

Multihoming without IP Identifiers

<[draft-nordmark-multi6-noid-00.txt](#)>

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet Draft expires April 20, 2004.

Abstract

This document outlines a potential solution to IPv6 multihoming in order to stimulate discussion.

This proposed solution relies on verification using the existing DNS to prevent redirection attacks, while allowing locator rewriting by (border) routers, with no per-packet overhead. The solution does not introduce a "stack name" type of identifier, instead it ensures that all upper layer protocols can operate unmodified in a multihomed setting while still seeing a stable IPv6 address.

DISCLAIMER: This work has been discussed extensively in a design team. The design team is still exploring multiple approaches and

INTERNET-DRAFT

Multihoming without IDs

Oct 20, 2003

this is an attempt to capture one such approach on paper. Because of this and due to lack of time to review the document one can not say that this is a product of the DT; errors and confusions should be attributed to the scribe and not to the DT.

INTERNET-DRAFT

Multihoming without IDs

Oct 20, 2003

Contents

1.	INTRODUCTION.....	4
1.1.	Non-Goals.....	4
1.2.	Assumptions.....	5
2.	TERMINOLOGY.....	5
2.1.	Notational Conventions.....	6
3.	PROTOCOL OVERVIEW.....	6
3.1.	Host Pair Context.....	10
4.	PROTOCOL WALKTHROUGH.....	11
4.1.	Initial Context Establishment.....	11
4.2.	Locator Change.....	13
4.3.	Handling Locator Failures.....	14
4.4.	Locator Set Changes.....	14
5.	HANDLING STATE LOSS.....	15
6.	ENCODING BITS IN THE IPv6 HEADER?.....	17
7.	COMPATIBILITY WITH STANDARD IPv6.....	18
8.	APPLICATION USAGE OF IDENTIFIERS.....	18
9.	CHECKSUM ISSUES.....	19
10.	IMPLICATIONS FOR PACKET FILTERING.....	19
11.	IPSEC INTERACTIONS.....	20
12.	SECURITY CONSIDERATIONS.....	20
13.	DESIGN ALTERNATIVES.....	20

14.	OPEN ISSUES.....	20
15.	ACKNOWLEDGEMENTS.....	21
16.	REFERENCES.....	21
16.1.	Normative References.....	21
16.2.	Informative References.....	22
	AUTHORS' ADDRESSES.....	23

[1.](#) INTRODUCTION

The goal of the IPv6 multihoming work is to allow a site to take advantage of multiple attachments to the global Internet without having a specific entry for the site visible in the global routing table. Specifically, a solution should allow users to use multiple attachments in parallel, or to switch between these attachment points dynamically in the case of failures, without an impact on the upper layer protocols.

This proposed solution uses existing DNS mechanisms to perform enough validation to prevent redirection attacks.

The goals for this proposed solution is to:

- o Have no impact on upper layer protocols in general and on transport protocols in particular.
- o Address the security threats in [[M6SEC](#)].
- o Allow routers rewriting the (source) locators as a means of quickly detecting which locator is likely to work for return traffic.
- o No per-packet overhead.
- o No extra roundtrip for setup.
- o Take advantage of multiple locators/addresses for load spreading.

[1.1.](#) Non-Goals

The assumption is that the problem we are trying to solve is site multihoming, with the ability to have the set of site locator prefixes change over time due to site renumbering. The assumption is that such changes to the set of locator prefixes can be relatively slow and managed; slow enough to allow updates to the DNS to propagate. Thus this proposal does not attempt to solve, perhaps related, problems such as host multihoming or host mobility.

This proposal also does not try to provide an IP identifier. Even though such a concept would be useful to ULPs and applications, especially if the management burden for such a name space was zero and there was an efficient yet secure mechanism to map from identifiers to locators, such a name space isn't necessary (and furthermore doesn't seem to help) when using the DNS to verify the

locator relationships.

[1.2.](#) Assumptions

The main technical assumptions this proposal makes is that the DNS infrastructure can be used for verification of the relationship between locators on both the initiator of communication and the responding peer. In particular, it assumes that getting DNS reverse maps (ip6.arpa) populated for the nodes that wish to take advantage of multihoming will not be a significant problem.

