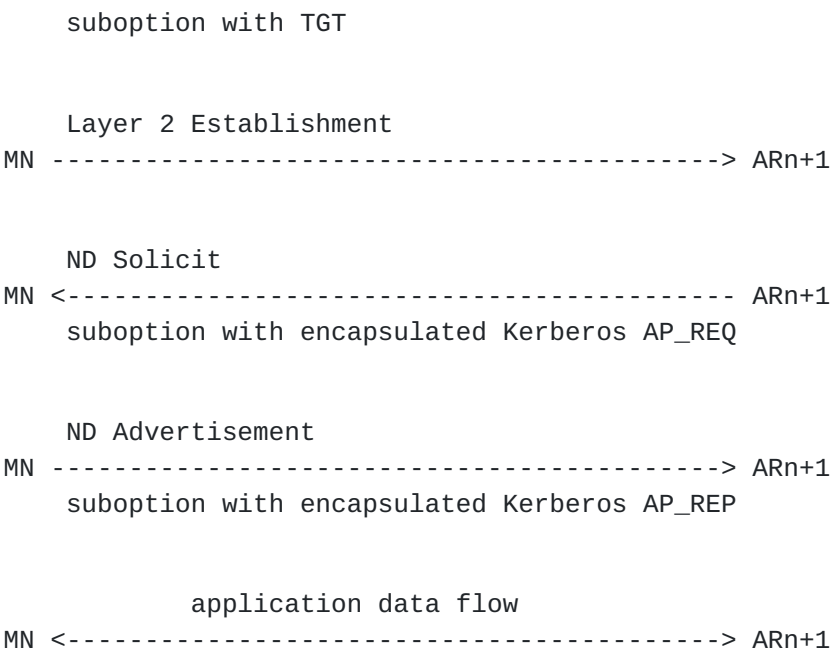

Figure 1: Global Authentication for Fast Handoffs

Before forwarding the AP_REQ to the MN, ARn+1 will send a message to its local AAA server with authorization data from the Kerberos service ticket extension. The reply from the local AAA server will contain the updated and massaged authorization data. ARn+1 will cache the authorization data and use it once the MN has authenticated to ARn+1. The protocol between ARn+1 and the local AAA server could be either Radius or Diameter. Upon receipt of the AP_REQ from the ARn+1, the MN will perform the normal Kerberos validation steps. In addition, the service principal identifier component of the client principal name in the ticket MUST be equal to mobileip. If not, the MN will fail the request by sending a KRB_ERROR message to the NR with the error code KRB_ERR_GENERIC.

Note that the AP_REQ will have the mutual authentication flag set. Thus, the MN is required to authenticate to the ARn+1 which it does by sending the AP_REP message to the ARn+1.

Subsequently, the MN sends any BCU's to its home agent (HA) and correspondent nodes (CN's). The MN may also send a router solicitation message to the ARn+1
with a new suboption asking for the reverse ticket (IAKERB [4]) and list of adjacent realms to be sent back in the router advertisement. The reverse ticket is constructed by the ARn+1 and is targeted at the ARn+1 from the MN. It allows fast subsequent authentications from the MN to the ARn+1. The list of adjacent realms can be used by the MN to precache crossrealm TGT's targeted at adjacent realms (as a background

task). This performance optimization allows subsequent global authentications for adjacent realms to skip exchanges with the remote user's KDC. Instead, a local KDC will be used.

Kerberos authentication uses the subprotocols from IAKERB [4]. In particular, the minimal messages subprotocol (which uses user-user authentication) is used for the global authentication exchange (illustrated in Figure 1). Precaching can use either standard Kerberos or standard IAKERB. Initial logon (which occurs before the first roam) should use standard IAKERB exchanges with an IAKERB proxy. As a result of initial logon, the MN user will obtain an initial TGT, as well as a service ticket targeted at the initial access router which acts as an IAKERB proxy. Initial logon can use public key certificates as described in pkinit [9].

For a break before make case, the MN may have to establish layer 2 with the ARn+1 without the benefit of the ND Redirect and Previous Router Notification messages. In this case, the security data from the ND Redirect message (in the above figure) will be placed in the router advertisement message sent from the ARn+1 to the MN in response to the MN's router solicitation message. The MN then sends a ND solicitation with the authentication flag and optionally, the TGT, if global authentication has been selected. In the global authentication case, the ARn+1 responds with a ND advertisement that includes the AP_REQ message. In this case the flow becomes:

