Automatic configuration of IPv6 addresses for MANET with multiple gateways (AMG)
draft-ruffino-manet-autoconf-multigw-03

Status of this Memo

Copyright Notice

Abstract

This document describes AMG, a mechanism for stateless autoconfiguration of IPv6 addresses for Mobile Ad-hoc Networks (MANETs), connected to the Internet by means of one or more gateways. Network prefixes are disseminated by Internet gateways and are used by nodes to configure a set of global IPv6 addresses. An algorithm is specified, by which nodes can choose the optimal address for data traffic. Configured global addresses are also advertised to other

MANET nodes, to minimize latencies in case of gateway failures, MANET partitions and mergers. The specified mechanism aims to be independent from any particular MANET routing protocol and to effectively exploit multiple gateways.

Table of Contents

1. Introduction

This document specifies AMG, a general-purpose stateless mechanism for automatic configuration of topologically correct, globally valid IPv6 addresses on nodes in a MANET connected to the Internet through one or more Internet gateways (see [I-D.singh-autoconf-adp]).

An overview of AMG can be given as follows: gateways periodically disseminate their IPv6 prefixes in the MANET; when nodes receive such prefixes, they build a set of global addresses, rank them with a best-prefix selection algorithm, start using one or more of them for data traffic and concurrently advertise them back to other MANET nodes.

An important feature of AMG is the use of a best-prefix selection algorithm, which enables MANET nodes to continuosly choose the best address to use, according to the MANET topology. In fact, a node can change the global address in use e.g. after the failure of the gateway announcing the prefix from which it derived its used global address or for performance reasons, e.g. to optimize downstream data traffic path. A second important feature is address advertisement: it enables nodes and gateways, to know all the addresses configured on each other nodes and to build routes towards them: in particular, packets arriving at the gateways from the Internet directed to any addresses of MANET nodes, can be routed with no delay, even after partitioning occur and many nodes may be forced to change addresses in use.

AMG aims to be independent from any particular MANET routing protocol; nevertheless we specify detailed operations in case the Optimized Link State Routing (OLSR) [RFC3626] is used. AMG uses the Generalized MANET Packet/Message Format [I-D.packetbb] for Prefix Advertisement messages.

This document is organized as follows: Section 2 describes the applicability scenarios; Section 3 exposes the problem of global address configuration in case of multiple gateways; Section 4 outlines the proposed solution. Logical data structures are detailed in Section 5 and operations are described in Section 6 and Section 7. Finally, Appendix B contains some considerations on using AMG with Mobile IPv6.

2. Applicability Scenarios

AMG can be used to configure valid IPv6 addresses on MANET nodes in different scenarios, both when the MANET is interconnected to the Internet and when it is standalone. The reference scenario is shown in Figure 1 (see also [I-D.ruffino-conn-scenarios]). In this scenario, MANETs are connected to the Internet by means of one or more gateways (GW1,GW2,GW3). Nodes (N1...N6), that are not directly linked to the external network, can use multi-hop paths to reach the gateways and forward outbound traffic to Internet hosts (H1).

```
                            H1
                             |
                             |
                      +---------------+
                      |   Internet    |**
                      +---------------+  *
                        *           *     *
                        *           *      *
                        *           *       *
                    GW1**           *       GW3
                      |         +--GW2       |
                      |         |   |        |
                   ---N1--------+   |        |
                  /      \          |       N5----N6
               N4         \        N2
                           N3-----/
```

Figure 1: MANET interconnected to an external network

An Internet gateway is a MANET node, equipped with a MANET interface, and a second interface, connected to the external network. Multiple gateways can improve reliability and fault tolerance, as there is no single point of failure in the network, and enable load balancing of traffic directed/coming to/from the Internet. Gateways, as well as nodes, can also be mobile devices, and, as such, can join and depart to/from the MANET at any time: the number of available gateways can therefore vary in time.

AMG can also be applied to scenarios where Internet connection is unavailable, i.e. when MANET is isolated or temporarily disconnected (see [I-D.ruffino-conn-scenarios] for a description of isolated MANET scenarios and applications).

3. Problem Statement and assumptions

Standard configuration methods for IPv6, i.e. stateful ([RFC3315]) and stateless ([RFC2462]) IPv6 autoconfiguration, cannot be applied to multi-link networks such as MANETs, as outlined by previous work ([I-D.singh-autoconf-adp], [I-D.perkins-manet-autoconf], [I-D.wakikawa-globalv6], [I-D.wakikawa-v6-support]). Standard methods have been designed for single-hop link (e.g. a single LAN segment), where all hosts and routers share the same link and don't address MANET intrinsic characteristics, such as multi-hop connections, partitions and mergers.