[2.](#) TERMINOLOGY

upper layer protocol (ULP)

- a protocol layer immediately above IP. Examples are transport protocols such as TCP and UDP, control protocols such as ICMP, routing protocols such as OSPF, and internet or lower-layer protocols being "tunneled" over (i.e., encapsulated in) IP such as

IPX, AppleTalk, or IP itself.

- interface - a node's attachment to a link.
- address - an IP layer name that contains both topological significance and acts as a unique identifier for an interface. 128 bits.
- locator - an IP layer topological name for an interface or a set of interfaces. 128 bits. The locators are carried in the IP address fields as the packets traverse the network.
- identifier - an IP layer identifier for an IP layer endpoint (stack name in [NSRG]). The transport endpoint is a function of the transport protocol and would typically include the IP identifier plus a port number. NOTE: This proposal does not contain any IP layer identifiers.

Application identifier (AID)

- an IP locator which has been selected for communication with a peer to be used by the upper layer protocol. 128 bits. This is used for pseudo-header checksum computation and connection identification in the ULP. Different sets of

communication to a node (e.g., different connections) might use different AIDs in order to enable load spreading.

address field

- the source and destination address fields in the IPv6 header. As IPv6 is currently specified this fields carry "addresses". If identifiers and locators are separated these fields will contain locators.

FQDN

- Fully Qualified Domain Name

[2.1.](#) Notational Conventions

A, B, and C are nodes.

FQDN(A) is the domain name for A.

Ls(A) is the locator set for A, which consists of L1(A), L2(A), ... Ln(A).

AID(A) is an application ID for A. AID(A) is always one member of Ls(A).

[3.](#) PROTOCOL OVERVIEW

In order to prevent redirection attacks this protocol relies on the DNS (for the nodes which support this protocol) having maintained and consistent forward and reverse maps. This allows any node, given one locator, to determine the corresponding FQDN and the set of locators for the node. Once those lookups have been performed, and the original locator is indeed part of the set, the node can happily allow any of those locators without being subject to redirection attacks. Keeping the FQDN around allows the solution to handle graceful renumbering by being able to redo the DNS (e.g., based on the TTL on the resource records).

DNS is also used to provide an indication of M6 capability of a node. The details of this is TBD but a simple example would be to introduce

a new M6 RR type in the DNS which has no RDATA; thus the mere existence of such a record at a FQDN implies that the node supports the M6 protocol.




Figure 1: Protocol stack

The proposal uses an M6 shim layer between IP and the ULPs as shown in figure 1, in order to provide ULP independence. Conceptually the M6 shim layer behaves as if it is an extension header, which would be ordered immediately after any hop-by-hop options in the packet. However, the amount of data that needs to be carried in an actual M6 extension header is close to zero. By using some encoding of the nexthdr value it is possible to carry the common protocols/extension headers without making the packets larger. The nexthdr encodings are discussed later in this document. We refer to packets that use this encoding to indicate to the receiver that M6 processing should be applied as "M6 packets" (analogous to "ESP packets" or "TCP packets").

Layering AH and ESP above the M6 shim means that IPsec can be made to be unaware of locator changes the same way that transport protocols can be unaware. Thus the IPsec security associations remain stable even though the locators are changing. Layering the fragmentation header above the M6 shim makes reassembly robust in the case that there is broken multi-path routing which causes different paths, hence potentially different source locators, for different fragments.

The proposal uses router rewriting of (source) locators as one way to determine which is the preferred (or only working) locator to use for return traffic. But not all packets can have their locators rewritten. In addition to existing IPv6 packets, the packets exchanged before M6 host-pair context state is established at the

receiver can not have their locators rewritten. Thus a simple

mechanism is needed to indicate to the routers on the path whether or not it is ok to rewrite the locators in the packet. Conceptually this is a single bit in the IPv6 header (we call it the "rewrite ok" bit) but there is no spare bit available. Later in the document we show how we solve this by allocating a range of next header values to denote this semantic bit.

Applications and upper layer protocols use AIDs which the M6 layer will map to/from different locators. The M6 layer maintains state, called host-pair context, in order to perform this mapping. The mapping is performed consistently at the sender and the receiver, thus from the perspective of the upper layer protocols packets appear to be sent using AIDs from end to end, even though the packets travel through the network containing locators in the IP address fields, and even though those locators might be rewritten in flight.