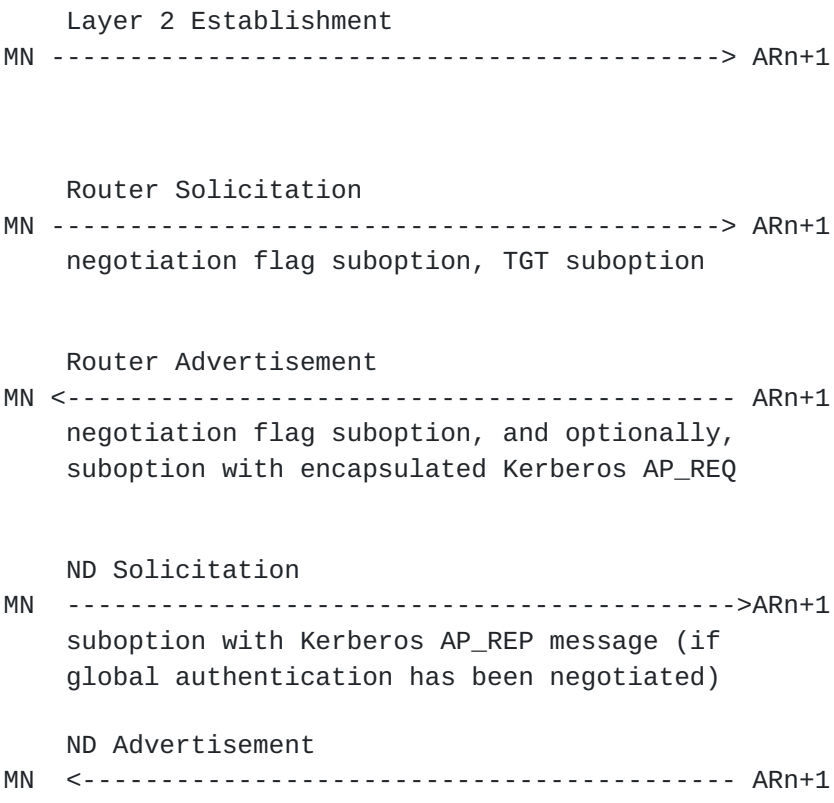

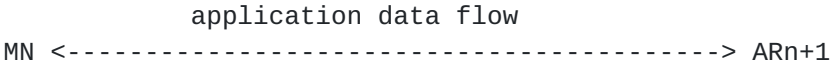

Figure 2: Global Authentication: Break Before Make

3. Neighbor Discovery Solicitation, Advertisement, Redirect Suboptions, and Previous Router Notification Suboption.

The following suboption contains the principal name and realm:

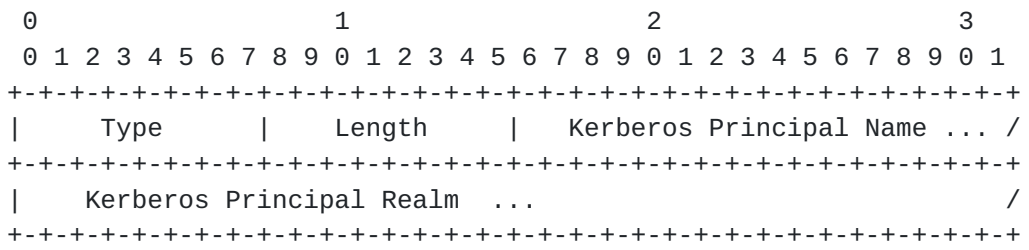

Fields:

Type	Message type. To be assigned.
Length	8-bit unsigned integer. The length in bytes of the option (including the type and length fields).
Principal Name	Kerberos Principal Name
Principal Realm	Kerberos Principal Realm

The following suboption contains the negotiation flag:

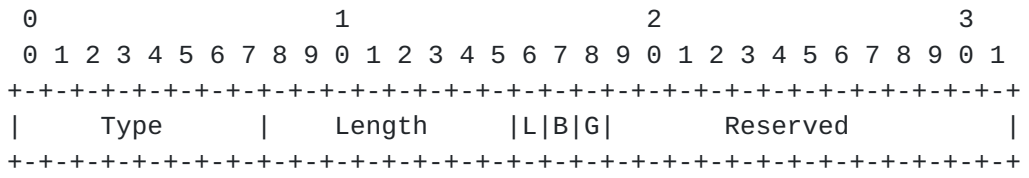

Fields:

Type	Message type. To be assigned.
Length	8-bit unsigned integer. The length in bytes of the option (including the type and length fields).
Flags	L - local authentication

B - local authentication followed by global
authentication
G - global authentication

The requestor sets one of the above flags and the responder sends the same suboption with either the same flag or a more strict flag. The flags increase in strictness from L to B to G. We note that an attacker can cause at most a denial of service attack by manipulating these flags in transit.

The following suboption contains a Kerberos protocol message:

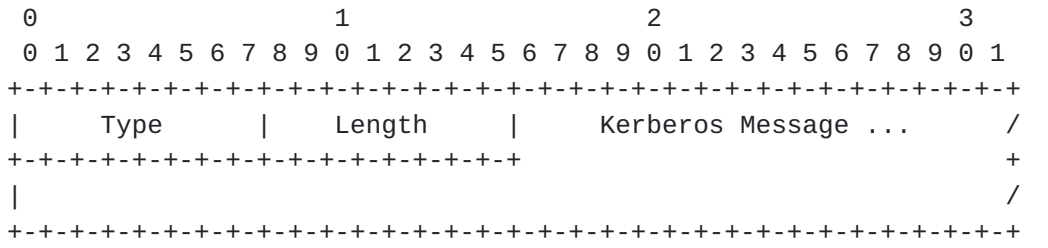

Fields:

Type	Message type. To be assigned.
Length	8-bit unsigned integer. The length in bytes of the option (including the type and length fields).
Kerberos Message	Kerberos protocol message as defined in [3]

Any of the above suboptions can be included in any of the neighbor discovery solicitation, advertisement, or redirect messages, as well as the previous router notification message. The principal name/realm suboption and the negotiation flag suboption can also occur in a router advertisement message.