This specification aims at solving the problems described in AUTOCONF Problem Statement [I-D.singh-autoconf-adp]. In particular, it is focused on global connectivity, i.e. its goals are to enable MANET nodes to build topologically correct global IPv6 addresses and to discover and exploit multiple Internet gateways, if present. It is designed to cope with partitions of the MANET and mergers of two or more MANETs.

It is worth noting that in the reference scenario (Section 2), different design choices imply different technical issues and requirements:

1. In case of multiple GWs announcing *one* network prefix, partitioning of the MANET may invalid routes in the Internet towards MANET nodes, compromising end-to-end reachability. E.g., if a MANET cloud breaks into two separate parts, each one containing a gateway, routers in the Internet cannot choose the correct gateway to deliver traffic for a MANET node. Recovery could be possible, e.g. using host routes, or using a signalling path through the Internet (if available), between partitioned gateways, but, currently, there is no consolidated way (i.e. IETF standard) to solve the problem. Moreover, the implications should be carefully studied, because it is quite likely such a mechanism would require additional protocols/mechanism to run on Internet routers, gateways and MANET nodes. This might limit the applicability of single-prefix autoconf to scenarios with no partitioning at all, e.g. small controlled ad-hoc networks, with very limited mobility, or static sensor networks.

2. In case of multiple gateways advertising *multiple* network prefixes, no coordination among gateways is needed, to preserve Internet routing consistency, even after partitions/mergers, since traffic is univocally routed towards the gateway owning one particular prefix. Using all the available prefixes, MANET nodes can build a "pool" of valid global IPv6 addresses. However, other problems may arise within the MANET itself.

In fact, traffic coming from the Internet is routed through the gateway which owns the prefix, used by nodes to build source address for data packets. Nodes' choice of which global address (and gateway) to use is critical, since it also directly affects the path followed by downstream data within the MANET: a node could choose a prefix announced by a very "far" gateway (in terms of routing metrics), while a "closer" one could exist. In this case, traffic could flow through a non-optimal path within the MANET. Therefore, MANET nodes should be able to choose, both at bootstrap and after major topological changes, the best gateway (e.g. the "closest" one in terms of routing metrics) to configure their address(es), in order to minimize latency and delay variation, maximize throughput and efficiently use radio resources. This can be summarized as the "best prefix selection" problem. Note that this problem is separate from the selection of a default gateway, which defines the default route for traffic directed to the Internet.

3. MANET nodes could reconfigure (frequently) their global address(es), due to partitioning, merging, movement and gateway failure. Following current version of [I-D.rfc2461bis], every unicast IPv6 address should be checked for uniqueness, prior to configuration. In a multiple-gateway/multiple-prefixes MANET, this could bring to a large amount of control signalling, especially if the ad-hoc network is very dynamic.

4. If nodes, involved in communications with Internet hosts, use only global addresses for route calculation (e.g. in OLSR, use of global address as originator address), existing MANET routes must be recomputed when connectivity to the gateway that assigned the prefix is lost. This, because nodes could choose to reestablish communications after the outage is detected, by exploiting another available gateway and therefore configuring a new global address.

3.1 Assumptions

It is therefore assumed that each gateway owns one or more topologically correct IPv6 network prefixes, which can be announced within the MANET. The mechanism by which gateways retreive this information is out of scope of this specification: it can be manually configured or dynamically set up, during link establishment towards the Internet, e.g. using DHCP with Prefix Delegation Option ([RFC3633]).

It is also assumed that different gateways advertise different prefixes, in order not to require special configuration both on

gateways themselves and on Internet routers. Moreover, as discussed above, use of MANET-local addresses to build in-MANET routes could be more efficient than use of global addresses, in case of frequent global address reconfiguration, especially if proactive protocols are used.

This specification explicitly does not address the problem of verifying uniqueness of a newly configured address. It is authors' belief that, due to the very low probability of an address conflict in IPv6 address space, an active Duplicate Address Detection mechanism may not be needed. A lightweight passive address conflict detection and repair mechanism could be used, instead of an try-and-wait approach (e.g. the model used in [RFC2462]), which can bring to signalling overhead and long latencies.

The specified autoconfiguration mechanism performs gateway discovery, but it is assumed that default route calculation is performed by routing protocol running on MANET nodes. This, because routing protocol is the only process responsible for calculating optimal routes in the MANET.

4. AMG Overview

AMG operates in 4 phases.

