

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 6, 2011

U. Herberg
T. Clausen
LIX, Ecole Polytechnique
July 5, 2010

Yet Another Autoconf Proposal (YAAP) for Mobile Ad hoc NETWORKS
draft-herberg-autoconf-yaap-00

Abstract

This document describes automatic configuration of MANET router interfaces, as well as prefix delegation to MANET routers. This autoconfiguration protocol is characterized by (i) adhering strictly to the Internet addressing architecture, (ii) being able to configure both MANET interface addresses and handle prefix delegation, and (iii) being able to configure both stand-alone MANETs, as well as MANETs connected to an infrastructure providing, e.g., globally scoped addresses/prefixes for use within the MANET.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 6, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
2.	Terminology	4
3.	Applicability Statement	5
4.	Protocol Overview and Functioning	5
5.	Protocol Parameters and Constants	6
5.1.	Protocol and Port Numbers	6
5.2.	Multicast Address	6
5.3.	IP Header	6
5.4.	Hold Times	6
5.5.	Time Granularity	7
5.6.	Message Intervals and Timeouts	7
5.7.	Jitter	8
6.	Information Sets	8
6.1.	Neighbor Set	8
6.2.	Initiator Selector Set	9
6.3.	Forwarded Set	9
7.	Unique Identifiers	10
8.	Packets and Messages	10
8.1.	Prefix Solicitation (PS) Message	10
8.1.1.	PS Message Specification	10
8.1.2.	PS Message Generation	11
8.1.3.	PS Message Jitter	11
8.1.4.	PS Message Processing	11
8.2.	Prefix Advertisement (PA) Message	11
8.2.1.	PA Message Specification	11
8.2.2.	PA Message Generation	12
8.2.3.	PA Message Jitter	12
8.2.4.	PA Message Processing	12
8.3.	Router Solicitation (RS) Message	13
8.3.1.	Local RS Message from Configuring Router to Initiator Router	13
8.3.2.	MANET-wide RS Message from Initiator Router	16
8.4.	Router Advertisement (RA) Message	17
8.4.1.	MANET-wide RA Message from Defensive Router	17
8.4.2.	Local RA Message from Initiator Router to Configuring Router	18
8.5.	Acknowledgement (AC) Message	19
8.5.1.	AC Message Specification	19
8.5.2.	AC Message Generation	19
8.5.3.	AC Message Jitter	19
8.5.4.	AC Message Processing	20

9. Message Considered for Forwarding	20
10. Router States	21
10.1. State Change	21
10.1.1. From Configuring State to Defensive State	21
10.1.2. From Defensive State to Configuring State	21
11. Initiator Selection	22
12. Prefix Conflict	22
13. Protocol Optimizations	22
13.1. Proxying	22
13.2. Unicast RA Messages	22
13.3. Optimized Broadcasting	22
14. Additional Considerations	23
14.1. Duplicate UUIDs	23
14.2. No initiator router exist	23
14.3. Initiator Proxying	23
15. Proposed Values for Parameters and Constants	23
16. IANA Considerations	23
16.1. Expert Review: Evaluation Guidelines	23
16.2. Message Types	23
16.3. Message TLV Types	24
16.4. Address Block TLV Types	25
17. Security Considerations	26
18. Normative References	26
Authors' Addresses	27

1. Introduction

This document adheres to the address architecture specified in [RFC5889], i.e.:

- o IP addresses configured on an interface are unique, at least within the routing domain, and,
- o No on-link subnet prefix is configured on this interface .

This specification allows to configure MANET-local and global prefixes to MANET routers. A router using the specified mechanism (i) acquires a MANET-wide unique prefix, and (ii) configures addresses from within this prefix for its interfaces as well as hosts attached to these routers (using any mechanism, such as SLAAC [RFC4861] or DHCPv6 [RFC3315]). This document does not stipulate how to configure link-local addresses or link-local prefixes.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

This document uses the terminology and notation defined in [RFC5444]. Additionally, it defines the following terminology:

- o Configuring Router

A Configuring Router is a router which has not yet configured a prefix, and which uses the protocol specified in this document in order to acquire a prefix.

- o Initiator Router

A Configuring Router selects one Initiator Router in the process of acquiring a prefix, which assists in validating the uniqueness of the chosen tentative prefix.

- o Defensive Router

A Defensive Router is a router which has already configured a prefix, and which can assist other, unconfigured routers (i.e. Configuring Routers) and thereby become an Initiator Router for unconfigured routers.

3. Applicability Statement

This autoconfiguration mechanism is applicable for MANET routers. The mechanism allows to configure MANET-wide and globally unique prefixes on a router, adhering to the address architecture specified in [RFC5889].

