IEEE P802.11s™/D1.01

Draft Amendment to STANDARD for Information Technology- - Telecommunications and information exchange between systems- - Local and metropolitan area networks- Specific requirements-

Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications:

Amendment <number>: ESS Mesh Networking

EDITORIAL NOTE—the amendment number will be inserted by IEEE-SA editorial staff during preparation for publication.

Prepared by the 802.11 Working Group of the IEEE 802 Committee

3 Park Avenue

New York, NY 10016-5997, USA

IEEE Standards Activities Department
Standards Licensing and Contracts
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA

Abstract: This amendment defines an IEEE 802.11 Wireless LAN (WLAN) Mesh using the IEEE 802.11 MAC/PHY layers that supports both broadcast/multicast individually addressed and unicast group addressed delivery over self-configuring multi-hop topologies.

Keywords: Wireless LAN, Medium Access Control, Mesh, Multi-hop

Introduction

(This introduction is not part of IEEE P802.11s/D1.01, Draft Amendment to Standard for Information Technology - Telecommunications and Information Exchange Between Systems - LAN/MAN Specific Requirements - Part 11: Wireless Medium Access Control (MAC) and physical layer (PHY) specifications: Amendment: ESS Mesh Networking.)

This amendment specifies enhancements to the following draft standard and draft amendments, in order to support mesh networking:

— — IEEE P802.11 REV-ma D8D9.0

— — IEEE P802.11k D4D7.10

— IEEE P802.11r D4.1

The networks described in this amendment make use of layer-2 mesh path selection and forwarding (that is, a mesh network that performs routing at the link layer). Mesh networks have advantageous properties in terms of robustness, range extension and density, but also have potential challenges such as power consumption and security. This amendment is specifically designed to address these challenges.

Notice to users

Errata

Errata, if any, for this and all other standards can be accessed at the following URL: http://standards.ieee.org/reading/ieee/updates/errata/index.html. Users are encouraged to check this URL for errata periodically.

Interpretations

Current interpretations can be accessed at the following URL: http://standards.ieee.org/reading/ieee/interp/index.html.

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents or patent applications for which a license may be required to implement an IEEE standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention. A patent holder or patent applicant has filed a statement of assurance that it will grant licenses under these rights without compensation or under reasonable rates and nondiscriminatory, reasonable terms and conditions to applicants desiring to obtain such licenses. The IEEE makes no representation as to the reasonableness of rates, terms, and conditions of the license agreements offered by patent holders or patent applicants. Further information may be obtained from the IEEE Standards Department.

Participants

At the time this draft amendment to standard was completed, the 802.11 Working Group had the following membership:

Stuart J. Kerry, *Chair*
Al Petrick and Harry Worstell, *Vice-chair*
Tim Godfrey, *Secretary*

EDITORIAL NOTE—a three column list of voting members of 802.11 on the day the draft was sent for sponsor ballot will be inserted

The following were officers of Task Group s:

Donald E. Eastlake 3rd, *Chair*
Stephen Rayment, *Secretary*
W. Steven Conner, *Technical Editor*

The following members of the balloting committee voted on this Standard. Balloters may have voted for approval, disapproval, or abstention.

EDITORIAL NOTE—a three-column list of responding sponsor ballot members will be inserted by IEEE staff

Editorial Notes

EDITORIAL NOTE—Two forms of editorial markup are used: Notes and Comments. Editorial Notes and Editorial Comments are not part of the amendment and will be removed before it is published, together with any other contents in this subclause. This paragraph is an example of how an Editorial Note is marked. Editorial Comments are marked (Ed:), and contain references to submissions or comment resolutions to track the origin of changes.

EDITORIAL NOTE—Headings with empty content or Headings preceding editing instructions that modify the contents of the referenced subclause are there to provide context to the reader of this document, they have no other significance.

EDITORIAL NOTE—Except when referring to tables and figures that exist in the baseline, figure and table numbers are preceded by "s" and are assigned sequentially. This will be changed prior to sponsor ballot.

EDITORIAL NOTE—The default IEEE-SA style for tables is to "float". This means that they be repositioned later, usually at the head of the next page, to avoid splitting the table and reduce the amount of blank space. The table can appear to move out of the subclause it is referenced first from, and can even split a paragraph. This is the intended IEEE-SA behavior, please do not report it as a defect in the draft.

