Michael Thomas
Cisco Systems
Oct 28, 2002

Analysis of Mobile IP and RSVP Interactions

draft-thomas-nsis-rsvp-analysis-00.txt

Status of this Memo

Abstract

IP Mobility along with RSVP makes my head hurt. Trying to guess the benefits (if any) of a proposed context transfer protocol is even more difficult to judge. This draft tries to make sense of some subset of the possible permutations in a possibly vain attempt have an overview of the problem space.

1. Introduction

The intersection of Mobile IP with RSVP is not well understood. This draft attempts to bring forth issues with their interaction, as well as diagram how a naive RSVP and MIP hosts would likely interact, assuming they implemented the protocols correctly. This version of the draft does not intend to be comprehensive as the vagueries of even simple cases raise plenty of questions. This draft does, however, try to speculate what the potential effects and/or advantages of Fast Mobility as well as the potential effects of a hypothetical context transfer protocol being worked on in the NSIS working group.

1.1 Terminology

MN mobile node

AR1 old access router

AR2 new access router

PDP policy decision point

AGG aggregation router; join point of reservations

CN correspondent node

X->Y means X sending traffic toward Y; ie, X is the source of the PATH and Y is the source of RESV's.

<-|-> change bar implies that the following messages can happen simultaneously. It should be noted that these are only suggestive of the concurrency and in many cases can be drawn differently given message arrival timing.

1.2 Assumptions

1.2.1 Handoff Trigger

As with all diagrams in this draft, we elide the mechanism by which AR1 decided that the handoff to AR2 was necessary.

1.2.2 Reservation Healing

In all of the diagrams, we expect that PATH's and RESV's are not propogated past the point needed to heal a reservation in place.

1.2.3 No Dog Leg Routing

The draft assumes that the reservation was put in place after the mobile node sending a binding update to the correspondent node and as such, the reservation is not being placed through the home agent.

1.2.4 All Parts of the PATH Must be Restored

In order to have a seamless handoff, all routers that are RSVP aware in the PATH after handoff must be alerted such that the PATH and RESV state can be installed on that router.

1.2.5 RSVP Session-ID

As currently formulated, RSVP generally uses the 5-tuple SRC,DST,PROTO,SPORT,DPORT as the index to which reservation a particular packet should be classified. It is also used as the means to refer to the reservation in subsequent RSVP signaling.

While adequate for many kinds of RSVP flows, using the 5-tuple as the session identifier runs into problems when the application desires admission of a certain class of traffic and expects to keep its traffic within a given Tspec, but would like to be able to change the filterspec mid-session. Consider, for example, the case of VoIP Call Waiting where the QoS envelope for the two calls is identical (say, 64kb/sec), but only one flow will be active at any one time. Unfortunately, RSVP as currently specified requires double booking of resources mainly because there is no way to assocatiate the new 5-tuple in the filterspec with the old reservation.

Mobility tickles the same problem, but in a different way. When a mobile node move from one router to another, it may change its care of address. It is assummed that the address used for the 5-tuple is the source address as seen in the IP header (as opposed to the actual home address). As such, instead of the destination of the reservation changing, the source address would be the variable part of the 5-

tuple instead of the destination, but the results yield a similar result: the reservation would need to be double booked, as well as the implication that any change of care of address would require a full round trip to the correspondent node.

For these reasons, this draft makes the wild and perhaps unfounded assumption that the ISSLL working group will step up to allowing RSVP to use a session identifier in addition to the normal 5-tuple to identify a reservation. Given the round trip implications, the author cannot see how seamless mobility with RSVP reservations can be achieved if this assumption is false.

1.2.6 Context Transfers

This draft dabbles in the "what-ifs" of some sort of means of transfering context between AR1 and AR2 on a handoff. The assumptions made here are:

- both the PATH and RESV state can be transfered
- authentication and authorization state can be transfered
- AR2 can, as is possible, act as if it were a set of interfaces on AR1 for the purposes of RSVP where the handoff looked like an ordinary topology change
- that the mobile may be made aware of the context context transfer and as such it would not need to consider AR2 to be completely naive with respect to reissuing PATH messages, or expecting RESV's.

  The intention here is not to define a mechanism nor suggest a particular solution. The intention here is only to illustrate whether context transfer provides any potential benefit to existing reservations.

  It should also be noted that there are a couple larger issue involved with context transfer which the RSVP portion of the problem space also needs to grapple with. Namely, context transfer is a posited means by which a router can restore the current state on another router for seamless handoffs, but it clearly does not answer the question of whether the move _ought_ to happen. A second issue arises as to what ought to happen when the QoS policies at one access router are different from the QoS policies at another

access router. For example, it is not hard to imagine a relatively congested access technology like cellular needing the admission control policies of RSVP, but another technology only requiring diffserv. Should the context transfer deal with QoS equivalencies? What ought be the mechanism to determine those equivalencies? And as above, ought the mobile node and/or previous access router have any say-so over such equivalencies from a policy or some other standpoint? Another issue is whether double booking of resources in a make-before-break kind of handoff is an issue. If so, how do you avoid it?

As stated earlier, this draft only raises these issues to hopefully get an eventual answer.

1.3 Initial Reservations

This draft assumes that there were two reservations set it place using the normal RSVP mechanisms between the mobile node and correspondent node.

MN->CN:

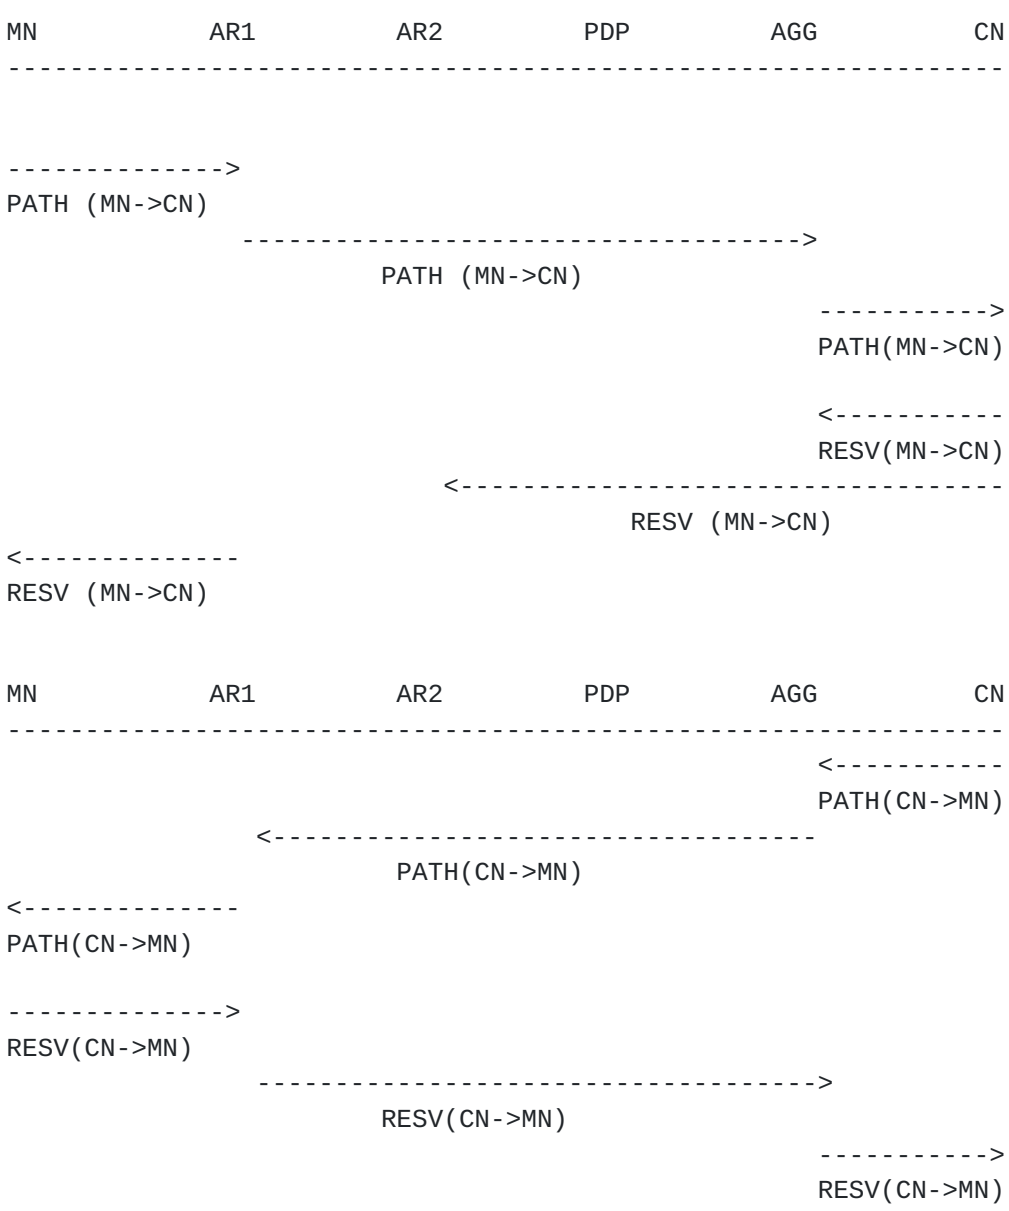

1.4 "Optimal" RSVP Handoff

This flow is fictional and is based on the incorrect assumption that if RSVP existed solely for the purpose of making seamless mobility work with established RSVP sessions, this is an optimal placement of reestablishment messages. This flow is intended to only be viewed as gauge in the successive flows as to how far they deviate from this self-proclaimed optimum. The assumption here is that through some unspecified means, AR2 came to know that the mobile was moving from AR1. It should be noted that AGG in this diagram may, in fact, be an

arbitrary number of hops -- both RSVP and non-RSVP -- away from AR2, and that it is just the common join point for the reservation. It is assumed that implementations will consider this in the actual network engineering.

MN->CN, CN->MN:

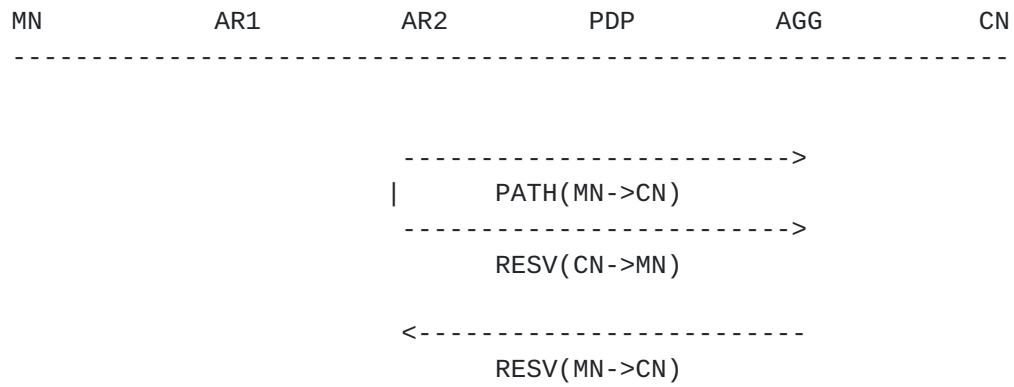

It should be pointed out that aside from the implication that a potential difference in the CoA will not hinder the reestablishment of reservations, there are several things wrong RSVP-wise with the "optimum". Namely:

1) It is unspecified how AR2 arrived with the PATH and RESV state required for both the MN->CN and CN->MN reservations. As we shall see later, MN itself would issue the PATH for the MN->CN reservation to reestablish PATH and RESV state.

2) The naked RESV from AR2 to AGG is illegitimate. There is no PATH state associated with it, and while we might be able to imagine AGG consing up the PATH state given the reservation on another interface, RSVP routers between AR2 and AGG would be completely clueless about the RESV. Indeed, since they have no PATH state, they wouldn't know where to forward the RESV for the next hop.

1.5 Additional Issues

1.5.1 Interdomain Trust Issues

While NSIS may not directly deal with interdomain trust issues, it

isn't clear how policy based admission interacts with topology changes. Is this well understood? In the flows below, what would go in the policy objects and what if anything would need to happen at the PDP's? What credentials need to be carried in RESV's of routers past the PEP's?

This memo tries to make an attempt to answer these questions. In no way does this memo intend to be exhaustive on this subject as a number of simplifying assumption have been made in the subsequent sections.

1.5.2 RSVP Aggregation

[2] defines a method of aggregating RSVP microflows into larger aggregates by tunneling individual RSVP microflows between an aggregator and deaggregator. The intended behavior is that the tunneled microflows will traverse the aggregated area of the network where it will receive the appropriate treatment by either installing a larger aggregated reservation, or perhaps using a diffserv service level agreement, etc.

It is quite possible that mobility will uncover further issues here. This draft does not consider aggregation, but future revisions of this draft or elsewhere ought to consider interactions with RSVP aggregation as well.

2. Same Care of Address

With the same care of address, we expect that the reservation will look the same and that this looks a whole lot like a topology change. The main deviation from the optimal flow is that there is the need for MN to reestablish the MN->CN reservation since AR2 is completely unaware fo any reservations that MN had established through AR1. In the CN->MN direction, AGG, upon seeing the host route, would under normal RSVP start sending PATH messages

2.1 MN->CN reservation

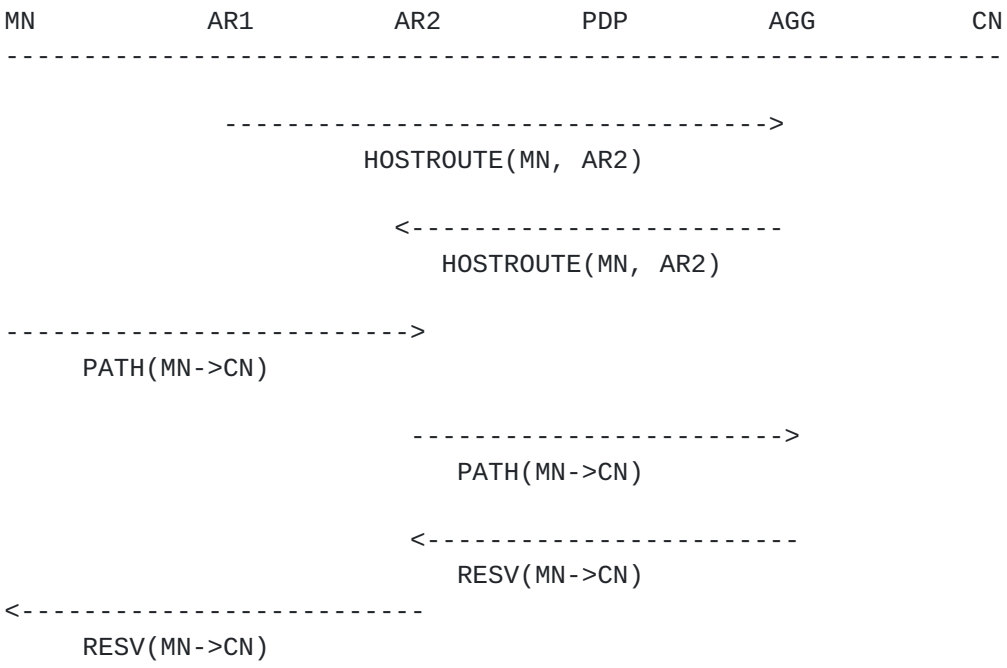

2.2 CN->MN reservation

The CN->MN reservation benefits from the fact that AGG is alerted to the topology change and can therefore initiate a PATH toward MN as with any other topology change. Since AGG is the join point, CN doesn't need to be informed of any changes as with standard RSVP.