This mechanism is applicable for IPv6.

4. Protocol Overview and Functioning

Each MANET router will acquire a prefix, constructed as `d:p:s::` with `d` being a prefix (possibly of size 0), common for the whole site (e.g., a global prefix assigned administratively to a given site, or provided by an Internet gateway), `p` being common for all routers in this MANET, and `s` being unique to a specific router.

The task for a router, appearing in a MANET, can thus be summarized as:

- o Acquiring `d` and `p` from the network;
- o Selecting `s`, unique within the MANET;
- o Configuring own MANET interfaces with addresses from within `d:p:s::` (and with a prefix length of /128).

A router, appearing in a network and wishing to be configured, is denoted a "Configuring router". Through a Prefix Solicitation (PS) / Prefix Advertisement (PA) message exchange, the router learns of the (already configured) routers in its vicinity, and selects one as initiator - the router which will assist in acquiring a valid configuration. The Configuring router extracts `d` and `p` from the PA received from the selected initiator, generates a tentative prefix `d:p:s::` (`s` being generated locally by the Configuring router, e.g., by a pseudorandom generator) and, by way of a Router Solicitation (RS) message requests from its initiator that this tentative prefix be verified unique within the MANET. The initiator then issues an RS to the entire MANET, containing the tentative address of its Configuring router, through the MANET. If the initiator does not (after due delay and retransmissions) receive a Router Advertisement (RA) indicating that the prefix is already in use, it will transmit an Autoconfiguration Confirmation (AC) message to the Configuring router, confirming that the Configuring router now "owns" `d:p:s::` and can become a Defensive router. If the initiator receives an RA indicating that the tentative prefix is already in use, it informs the Configuring router by issuing an RA.

A Defensive router has two tasks. First, if receiving a RS containing its own d:p:s::, it must respond by issuing an RA. Second, if receiving a PS, it must respond by a PA, thus accept becoming initiator and act as described above.

The initiator and Configuring routers communicate using link-local multicast to LL-MANET-Routers [RFC5498], with the configuring router using the Unspecified Address as source [RFC3513]. These two routers identify traffic destined to each other by way of UUIDsS. [RFC4122], embedded in all messages exchanged between them. UUIDs are 16 octets long, and as they are exchanged in messages only between neighboring routers, they need only be locally unique. Network-wide messages (RS/RA) are "proxyed" by the initiator, which is already configured and which thereby has a MANET-wide unique address.

5. Protocol Parameters and Constants

This specification uses the parameters and constants described in this section.

5.1. Protocol and Port Numbers

This protocol specifies PS, PA, RS, RA, and AC messages, which are included in packets as defined by [RFC5444]. These packets may be sent either using the "manet" protocol number or the "manet" well-known UDP port number, as specified in [RFC5498].

5.2. Multicast Address

The packets (as defined by [RFC5444]) which contain the messages specified by this document may be locally transmitted using LL-MANET-Routers, as specified in [RFC5498].

5.3. IP Header

If the router is in the Configuring state, the source address in the IP header containing control messages is set to the Unspecified Address, as specified in [RFC3513]. For a router in any other state, the source address is set to any (non link-local) IP address of the interface transmitting the packet.

5.4. Hold Times

- o N_HOLD_TIME - determines how long a Neighbor Tuple is kept in the Neighbor Set in milliseconds.

- o I_HOLD_TIME - determines the maximum time duration that a Initiator Selector Tuple is kept in the Initiator Selector Set in milliseconds.
- o F_HOLD_TIME - is the period after receipt of a message that is forwarded by this router for which that information is recorded, in order that the message is not forwarded again if received again.

The following constraints applies to these router parameters:

- o N_HOLD_TIME >= 0
- o I_HOLD_TIME >= 0
- o F_HOLD_TIME >= 0

5.5. Time Granularity

The constant C (time granularity) is used as specified in [RFC5497].

5.6. Message Intervals and Timeouts

- o PS_INTERVAL - specifies the time interval between two successive PS messages in milliseconds.
- o PS_TIMEOUT - specifies the maximum time, in milliseconds, after the first transmitted PS which a router waits for an incoming PA message. If no PA message has been received, this router switches from Configuring State to Defensive state as described in Section 10.1.1.
- o RS_INTERVAL - specifies the time interval between two successive local RS messages in milliseconds.
- o RS_TIMEOUT - specifies the maximum time, in milliseconds, after the first transmitted local RS message which a router waits for an incoming RA or AC message. If no such message has been received, this router switches from Configuring State to Defensive state.
- o RA_INTERVAL - specifies the time interval between two successive RA messages in milliseconds.
- o RA_TIMEOUT - specifies the maximum time, in milliseconds, after the first transmitted RA during which this router transmits RA messages.