1) MANET-local address configuration
   At bootstrap, nodes configure a permanent MANET-local address and use it as identifier in routing protocol messages (e.g. as main address in OLSR). This specification recommends the use of IPv6 Unique Local Addresses [RFC4193] to configure MANET-local addresses. As explained in Section 3, such addresses are not verified for uniqueness prior to configuration; instead, an address conflict detection mechanism may be started, that monitors routing packets and other events, reacting only when a conflict is detected (out of scope of this specification). When this step completes, MANET nodes (highly likely) own an unique IPv6 address, which can be used for communications within the MANET.

2) Prefix Advertisement
   Internet gateways periodically advertise global network prefixes by means of a message, named Prefix Advertisement (PA). It is assumed that a forwarding engine is available in the MANET. The current version of this specification assumes the use of SMF ([I-D.SMF]); future versions may define a specialized forwarding procedure for PAs;

3) Best prefix selection and global address configuration
   MANET nodes receive Prefix Advertisement (PA) messages from the gateways. They use the prefixes stored in PAs to configure a set of global IPv6 addresses, built according to [RFC4291]: at least, they build an address from each received prefix. Nodes apply an algorithm, named Best Prefix Selection (see Section 7.2.1), using routing metrics of Internet gateways, if available, to rank the received prefixes. Nodes can use the ranking to select appropriate source address for Internet traffic.

4) Advertisement of global addresses
   Nodes start broadcasting the configured global addresses (in step 3.) to other MANET nodes: this operation enables all nodes to bind MANET-local addresses, used by routing protocol for path calculations, to global addresses. As a result, they can use existing routes to deliver data traffic coming from the Internet and directed to any global address in the MANET. Routing protocols can be used for address advertisement: OLSRv1 (with HNA or MID messages), OLSRv2 (with TC messages) and DYMO already have the capability to advertise multiple addresses along with the main address of each node.

Nodes periodically reapply Best Prefix Selection algorithm (step 3.

above), inspecting the routing table. The selection can also reactively triggered, to re-order prefixes after a significant event occurs, for example:

- o one ore more gateways, which advertises prefixes used by the node, become unreachable.

- o node detects a significant topological change, after which the prefix, previously ranked as the best one, corresponds to an non-optimal path (see Section 3).

The main advantages of the proposed solution are the following:

- o in case of gateway failure or MANET partitioning, nodes can immediately use another valid global address to start data communications with other nodes and Internet hosts: no significant delay due to routing protocol operations is experienced. This, because the local topological information is bound to a MANET-local address and is independent from the global address currently used. As described above (step 4.), all other MANET nodes already know the correct path to reach the node by this address.

- o the path followed by downstream traffic coming from the Internet can be optimized, with respect to the routing metric. This can be achieved, using a Best Prefix Selection algorithm (Default Gateway method, described in Section 7.2.1) that assigns the highest rank to the address derived from a prefix announced by the default gateway (i.e., the best one routing protocol chooses).

- o a gateway which becomes a node, e.g. as the result of losing connectivity towards the external network, can immediately receive downlink traffic by using another active gateway.

5. Logical data structures

In this section, two data structures used by AMG are defined: the Prefix Information Base (PIB), which stores the received prefixes, and the Global Addresses Information Base (GAIB), which stores global addresses built by the node. Best prefix Selection algorithm is applied on GAIB entries.

5.1 Prefix Information Base

The Prefix Information Base (PIB) contains the delegated prefixes announced by gateways within the MANET and it is filled processing Prefix Advertisements. It is maintained by each node and gateway. If a gateway advertises multiple prefixes, PIB MUST be filled with an entry for each advertised prefix.

Entries of the PIB have the following structure (length of each field is expressed in octets):

Field	Data
P_address (16)	MANET-local address of the gateway which sent the PA
P_network (16)	A prefix owned by the gateway whose MANET-local address is P_address
P_prefix_length (1)	Length of the prefix contained in P_network field
P_time (1)	Validity time

Table 1: Prefix Information Base (PIB)

5.2 Global Addresses Information Base (GAIB)

The Global Addresses Information Base (GAIB) stores the set of the global addresses built by a node, after processing Prefix Advertisements carrying global prefixes, i.e. using global prefixes contained into PIB. It is maintained on each node and gateway. The refresh of GAIB entries tightly depends on the state of the entries of PIB, as the validity of a global Address is bound to the validity of the global prefix from which the global Address has been derived.

Best Prefix Selection algorithm is applied on entries of GAIB, which

are re-ordered accordingly.

Entries of the Global Addresses Information Base have the following structure:

Field	Data
G_address (16)	A valid global IPv6 address, owned by the node
G_prefix_length (1)	Length of the prefix in G_address
G_metric (1)	Routing metric of the gateway, owner of the prefix used to build G_address.

Table 2: Global Addresses Information Base

Entries in GAIB are used to fill routing protocol messages, responsible for advertising global addresses of each node to other MANET nodes (e.g. HNA and MID in OLSRv1), as explained in Section 7.3. Optimisations, such as advertising only a subset of the GAIB, to decrease overhead in the network, may be specified in future versions of this specification.