4. Router Solicitation and Advertisement Options

The following router solicitation message suboption is used to request

that the access router return a list of adjacent realms in its router advertisement. The access router may also return a Kerberos message suboption with a reverse ticket.

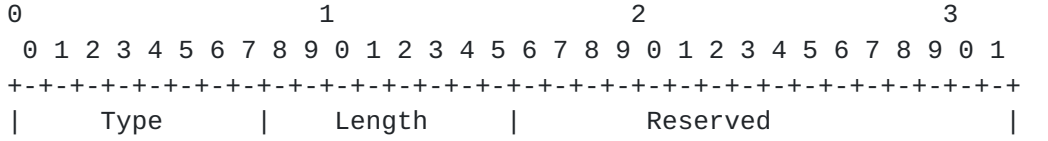

Fields:

Type	Message type. To be assigned.
Length	8-bit unsigned integer. The length in bytes of the option (including the type and length fields).

The following router advertisement message suboption is used to return a list of adjacent realms in the router advertisement:

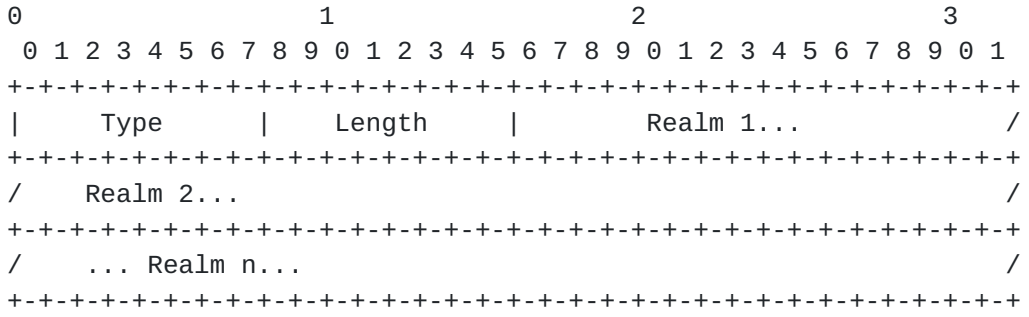

Fields:

Type	Message type. To be assigned.
Length	8-bit unsigned integer. The length in bytes of the option (including the type and length fields).
Realm i	The realm name of the ith geographically adjacent Kerberos realm.

The following router solicitation message suboption is used to request that the access router return the principal name/realm suboption and the negotiation flag suboption (as described above). It is used in break before make handoff cases.

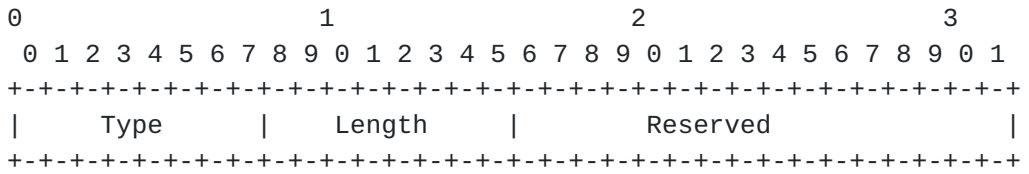

Fields:

Type	Message type. To be assigned.

Length 8-bit unsigned integer. The length in bytes of the option (including the type and length fields).

5. Authorization

We propose the following approach to authorization for the global case (in the local case, authorization state is passed from the previous router to the new router).

The NR's KDC will map authorization data in the user's Kerberos ticket to local authorization attributes which will be placed in a ticket extension of the issued ticket.

Here we propose a new Kerberos authorization data type:

AD-MOBILEIP-ATTRIBUTES TBD

The data is the ASN.1 OCTET STRING encoding of the following:

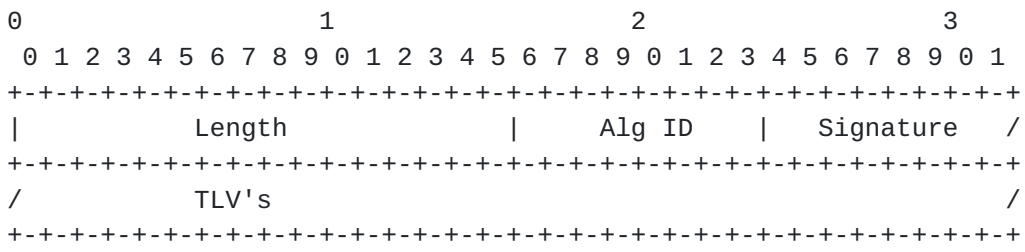

Fields:

Length 8-bit unsigned integer. The length in bytes of the option (including the type and length fields).

AlgID Algorithm Identifier

Signature Digital signature over TLV's.

TLV's Authorization attributes in TLV form

We also propose a new ticket extension type:

TE-MOBILEIP-ATTRIBUTES 9

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Deering, S., Hinden, R., "Internet Protocol, Version 6 (IPv6)

Specification", RFC 2460, December 1998.

[3] Kohl, J., Neuman, C., "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.