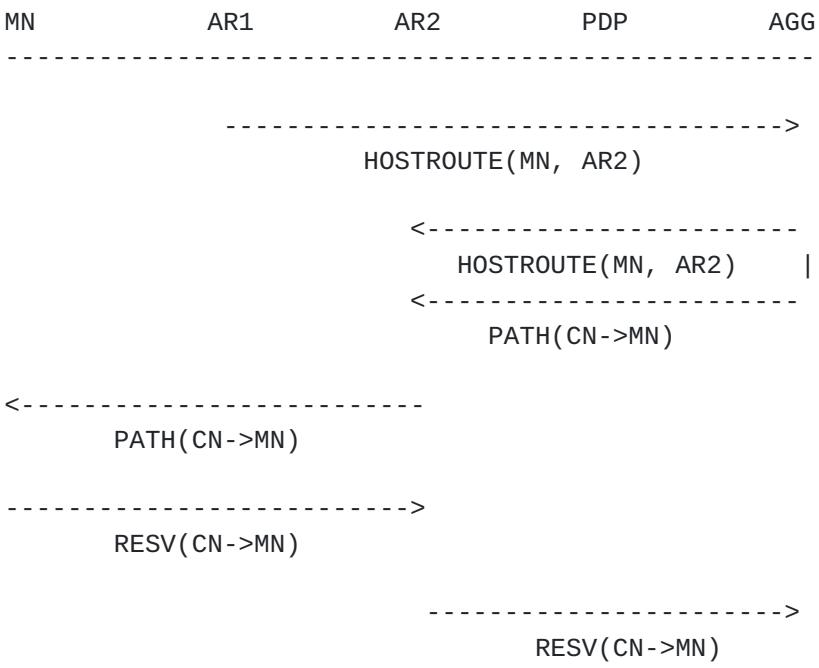

2.3 CN->MN & MN->CN reservations

With concurrent reservations, we can draw the flow as follows. Note that in both flows above, a round trip between MN and AR2 are required to either reestablish state, or conform to standard RSVP.

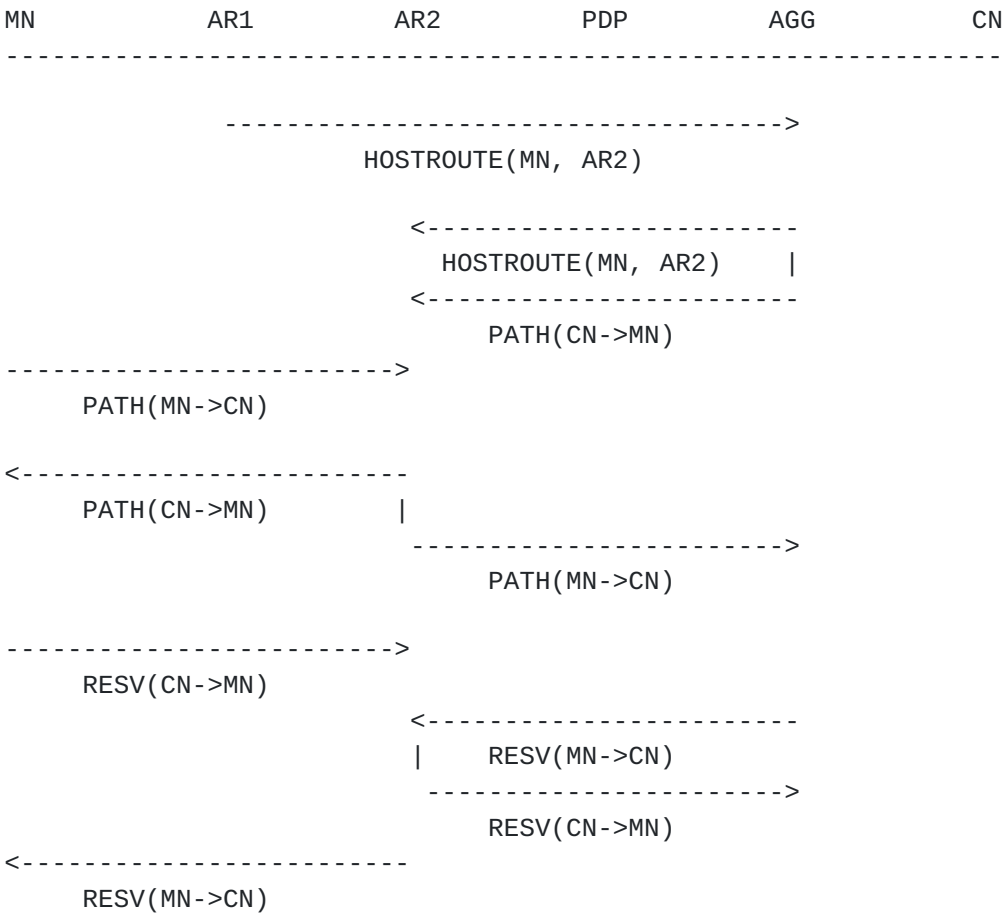

2.4 CN->MN & MN->CN Reservations with Context Transfer

In the context transfer case, we can see that if we assume that AR2 would act as AR1 on a topology change, it would immediately send a PATH message for the MN->CN reservation upon installing the PATH state locally. Assuming that MN is aware of the context transfer, it would need no RSVP signaling to restore the reservation on its new interface.

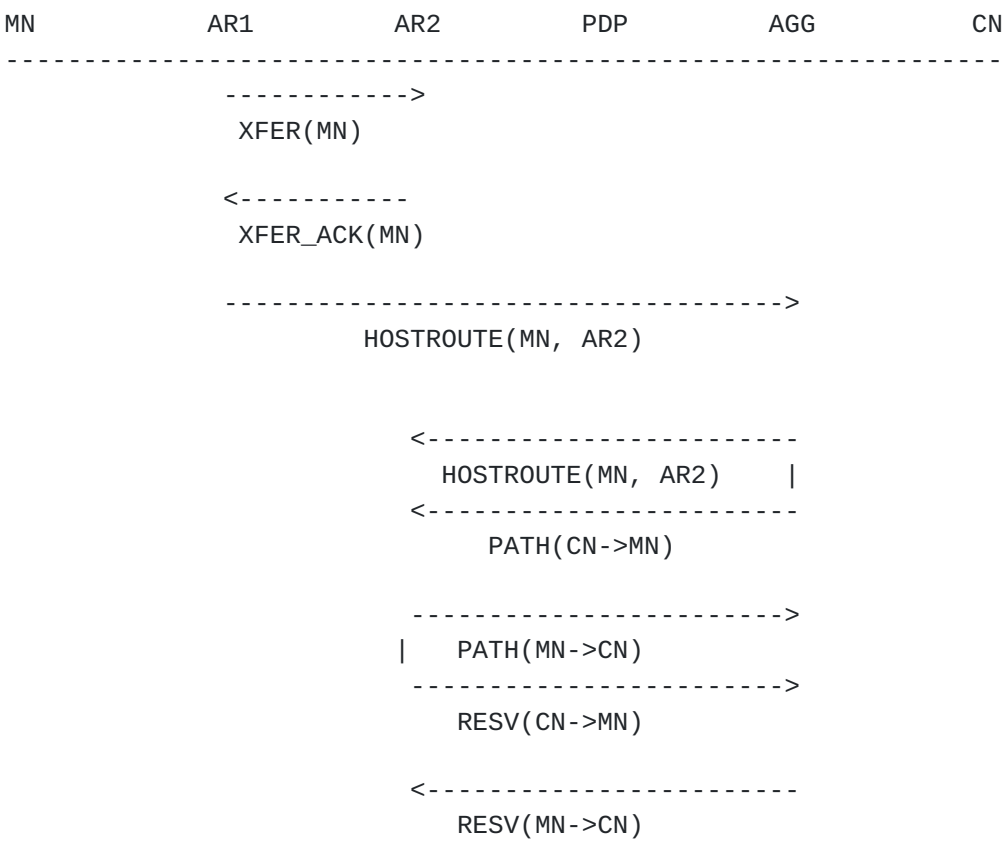

3. Different Care of Address

Here we expect to be able to use a shared reservation for the new CoA, but no other assumptions are made beyond that.

3.1 MN->CN

As we can see, the MN->CN with new CoA require that MN launches a PATH to AR2 which is unware of the existing reservations on AR1. These are aggregated as a topology change at AGG and need not traverse back to CN.

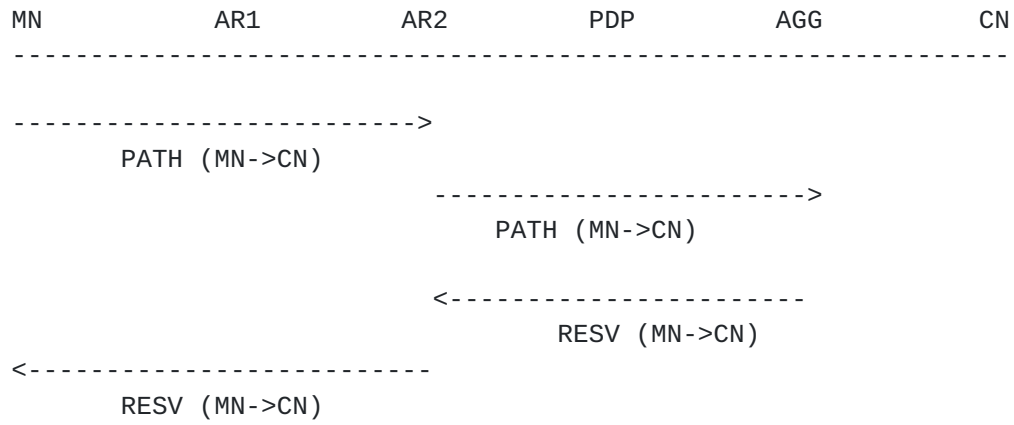

3.2 CN->MN

Since PATH's originate from the sender and only MN is aware of the topology change, the only means of tracing the new path is for CN to launch a new PATH. One obvious trigger might be a binding update back to CN which would alert CN that the reservation's topology has changed.

As can be readily seen, the reservation cannot be reestablished to CN until it receives a binding update.

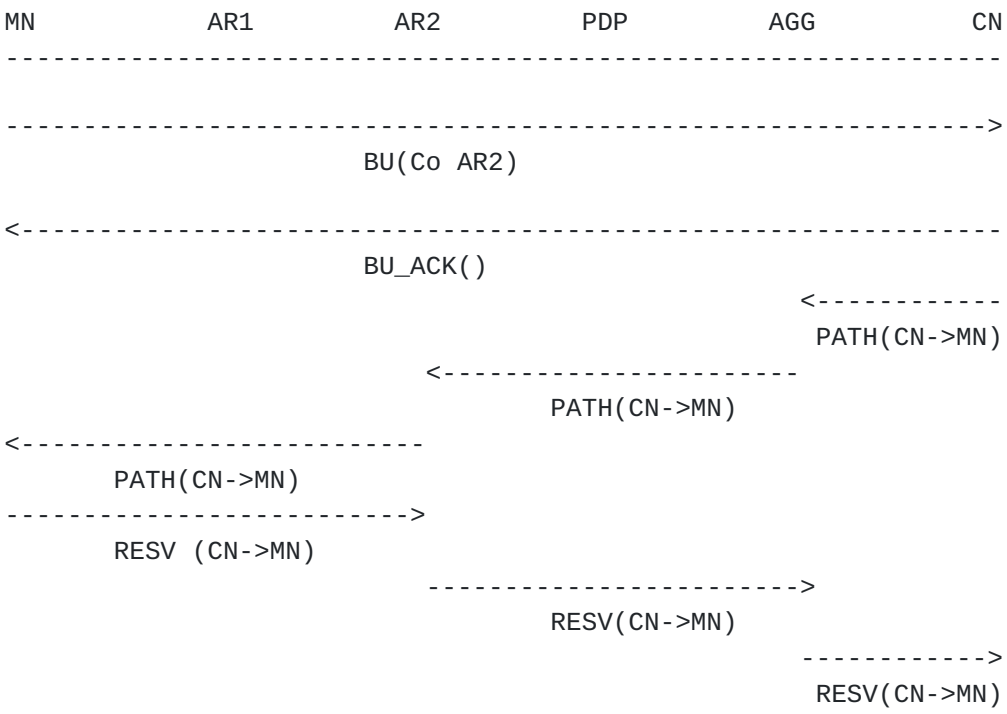

3.3 MN->CN, CN->MN Reservations

Combining the two flows yields the following. Note that there is very little overlap as well as the six messages to/from the mobile node. This is a far cry from our "optimal" flow.