5.2.1 Global Addresses Information Base Management

For each (valid) prefix contained into Prefix Information Base, a node creates a new entry as follows:

1. it creates a IPv6 global address as described in [RFC4291] and stores it in the G_address field. It stores the prefix length in G_prefix_length.

2. it looks-up the routing table and retreives the metric of the gateway that advertised the network prefix used to build G_address, if available. It stores the metric in G_metric field. It can retreive MANET-local address of the gateway inspecting the PIB: the node performs a search using network prefix as key. If the look-up fails, G_metric is left blank.

GAIB maintanance is periodically executed: i.e. routing table look-up is periodically looked-up, to update values in G_metric field of all entries.

If an entry stored into Prefix Information Base is removed, e.g.

after P_time expiration, all the addresses derived from the expired prefix must be removed from GAIB as well.

6. Prefix Advertisement

This section details format and processing of Prefix Advertisement (PA) messages in AMG. PAs conform to the generalized message format, as specified in [I-D.packetbb].

6.1 Prefix Advertisement (PA) messages format

The general PA message structure is shown in Figure 2.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   msg-type    |  RSRV |1|0|1|0|          msg-size             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                                                               |
|                                                               |
|                     originator-address                        |
|                                                               |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        msg-seq-number         |    msg-tlv-block-length=0     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  head-length  |                   head                        :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:               head                            |   num-tails   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             tail                              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                             tail                              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:             tail              |              ...              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             tail                              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                             tail                              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:             tail              |       tlv-block-length        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   tlv-type-1  | tlv-semantics |     length    |     value     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      ...      |     value     |   tlv-type-2  | tlv-semantics |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     length    |     value     |      ...      |     value     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Format of Prefix Advertisement messages

Where each field of the message has the following meaning:

- <msg-type> = Prefix Advertisement (T.B.D.)

- <msg-semantics>:

  * <originator-address> and <msg-seq-number> are included by setting bit 1

  * bit 3 indicates that the message should be intended for forwarding even if the message type is unknown to the recipient

- <msg-header-info>

  * <originator-address> is filled with the main address of the gateway sending PA messages

  * PA contains <msg-seq-number>

- <addr-block> contains the delegated prefixes owned by the gateway, which generated the message.

- <tlv-block>

  * <tlv-type-1> = VALIDITY_TIME

    + <tlv-semantics>: bit 2 and bit 3 (noindex e multivalue) can be set to support multiple addresses with different values

    + <value>: validity time of the information stored in the PA message. It is set to PA_HOLD_TIME. Validity times are expressed in mantissa and exponent format

  * <tlv-type-2> = PREFIX_LENGTH

    + <tlv-semantics>: bit 2 and bit 3 (noindex e multivalue) can be set to support multiple addresses with different values

    + <value>: prefix length, expressed in octets, of the network prefix carried in PA message.

- final padding to 32 bit boundary is optional, and is not included in message length.

This specification assumes use of SMF [I-D.SMF] to broadcast PAs within the MANET. Using SMF, "hop-count" field, commonly used in

MANET protocols, cannot be incremented by nodes, upon forwarding the message and therefore it is not included in PAs. However, in a MANET, hop-count between two nodes in a MANET can be generally different from the metric associated with the route between the same two nodes, as calculated by routing protocol. Thus, AMG retreives gateways' metric from the routing table, both to fill GAIB (Section 5.2.1) and to apply Best Prefix Selection algorithm (Section 7.2.1).

6.2 PA message generation

PAs are originated by gateways every PA_INTERVAL seconds. Every PA contains all the prefixes owned by the gateway, along with associated validity time. The list of prefixes can be partial in each PA message (e.g., due to message size limitations, imposed by the network, or other policies), but parsing of all PA messages describing the interface set from a node MUST be complete within a certain refreshing period (PA_INTERVAL). The information contained in the PA messages is used by the nodes to build their Global Addresses.

6.3 PA message processing and PIB management

PAs are received by all MANET nodes and gateways, which MUST process every received PA according to this Section. Upon processing a PA message, the P_time MUST be computed from the VALIDITY_TIME tlv of the message header (see [RFC3626]). The PIB SHOULD then be updated as follows; for each network prefix (Network Address, Prefix Length) in the message:

1. if an entry in the association set already exists, where:

   P_addr == Originator Address

   P_network_addr == Network Address

   P_prefix_length == Prefix Length

   then the holding time for that entry MUST be set to:

   P_time = current time + validity time

2. otherwise, a new entry MUST be recorded with:

   P_gateway_addr = Originator Address

P_network_addr = Network Address

P_prefix_length = Prefix Length

P_time = current time + validity time

7. Address configuration mechanism

This section describes the configuration operations for MANET nodes and gateways.