[4] Swift, M., Trostle, J., "Initial Authentication and Pass Through Authentication Using Kerberos V5 and the GSS-API (IAKERB)", Internet draft (work in progress), draft-ietf-cat-iakerb-04.txt, July 2000.

[5] Johnson, D. and Perkins, C., "Mobility Support in IPv6", Internet draft (work in progress), draft-ietf-mobileip-ipv6-12.txt, April 2000.

[8] T. Narten, E. Nordmark, and W. Simpson. Neighbor Discovery for IP Version 6 (IPv6). Request for Comments (Draft Standard) 2461, Internet Engineering Task Force, December 1998.

[9] Tung, B., Neuman, C., Hur, M., Medvinsky, A., Medvinsky, S., Wray, J., Trostle, J., "Public Key Cryptography for Initial Authentication in Kerberos", Internet draft (work in progress), draft-ietf-cat-kerberos-pk-init-11.txt, April 2000.

Appendix C

This appendix shows a fully integrated RSVP/COPS/Kerberos solution. It is very similar to Appendix A, and will be merged at a later date.

Abbreviations:

PO: RSVP/COPS policy object
TR: Ticket Relay Object: an object which relays a shared secret back to the initiator and relay initiator. This subject deserves a great deal more discussion, and will be added in future versions of this draft.
HO: Handoff Object: an object which requests a handoff to occur between AR1 and AR2 with optional forward tunneling from AR1 to AR2. If handoff objects desire forward tunneling, they must contain an embedded handoff object
EBU: Embedded Binding Update: an embedded binding update is a normal IPv6 mobility binding update message which is encapsulated in a wrapper which proves MN's identity as well as proving freshness, etc. EBU's may be embedded in RSVP and COPS messages.
tick[XXX]: a Kerberos ticket for which the server half of the ticket is encrypted in XXX's key. MN is always assumed to be in the client half of the ticket.

Initial Access

When the mobile node initially desires access, it must go through a full initial access flow. This flow assumes that MN does not have any valid Kerberos tickets cached. If it does, it may omit the interactions with the appropriate KDC's.

The goal of this flow is to just obtain best effort access from AR1. As such, the RSVP request is directed at AR1.