- o AC_INTERVAL - specifies the time interval between two successive AC messages in milliseconds.
- o AC_TIMEOUT - specifies the maximum time, in milliseconds, after the first transmitted AC during which this router transmits AC messages.

5.7. Jitter

- o PS_MAXJITTER - represents the value of MAXJITTER used in [RFC5148] for generated PS messages on this MANET interface.
- o PA_MAXJITTER - represents the value of MAXJITTER used in [RFC5148] for generated PA messages on this MANET interface.
- o F_MAXJITTER - represents the default value of MAXJITTER used in [RFC5148] for messages forwarded by this router.

6. Information Sets

6.1. Neighbor Set

A router's Neighbor Set records site prefixes and MANET prefixes of each of its 1-hop Defensive neighbor routers. It consists of Neighbor Tuples, each representing a single such 1-hop neighbor:

(N_UUID, N_is_initiator, N_site_prefix, N_site_prefix_length,
N_manet_prefix, N_manet_prefix_length, N_iface, N_time)

where:

N_UUID - is the unique identifier of the neighbor, as received in a PA message.

N_is_initiator - is a boolean flag that determines whether this neighbor is selected as initiator.

N_site_prefix - is the site prefix d as indicated by a PA message.

N_site_prefix_length - is the length of N_site_prefix in number of bits.

N_manet_prefix - is the MANET prefix p as indicated by a PA message.

N_manet_prefix_length - is the length of N_manet_prefix in number of bits.

N_iface - is the identifier of the network interface on which the neighbor can be reached.

N_time - specifies when this Tuple expires and MUST be removed.

6.2. Initiator Selector Set

A router's Initiator Selector Set records UUIDs and tentative prefixes of each 1-hop Configuring neighbor router, which has selected this router as its initiator. It consists of Initiator Selector Tuples, each representing a single 1-hop neighbor:

(I_UUID, I_prefix, I_prefix_length, I_iface, I_time)

where:

I_UUID - is the unique identifier of the initiator selector, as received in a local RS message.

I_prefix - is the tentative prefix s as indicated by the initiator selector in an RS message.

I_prefix_length - is the length of I_prefix in number of bits.

I_iface - is the identifier of the network interface on which the neighbor can be reached.

I_time - specifies when this Tuple expires and MUST be removed.

6.3. Forwarded Set

A router has a single Forwarded Set which records signatures of messages which have been forwarded by the router. It consists of Forwarded Tuples:

(F_type, F_orig_addr, F_seq_number, F_time)

where:

F_type - is the forwarded Message Type;

F_orig_addr - is the originator address of the forwarded message;

F_seq_number - is the message sequence number of the forwarded message;

F_time - specifies the time at which this Tuple expires and MUST be removed.

7. Unique Identifiers

Throughout this document, it is supposed that routers have a unique identifier (UUID) of length of 16 octets, according to [RFC4122]. This identifier can be created by a pseudo-random number generator or any other means (such as calculated from the MAC address of a network interface). The identifier serves to discern routers without configured IP addresses in a 1-hop neighborhood. They SHOULD be unique within a link-local scope. If they are not unique within this scope, the mechanism will still work, however, as described in Section 14.1.

8. Packets and Messages

The packet and message format used by this protocol is defined in [RFC5444], which is used with the following considerations:

- o This protocol specifies five Message Types, the Prefix Solicitation (PS) message, the Prefix Advertisement (PA) message, the Router Solicitation (RS) message, the Router Advertisement (RA) message, and the Acknowledgment (AC) message.
- o PS, PA, and AC messages MUST NOT be forwarded, i.e., a <msg-hop-limit>, if present, MUST have the value 1.
- o All five message types MAY be included in multi-message packets as specified in [RFC5444].
- o Received messages MUST be parsed in accordance with [RFC5444]. A message which is not in conformance with [RFC5444] MUST be discarded without being processed.
- o This protocol specifies two Address Block TLVs. These TLV Types are all defined only with Type Extension = 0. TLVs of other types, and of these types but without Type Extension = 0, are ignored by this protocol.

8.1. Prefix Solicitation (PS) Message

A PS message is broadcast by a Configuring router to its 1-hop neighbors in order to discover already configured routers, and if such exist, to acquire their site prefix and MANET prefix.

8.1.1. PS Message Specification

The <msg-addr-length> field is set to 15. The PS message MUST not contain a <msg-orig-addr> header field. <msg-hop-limit> SHOULD be set

to 1.

The PS message does not contain any addresses, Address Block TLVs, or Message TLVs (but MAY so by an extension of this protocol).