7.1 MANET-local Address configuration

At bootstrap, each node and gateway builds one address following [RFC4193] and configures it on one of its interface partecipating to MANET routing. ULA addresses MAY be used for multi-hop transmissions local to the MANET (e.g. [I-D.jelger-mla]). Other MANET interfaces can be configured with ULA as well, but nodes must choose one of its MANET-local addresses as main address and activate the SMF process. MANET-local address SHOULD be used as originator address in routing protocol messages. A passive address conflict detection mechanism, such as [I-D.laouiti], MAY be started after configuration, as explained in Section 3.1.

7.2 Global addresses configuration

As a configuring node joins the appropriate multicast group and activates SMF, it starts receiving PAs from available Internet gateways. It processes PAs and updates PIB accordingly, as described in Section 6.3. It concurrently fills GAIB, as described in Section 5.2.1, building a set of valid global IPv6 addresses. The node can configure one or more of the global addresses on its interfaces and start using them as source addresses in data packets.

To choose the optimal source address to use, the nodes executes the Best Prefix Selection algorithm on entries of GAIB. Two situations can arise:

- o node configures only one global address: in this case, the selection algorithm can help the node to choose the best address to use;

- o node configures multiple global addresses: in this case, [RFC3484] controls the source address selection. The Best Prefix Selection algorithm could nevertheless be use as an aid to RFC3484, in order to select the most convenient address to use.

7.2.1 Best Prefix Selection Algorithm

This section features two Best Prefix Selection algorithms. In the current version of this document only algorithm traces are included.

The best prefix selection algorithm must take into account aspects

related to MANET topology, e.g. the routing metrics of the gateways and external aspects, e.g. the number and type of active data sessions. It is assumed that a node, inspecting the routing table, monitors the current metric value of every reachable gateway generating PA messages and always knows which is the current default gateway, chosen by routing protocol. In this section two alternative algorithms are proposed.

1. Default Gateway Method: nodes always rank the prefix announced by the current Default Gateway as the highest possible and sort the remaining prefixes by descending value of routing metrics (i.e. worse metric means lower ranking).

   In case of OLSR and DYMO, the default gateway is the closest gateway in terms of number of hops. This algorithm solves the downlink path optimization problem described in Section 3. In fact, if the node uses a global IPv6 address derived from the prefix announced by the default gateway, traffic to and from the external network flows through the same gateway. As a drawback, if MANET topology frequently changes, nodes which configured only one global address on MANET interface, may experience frequent address changes, which can cause disruption of data sessions.

2. Threshold Method: nodes calculate the difference 'D' between the metric of the gateway announcing the prefix currently in use and the metric of each other known gateway. Then, they discard gateways whose D value is below a predefined threshold. Finally, if the result is not an empty set, they re-sort the prefixes of the non-discarded gateways by descending value of routing metrics.

   Threshold method accounts for situations when, although nodes use non-optimal prefixes (and traffic flows on non-optimal paths), they prefer to keep current address preferences unchanged, because the benefits provided by a prefix change are not worth the overhead required by the change itself. E.g. when a node has one single configured address and many active data sessions, a prefix change may bring to severe loss of data. Choosing one single value of the threshold for many deployment environments can be difficult: it highly depends on the metric and other factors, which do not strictly depend on routing, e.g. Quality of Service required by applications and number of active data sessions.

The selection algorithm SHOULD be executed at bootstrap time, after a node receives the first set of global prefixes. Nevertheless, this

operation SHOULD also be executed when particular events trigger a topological change in the MANET. Such events have been cited in Section 4 and can be further detailed as follows:

1. Failure of the gateway owning the chosen prefix;

2. A partition, after which the node and the gateway, owning the chosen prefix, belong to two different MANETs;

3. The gateway, which announces the chosen prefix, becomes a node, e.g. after shutting down the interface connecting it to the external network, and stops announcing prefixes;

4. After a movement of one or more MANET devices, a gateway has a better metric than the gateway announcing the chosen prefix;

5. A merging occurs, after which a gateway previously not connected to the MANET may have the best metric value.

In case of events 1., 2. and 3. PIB entries expire, because the global prefix, which the global addresses in use are derived from, is no more listed into PA messages: the node MUST also change its global address, choosing one of the prefixes announced by active gateways. In case of 4. and 5., the node determines that its global addresses in use are derived from a non-optimal prefix, according to the the topological information it owns: in this cases, the node MAY use other valid global addresses, derived from optimal gateways. Appendix B elaborates on address changes on MANET nodes.