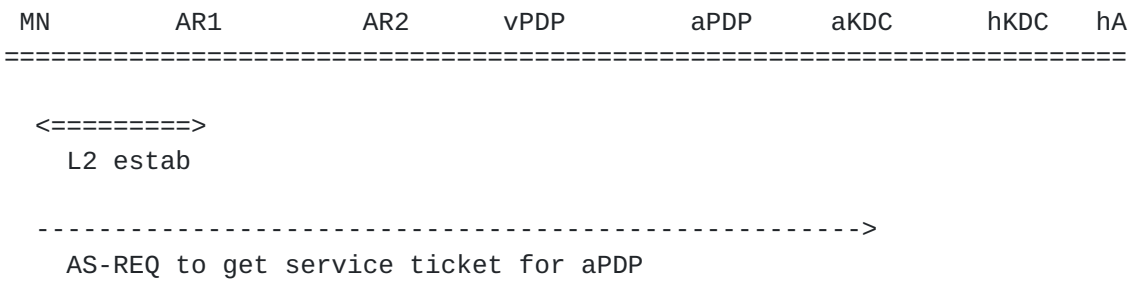

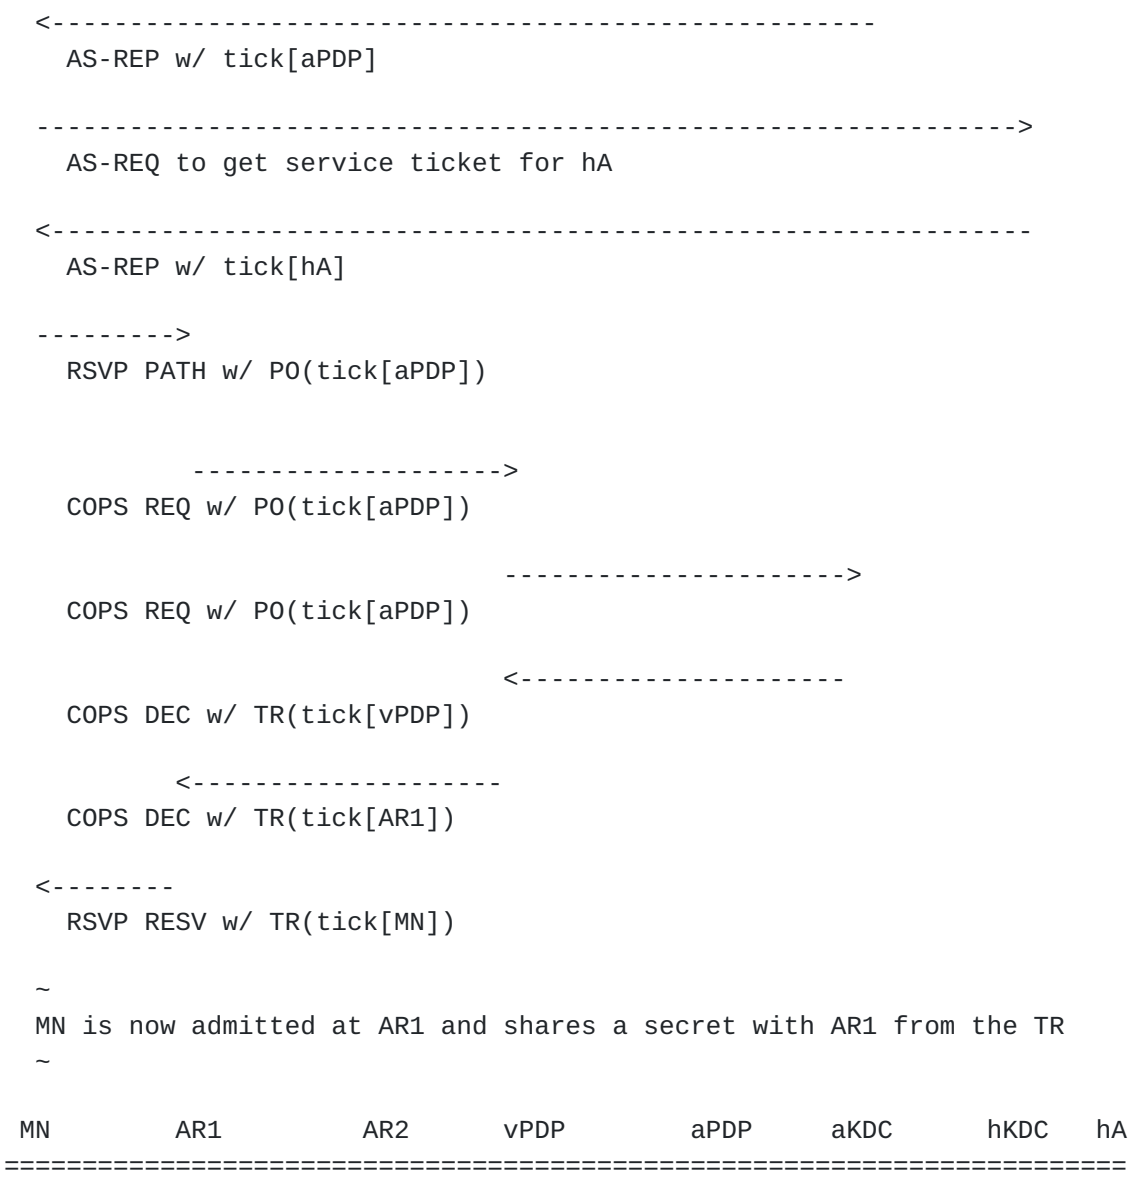

~
MN is now admitted at AR1 and shares a secret with AR1 from the TR
~

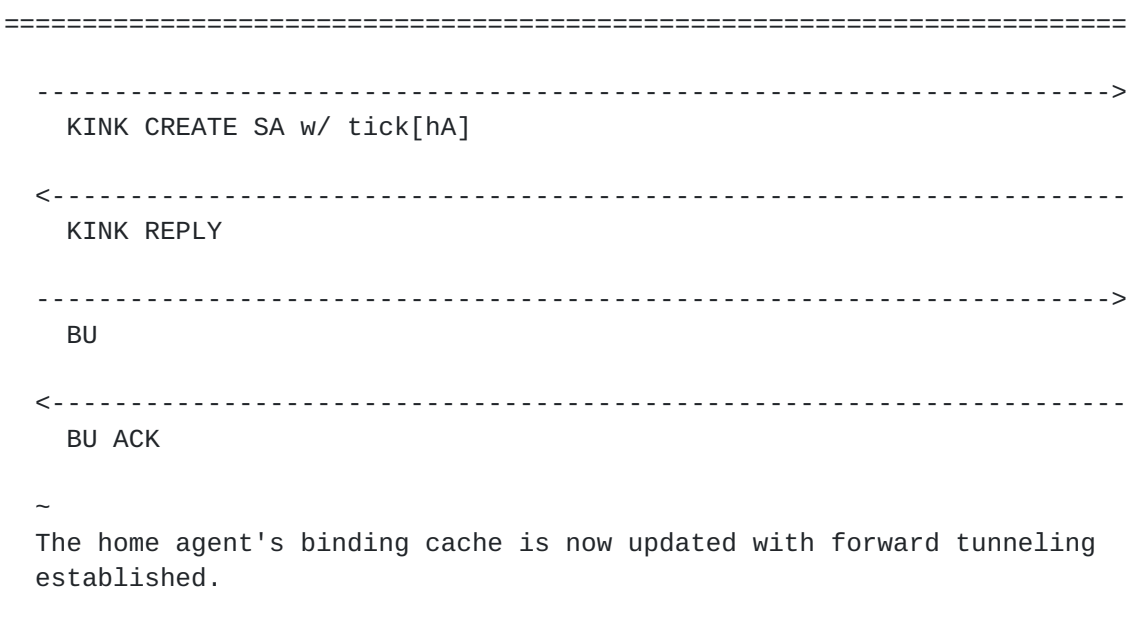

~
The home agent's binding cache is now updated with forward tunneling established.