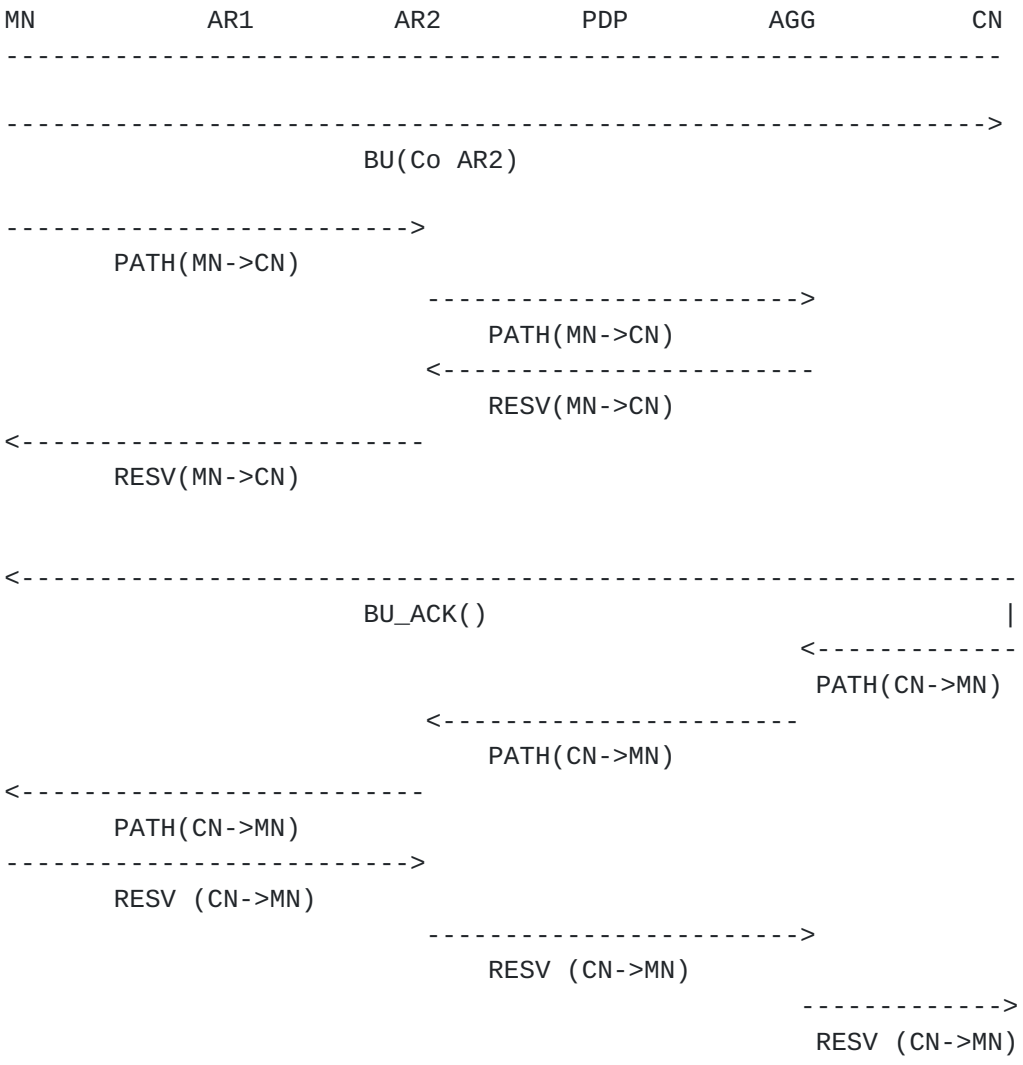

3.4 MN->CN, CN->MN Reservations with Context Transfer

The context tranfer scenario here is better than without context transfer, but still has many problems. For the CN->MN flow, there is still the dependency to use the binding update to CN to reinitiate the PATH message. However, while there were two reservations missing in the previous flow (AGG->AR2, AR2->MN), it is possible that since AR2 had the PATH and RESV state from AR1 for the hop to MN, it could proactively reestablish that portion of the reservation before it actually sees the RESV coming back from AGG.

This seems on its face -- as above -- to be a rather unfortunate and undesirable outcome. The main problem is that unlike the same CoA context transfer, the host route update performed the job of alerting the routers of the topology change. There seems to be no obvious mechanism to do this with RSVP itself.

Note that this flow is silent about any fast handoff between AR1->AR2 tunnel. Section 4 will explore the interaction between fast handoffs, context transfers and mobility, and reservations.

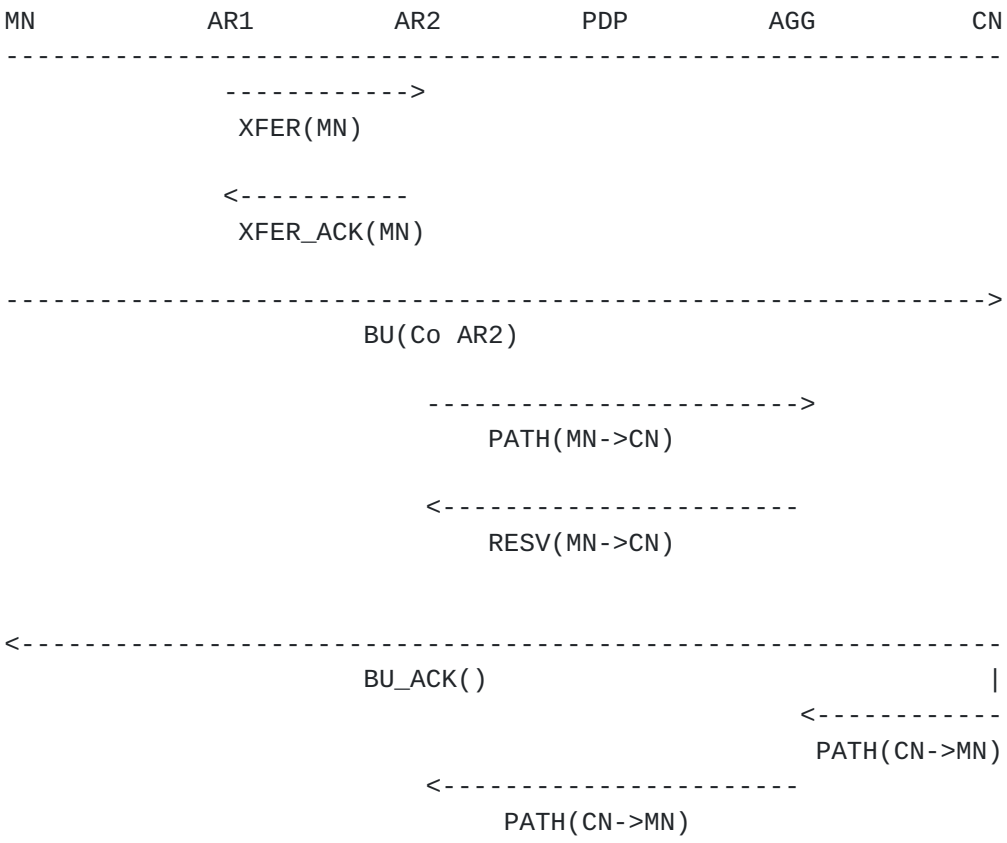

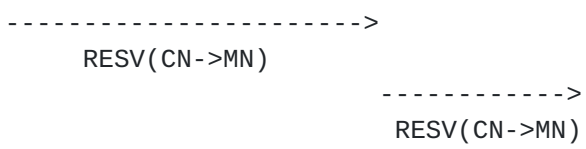

4. Mobile IP Fast Handoffs

It is expected that when a mobile node desires a seamless handoff that it will use the fast handoff work being currently addressed in the Mobile IP working group [refxxx]. While most of the details are unimportant, the pertinent part of the work is that there will be a forward tunnel established for flows originating from the correspondent node to the mobile node. This forward tunnel is identical to the forward tunnel between the mobile node's home agent and the mobile node itself. While the tunnel may in fact be transient between the old access router and the new, there may still be high motivation to make certain that the so-called side-haul traffic has admission control as well.

Because this traffic is only in the direction of the correspondent node to the mobile node and because it is definitionally the different CoA case, we only need to consider two additional cases: RSVP when there is no context transfer, and our hypothetical router context transfer protocol. In this section, we denote the forward tunnel establishment message as a binding update.

4.1 Fast Handoffs without Context Transfer

Fast handoffs as currently being formulated in the MIP working group use a forward tunnel from a mobile IP router that is topologically/geographically close to the mobile node on the old access link (typically the first hop router) toward the mobile node. Since this is a tunnel, RSVP and specifically [RFC2746] treats the as a new interface on the old access router for which the old access router and mobile node must create a new reservation for the forward tunnel in order to preserve the quality of service that the reservation on the direct link was previously entitled to. Note that fast handoff is not involved with how the reservation is reestablish in the direction of CN->MN.

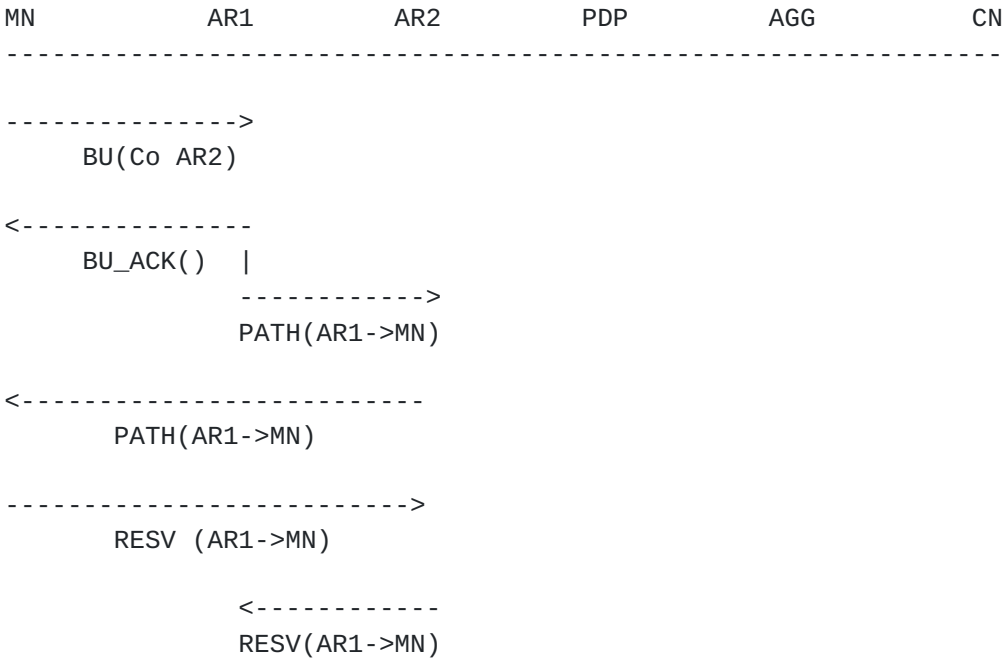

4.2 Fast Handoffs with Context Transfer

With a context transfer, again assuming the mobile is aware of the transfer so that it does not add unnecessary signling. The net effect here is that the mobile only needs to initiate the fast handoff by signaling AR1. As above, AR1 will initiate a PATH message, but since AR2 already knows about the reservation ahead of time due to the context transfer, it knows that the next hop has the same PATH and RESV state. As such AR2 is the join point and does not need to forward the PATH to MN and instead immediately sends the RESV back to AR1.

This has the nice effect that there is minimal signaling between the mobile node and the access routers.

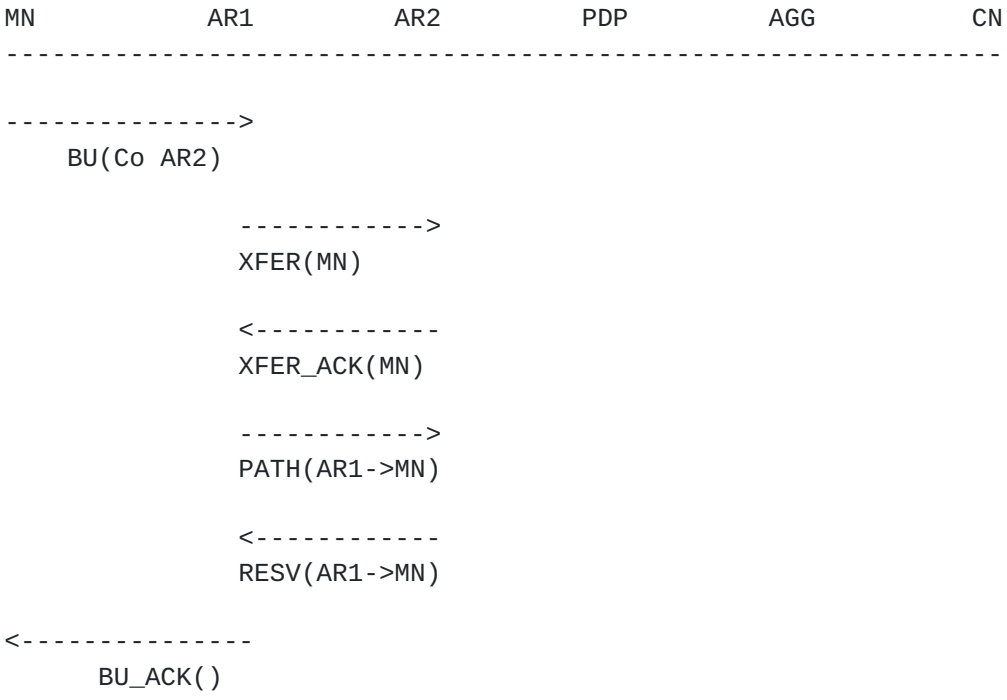

4.3 MN->CN, CN->MN with Context Transfer and Fast Handoff

This combined flow shows the effects of context transfer, fast handoff with both flows. Note that the mobile node need only initiate the transfer, but not do any other signaling on the last hop interface that is in the critical path.