7.3 Global addresses broadcasting

After nodes have built and configured one or more global addresses, they advertise them within the MANET. Doing this, other MANET nodes bind each other node's MANET-local address with the global addresses owned by each node. Since DYMO, OLSRv1 and OLSRv2 can already support advertisement of multiple addresses, belonging to a single node, this specification does not define any new transport mechanism. In the following sections, implementation guidelines are given for the mentioned protocols.

7.3.1 OLSRv1 operations

OLSRv1 operations are detailed in Appendix D. Here, MID messages are used to disseminate global addresses. HNAs messages can be used as well.

7.3.2 OLSRv2 operations

In the current version of OLSRv2, HNA and MID functionalities have been incorporated in TC messages. Procedures defined for OLSRv1 can be modified to work with v2 processing.

7.3.3 DYMO operations

Every DYMO node may insert its addresses in the routing messages (RREP, RREQ, RERR) it forwards to other nodes. Thus, when a node generates a RREQ to create a route to a destination, every node in the path will have information (in particular, the metric, i.e. the hop-count, and a valid route) about IP addresses owned by other nodes in the path. AMG can exploit this feature to advertise configured global addresses to other MANET nodes, and in particular, to the gateways.

In order to apply Best Prefix Selection algorithm, AMG needs to know metrics associated with gateways. In DYMO, this could be accomplished by having the gateways periodically generating a special gratuitous RREQ message with IPDestinationAddress = MANETcastAddress and Target.Address = ALL_ZERO. Nodes receiving this RREQ MUST NOT generate a RREP to gateways, but must create or update routes and metrics towards them and intermediate nodes.

7.4 Gateway operations

Internet gateways have at least one global IPv6 address, belonging to the external network and used on the external interface. While the mechanism, by which such address is acquired, is out of scope of this specification, the configuration of the global address used on the MANET interface is described in this section.

Gateways MUST configure the global IPv6 address of their MANET interface using the mechanism specified in Section 5 and Section 7: a gateway MUST execute the operations described in these sections for MANET nodes. Gateways MUST always select the prefix they own, to configure global address they will use on MANET interface. Finally, gateways MUST process PAs received from other gateways, build global addresses and disseminate them as described in Section 7.3.

As described in Section 4, a gateway can change its mode of operations, becoming a node, for a number of reasons, e.g. because it has lost connectivity with the external network or because of its Internet interface failure. When a gateway is active, but it has indications that the Internet is unreachable, it stops generating PA messages and executes the operations described in Section 7.2 and Section 7.3.

Since the gateway has already disseminated its new global address (after it first received prefixes from other gateways), it can start communicating with the hosts located outside the MANET with negligible latency.

8. Security Considerations

T.B.D.

9. IANA Considerations

This document has no actions for IANA.

10. References

10.1 Normative references

[RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[RFC3626] Clausen, T. and P. Jacquet, "Optimized Link State Routing Protocol (OLSR)", RFC 3626, October 2003.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC4291] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 4291, February 2006.

10.2 Informative References

[I-D.DYMO] Chakeres, I., Belding-Royer, E., and C. Perkins, "Dynamic MANET On-demand (DYMO) Routing", draft-ietf-manet-dymo-04 (work in progress), October 2005.

[I-D.OLSRv2] Clausen, T. and C. Dearlove, "The Optimized Link-State Routing Protocol version 2", draft-ietf-manet-olsrv2-01 (work in progress), March 2006.

[I-D.SMF] Macker, J., "Simplified Multicast Forwarding for MANET", draft-ietf-manet-smf-02 (work in progress), March 2006.

[I-D.jelger-mla] Jelger, C., "MANET Local IPv6 Addresses", draft-jelger-autoconf-mla-00 (work in progress), April 2006.

[I-D.laouiti] Adjih, C., Laouiti, A., and K. Mase, "Address autoconfiguration in Optimized Link State Routing Protocol", draft-laouiti-manet-olsr-address-autoconf-01

(work in progress), July 2005.

[I-D.packetbb]
Clausen, T., Dearlove, C., and J. Dean, "Generalized MANET Packet/Message Format", draft-ietf-manet-packetbb-00 (work in progress), February 2006.

[I-D.perkins-manet-autoconf]
Perkins, C., Malinen, J., Wakikawa, R., and E. Belding-Royer, "IP Address Autoconfiguration for Ad Hoc Networks", I-D draft-perkins-manet-autoconf-01.txt, November 2001.

[I-D.rfc2461bis]
Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", draft-ietf-ipv6-2461bis-06 (work in progress), October 2005.