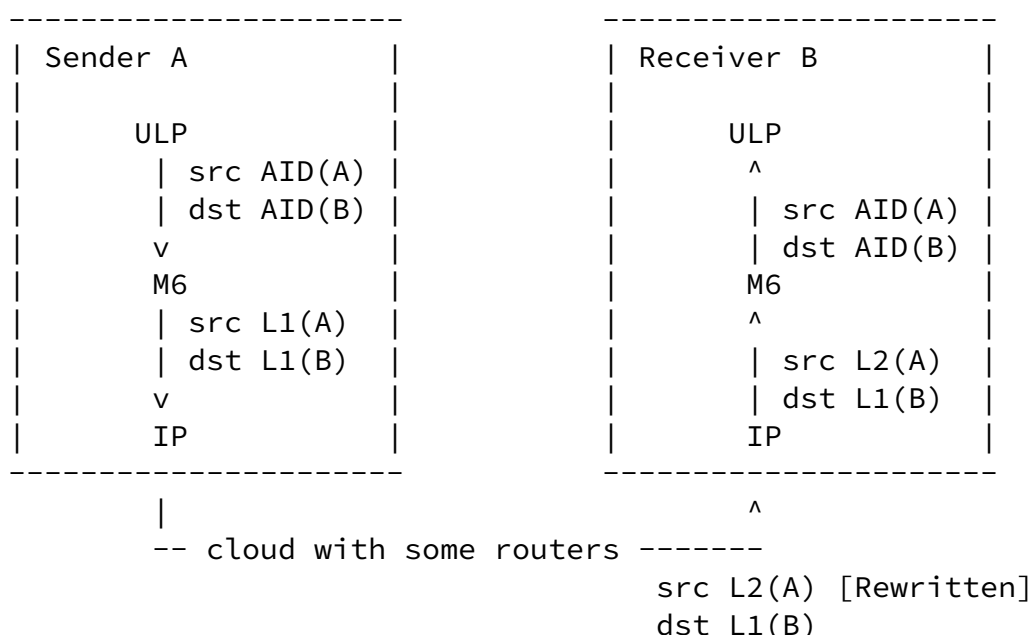


Figure 2: Mapping with router rewriting of locators.

The result of this consistent mapping is that there is no impact on the ULPs. In particular, there is no impact on pseudo-header checksums and connection identification.

Conceptually one could view this approach as if both AIDs and locators being present in every packet, but with a header compression mechanism applied that removes the need for the AIDs once the state has been established. As we will see below the flowid will be used akin to a "compression tag" i.e., to indicate the correct context to use for decompression.

The need for some "compression tag" is because the desire to allow

load spreading and handle site renumbering. Without those desires it could have been possible to e.g. designate one fixed locator as the AID for a node and storing that in the DNS. But instead different connections between two nodes are allowed to use different AIDs and on reception of a M6 packet the correct AIDs may be inserted into the IP address fields before passing the packet to the ULP. The flowid serves as a convenient "compression tag" without increasing the packet size, and this usage doesn't conflict with other flowid usage.

In addition to the zero overhead data messages, there are three different M6 message types introduced (which could be defined as new ICMPv6 messages). They are used to perform a 3-way handshake to create state at both endpoints without creating state on the first received packet (which would introduce a memory consumption DoS attack). The three message types are called:

- o Context request message; first message of the 3-way context establishment. Sent by the responder when a data packet arrives with no context state. An ULP packet can be piggybacked on this message.
- o Context response message; second message of the 3-way context establishment. Sent in response to a context request. An ULP packet can be piggybacked on this message.
- o Context confirm message; third message of the 3-way context establishment. Sent in response to a context response. An ULP packet can be piggybacked on this message.

Similar to MAST [[MAST](#)] the above exchange can be performed asynchronously with data packets flowing between the two nodes; until context state has been established at both ends the packets would flow without allowing router rewriting of locators and without the ability for the hosts to switch locators.

Once the 3-way state creation exchange has completed there is host-pair context state at both hosts. At that point in time the responder (which didn't use DNS before the setup) can asynchronously start retrieving and verifying additional locators using the DNS. Once a peer locator has been verified the node will accept packets from that locator as well as dynamically switch to using the last received source locator (that is already verified) as the destination locator for return traffic.