8.1.2. PS Message Generation

A PS message is generated per router and sent over all MANET interfaces when the router is not yet configured and desires to acquire a MANET-wide unique prefix.

If no PA message has been received on any interface after PS_INTERVAL, the router generates a new PS message.

If no PA message has been received within PS_TIMEOUT after the first generated PS message, the router assumes that it is the first router of the MANET and switches to the "Defensive state" (as described in Section 10.1.1).

8.1.3. PS Message Jitter

If using data link and physical layers which are subject to packet loss due to collisions, PS messages SHOULD be jittered as described in [RFC5148]. PS messages MUST NOT be forwarded, and thus message forwarding jitter does not apply to PS messages. Each form of jitter described in [RFC5148] requires a parameter MAXJITTER. These parameters may be dynamic, and are specified by PS_MAXJITTER.

8.1.4. PS Message Processing

A router in Configuring state MUST ignore incoming PS messages.

A router in the Defensive state MAY trigger a PA message upon reception of an incoming PS message. The PA message is transmitted on the same interface the PS message has been received as described in Section 8.2.

8.2. Prefix Advertisement (PA) Message

A PA message is broadcast by a Defensive router to its 1-hop neighbors, as a response to an incoming PS message.

8.2.1. PA Message Specification

The <msg-addr-length> field is set 15. <msg-hop-limit> SHOULD be set to 1.

The PA message MUST contain a UUID Message TLV with the router's

Unique identifier (UUID) as TLV value and length 16 and type extension 1 (which indicates that it is a source UUID).

The PA message MUST contain the MANET prefix, contained in an Address Block and associated with a MANET_PREFIX TLV. The type extension of the MANET_PREFIX TLV determines the length of the prefix in number of bits.

The PA message MAY contain the site prefix, contained in an Address Block and associated with a SITE_PREFIX TLV. The type extension of the SITE_PREFIX TLV determines the length of the prefix in number of bits.

8.2.2. PA Message Generation

A PA message is generated by a Defensive router as a response to a received PS message, and transmitted over the same interface as the one on which the PS message was received.

8.2.3. PA Message Jitter

If using data link and physical layers which are subject to packet loss due to collisions, PA messages SHOULD be jittered as described in [RFC5148]. PA messages MUST NOT be forwarded, and thus message forwarding jitter does not apply to PA messages. Each form of jitter described in [RFC5148] requires a parameter MAXJITTER. These parameters may be dynamic, and are specified by PA_MAXJITTER.

8.2.4. PA Message Processing

If the router is in the Configuring state, the message is processed, otherwise it MUST be ignored.

If the message is processed, the following steps MUST be performed:

1. Find the Neighbor Tuple which corresponds to the UUID contained in the UUID Message TLV of the message (henceforth current tuple).
2. If no such tuple exists, create a new one (henceforth current tuple) and add it to the Neighbor Set.
3. The values of the current tuple are set to:
 - * N_UUID: the unique identifier of the neighbor, as acquired from the UUID Message TLV in the PA message.

- * `N_is_initiator`: false
- * `N_site_prefix`: is set to the prefix associated with the `SITE_PREFIX` Address Block TLV in the previously received PS message that triggered this PA transmission; this value may be NULL if no such TLV has been contained in the PS message.
- * `N_site_prefix`: is set to the length of the `N_site_prefix` in number of bits, as acquired from the type extension of the `SITE_PREFIX` TLV or 0 if no `SITE_PREFIX` TLV has been contained in the PS message.
- * `N_manet_prefix`: is set to the MANET prefix associated with the `MANET_PREFIX` Address Block TLV in the previously received PS message that triggered this PA transmission.
- * `N_manet_prefix_length`: is set to the length of `N_manet_prefix` in number of bits, as acquired from the type extension of the `MANET_PREFIX` Address Block TLV.
- * `N_iface`: is set to the interface identifier on which the PA message has been received.
- * `N_time`: current time + `N_HOLD_TIME` (both in milliseconds)

8.3. Router Solicitation (RS) Message

A Configuring router sends a local Router Solicitation (RS) Message to its initiator (using link-local multicast and the UUID of the initiator in a Message TLV). Upon reception of this local RS, the initiator broadcasts a Router Solicitation (RS) Message throughout the MANET, on behalf of the Configuring router. In the following, these two different cases (RS between Configuring router and initiator router, and RS flooded by the initiator) are discerned. In both cases, the same message type (RS) is used, as only the link-local RS contains the UUID TLV, which allows to differentiate between the two different message scopes.