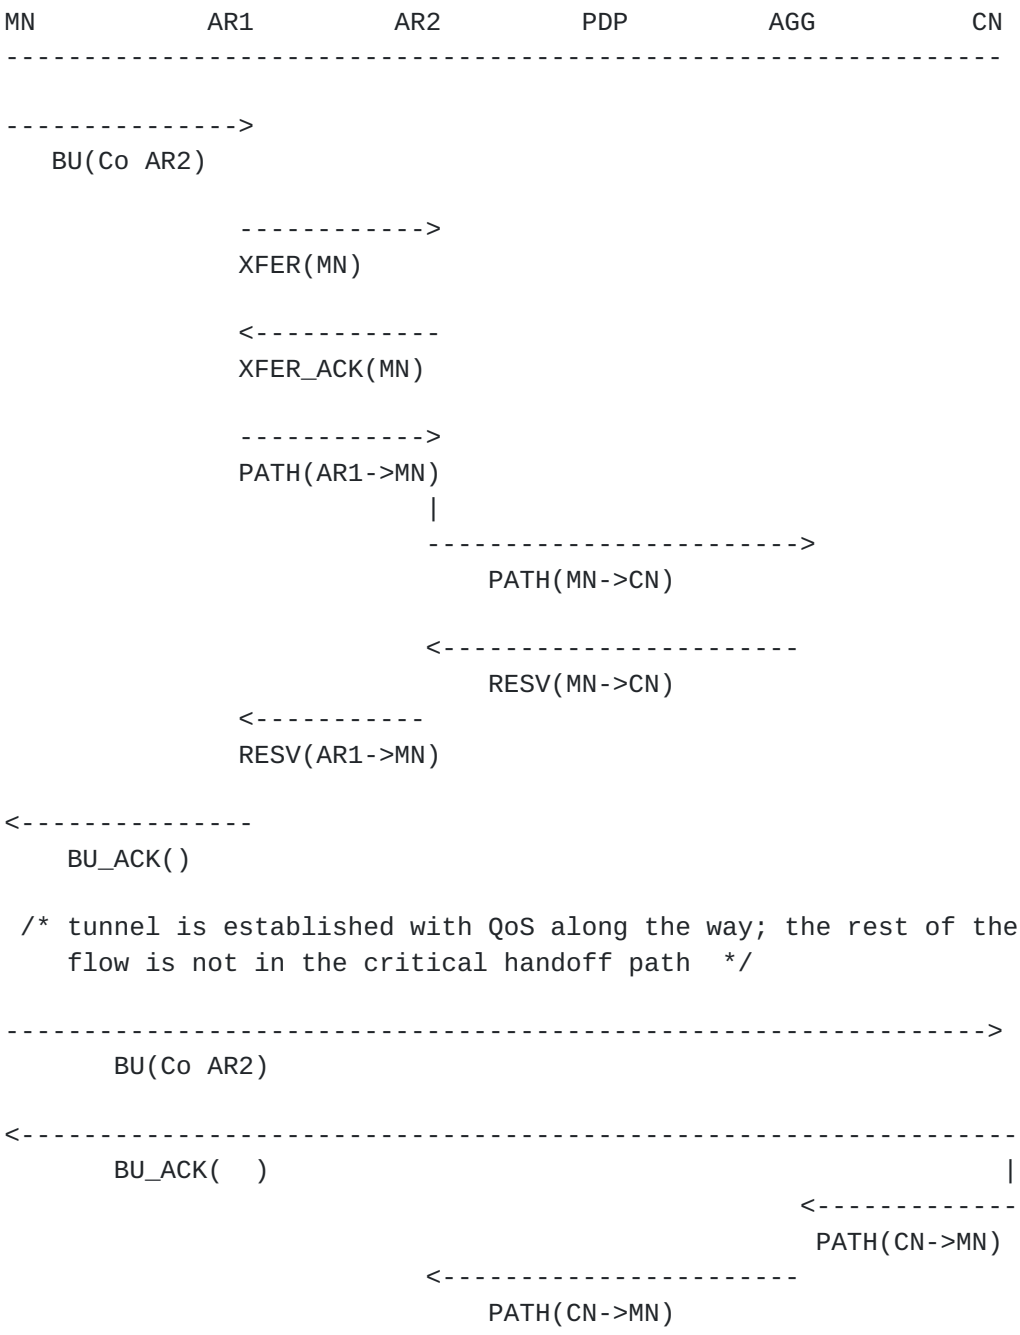

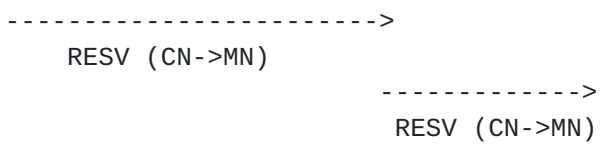

4.4 Reservations Through the Home Agent

In some cases, the correspondent node may be unable or unwilling to perform a binding update. In this case, the reservation will always flow through the home agent in the CN->MN direction. It should be noted that this flow is identical to the Fast Handoff flow, aside from the details of which routers are also affected during reestablishment. This should not be surprising since Fast Handoff is just using the old access router as a "local" home agent to establish the identical kind of forward tunnel that the mobile node's real home agent would create.

Since the tunnel between MN and HA is effectively a virtual interface between the two, the same considerations apply as in [RFC2746], which is to say that in order to restore the reservation on the new acces router between the home agent and the mobile, all of the same considerations apply as in the fast handoff case. Since Fast Mobility is not dependent on how the packets arrive at AR1 (tunneled, not tunneled), packets arriving through the home agent are treated like any other packets destined for the mobile node.

Therefore, fast mobility can be achieved as usual by tunneling the already tunneled packets from AR1 to the mobile node. The only thing that changes is that the filterspec will need to be adjusted to reflect encapsulated data from the Home Agent instead of the actual reservation to the correspondent node, but this is inherent in [RFC2746].

Of course, we end up with two levels of tunnel for the duration of the fast handoff tunnel, but this seems inherent with the design of fast handoffs. Whether the implied bandwidth overhead is too burdensome is beyond the scope of this draft.

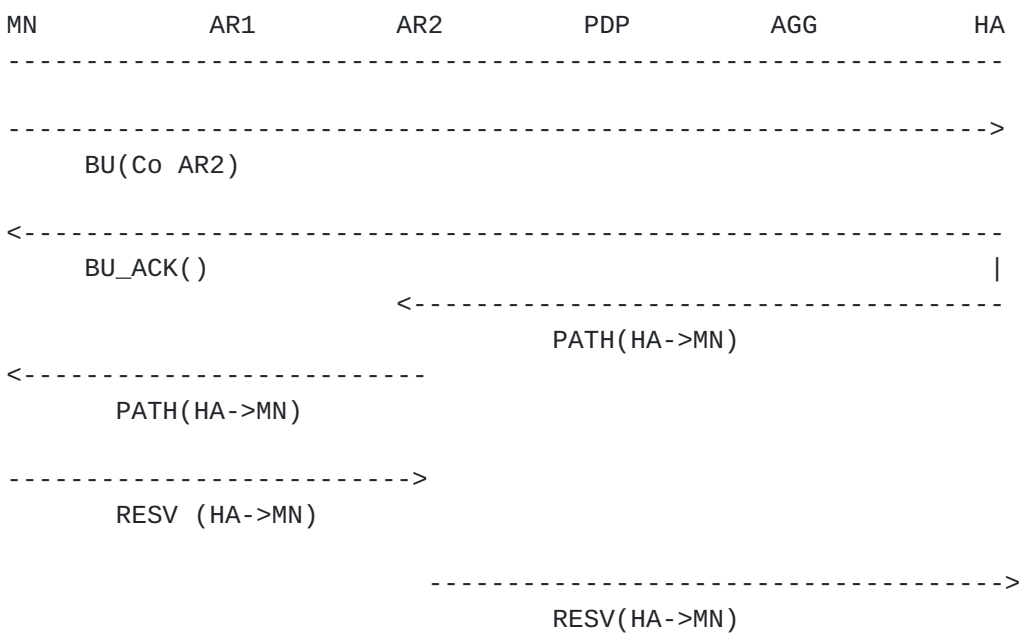

5 Analysis

At first glance comparing the flow in section 2.4 and 4.3 the with the "optimal" flow, seems anything but optimum. However the "optimal" only gave a minimum amount of signaling which would be required; it did not assign any metrics to deviations from optimal. In this section, we shall try to assign some numeric quantities to determine how far each flow is from optimum.

To do this, we assign a numeric times to each link, These numbers are arbitrarily chosen here and attempt to express the relative cost of sending a message and its associated latency in milliseconds. For any particular technology these numbers should be adjusted. Since the end goal is seamless mobility, we define the total numeric cost as the time between the initiation of a handoff and the when a functional bidirectional reservation is in place between the mobile node and the correspondent node care of the new access router.

Costs:

- Tlast = MN<->AR: 20ms. Last mile is always s l o w.
- Tagg = AR<->AGG: 2ms. Assumedly fast ethernet, etc
- Tarar = AR<->AR: 5ms. perhaps slower than Tagg
- Tcn = MN<->CN: 100ms. This is a very arbitrary number; we assume

CN is very far away

5.1 Optimal Flow

In our fictional flow, we only have 3 messages between AR and AGG. Therefore our calculation is:

Tagg * 3 = 2*3 = 6ms.

5.2 Optimized Real Flows

Section 2.4 is the same CoA case with context transfer. It's calculation is:

Tarar*2 + Tagg*6 = 5*2 + 2*6 = 22

Section 3.4 is the different CoA case with context transfer, but no fast handoff:

Tarar*2 + Tcn*2 + Tagg*4 + Tlast*2 = 5*2 + 2*100 + 2*4 + 2*20 = 258

Section 4.3 is the different CoA with context transfer, but with fast handoffs. Note that we do not count the cost of the round trip to the CN and associated messaging since it is not in the critical path:

Tlast*2 + Tarar*4 + Tagg*2 = 40 + 20 + 4 = 64

5.2.1 Optimized Conclusions (or v.v.)

The general results should not be completely surprising. The fictional is always best as fiction is wont to be. Keeping the same care of address is the next best thing though as the amount of time to propogate the host route is relatively small and it doesn't involve any communication with the mobile or correspondent nodes. The cost of even using context transfer without fast handoff is unsuprisingly bad as it requires a potentially expensive round trip back to the correspondent node. For real time mobility, this will undoubtably be unacceptiable. For whole enchilada, we find a middle ground. The main difference between the same and different CoA cases is actually in the need for a round trip on the slow last mile interface. Assuming that not many more of these slow round trips are not necessary, however, this may prove a fast enough for most cases. It should be noted that proactive handoffs from AR1 to AR2 are not accounted for here specifically, but ought to produce somewhat better results in the different CoA case.

5.3 Other Flows

For completeness sake, the rest of the flows are listed here.

/* some day */

6 Mobility and Interdomain Policy

[RFC2749] and [RFC2753] provide a framework for policy based admission control for RSVP. The general case is that there may be many policy enforcement points (PEP) as well as many policy decision points (PDP). In general, it is envisioned that PEP's will be strategically placed at the borders of RSVP enabled networks to enforce policy based admission control, though there is an implementation detail. Although a reservation may flow through any number of differently administered RSVP domains, for simplicity's sake we will limit the number of administrative domains each reservation traverses to two for policy decisions as well as the domains that the reservation crosses.

We define a cross realm reservation as one where the realm of the current set of policy objects will either require adjustment by the sender to procure the proper policy objects, or expect that one of the upstream routers supply a new policy object which will allow the reservation to pass the next PEP.

To complete the scenario for mobility, we must consider at least one other possibility: that the mobile node may move from one visited network administrative domain to another. For simplicity, we consider that the PDP that the mobile node is authorized by is in the same domain as its home agent, though in reality there need be no such linkage.

The following diagram illustrates the various networks and their associated PDP's. We will consider all of the reservations this draft in turn.

The following conventions are used in this diagram and the rest of this section:

box Boxes denote various networks which are administered separately.

AR an access router which is a PEP on the ingress side for reservations. That is to say, it is a PEP only in the direction coming from the untrusted interface, eg the mobile node.

BR a border router which is a PEP on the ingress from another administrative domain typically toward the big I internet.

PDP a policy decision point for a given administrative domain.

MN a mobile node which is in the same administrative domain as the home network.

CN a node which is in the same administrative domain as the correspondent network.

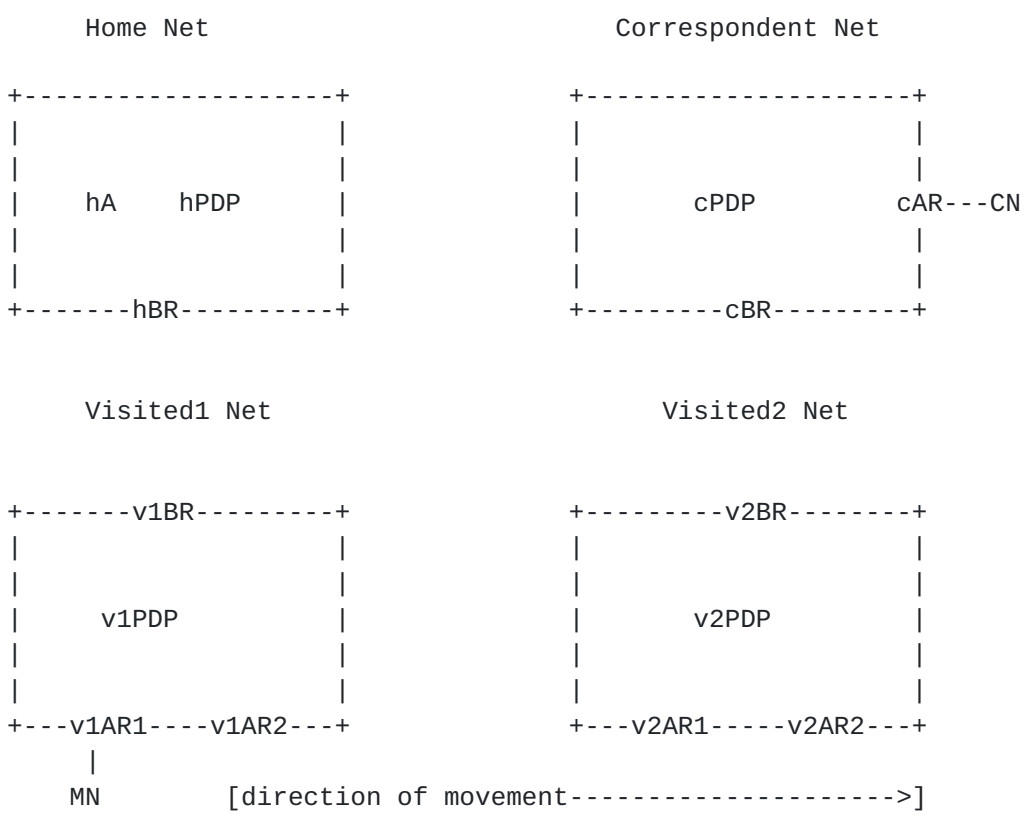

6.1 Bilateral vs Hierarchical Cross Realm Policy

[RFC2753] describes two possible approaches to interrealm policy based admission control. The first method requires hopwise bilateral arrangements between the administrative domains. That is, as a

reservation's RESV crosses administrative boundaries, the appropriate policy object which should grants admission is placed in the RESV. An RSVP router may add to, replace, or delete policy objects as the reservation passes by. Using this method, routers would normally rewrite a new policy object acceptible for the next upstream PEP.

The second method -- somewhat obliquely refered to in section 5.4 of [RFC2753] -- uses cascaded PDP's such that if a PDP cannot make a policy decision locally because it is not in the same realm as the current policy object, it can relay the policy request to a PDP in the realm of the policy object for a decision. This method also allows for the establishment of a hierarchy of PDP's which eventually relay the policy request to the final PDP for a decision.

The following section will describe the tradeoffs at work, and what the specific problems are with respect to mobility.

A specific example is in order here. Suppose that there exists a mobile node which is serviced by AT&T and has a certificate to prove that it is mikes-phone@attwireless.net. mikes-phone@attwireless.net then attaches itself to MCI's network and want to instantiate a reservation -- say for a phone call. The only certificate that the phone posseses is the one that AT&T installed as part of the service agreement. The phone will therefore, place that certificate with a digital signature into a policy object when it needs to send the RESV. While it is possible for the phone to have multiple different providers (and hence) multiple certificates, this is not scalable and therefore the general case is where the phone does not directly possess a credential to pass each PEP/PDP it might encounter for the reservation.

As normal, MCI's PEP will forward the RESV to its local PDP for a decision. While it's not too far a stretch of the imagination to believe that MCI's PDP could authenticate mikes-phone@attwireless.net -- assuming it shared and trusted AT&T's root certificate -- there still remains one very large problem: MCI's PDP is completely clueless about what mikes-phone@attwireless.net is _authorized_ for. This brings up an interesting question about what MCI's PDP does when it sees a cross realm request from its PEP. MCI's PDP could:

1) Expect AT&T's subscriber database and policies are replicated at MCI's PDP. 2) Authorize anything it can authenticate 3) Expect that there is policy/authorization information in the certificate that AT&T gave me and that the PDP knows how to act on it 4) Relay that (now authenticated) request back to something that does understand the policy

(1) is, of course, a joke that has no way of scaling, and would most likely not be be acceptible even if it were scalable.

Clearly, (2) is pretty suboptimal if for no other reason than the fact that AT&T has no way to revoke the certificate in anything approaching real time if the service is canceled. Also, any kind of restrictions AT&T wishes to place on the phone's service would be unknown to MCI's PDP.