INTERNET-DRAFT

Multihoming without IDs

Oct 20, 2003

[3.1.](#) Host Pair Context

The host pair context is established on the initiator of communication based on information learned from the DNS (either by starting with a FQDN or with an IP address -> FQDN lookup). The responder will establish some initial state using the context 3-way handshake and later discover and verify the peer's locators using the DNS.

The context state contains the following information:

- the peer locator which the ULP uses as ID; AID(peer)
- the local locator which the ULP uses as ID; AID(local)
- the set of peer locators; Ls(peer)
- for each peer locator, a bit whether it has been verified with the DNS (by doing reverse + forward lookup)
- the preferred peer locator - used as destination; Lp(peer)
- the set of local locators; Ls(local)
- the preferred local locator - used as source; Lp(local)
- the flowid used to transmit packets; F(local)
- the flowid to expect in receive packets; F(peer)
- the fqdn for the peer; FQDN(peer)
- State about peer locators that are in the process of being verified in the DNS

This state is accessed differently in the transmit and receive paths. In the transmit path when the ULP passes down a packet the key to the context state is the tuple <AID(local), AID(peer)>; this key must identify at most one state record. In the receive path it is the

F(peer) plus one of the locators in each of Ls(local) and Ls(peer) that are used to identify at most one state record.

These uniqueness requirements imposed by those lookup keys uniquely identifying one state record means that one can not create multiple records (e.g. with different FQDN or locator sets) that have the same AID pair, and the peer must pick a flowid so that host pair contexts which have at least one common members in Ls(local) and in the Ls(peer) sets, but with different AID pair, gets a different

F(local). The context state at both ends must be consistent for this to be completely robust. One way of ensuring this is to have each node perform a periodic DNS lookup of its own FQDN in order to have a current Ls(local) that is the same as the Ls(peer) that the peer would find in the DNS.

Note that the flowids could be selected to be finer grain than above; for instance having a different flowid for each connection. Doing so requires some efficient datastructure organization at the receiver to map multiple F(peer) to the same context.

[4.](#) PROTOCOL WALKTHROUGH

[4.1.](#) Initial Context Establishment

Here is the sequence of events when A starts talking to B:

1. A looks up FQDN(B) in the DNS which returns Ls(B) plus "B is M6 capable" One locator is selected to be returned to the application: AID(B) = L1(B) The others are installed in the M6 layer on the host with AID(B) being the key to find that state. To make sure that the lookup from AID(B) returns a single state record it appears that one needs to do a reverse lookup AID(B)->fqdn and check that the result is FQDN(B). Whether this check

can be deferred until two entities try to use the same AID(B) for a different Ls is for further study. Always doing the reverse lookup would be more predictable in any case.

2. The ULP creates "connection" state between AID(A)=L1(A) and AID(B) and sends the first packet. L1(A) was picked using regular source address selection mechanisms.
3. The M6 layer matches on AID(B) and finds the proto-context state (setup in step #1) with Ls(B). The existence of that state will make the M6 layer send a M6 packet. The M6 layer selects a flowid F(local) consistent with the above uniqueness requirements (which ensure that the receiver will map to the correct AID pair).
4. The packet (TCP SYN or whatever) is sent to peer with locators

L1(A) to L1(B) i.e., the same as the AIDs. Since B doesn't have any context state yet, A will not set the "rewrite ok" bit in the header.

5. Node B receives packet and sees it is a "M6 packet". Passes the packet to the M6 shim layer. The M6 layer can't create state on the first packet, but since the rewrite bit is not set in the packet it can pass the packet unmodified to the ULP. The ULP sees a packet identified by AID(A), AID(B).

The M6 layer initiates a state creation 3-way exchange by forming a context request message. The same technique as in [\[MIPv6\]](#) can be used to securely do this exchange without any local state; use a local key which is never shared with anybody and pass the context state, a timestamp, and the keyed hash of the state+timestamp in the context request packet. When the state, timestamp, and keyed hash value is returned in the context response message, the hash is used to verify that the state hasn't been modified.

The 3-way exchange is done asynchronously with ULP packets, but it is possible (assuming the MTU allows) to piggyback ULP packets on this exchange.