8.3.1. Local RS Message from Configuring Router to Initiator Router

8.3.1.1. Local RS Message Specification

The `<msg-addr-length>` field is set to 15. The local RS message MUST NOT contain a `<msg-orig-addr>` header field. `<msg-hop-limit>` SHOULD be set to 1.

The local RS message MUST contain a UUID Message TLV with the UUID of the router's selected initiator as value and type extension 0 (i.e.

indicating a destination UUID).

The local RS message MUST contain a UUID Message TLV with this router's UUID as value and type extension 1 (i.e. indicating a source UUID).

The local RS message MUST contain a TIMEOUT Message TLV with the value being:

$$\text{time_of_first_sent_RS} - \text{current_time} + \text{RS_TIMEOUT}.$$

where

- o `time_of_first_sent_RS` is the time of the first RS that has been sent for the tentative prefix, or the current time if this message is the first one.

The value is encoded using the format specified in [RFC5497].

The local RS message MUST contain the tentative router prefix `s` in an address block associated with a `TENTATIVE_PREFIX Address Block TLV` with the value hereof being the length of the prefix in number of bits.

8.3.1.2. Local RS Message Generation

An RS message is generated after the router has selected an initiator (as specified in Section 11). The message is sent to the MANET interface on which the initiator can be reached (determined from `N_iface` in the Neighbor Set).

If no RA or AC message has been received on the `N_iface` interface from the selected initiator (determined from the Neighbor Tuple of the initiator) after `RS_INTERVAL` milliseconds, the router generates a new local RS message. This process is repeated until an RA or AC message has been received, or otherwise `RS_TIMEOUT` milliseconds have passed since the first generated local RS message for this prefix.

If then still no RA or AC message has been received, the router SHOULD restart the autoconfiguration process and start transmitting PS messages again (refer to Section 8.1) or select a new initiator router.

8.3.1.3. Local RS Message Jitter

No jitter is required for local RS messages, since only one router - the selected initiator - will process them.

8.3.1.4. Local RS Message Processing

If the message does not contain any UUID Message TLVs, it is considered a MANET-wide RS Message and processed accordingly (refer to Section 8.3.2).

If the router receiving an RS is in Configuring state, the message MUST be ignored.

If the message does not contain a UUID Message TLV with type extension 0, it MUST be ignored. Otherwise, if the value of that UUID Message TLV does not correspond to this router's UUID, the message MUST be ignored.

If the message does not contain an address associated with a TENTATIVE_PREFIX Address Block TLV, the message MUST be ignored.

If a prefix conflict between the router's prefix s and the address associated with a TENTATIVE_PREFIX Address Block TLV has been detected, a local RA message MUST be triggered (as specified in Section 8.4.2).

If the message is not ignored, the following steps MUST be performed:

1. If no Initiator Selector Tuple exists, which corresponds to the UUID contained in the UUID Message TLV with type extension 0 in the message, create a new tuple and add it to the Initiator Selector Set. The values of the newly created tuple are:
 - * I_UUID: the unique identifier of the neighbor, as acquired from the UUID Message TLV in the RS message.
 - * I_prefix: is the prefix associated with the TENTATIVE_PREFIX Address Block TLV in the RS message.
 - * I_prefix_length is the length of the I_prefix in number of bits, as acquired from the value of the TENTATIVE_PREFIX Address Block TLV.
 - * I_iface: is set to the interface identifier on which the RS message has been received.
 - * I_time: $\text{current_time} + \max(\text{I_HOLD_TIME}, \text{value_of_the_TIMEOUT_Message_TLV})$ # in milliseconds
2. A MANET-wide RS message is generated (refer to Section 8.3.2.2).

8.3.2. MANET-wide RS Message from Initiator Router

8.3.2.1. MANET-wide RS Message Specification

The <msg-addr-length> field is set to 15. The MANET-wide RS message MUST contain a <msg-orig-addr> header field, which contains the prefix s of this router. <msg-hop-limit> SHOULD set to 255. The MANET-wide RS message MUST contain a <msg-seq-num> field.

The MANET-wide RS message MUST NOT contain any UUID Message TLVs.

The MANET-wide RS message MUST contain the tentative prefix in an address block associated with a TENTATIVE_PREFIX Address Block TLV with the value being the length of the prefix in number of bits. The prefix is acquired from the I_prefix field in the Initiator Selector Tuple that corresponds to the initiator selector that has sent the local RS message, which triggered this MANET-wide RS message.

8.3.2.2. MANET-wide RS Message Generation

Whenever a router has received a local RS message from a router that has selected it as its initiator (i.e. for which a corresponding Initiator Selector tuple exists), it generates a MANET-wide RS message and sends it to all MANET interfaces.