(3) has two main problems: first, it would require a common policy scheme that is codified in the certificate which both providers understand. This leads to a least common denominator problem, and would most likely lead to a huge interoperability problem as the bilateral agreements end up special casing various policies amongst themselves.

Secondly -- and probably more serious -- is that policy as captured in the certificate is effectively static. Any changes of policy would require that the old certificate be revoked and a new certificate issued. This may be OK when the policy is relatively static and simple but if we want to allow a dynamic policy, (eg, calling cards) a static policy mechanism is not going to be well suited.

(4) provides the most dynamic and fine grained policy control. Revocation isn't much of an issue because authentication doesn't automatically translate to authorization for any services. Also: since the policy logic is self contained within AT&T's PDP and it has access to all of the necessary subscriber information, doesn't rely on any krufty less-than-grand-unified policy schema. Another nice side effect is that the authentication relationship needed is between the phone and AT&T's PDP which are not coicidentally in the same realm.

(4)'s tradeoff, of course, is that it requires a potentially expensive round trip to AT&T's PDP for decisions. One mitigation is that AT&T's PDP could allow the decision to be cached at MCI's PDP, but the problem never completely goes away. However, in the specific case of mobility and fast handoffs, as we shall see in subsequent sections, the round trip penalty may not be a catastrophic penalty. The reason is that there is other messaging as is evident in the previous sections which make a very strong case for Fast Handoffs and Context Transfer.

Thus to simplify this already too-weighty draft, we will focus on (4) as it seems like it provides the proper framework for mobile nodes as they are most likely to be deployed. This is not to say that any of possibilities are not viable in, perhaps, more limited situations. This draft does not attempt to pass judgement on their feasibility elsewhere.

6.2 Assumptions

The following assumptions are made here:

- All policy based admission control must use some form of proof of identity

- The crypto used to assert identity in the RSVP RESV requests from the MN and CN are digitally signed using a X.509 certificate with which the node and PDP share a common root CA. Other schemes such as Kerberos [RFC1510] and plain text passwords may be used as well, though it is not guaranteed they will produce identical results here.

- That the PATH, PEP and PDP's in the PATH are not known in advance.

- Original policy objects in the RESV's may be deleted and/or new policy objects may be added to the RESV.

- PDP's have pre-existing relationships with one another and function as the settlement point for accounting by either using bilateral agreements, or some clearinghouse arrangement ala [OSP]

6.3 Initial Case

The initial case for policy based admission is where the mobile node is on Visited Network 1 with a bidirectional reservation with the correspondent node CN.

6.3.1 MN->CN

In this case, the correspondent node supplies the policy objects to gain admission. Note the explicit use of cascaded PDP's to defer the decision back to the PDP in the originating administrative domain.

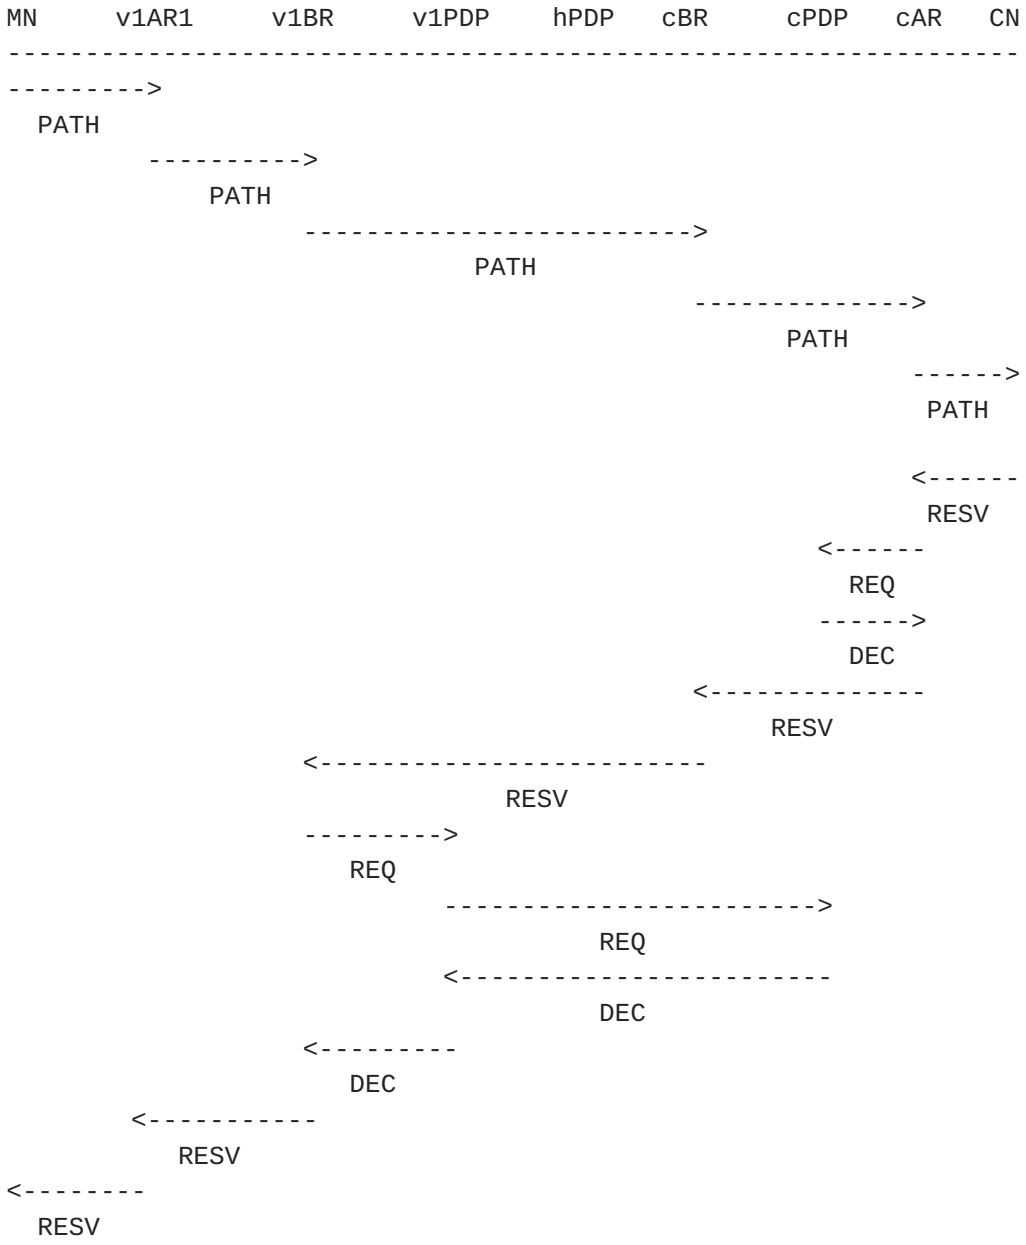

6.3.2 CN->MN

In the CN to MN case, the MN is the one that needs to provide credentials to the PDP's,

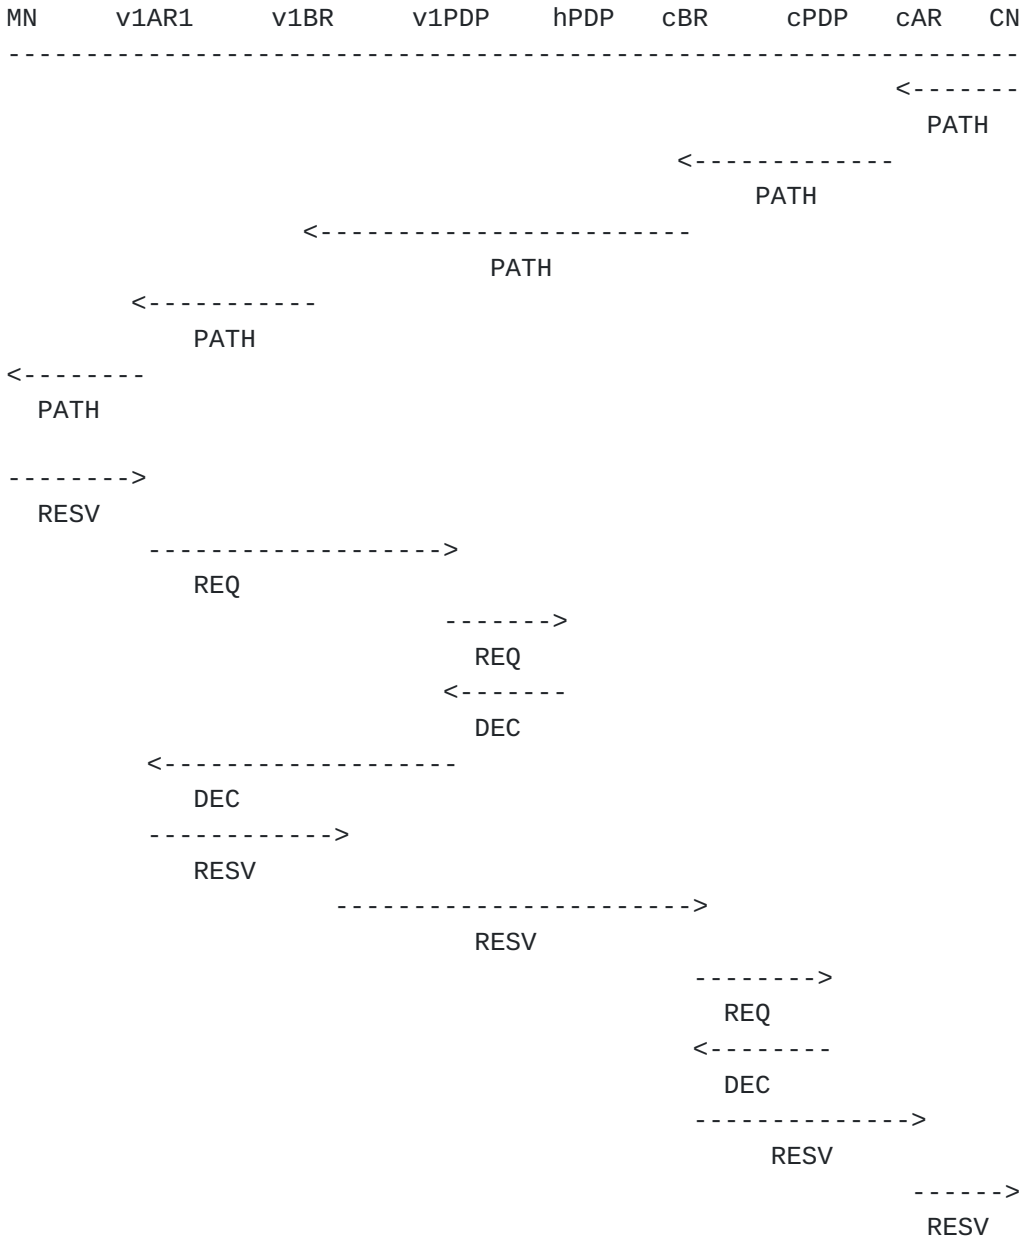

6.4.1 MN->v1AR2, MN->CN Direction

If the mobile moves from v1AR1 to v1AR2 in the normal mobile case, it must initiate a PATH message to trace the new route back to CN. Note that this is no different than the flow in 6.2.1 with the exception of v1AR1 being replaced by v1AR2; no binding updates are necessary.

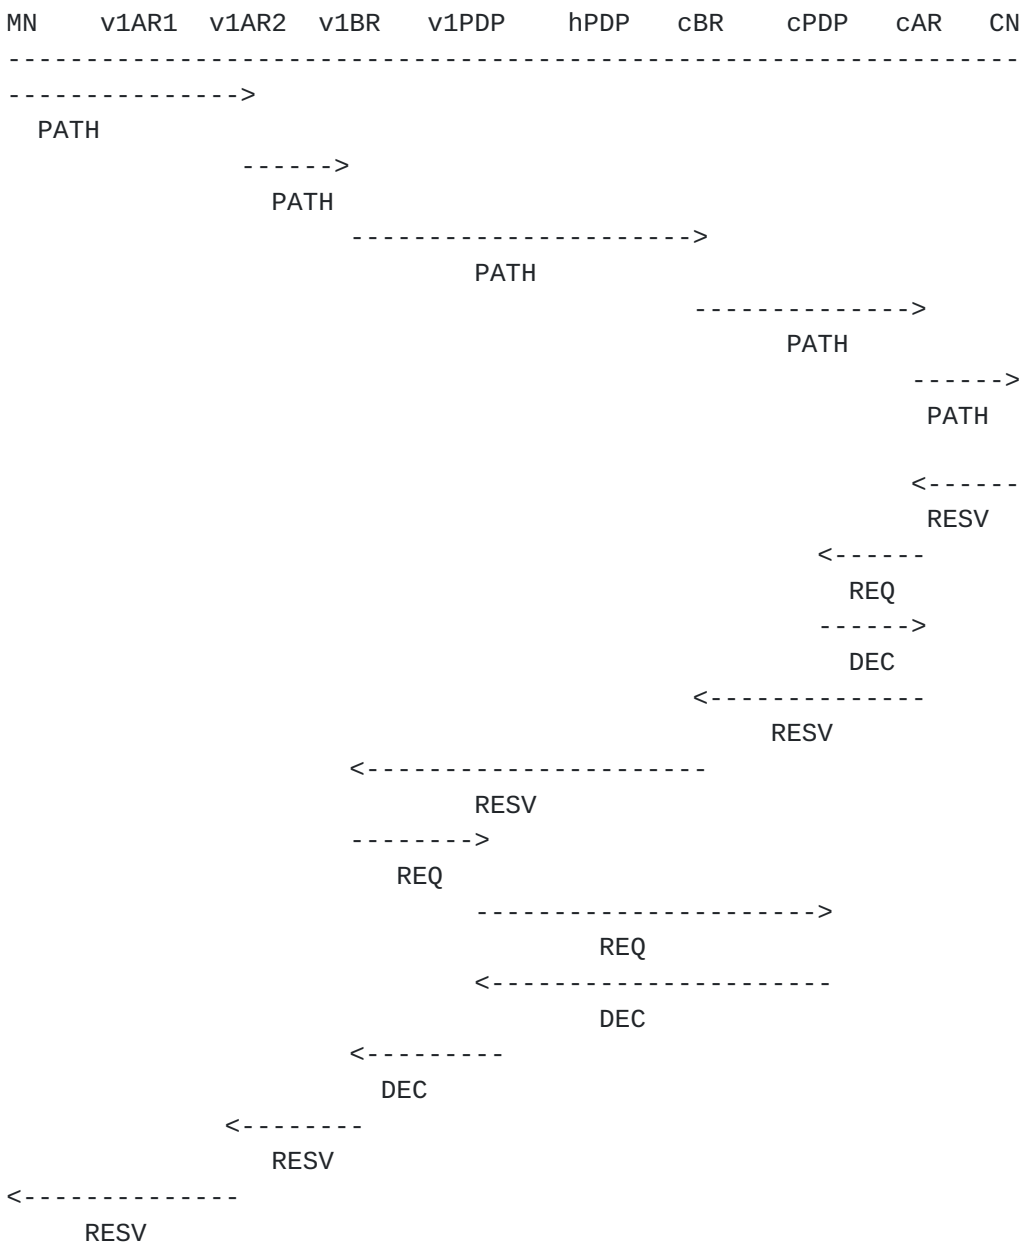

6.4.2 MN->v1AR2, CN->MN Direction

The only difference with this flow and the flow in 6.2.2 is that a binding update to CN is need to stimulate the CN to retrace the PATH for the reservation