Should ULP packets be passed down to the M6 layer on B before

the context response message has been received there will be no context state and no state installed as a result of a DNS lookup (as on A). This will indicate that the ULP message should be passed as-is (not as an M6 message) to the peer. Thus during the 3-way exchange packets can flow in both directions using the original locators=AIDs.

6. Node A receives the context request message. It verifies that the message is related to something it sent by looking at the locators (should match the AIDs) and the flowid it sent (which is in the state in the context request message).

If a ULP packet was piggybacked A will pass that to the ULP.

Then A sends a content response which has the same information as the context request plus a nonce/timestamp that A selected.

7. Node B receives the context response message. It verifies that the hash of the state is correct using its per-node key. At this point in time it knows that A is at least not just blasting out packets as a DoS - A is also responding to context request messages. Thus B goes ahead and allocates state at this point in time using the state that is in the context response message.

The M6 layer selects a flowid F(B) consistent with the above uniqueness requirements (which ensure that the receiver will map to the correct AID pair). At this point in time B has enough information to handle M6 packets from A, even though it hasn't yet determined and verified any additional peer locators from the DNS. It has also the state (F(B) mainly) necessary send data packets to A with "rewrite ok" set. Thus B sends a context confirm message to A which contains A's nonce/timestamp from the context response and F(B).

If a ULP packet was piggybacked on the context response B will pass that to the ULP.

At this point in time B can start asynchronously and incrementally to extract and verify Ls(A) from the DNS. The first lookup consists of finding L1(A)=AID(A) in ip6.arpa to get the FQDN and record it, and lookup the AAAA RR set for that FQDN to get Ls(A). Then verify (also incrementally) that each member

of Ls(A) is indeed assigned to A by doing a reverse lookup of each one (except L1(A) which was already looked up). Only when the reverse lookup of a given peer locator has completed is that locator marked as verified.

8. Node A receives the context confirm message, verifies the nonce/timestamp, and records F(peer) from the packet.

If a ULP packet was piggybacked on the context confirm A will pass that to the ULP.

At this point in time A knows that B has context state, thus it can start sending packets with "rewrite ok" set.

[4.2.](#) Locator Change

This is the sequence of events when B receives a packet with a previously unused source locator for A, for instance L2(A).

Node B receives M6 packet with source L2(A) and destination L1(B). Looks up context state using the flowid and the locators. If this lookup succeeds then the locator is acceptable for incoming packets (even though it might not have been verified for use as return traffic) and the packet is rewritten to contain the AIDs from the context state and passed to the ULP.

If L2(A) has not been verified then it would make sense for B to put that first in the list of asynchronous DNS verifications that are

needed. If/once L2(A) has been verified B can make it the preferred peer locator for use when sending packets to AID(A).

The verification needs to complete before using the locator as a destination in order to prevent 3rd party DoS attacks [[M6SEC](#)].

If a node receives a packet with a known flowid but where the locators (src and dst) are not part of the sets it drops the packet. [TBD: Useful vs. harmful to send ICMP error?]

[4.3.](#) Handling Locator Failures

Should not all locators be working when the communication is initiated some extra complexity arises, because the ULP has already been told which AIDs to use. If the locators that were selected to be AIDs are not working it isn't possible to send a zero-overhead initial packet from A to B. Instead both the AIDs and the working locators need to be conveyed. This could be done by either reusing IP-in-IP encapsulating or defining another M6 message type which carries both.

After context setup the sender can use retransmit hints from the ULP to get the M6 layer to try a different verified locator.

If one outbound path from the site fails and the border routers rewrite source locators then the peer in another site will see packets with the working source locators. Once that has locator been verified, the return path will switch to use the working locator. As long as both ends are transmitting packets this will relatively quickly switch to working locators except when both nodes experience a failing locator at the same time.

Without locator rewriting one would need to add some notification e.g., by defining a new bit in the router advertisement prefixes (IMHO this is semantically different than the preferred vs. deprecated stuff), but we also need some mechanism to carry this info from the border routers to the routers on each subnet.

[4.4.](#) Locator Set Changes

Due to events like site renumbering the set of locators assigned to a node might change at a slow rate. Since this proposal uses the locators in the DNS as the credible source for which locators are assigned there is some coordination necessary to ensure that before a

host, or the border routers for a site doing rewriting, start using a new source locator, that locator has propagated through the DNS so that the peer could have discovered it.