8.3.2.3. MANET-wide RS Message Jitter

No jitter is required for generating MANET-wide RS messages. However, for forwarded RS messages, jitter considerations as specified in Section 9 apply.

8.3.2.4. MANET-wide RS Message Processing

If the message contains a UUID Message TLV, it is considered as local RS Message and processed accordingly (refer to Section 8.3.1.4).

If this router is in the Configuring state, the message MUST be ignored.

If the message does not contain an address associated with a TENTATIVE_PREFIX Address Block TLV, the message MUST be ignored.

If a prefix conflict between the router's prefix s and the address associated with a TENTATIVE_PREFIX Address Block TLV has been detected, a MANET-wide RA message MUST be generated, specified in Section 8.4.1.

If no conflict has occurred, the message is considered for forwarding

as specified in Section 9.

8.4. Router Advertisement (RA) Message

A Defensive router floods a MANET-wide Router Advertisement (RA) Message throughout the MANET if it has detected a conflicting prefix (as specified in Section 12) advertised in an RS message. When the initiator router for the conflicting prefix receives this RA, it generates a local RA message to its initiator selector (using link-local multicast and the UUID of the initiator in a Message TLV). In the following, these two different cases (RA flooded by the conflicting router, and RA between initiator router and Configuring router) are discerned. In both cases, the same message type (RA) is used, as only the link-local RA contains the UUID TLV, which allows to differentiate between the two different message scopes.

8.4.1. MANET-wide RA Message from Defensive Router

8.4.1.1. MANET-wide RA Message Specification

The <msg-addr-length> field is set to 15. The MANET-wide RA message MUST contain a <msg-orig-addr> header field. <msg-hop-limit> SHOULD be set to 255. The MANET-wide RA message MUST contain a <msg-seq-num> field.

The MANET-wide RA message MUST NOT contain a UUID Message TLV.

The MANET-wide RA message MUST put its prefix (i.e. the prefix which is conflicting) in an address block and associate with a TENTATIVE_PREFIX Address Block TLV.

8.4.1.2. MANET-wide RA Message Generation

A MANET-wide RA message is generated as a response to an incoming MANET-wide RS message if a prefix conflict has been detected (as specified in Section 12). The message is sent to all MANET interfaces.

The router generates a new RA message RA_INTERVAL milliseconds after the last generated RA message. This process is repeated until RA_TIMEOUT milliseconds have passed since the first generated RA message for this prefix.

8.4.1.3. MANET-wide RA Message Jitter

No jitter is applied to generated RA messages. When forwarding an RA message, jitter is applied as specified in Section 9.

8.4.1.4. MANET-wide RA Message Processing

If the message contains a UUID TLV, then:

- o if this router is a Configuring router, the message is considered as local RA Message and processed accordingly (refer to Section 8.4.2).
- o or otherwise the message MUST be ignored.

If the message does not contain a UUID TLV, then:

- o if the message contains an address associated with a TENTATIVE_PREFIX TLV, and the address is contained in the I_prefix field of an Initiator Selector Tuple, then the router generates a local RA as specified in Section 8.4.2.
- o otherwise, the message is considered for forwarding as specified in Section 9.

8.4.2. Local RA Message from Initiator Router to Configuring Router

8.4.2.1. Local RA Message Specification

The <msg-addr-length> field is set to 15. <msg-hop-limit> SHOULD be set to 1.

The local RA message MUST contain a UUID Message TLV with type extension 0 with the value being the UUID of the initiator selector. The UUID is determined from the I_UUID field of the Initiator Selector Tuple whose I_prefix is equivalent to the address of the MANET-wide RA which is associated with an TENTATIVE_PREFIX Address Block TLV.

8.4.2.2. Local RA Message Generation

A local RA message is generated as a response to an incoming MANET-wide RA message, if the message contains a tentative prefix, which is equivalent to an I_prefix entry of an Initiator Selector Tuple. The message is sent to the I_iface interface of the initiator selector.

8.4.2.3. Local RA Message Jitter

No jitter is applied.

8.4.2.4. Local RA Message Processing

If the message does not contain a UUID Message TLV, then the message is processed as specified in Section 8.4.1.4.

If the message contains a UUID TLV, then:

- o if this router is a Configuring router and if the TLV value from the UUID Message TLV corresponds to this router's UUID, then this router MUST:
 - * Choose another tentative prefix *s*, and
 - * Restart the autoconfiguration process by generating RS messages with this new tentative prefix.
- o or otherwise the message MUST be ignored.

8.5. Acknowledgement (AC) Message

8.5.1. AC Message Specification

The <msg-addr-length> field is set to 15. <msg-hop-limit> SHOULD be set to 1.