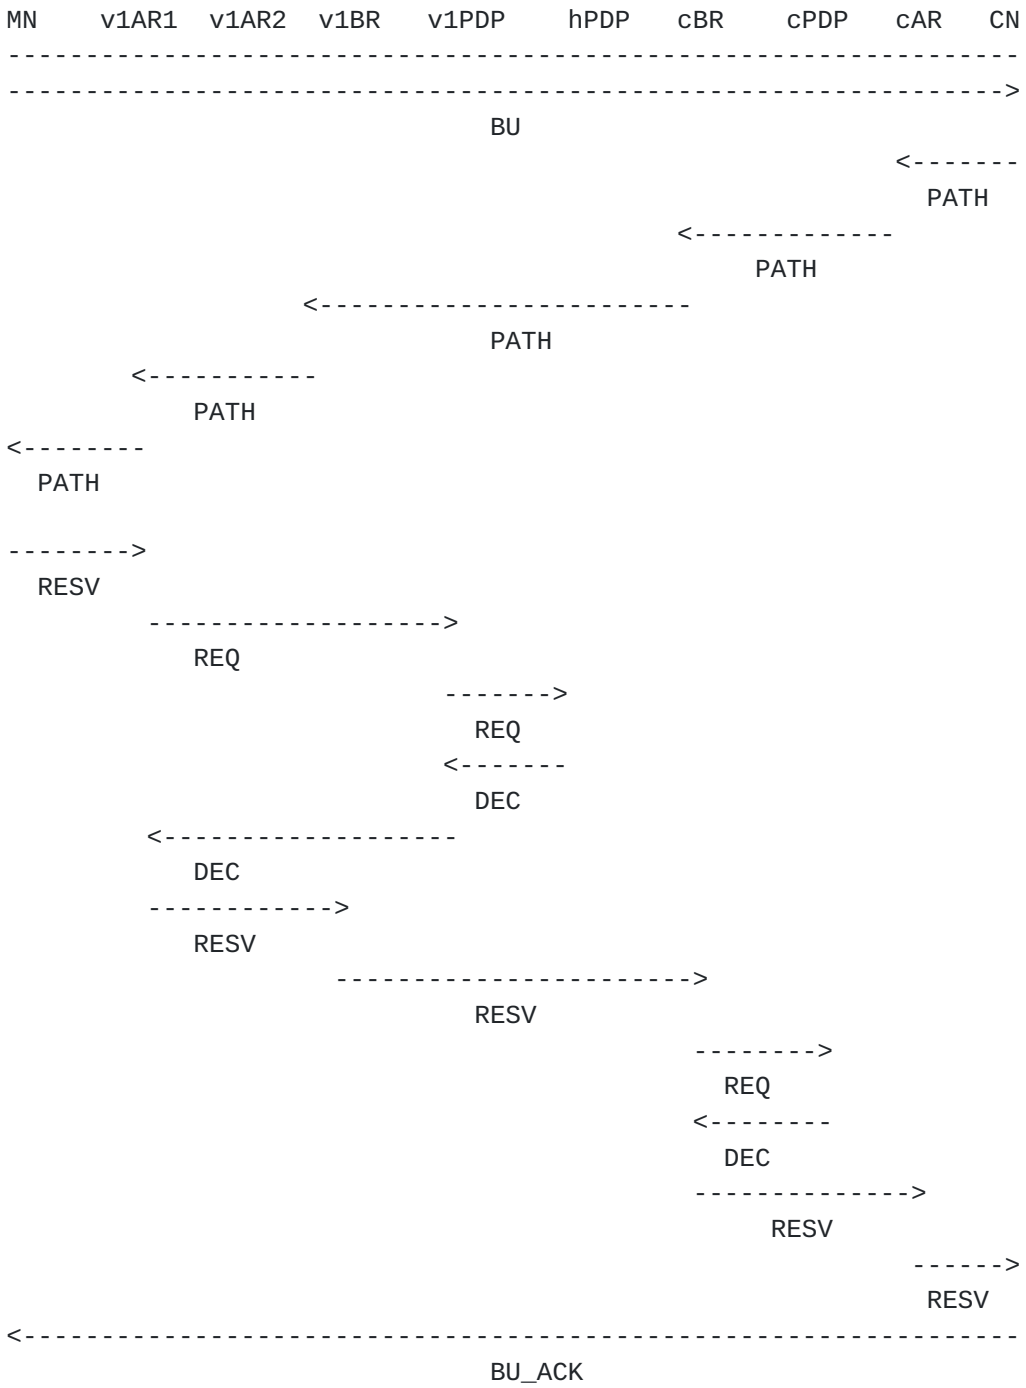

7 Same Care of Address

With policy, the same care of address case has fewer scenarios because it does not deal with the case of the mobile node traversing administrative domains because we expect that that will require a different care of address by definition. For the sake of simplicity, v1BR is colocated with the AGG functionality in previous sections.

7.1 MN->v1AR2, MN->CN Direction

Note that in this flow the routing change is completely localized within the trusted part of the visited network, therefore none of the PDP's need to be consulted.

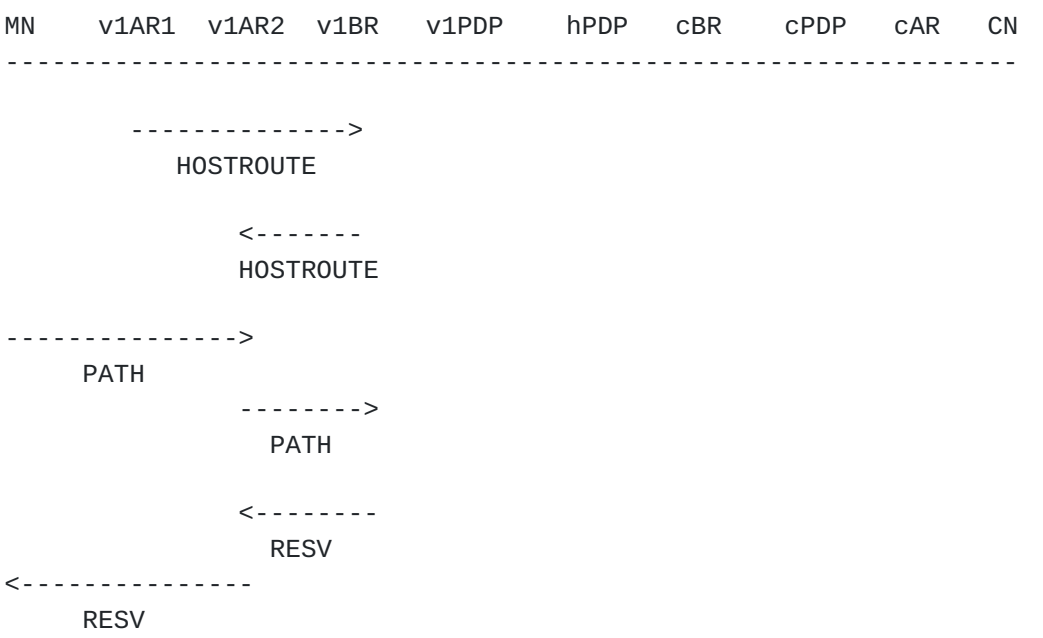

7.2 MN->v1AR2, CN->MN Direction

In this flow the mobile node must present a policy object to gain admission since it must pass the v1AR1 PEP. This is clearly suboptimal as it requires a potentially costly round trip to MN's home PDP.

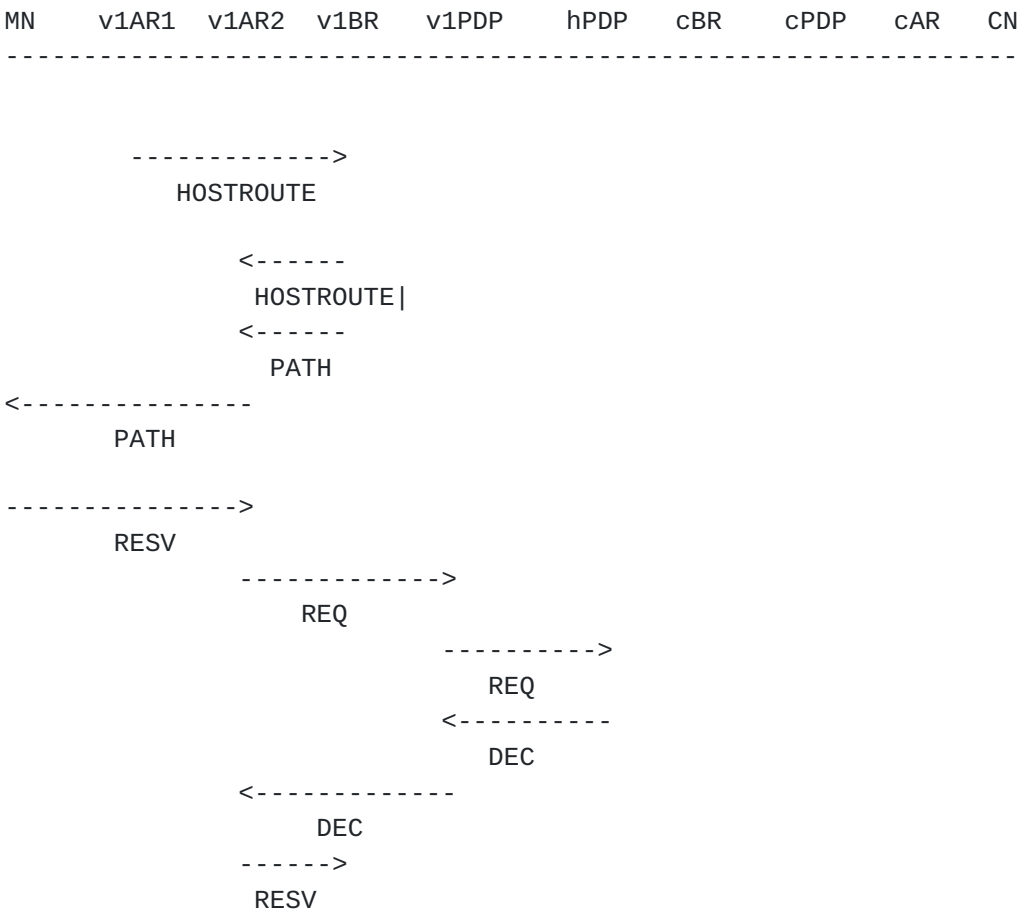

7.2 MN->v1AR2 with Context Transfer

If the context were to be transfered from v1AR1 to v1AR2, we should be able to forego the mobile node from having to submit its credentials for hPDP since it was already authorized by hPDP at v1AR1 which is within the same administrative domain. Note that further RSVP messages from MN will require that v1AR2 share a key with MN to sign the Integrity Object, but assuming that a shared key for v1AR1 was transfered the same key could be used for both access routers. How MN and v1AR1 came to share a key is outside of the scope of this draft.

Note that this flow is, in fact, identical to the flow in 2.4, which of course is a good thing.

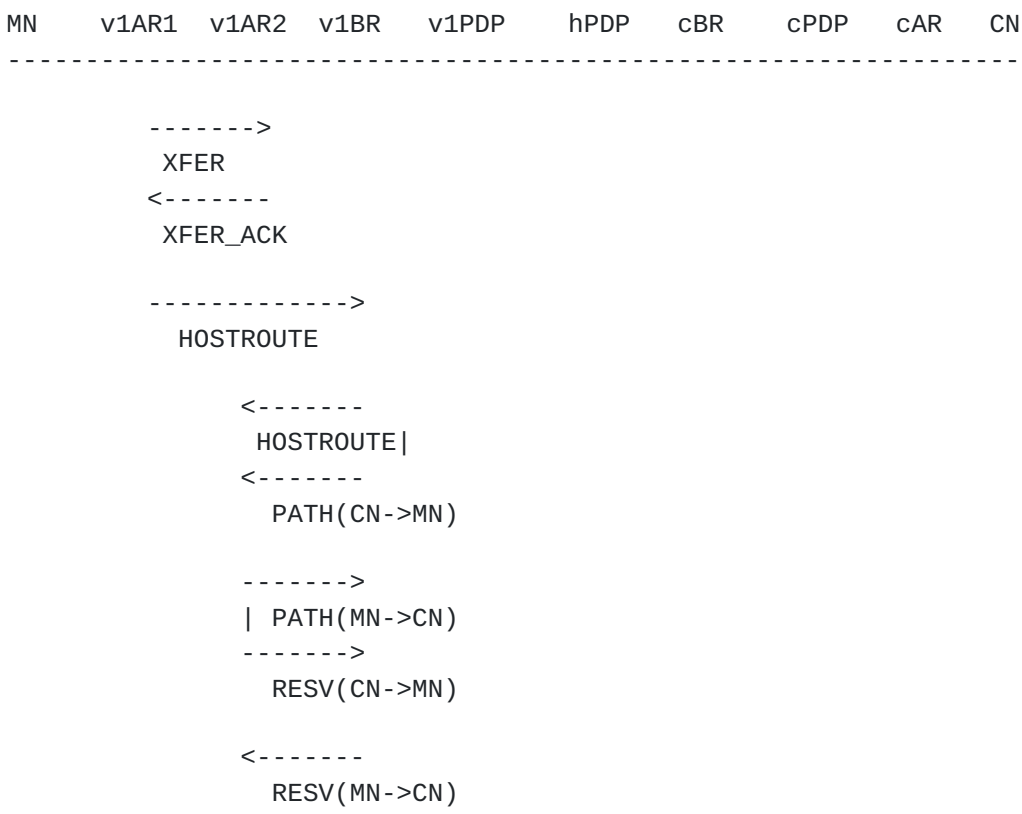

8 Change of Care of Address

With normal mobile IP, reestablishment of the reservation isn't any different than the initial case. That is to say, the reservation must be fully reauthorized as if it were a brand new reservation. As such we can combine the case where we cross administrative boundaries since the flows are identical.

8.1 MN->v1AR2, MN->CN Direction

If the mobile moves from v1AR2 to v2AR1 in the normal mobile case, it must initiate a PATH message to trace the new route back to CN. Note that this is no different than the flow in 6.2.1 with the proper transposition to v1AR2 to v2AR1.