Due to concerns about having packets with unknown, hence potentially bogus, source locators triggering DNS lookups this proposal instead uses the DNS TTL as an indication that the set of locators need to be refreshed. One could also envision a combination of receiving a packet *and* the DNS TTL having expired as the trigger to redo the DNS lookups.

When DNS TTL expires on either node it performs a new fqdn->Ls lookup to get the new set of locators. (Presumably failures to redo the lookup shouldn't have a negative effect.)

5. HANDLING STATE LOSS

The protocol needs to handle two forms of state loss:

- a peer loosing all state,
- the M6 layer garbage collecting state too early due to not being aware of what all ULPs do.

The first case is the already existing case of a node crashing and "rebooting" and as a result loosing transport and application state. In this case there are some added complications from the M6 layer since a peer will continue to send packets assuming the context still exists and due to the loss of state on the receiver it isn't even able to pass the correct packet up to the ULP (e.g., to be able to get TCP to generate a reset packet) since it doesn't know what AIDs to use when replacing the locators.

The second case is a bit more subtle. Ideally an implementation shouldn't discard the context state when there is some ULP that still depends on this state. While this might be possible for some implementations with a fixed set of applications, it doesn't appear to be possible for implementations which provide the socket API; there can be things like user-level "connections" on top of UDP as well as application layer "session" above TCP which retain the identifiers from some previous communication and expect to use those identifiers at a later date. But the M6 layer has no ability to be aware of this.

Thus an implementation shouldn't discard context state when it knows it has ULP connection state (which can be checked in e.g., Unix for TCP), or when there is active communication (UDP packets being sent to AID(A) recently), but when there is an infrequently communicating user-level "connection" over UDP or "session" over TCP the context state might be garbage collected even though it shouldn't.

For instance if B lost all state and A retransmits a packet with flowid, L3(B), L2(A) then what is needed is a packet to L1(B) from L1(A) passed to the ULP so that the ULP can send an error (such as a TCP reset). But B has no matching state thus it needs to send an error back. Should the packet not have "rewrite ok" set it can pass it to the ULP since it knows that such packets contain locators that are AIDs. Thus the protocol needs a "Unknown context" error message - probably akin to the one defined in [\[CB128\]](#).

Local loss of context state (on B in this example) has slightly different issues since without any context state the M6 layer on B can not determine whether a peer is a standard IPv6 node or a node which supports multihoming. B can determine this by doing a reverse lookup of AID(A)->FQDN(A) followed by a FQDN(A) lookup to see if there is an M6 record (and get the locator set of A as well).

But if B is communicating with both standard IPv6 nodes and nodes which support multihoming then it has to avoid doing these DNS lookups for every packet sent to a standard IPv6 node. Implementation tricks (such as "has this socket ever used M6" flag at the socket layer, and "negative caching" of peers that do not support M6) can be useful to avoid performance overhead.

If as part of this B determines that A is M6 capable it has the same information as the initiator during the initial context establishment thus it can follow that procedure. If A didn't garbage collect its end of the state this will require some extra work to come up with a single host-pair context for a pair of AIDs at both ends with consistent flowids in the two nodes (i.e., F(local) needs to match F(peer) at the other node). Specifying this is for further study.

INTERNET-DRAFT

Multihoming without IDs

Oct 20, 2003

6. ENCODING BITS IN THE IPv6 HEADER?

The idea is to pick extra IP protocol values for common combinations, and have a designated protocol value to capture the uncommon IP protocols which might use M6.

We pick two unused ranges of IP protocol values which 8 numbers each (assuming we will not need more than 7 common transport protocols).

The ranges start at P1 and P2, respectively:

P1 TCP over M6 - rewrite ok

P1+1 UDP over M6 - rewrite ok

P1+2 SCTP over M6 - rewrite ok

P1+3 RDDP over M6 - rewrite ok

P1+4 ESP over M6 - rewrite ok

P1+7 escape - any protocol over M6 - rewrite ok

In this case we spend another 8 bytes (minimum IPv6 extension header size due to alignment rule) to carry the actual IP protocol. This causes some mtu concerns for those protocols, but they aren't very likely to be used with M6?