[I-D.ruffino-conn-scenarios]
Ruffino, S., Stupar, P., Clausen, T., and S. Singh, "Connectivity Scenarios for MANET", draft-ruffino-conn-scenarios-00 (work in progress), January 2006.

[I-D.singh-autoconf-adp]
Singh, S., Clausen, T., Perkins, C., and P. Ruiz, "Ad hoc network autoconfiguration: definition and problem statement", draft-singh-autoconf-adp-03 (work in progress), October 2005.

[I-D.wakikawa-globalv6]
Wakikawa, R., Malinen, J., Perkins, C., Nilsson, A., and A. Tuominen, "Global connectivity for IPv6 Mobile Ad Hoc Networks", I-D draft-wakikawa-manet-globalv6-04.txt, October 2003.

[I-D.wakikawa-v6-support]
Wakikawa, R., "IPv6 Support on Mobile Ad-hoc Network", draft-wakikawa-manet-ipv6-support-00 (work in progress), February 2005.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

Authors' Addresses

Simone Ruffino
Telecom Italia
Via G.Reiss Romoli 274
Torino 10148
Italy

Phone: +39 011 228 7566
Email: simone.ruffino@telecomitalia.it

Patrick Stupar
Telecom Italia
Via G.Reiss Romoli 274
Torino 10148
Italy

Phone: +39 011 228 5727
Email: patrick.stupar@telecomitalia.it

Appendix A. Acknowledgments

The authors would like to thank Ivano Guardini and Ivano Verzola for their valuable comments.

Appendix B. Considerations on address changes

As explained in Section 7.2, nodes and gateways can configure multiple addresses on their MANET interfaces, if multiple gateways and/or prefixes are available. [RFC3484]is used to choose among configured global addresses the source addresses for data communications. The output of Best Prefix Selection algorithm, i.e. prefix ranking, can be used as an hint, to enable nodes to choose always the optimal prefix and address.

Configuring multiple addresses enable nodes to effectively exploit multiple egress points in the MANET and to smoothly change source addresses for data traffic, following MANET topological changes. If a MANET is very dynamic, the best gateway for a node can change very quickly: if multiple address are available, node can simply choose one of them as source address to set-up new data communications.

If only one address is configured on MANET interface (e.g. due to user preferences, configuration or other policies), then, according AMG (Section 7.2.1), node could experience frequent address changes, e.g. in order to optimize downlink traffic coming from external hosts. Generally, such address changes imply active data sessions interruption.

In order to cope with this, Mobile IPv6 [RFC3775] can be used. It is worth noting that AMG, in particular using mechanism explained in Section 7.3, reduces (ideally to zero) the latency introduced by a global address change, granting better performances when MANET nodes use MIPv6. In fact, if a node experiments a change from a first gateway to a second gateway, then it chooses a new global address, associated to the second gateway, and it sends a Binding Update message, registering the new chosen address as the new Care-of Address. When the Binding Acknowledge message from the Home Agent arrives at the gateway, immediately a route to the node will be available, because the new Care-of Address was announced in the MANET (as described in Section 7.3). Therefore handover latency is reduced to the time needed to send a Binding Update message and receive the correspondent Binding Acknowledge message routing latency is negligible.

Appendix C. Prefix Advertisement Examples

Figure 3 and Figure 3 depict two example of PA messages.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   msg-type    |0|0|0|0|1|0|1|0|           msg-size            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                                                               |
|                                                               |
|                     originator-address                        |
|                                                               |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        msg-seq-number         |0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  head-length  |                   head                        :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:               head                            |0|0|0|0|0|0|1|0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            tail 1                             :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                            tail 1                             :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:          tail 1               |            tail 2             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            tail 2                             :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                            tail 2                             :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|0|0|0|0|0|0|0|0|0|0|0|1|0|1|0| VALIDITY_TIME |0|0|0|0|1|1|0|0|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|0|0|0|0|0|1|0|    value_1    |    value_2    | PREFIX_LENGTH |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|0|0|0|1|1|0|0|0|0|0|0|0|0|1|0|    length_1   |    length_2   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: PA containing two different prefixes, with different validity times and prefix lengths

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |   msg-type    |0|0|0|0|1|0|1|0|           msg-size            |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   +                                                               |
   |                                                               |
   +                    originator-address                         |
   |                                                               |
   +                                                               |
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |        msg-seq-number         |0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  head-length  |                                               |
   +-+-+-+-+-+-+-+-+                                               +
   |                                                               |
   +                           head                                +
   |                                                               |
   +                                                               +
   |                                                               |
   +               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |               |0|0|0|0|0|0|0|1|0|0|0|0|0|0|0|0|0|0|0|0|1|0|0|0
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | VALIDITY_TIME |0|0|0|0|0|1|0|0|0|0|0|0|0|0|0|1|    value_1    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | PREFIX_LENGTH |0|0|0|0|0|1|0|0|0|0|0|0|0|0|0|1|    length_1   |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: PA containing one single prefix

Appendix D. OLSRv1 operations

This section details the global address advertising procedure, described in Section 7.3, using OLSRv1 MID messages.