The AC message MUST contain a UUID Message TLV with type extension 0 with the UUID of the Configuring router as TLV value.

The AC message MUST contain the tentative prefix in an address block associated with an TENTATIVE_PREFIX Address Block TLV with the length of the prefix in number of bits as TLV value.

The AC message MAY contain any other TLV as specified by extensions to this protocol.

8.5.2. AC Message Generation

An AC message is generated on a Defensive router when no RA message has been received for a tentative prefix of an initiator at the time the Initiator Selector Tuple expires (i.e. at *I_time*). The message is sent to the *I_iface* interface of the initiator selector.

8.5.3. AC Message Jitter

No jitter is applied.

8.5.4. AC Message Processing

If the receiving router is not a Configuring router, then the message MUST be ignored.

Otherwise, if:

- o the message contains a UUID TLV and the value from the UUID Message TLV corresponds to this router's UUID and
- o the message contains an address associated with a TENTATIVE_PREFIX Address Block TLV, and the address is equivalent to this router's tentative prefix, then this router MUST:
 - * Select this prefix as permanent.
 - * Switch in the Defensive state (as specified in Section 10.1.1)
- o or otherwise the message MUST be ignored.

9. Message Considered for Forwarding

If a message (the "current message") is considered for forwarding, then the following tasks MUST be performed:

1. If a Forwarded Tuple exists with:
 - * F_type = the Message Type of the current message, AND;
 - * F_orig_addr = the originator address of the current message, AND;
 - * F_seq_number = the sequence number of the current message.then the current message MUST be silently discarded.
2. Otherwise:
 1. The Message Header of the current message is modified by:
 - + if present, decrement <msg-hop-limit> in the Message Header by 1, AND;
 - + if present, increment <msg-hop-count> in the Message Header by 1.

2. The message is transmitted over all interfaces.

If using data link and physical layers which are subject to packet loss due to collisions, forwarded messages SHOULD be jittered as described in [RFC5148]. PS messages MUST NOT be forwarded, and thus message forwarding jitter does not apply to PS messages. Each form of jitter described in [RFC5148] requires a parameter MAXJITTER. These parameters may be dynamic, and are specified by F_MAXJITTER.

10. Router States

A router is in either of the following states:

- o Configuring
- o Defensive

10.1. State Change

10.1.1. From Configuring State to Defensive State

A router can change from Configuring state to Defensive state (i.e. complete the autoconfiguration process) in two different ways:

1. If no PA message has been received within PS_TIMEOUT milliseconds after the first generated PS message of a Configuring router, the router assumes that it is the first router of the MANET. It then uses a predefined site-local prefix part "d", creates a random MANET prefix part "p" of length PREFIX_LENGTH using a pseudo-random number generator and concatenates a random router prefix part "s". The Neighbor Set can be cleared (i.e. all Neighbor Tuples deleted). The router switches to the "Defensive state".
2. If the Configuring router receives an AC message from its initiator router for the tentative prefix, the Configuring router uses a predefined site-local prefix part "d" and the MANET prefix part "p" that it has acquired by PA messages from the Initiator. It then concatenates a the tentative router prefix part "s", which it has verified to be unique within the MANET. The Neighbor Set can be cleared (i.e. all Neighbor Tuples deleted). The router switches to the "Defensive state".

10.1.2. From Defensive State to Configuring State

TBD: Router releases the prefix and needs a new prefix.

11. Initiator Selection

After a Configuring router has received at least one PA message, it may at any time select its initiator router among the neighbors from which it has received a PA message. As soon as it has selected the initiator router, it can start generating RS messages. This document does not specify how to select the initiator router. For example, a router could select the first router that from which it has received a PA message. PA messages MAY also include additional information by means of TLVs that can be used to determine the best initiator router.

Examples of such information may be:

- o number of routers in the MANET;
- o density of the routers in the MANET;
- o distance to an Internet gateway;
- o distance to a certain address prefix.

12. Prefix Conflict

A prefix conflict occurs when a requested prefix (by means of an RS message) is equivalent to the router prefix of this router.

TBD

13. Protocol Optimizations

The following protocol optimizations MAY be applied to the base protocol in order to increase performance.

13.1. Proxying

TBD

13.2. Unicast RA Messages

TBD

13.3. Optimized Broadcasting

TBD

14. Additional Considerations

14.1. Duplicate UUIDs

TBD

14.2. No initiator router exist

TBD

14.3. Initiator Proxying

TBD

15. Proposed Values for Parameters and Constants

TBD

16. IANA Considerations

This specification defines five Message Types, which must be allocated from the "Message Types" repository of [RFC5444], four Message TLV Types, which must be allocated from the "Message TLV Types" repository of [RFC5444], and one Address Block TLV Type, which must be allocated from the "Address Block TLV Types" repository of [RFC5444]

16.1. Expert Review: Evaluation Guidelines

For the registries where an Expert Review is required, the designated expert SHOULD take the same general recommendations into consideration as are specified by [RFC5444].