P2 TCP over M6 - no rewrite

P2+1 UDP over M6 - no rewrite

P2+2 SCTP over M6 - no rewrite

P2+3 RDDP over M6 - no rewrite

P2+4 ESP over M6 - no rewrite

P2+7 escape - any protocol over M6 - no rewrite

In this case we spend another 8 bytes (minimum IPv6 extension header size due to alignment rule) to carry the actual IP protocol. This causes some mtu concerns for those protocols, but they aren't very likely to be used with M6?

Thus a router would check if the protocol is in the P1 range and if so, it can rewrite the locator(s). A host would check a received packet against both P1 and P2 ranges and if so pass it to the M6 shim layer.

Some possible alternatives to the above encoding is to:

- use something combination of the universal/local and group bit in the interface id of the source address to indicate "rewrite ok".

- steal the ECN bits from the traffic class before ECN becomes a proposed standard? Don't think this will be popular!
- always have a shim header - adds 8 bytes overhead per packet.

7. COMPATIBILITY WITH STANDARD IPv6

A node can easily implement M6 in a way that interoperates with current IPv6 as follows.

When the DNS lookup routines do not find an M6 record for the peer they will return the AAAA resource record set to the application; those would be the IPv6 addresses. When the ULP passes down these addresses the M6 layer will not have any state generated by the DNS lookup code, thus no M6 processing will take place on the sender. (Note that this interferes with the M6 layer state recovery concerns above.)

The receive side handles both standard IPv6 and M6 since it demultiplexing on whether a packet is an M6 packet.

8. APPLICATION USAGE OF IDENTIFIERS

The upper level protocols will operate on AIDs which are mere locators. Thus as long as a site hasn't renumbered the AID can be used to either send packets to the node, or (e.g. if that locator isn't working), it is possible for an application to do a reverse lookup plus forward lookup of the AID to get the set of locators for the peer.

Once a site has been renumbered the AIDs which contain the old prefix will no longer be useful. Hence applications must try to honor the DNS TTL somehow.

Applications which use to map the peer's IP address to a domain name today perform a reverse lookup in the DNS (e.g., using the `getnameinfo()` API). This proposal doesn't add or subtract to the benefits of performing such reverse lookups.

9. CHECKSUM ISSUES

The IPv6 header does not have a checksum field; the IPv6 address fields are assumed to be protected by the ULP pseudo-header checksum. The general approach of an M6 shim which replaces locators with identifiers (where only the identifiers are covered by ULP checksum) raises the potential issue of robustly handling bit errors in the address fields.

With the definition of the M6 shim there can be undetectable bit errors in the flowid field or the nexthdr field which might adversely affect the operation of the protocol. And since the AIDs are what's covered by the ULP's pseudo-header checksum the locators in the address fields are without checksum protection. Since the locators are verified against the set of locators determined from the DNS, the worst effect of an undetected bit error in the locators is an error message being sent back [TBD]

Undetected bit errors in the flowid can have the following effects [TBD]

10. IMPLICATIONS FOR PACKET FILTERING

Ingress filtering should be replaced by locator rewrite when the "rewrite ok" bit is set.

Locator rewriting (when the bit is set) can be applied at places where ingress filtering isn't currently performed (e.g., due to multihoming issues).

Firewall filtering potentially require modifications to be aware of M6. All the packets contain locator thus a firewall would need to be aware of the context state to let the correct packets through. Such firewalls could optionally perform their own verification by issuing DNS lookups the same way as the endpoint. However, the firewalls probably has to be more careful not exposing themselves to DoS attacks by doing too much DNS lookups.

11. IPSEC INTERACTIONS

As specified all of ESP, AH, and key management is layered above the M6 layer. Thus they benefit from the stable identifiers provided above the M6 layer. This means the IPsec security associations are unaffected by switching locators.

The alternative would be to layer M6 above IPsec, but that doesn't seem to provide any benefits. Since we want to allow routers performing locator rewriting it isn't possible to take advantage of for instance AH to protect the integrity of the IP headers.

12. SECURITY CONSIDERATIONS

Early analysis indicates this addresses the issues in [[M6SEC](#)].