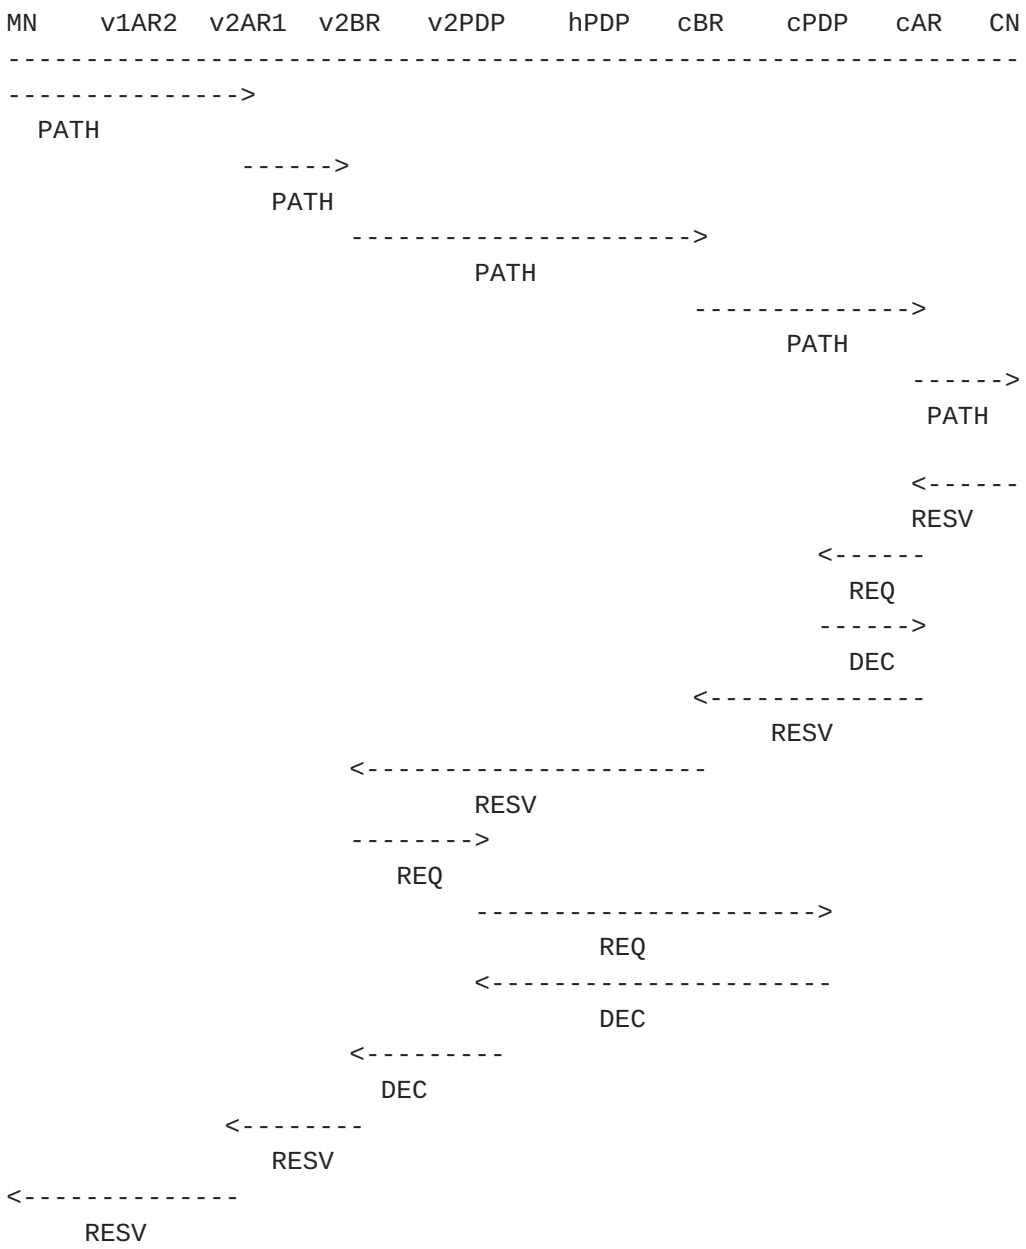

8.2 MN->v2AR1, CN->MN Direction

Likewise as in 8.1, the only difference with this flow and the flow in 6.2.2 is that a binding update to CN is need to stimulate the CN to retrace the PATH for the reservation as well as the same transposition of v1AR2 to v2AR1 as above.

Obviously, this flow leaves a lot to be desired since it must incur serveral expensive round trips to distant PDP's.

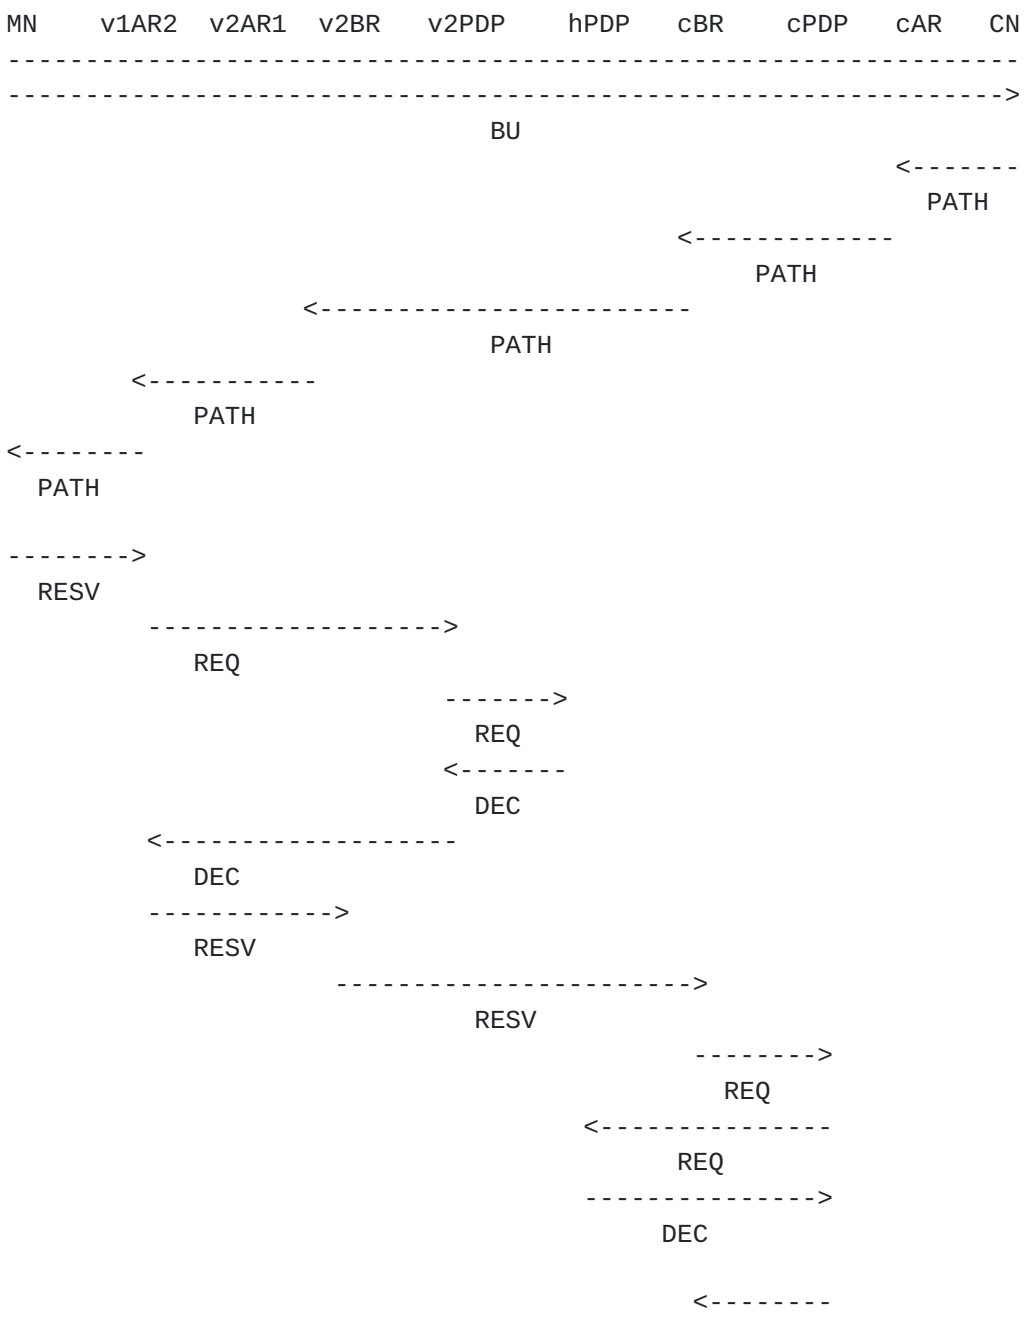

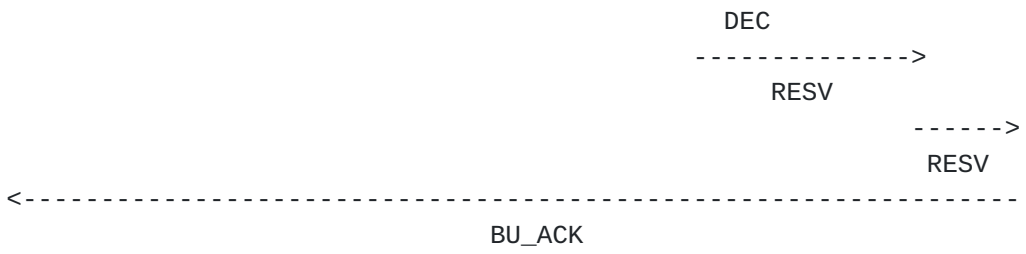
DEC
RESV
RESV
BU_ACK

8.3 MN->v1AR2 with Context Transfer

In a care of address change, context transfer may help the situation somewhat. Since it isn't clear whether context can be transfered across administrative domains, this flow shows the case where context transfer ought to work.

As with other flows that do not use fast handoff in the preceding sections, this flow shows that the MN->CN direction benefits from the transfer, whereas the CN->MN direction is still bound by a round trip to the CN as well as round trips to the home PDP since the correspondent node's PEP is unaware of the context transfer and will demand credentials as usual at its border, though not quite as bad as 8.2 since the v1AR2 will be able to act as a local PDP since it has the context from v1AR1.

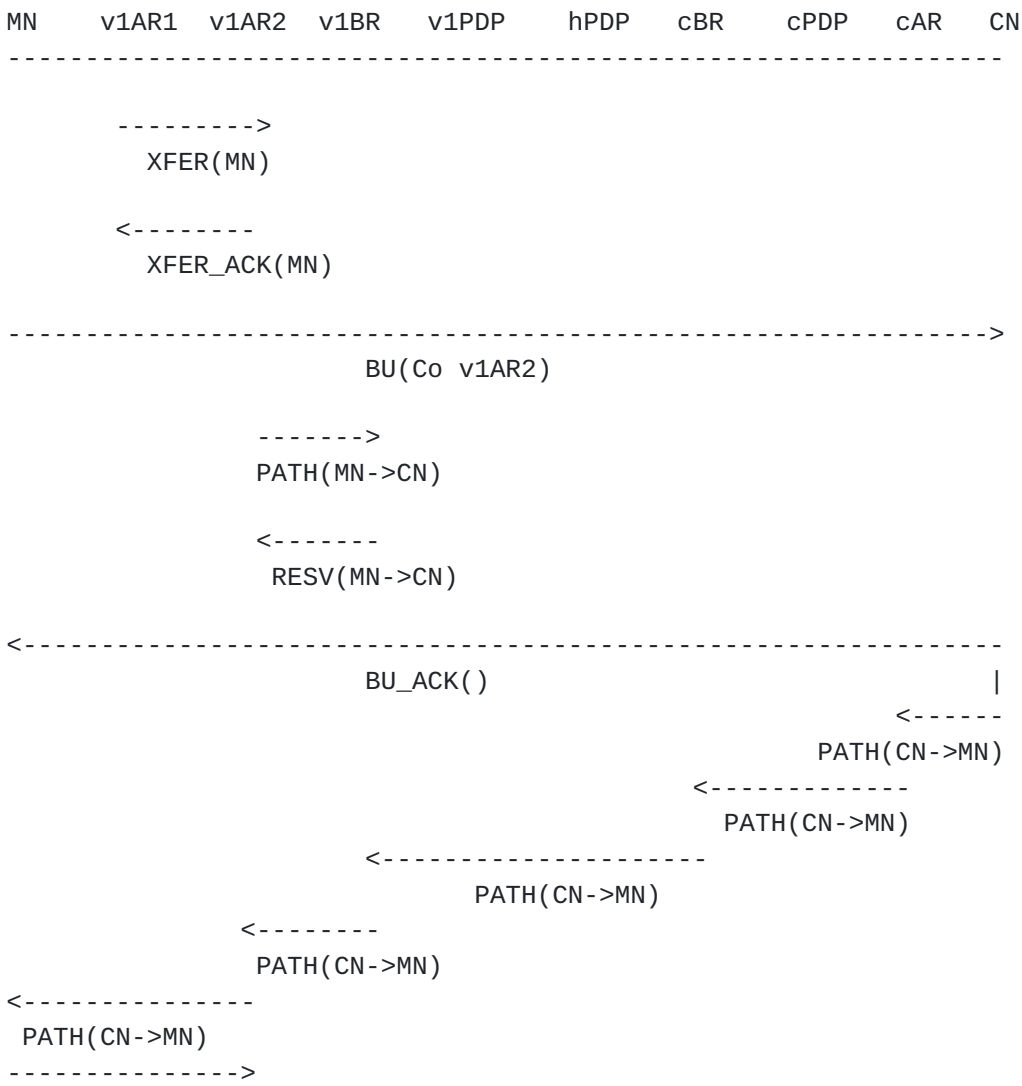

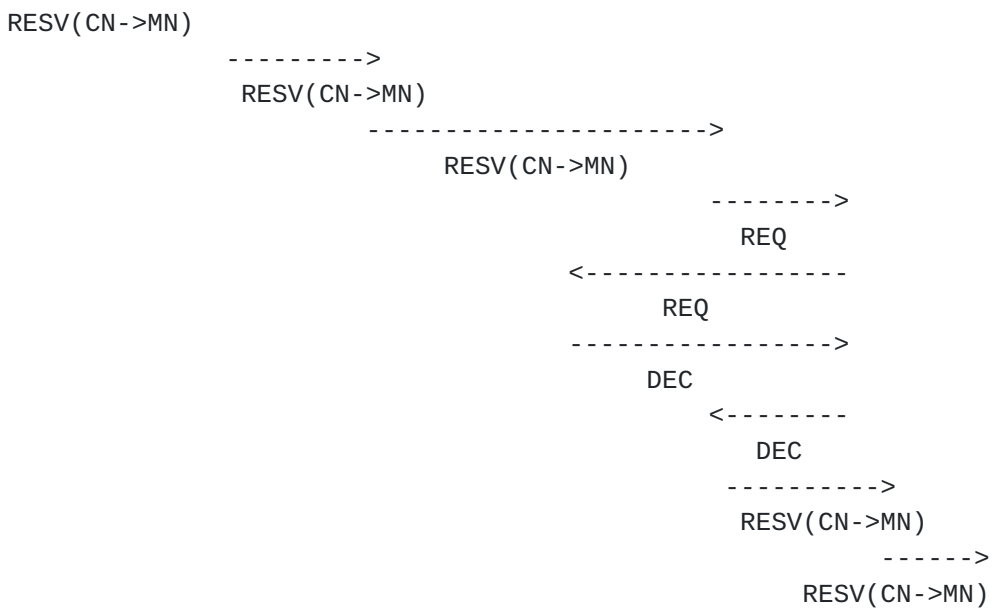

9 Fast Handoffs with Policy Based Admission

Fast handoff is clearly need to bridge the latency of the flow in section 8.3; policy based admission only exacerbates an existing dreadful situation. The following sections examine how Fast Handoffs interact with policy based admission along with some conjecture as to how Context Transfer may improve the situation.