Note: if the MN is already in possession of valid tickets to the aPDP

or hA, the AS-REQ and AS-REP's in the flow can be omitted.
~

Intra Trust Domain Fast/Smooth Handoff

It is assumed that AR1 and AR2 have a pre-existing security association.
If they do not, a security association between the two would need to be established using normal IPsec keying mechanisms.

This flow assumes that the MN is capable of transmitting and receiving on AR2. It may or may not still have connectivity to AR1. It is also assumed that AR1 and AR2 are with in the same trust domain where AR1 can act as a PDP for AR2.

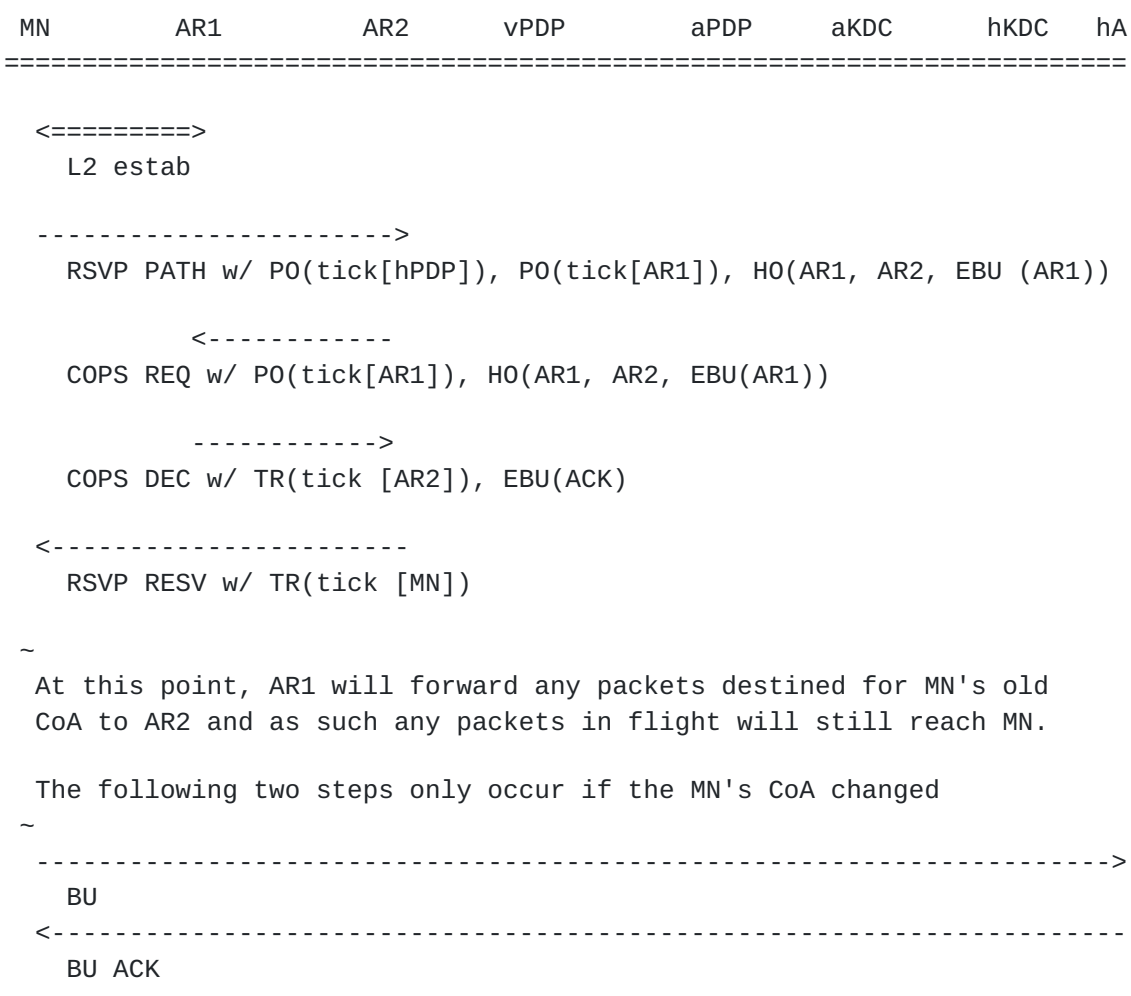

Intra Trust Domain Fast/Smooth Handoff to a cNode with QoS

This flow describes a fast/smooth handoff with a cNode with an established QoS flow. It is assumed that cNode and MN already

have established keys between them. How they are initially established is outside of the scope of this document.

It is assumed that AR1 and AR2 have a pre-existing security association.
If they do not, a security association between the two would need to be established using normal IPsec keying mechanisms.

This flow assumes that the MN is capable of transmitting and receiving on AR2. It may or may not still have connectivity to AR1. It is also assumed that AR1 and AR2 are with in the same trust domain where AR1 can act as a PDP for AR2.