[13.](#) DESIGN ALTERNATIVES

Use an actual extension header for M6 and place the context tag in that header instead of using the flowid. This would make the packets 8 bytes larger since the minimum extension header size is 8 bytes due to the alignment rules for extension headers in IPv6.

[14.](#) OPEN ISSUES

DNS lookup fails or times out on the receiver; what should one do? Send error?

Is it possible to facilitate transition to M6 using some "M6 proxy" at site boundaries until all important hosts in a site have been upgraded to support M6? Would would be the properties of such a proxy? Would it place any additional requirements on the protocol itself?

One of the issues with FQDNs mapping to AAAA records is that in some cases multiple AAAA records mean a multihomed host and in other cases it means multiple hosts providing the same service. If we need to introduce a new RRtype for M6, would it be useful to try to make this host/service distinction more clear at the same time? An example solution would be that the M6 record would by its existence indicate the M6 capability, and its RDATA would contain a list of host names which would be used to resolve the AAAA records for each host implementing the service.

Would destination locator rewriting be a useful way for the routing system to pass some information to the node? Or is source locator

rewriting sufficient?

[15.](#) ACKNOWLEDGEMENTS

This document is the result of discussions in a MULTI6 design team but is not the "product" of that design team. The scribe wishes to acknowledge the contributions of (in alphabetical order): Iljitsch van Beijnum, Brian Carpenter, Tony Li, Mike O'Dell, and Pekka Savola.

The idea to allow locator rewriting by routers was first presented by Mike O'Dell [[ODELL96](#)]. The techniques for avoiding state DoS attacks on the first packet are patterned after [[MIPv6](#)].

[16.](#) REFERENCES

[16.1.](#) Normative References

[M6SEC] Nordmark, E., and T. Li, "Threats relating to IPv6 multihoming solutions", [draft-nordmark-multi6-threats-00.txt](#), October 2003.

[ADDR-ARCH] S. Deering, R. Hinden, Editors, "IP Version 6 Addressing Architecture", [RFC 3513](#), April 2003.

[IPv6] S. Deering, R. Hinden, Editors, "Internet Protocol, Version

[draft-nordmark-multi6-noid-00.txt](#)

[Page 21]

[16.2.](#) Informative References

[NSRG] Lear, E., and R. Droms, "What's In A Name: Thoughts from the

NSRG", [draft-irtf-nsrg-report-09.txt](#) (work in progress), March 2003.

[MIPv6] Johnson, D., C. Perkins, and J. Arkko, "Mobility Support in IPv6", [draft-ietf-mobileip-ipv6-24.txt](#) (work in progress), June 2003.

[AURA02] Aura, T. and J. Arkko, "MIPv6 BU Attacks and Defenses", [draft-aura-mipv6-bu-attacks-01](#) (work in progress), March 2002.

[NIKANDER03] Nikander, P., T. Aura, J. Arkko, G. Montenegro, and E. Nordmark, "Mobile IP version 6 Route Optimization Security Design Background", [draft-nikander-mobileip-v6-ro-sec-01](#) (work in progress), June 2003.

[ODELL96] O'Dell M., "8+8 - An Alternate Addressing Architecture for IPv6", [draft-odell-8+8-00.txt](#), October 1996,

[MAST] D. Crocker, "MULTIPLE ADDRESS SERVICE FOR TRANSPORT (MAST): AN EXTENDED PROPOSAL", [draft-crocker-mast-protocol-01.txt](#), October, 2003.

[CB128] E. Nordmark, "Strong Identity Multihoming using 128 bit Identifiers (SIM/CBID128)", [draft-nordmark-multi6-sim-00.txt](#), October 2003.

[IPv6-SA] R. Atkinson. "Security Architecture for the Internet Protocol". [RFC 2401](#), November 1998.

[IPv6-AUTH] R. Atkinson. "IP Authentication Header", [RFC 2402](#), November 1998.

[IPv6-ESP] R. Atkinson. "IP Encapsulating Security Payload (ESP)", [RFC 2406](#), November 1998.

AUTHORS' ADDRESSES

Erik Nordmark
Sun Microsystems, Inc.
17 Network Circle
Mountain View, CA
USA

phone: +1 650 786 2921
fax: +1 650 786 5896
email: erik.nordmark@sun.com

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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