9.1 Fast Handoff with Policy

For a fast hand we establish a forward tunnel between v1AR1 and v1AR2 which requires that MN reauthorize with v1PDP and hPDP. This is potentially a long round trip which is in the critical path of setting up the tunnel which will field the traffic from the correspondent nodes during the transition.

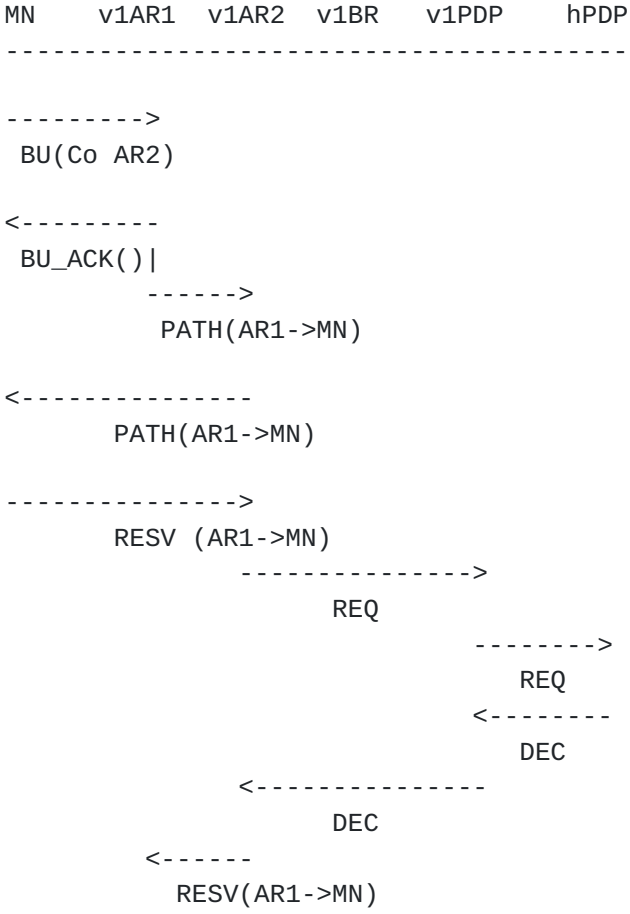

9.2 Fast Handoffs with Context Transfer

There is a _major_ assumption here that existence of the reservation from v1AR1 to v1AR2 would also allow v1AR1 to authorize the reservation for the forward tunnel to v1AR2. This is, perhaps, a dubious assumption and does not seem like the way that the COPS policy framework is intended to work, so this is all arguable. With a stateful v1PDP, or v1AR1 acting as a local PDP it may be acceptible as a local policy to give a grace period to mobile nodes so long as their other reservations were admitted.

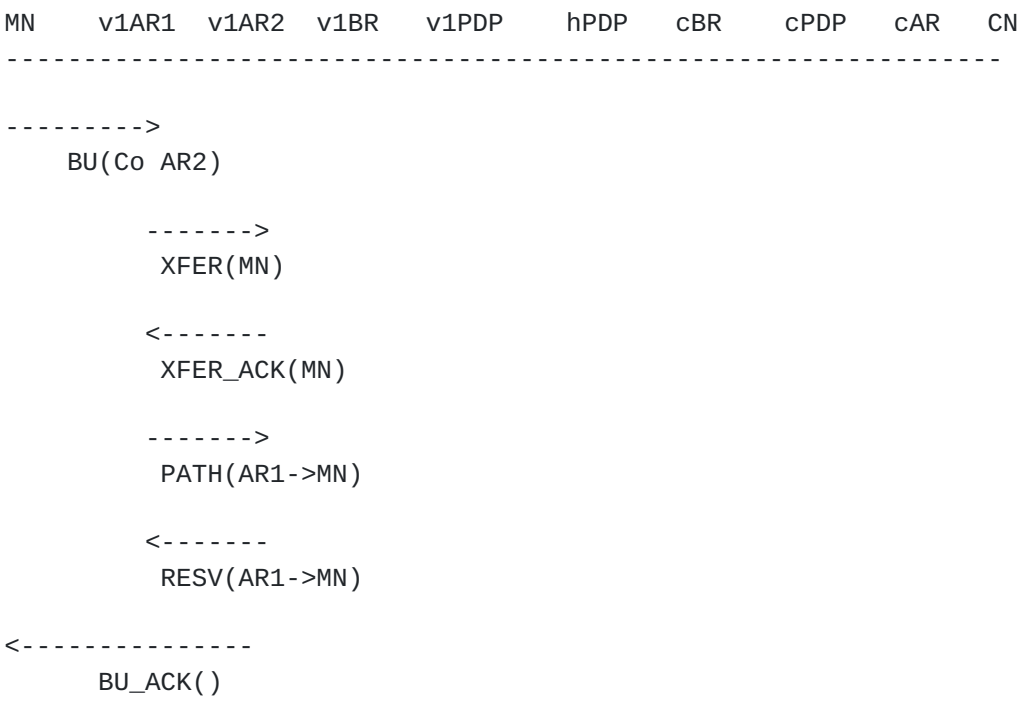

9.3 MN->CN, CN->MN with Context Transfer and Fast Handoff

This combined flow shows the effects of context transfer fast handoff, and policy based admission control with both flows. Note that the mobile node need only initiate the transfer, but not do any other signaling on the last hop interface that is in the critical path. Other than the addition of PDP traffic on the reestablishment of the CN->MN path after the fast handoff was made, this flow is otherwise identical to the flow in section 4.3.

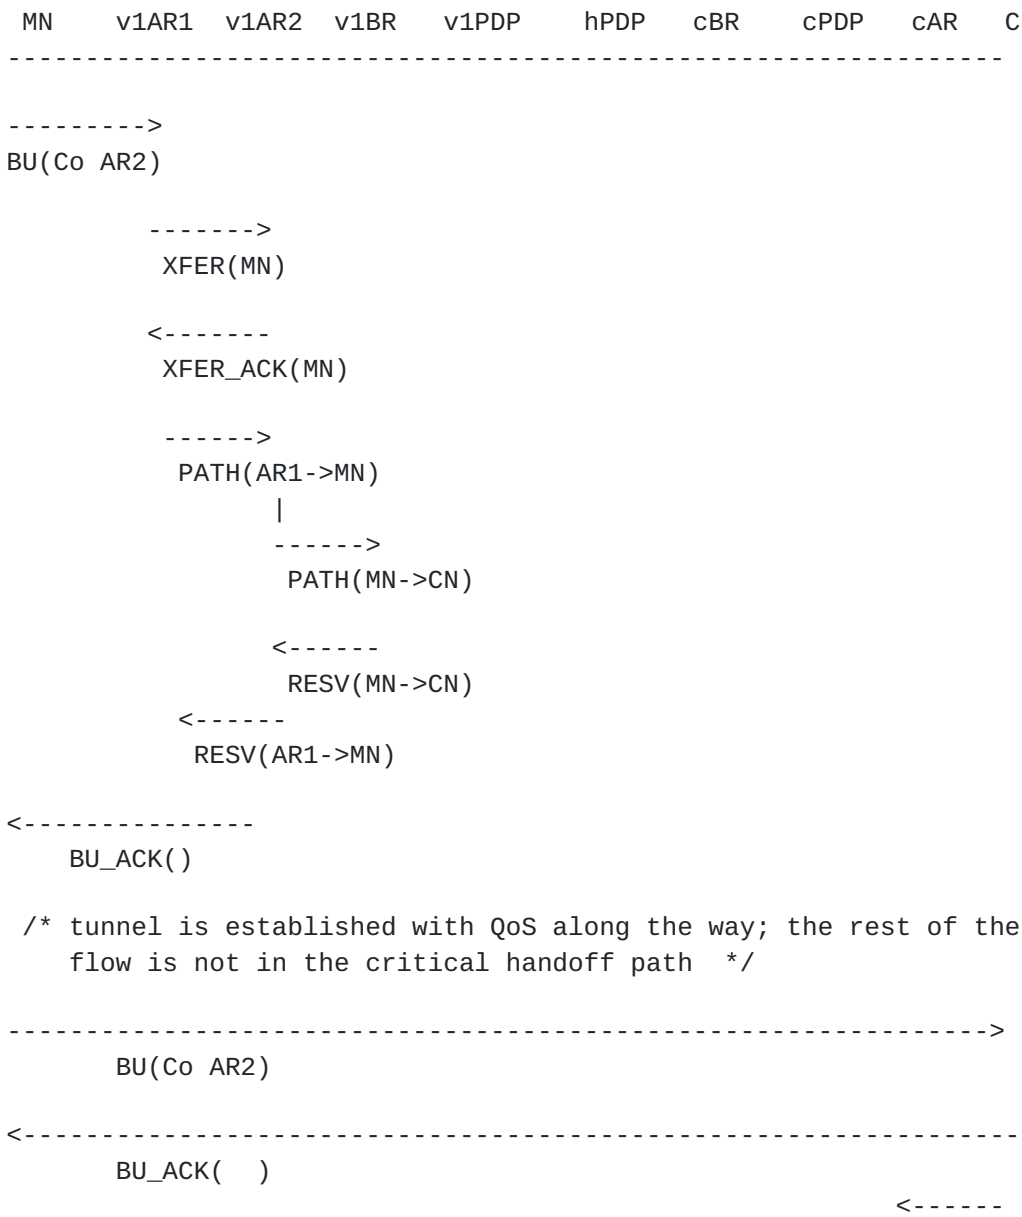

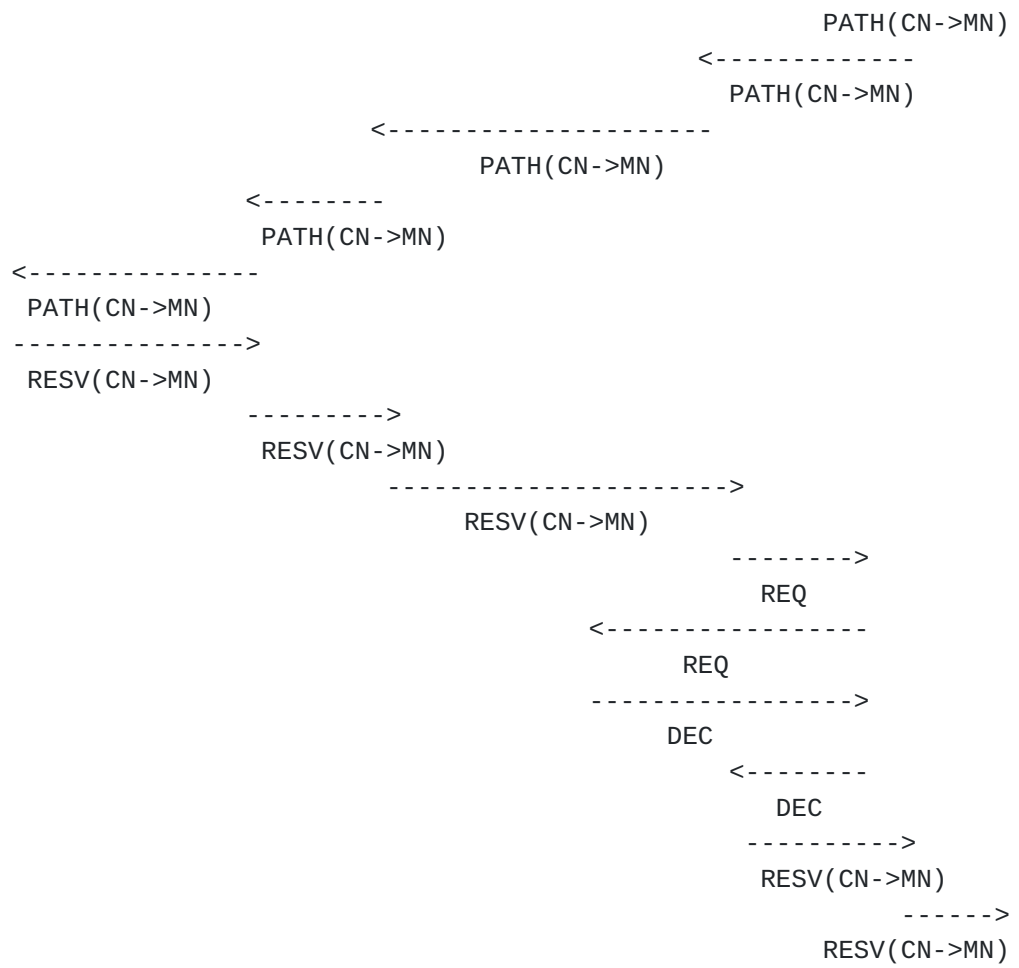

[need to work on the analysis section here...]

References

[RFC2205]
Network Working Group, R. Braden, Ed, "Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification"

[RFC2746]
Network Working Group, A. Terzis, et al, "RSVP Operation Over IP Tunnels"

[RFC2749]
Network Working Group, S. Herzog, Ed "COPS Usage for RSVP"

[RFC2752]
Network Working Group, S. Yadav, et al, "Identity Representation for RSVP"

[RFC2753]

Network Working Group, R. Yavatkar, et al, "A Framework for Policy-based Admission Control "

[RFC2747]
Network Working Group, F. Baker, et al, "RSVP Cryptographic Authentication"

[1] Mobile IP Working Group, D. Johnson, C. Perkins, "Mobility Support in IPv6" draft-ietf-mobileip-rfc2002bis-03.txt

[2] ISSLL Working Group, F. Baker, et al. "RSVP Reservation Aggregation" draft-ietf-issll-rsvp-aggr-02.txt

Acknowledgments

I'd like to thank Dave Oran, Bruce Davie and Ron Cohen who took the time to review this draft, and ponder many of my seemingly innocuous questions.

Author's Address

Michael Thomas
Cisco Systems
375 E Tasman Rd
San Jose, Ca, 95134, USA
Tel: +1 408-525-5386
email: mat@cisco.com