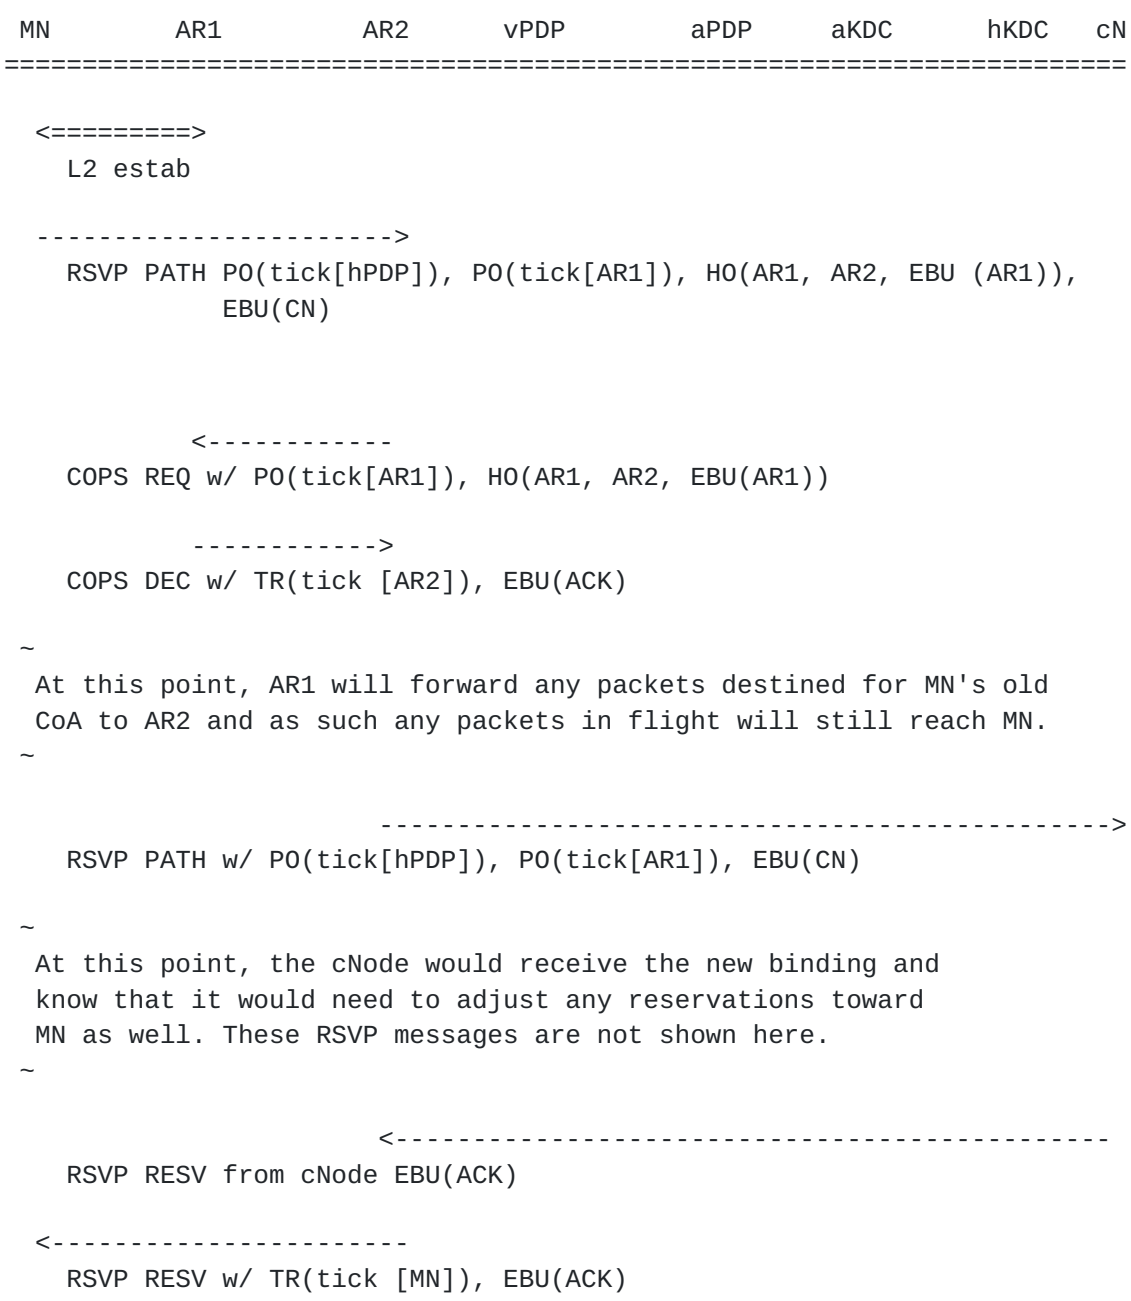

Inter Trust Domain Handoff

In the case where AR1 and AR2 do not trust another domain's cached policy decisions, it will necessitate another authoritative policy decision. Note that the MN is at liberty to send as many policy objects as it feels may be needed to expedite the decision. Also: if the MN has local credentials it may send them as well which would also result in a fast though not necessarily smooth handoff.

Note that this flow is nearly identical to the initial access flow. The only difference is that the MN sends a PO and HO for AR1. Because AR2 and AR1 are not within the same trust domain, AR2 ignores that policy object since it is not a trusted PDP.