16.2. Message Types

This specification defines five Message Types, to be allocated from the 0-223 range of the "Message Types" namespace defined in [RFC5444], as specified in Table 1.

Type	Description
TBD1	PS: Prefix Solicitation message
TBD2	PA: Prefix Advertisement message
TBD3	RS: Router Solicitation message
TBD4	RA: Router Advertisement message
TBD5	AC: Acknowledgment message

Table 1: Message Type Assignment

16.3. Message TLV Types

This specification defines four Message TLV Types, which must be allocated from the "Message TLV Types" namespace defined in [RFC5444]. IANA is requested to make allocations in the 0-127 range for these types. This will create four new Type Extension registries with assignments as specified in Table 2, Table 3, Table 4, and Table 5. Specifications of these TLVs are in Section xxx and Section xxx, respectively. Each of these TLVs MUST NOT be included more than once in a Message TLV Block.

Name	Type	Type Extension	Description	Allocation Policy
UUID	TBD6	0	Specifies the destination UUID of a neighbor router	Expert Review
UUID	TBD7	1	Specifies the source UUID of this router	
UUID	TBD8	2-255	Unassigned	

Table 2: Message TLV Type assignment: UUID

Name	Type	Type Extension	Description	Allocation Policy
MANET_PREFIX	TBD9	0-128	Specifies the MANET prefix part with the type extension being the length of the prefix in number of bits.	

MANET_PREFIX	TBD10	1-255	Unassigned	Expert Review
--------------	-------	-------	------------	---------------

Table 3: Message TLV Type assignment: MANET_PREFIX

Name	Type	Type Extension	Description	Allocation Policy
SITE_PREFIX	TBD11	0-128	Specifies the SITE prefix part with the type extension being the length of the prefix in number of bits.	
SITE_PREFIX	TBD12	129-255	Unassigned	Expert Review

Table 4: Message TLV Type assignment: SITE_PREFIX

Name	Type	Type Extension	Description	Allocation Policy
TIMEOUT	TBD13	0	Specifies the timeout of the Configuring router which are waiting for AC or RA messages.	
TIMEOUT	TBD14	1-255	Unassigned	Expert Review

Table 5: Message TLV Type assignment: TIMEOUT

16.4. Address Block TLV Types

This specification defines one Address Block TLV Type, which must be allocated from the "Address Block TLV Types" namespace defined in [RFC5444]. IANA are requested to make allocations in the 8-127 range for this type. This will create a new Type Extension registry with assignments as specified in Table 6. Specification of this TLV is in Section xxx.

Name	Type	Type Extension	Description	Allocation Policy
TENTATIVE_PREFIX	TBD15	0	Specifies the tentative prefix that is checked for uniqueness.	
TENTATIVE_PREFIX	TBD16	1-255	Unassigned	Expert Review

Table 6: Address Block TLV Type assignment: TENTATIVE_PREFIX

17. Security Considerations

TBD

18. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, BCP 14, March 1997.
- [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.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique Identifier (UUID) URN Namespace", RFC 4122, July 2005.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC5148] Clausen, T., Dearlove, C., and B. Brian, "Jitter Considerations in Mobile Ad Hoc Networks (MANETs)", RFC 5148, February 2008.
- [RFC5444] Clausen, T., Dearlove, C., Dean, J., and C. Adjih, "Generalized MANET Packet/Message Format", RFC 5444,

February 2009.

- [RFC5497] Clausen, T. and C. Dearlove, "Representing multi-value time in MANETs", RFC 5497, March 2009.
- [RFC5498] Chakeres, I., "IANA Allocations for Mobile Ad Hoc Network (MANET) Protocols", RFC 5498, March 2009.
- [RFC5889] Baccelli, E. and M. Townsley, "IP Addressing Model in Ad Hoc Networks", RFC 5889, July 2010.

Authors' Addresses

Ulrich Herberg
LIX, Ecole Polytechnique
91128 Palaiseau Cedex,
France

Phone: +33-1-6933-4126
Email: ulrich@herberg.name
URI: <http://www.herberg.name/>

Thomas Clausen
LIX, Ecole Polytechnique
91128 Palaiseau Cedex,
France

Phone: +33-6-6058-9349
Email: T.Clausen@computer.org
URI: <http://www.thomasclausen.org>