D.1 Data structures: MID Information Base

The Multiple Interface Association Information Base, is defined in [RFC3626]: the present specification modifies the semantics of one of its fields. Multiple Interface Association Information Base is defined in [RFC3626] and is filled processing MID messages. [RFC3626] mandates that these messages are generated by a MANET node only when it is equipped with multiple physical interfaces, through which it is connected to the MANET and participates to OLSR. MIDs contain the addresses configured on the node's physical interfaces. The node is identified by multiple valid IPv6 addresses, one for each interface connected to the MANET: Multiple Interface Association Information Base contains bindings between such addresses and the main address of the node. Using this table, MANET nodes can set-up routes not only towards main address of other nodes, but also towards multiple interface addresses associated to main address. Following [RFC3626], a node connected to the MANET by means of a single interface MUST NOT generate MIDs.

In this specification, as described in Appendix D.2, MID messages generated by a node contain the list of the Global Addresses, i.e. the list of all the global addresses the node may configure on its MANET interface. The Multiple Interface Association Information Base is maintained by each node and gateway and is used to store the bindings between the Global Addresses and MANET-local Addresses of other nodes.

It is worth noting that the semantics of the entries in Multiple Interface Association Information Base, as well as of MID messages, is changed by this specification, since multiple Global Addresses can be configured on a single interface. This semantic change has no effect on the processing of MID messages and it is completely backward-compatible: in fact, from a node's perspective, addresses announced in MID messages can be single addresses configured on multiple interfaces, or multiple addresses configured on a single interface. Routing table construction rules are not changed: nodes build necessary routes to both primary and secondary addresses following [RFC3626].

Hence, Multiple Interface Association Information Base entries have the following semantics (same as specified in [RFC3626]):

Field	Data
I_iface_addr	a Global Address built (and possibly configured) by a node
I_main_addr	the main address of the node which has built the Global Address contained into I_iface_addr field
I_time	Validity time

Table 3: Multiple Interface Association Information Base

D.2 MID messages

By means of standard MID messages processing, when OLSR eventually converges, the node is reachable at any of its Global Addresses : MANET nodes' routing tables contain a route for each secondary address listed into MID messages. A packet whose destination is one of the secondary addresses of a node (e.g. traffic from external hosts to MANET nodes) can therefore be routed within the MANET. Return traffic will be destined to such secondary address and will be routed within the MANET by means of the topological information inserted into MID messages.

D.2.1 MID message generation

A MID message is sent by a node in the network to announce its Global Addresses. I.e., the MID message contains the list of the Global Addresses which have been built by it and inserted into GAIB. The list of Addresses can be partial in each MID message (e.g., due to message size limitations, imposed by the network), but parsing of all MID messages describing the Global Information Base of a node MUST be complete within a certain refreshing period (MID_INTERVAL). The information contained in the MID messages is used by the nodes to route packets, which may be destined to one (or more) of the Global Addresses, chosen by a node to communicate with hosts located outside the MANET.

D.2.2 MID message forwarding

Upon receiving a MID message, following the rules of section 3 of [RFC3626], the message MUST be forwarded according to section 3.4 of

[RFC3626].

D.2.3 MID message processing

MID messages are processed as described in [RFC3626]. The tuples in the multiple interface association set are recorded with the information that is exchanged through MID messages. Upon receiving a MID message, the "validity time" MUST be computed from the Vtime field of the message header (as described in Section 3.3.2 of [RFC3626]). The Multiple Interface Association Information Base SHOULD then be updated as follows:

1. If the sender interface (note: not the originator) of this message is not in the symmetric 1-hop neighborhood of this node, the message MUST be discarded.

2. For each Global Address listed in the MID message:

    1. If there exist some tuple in the interface association set where:

        I_iface_addr == Global Address, AND

        I_main_addr == Originator Address,

       then the holding time of that tuple is set to:

        I_time = current time + validity time.

    2. Otherwise, a new tuple is recorded in the interface association set where:

        I_iface_addr = Global Address,

        I_main_addr = Originator Address,

        I_time = current time + validity time.

Appendix E. Proposed Values for Constants

E.1 Emission Intervals

PA_INTERVAL   = 5 seconds

MID_INTERVAL   = 5 seconds

E.2 Holding Time

PA_HOLD_TIME   = 3 x PA_INTERVAL

MID_HOLD_TIME   = 3 x MID_INTERVAL

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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