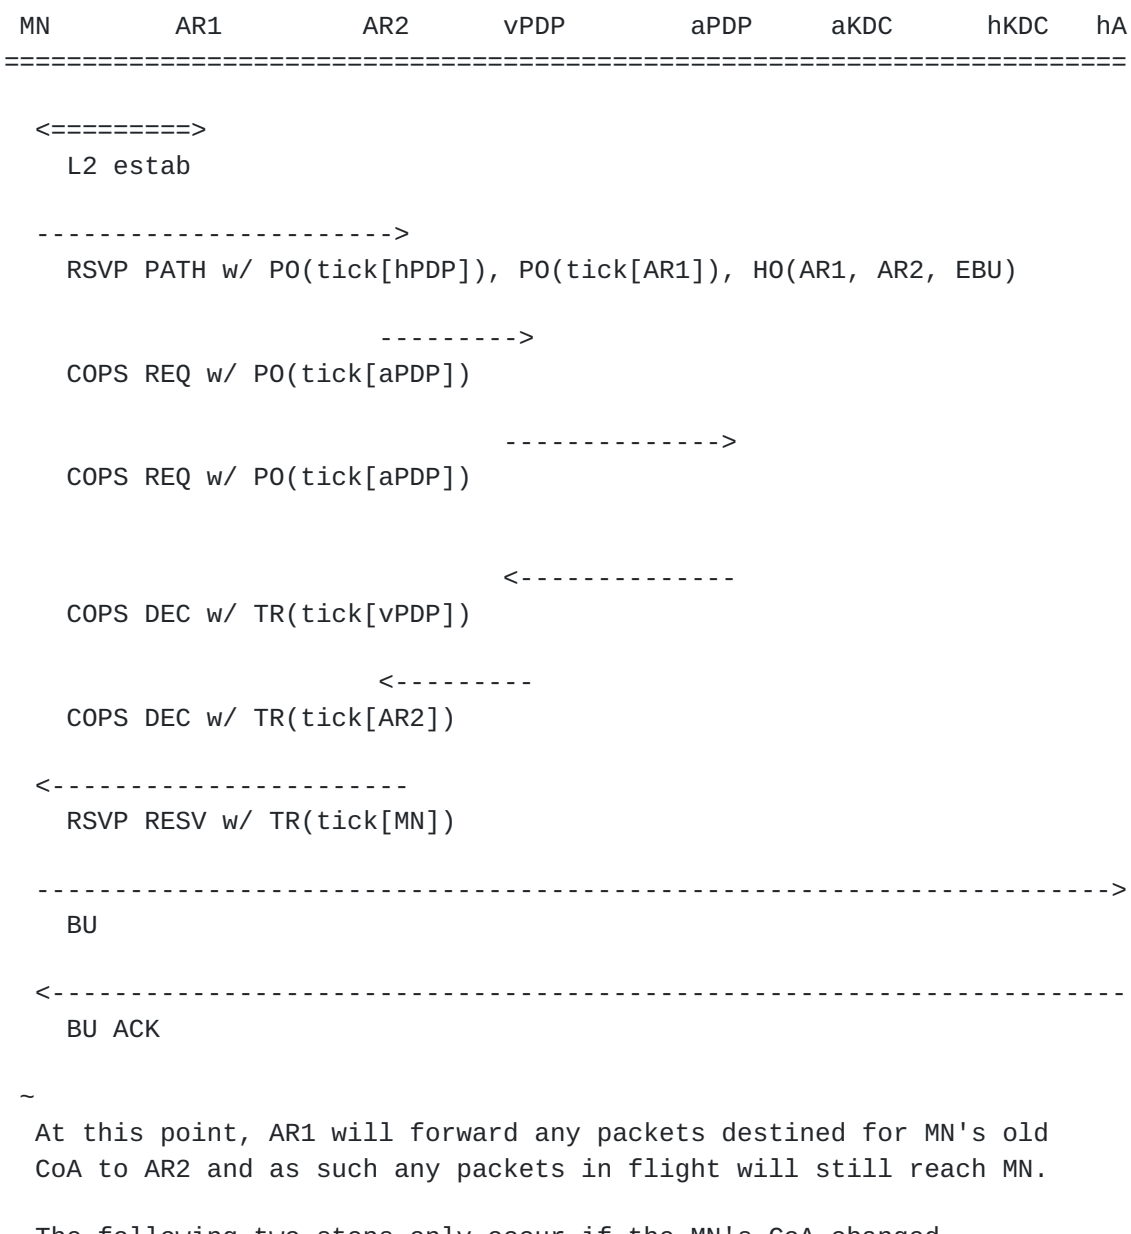

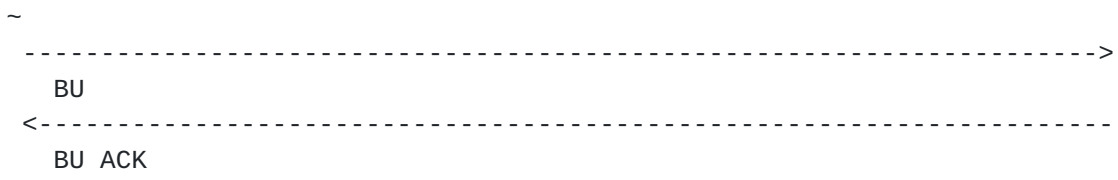
~
BU
BU ACK