EDITORIAL NOTE—Line numbering is only approximate. This is a limitation of the FrameMaker tool. Whitespace between paragraphs is part of the IEEE-SA style, as defined in their templates. The combination of these two facts leads to the appearance of blank lines in the draft between every paragraph. Please do not report this as an editorial defect as it is the unavoidable behavior.

EDITORIAL NOTE—New subclauses are generally introduced by an editorial instruction "insert the following new subclause". New subclause headings are generally introduced by an editorial instruction "insert the following new subclause heading". Each new heading or subclause has its own editorial instruction. The instruction intentionally does not include where to insert the subclause because that is determined uniquely by the subclause number.

EDITORIAL NOTE—Pronunciation. It is assumed that while reading the spec aloud, a reader will read "MP" as "emm pea" rather than read it as "mesh point". This determines the spelling of the indefinite article to be "an" rather than "a".

Status of this document

Draft	Date	Changes
D1.01	2007-03-09	Conversion of draft from Word format to FrameMaker format. Implementation of most comment resolutions marked as Accept/Counter in 11-07/23r5 and 11-07/23r6 adopted by motions during January 2007 meeting. Implemented resolutions are marked with "D1.01" in "Edited in Draft" column of 11-07/23r20.

Table of Contents

List of Figures

List of Tables

IEEE P802.11s™/D1.01

Draft STANDARD for Information Technology- Telecommunications and information exchange between systems- Local and metropolitan area networks- Specific requirements-

Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications:

Amendment <number>: ESS Mesh Networking

EDITORIAL NOTE—the amendment number will be inserted by IEEE-SA editorial staff in the publication preparation phase.

The editing instructions are shown in bold italic. Four editing instructions are used: change, delete, insert, and replace. Change is used to make corrections in existing text or tables. The editing instruction specifies the location of the change and describes what is being changed by using strikethrough (to remove old material) and underscore (to add new material). Delete removes existing material. Insert adds new material without disturbing the existing material. Insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. Replace is used to make changes in figures or equations by removing the existing figure or equation and replacing it with a new one. Editorial notes will not be carried over into future editions because the changes will be incorporated into the base standard.

2. Normative references

Insert the following citations at the appropriate locations in Clause 2:

IETF RFC 3561, "Ad hoc On-Demand Distance Vector (AODV) Routing", C. Perkins, E. Belding-Royer, S. Das, July 2003.

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

3. Definitions

Insert the following new definitions alphabetically, renumbering as necessary:

3.1 WLAN Mesh: A WLAN Mesh is an IEEE 802.11-based WDS which is part of a DS, consisting of a set of two or more Mesh Points interconnected via IEEE 802.11 links and communicating via the WLAN Mesh Services. A WLAN Mesh may support zero or more entry points (Mesh Portals), automatic topology learning and dynamic path selection (including multiple hop paths).

3.2 WLAN Mesh Services: The set of services provided by the WLAN Mesh that support the control, management, and operation of the WLAN Mesh, including the transport of MSDUs between Mesh Points within the WLAN Mesh. WLAN Mesh Services supplement DSS (Distribution System Services).

3.s3 Mesh: A Mesh is network consisting of two ore more Mesh Points communicating via Mesh Services..

3.s4 Mesh AP (MAP): Any Mesh Point that is also an Access Point.

3.s5 Mesh Broadcast: A frame forwarding mechanism for transporting MSDUs to all Mesh Points within a Mesh.

3.s6 Mesh Deterministic Access Opportunity (MDAOP): MDAOP is a period of time within every Mesh DTIM interval that is set up between a transmitter and a receiver (see clause 9.14.1)receiver.

3.7 Mesh Point (MP): Any IEEE 802.11 entity that contains an IEEE 802.11-conformant Medium Access Control (MAC) and Physical Layer (PHY) interface to the Wireless Medium (WM), that is within a WLAN Mesh, and that supports WLAN Mesh Services.

3.8 Mesh AP (MAP): Any Mesh Point that is also an Access Point.

3.9 Mesh Portal: A point at which MSDUs exit and enter a WLAN Mesh to and from other parts of a DS or to and from a non-802.11 network. A Mesh Portal can be collocated with an IEEE 802.11 portal.

3.s10 Mesh Link: A bidirectional IEEE 802.11 data link between two associated Mesh Points.

3.s11 Link Metric: A criterion used to characterize the performance/quality/eligibility of a mesh link as a member of a mesh path. A mesh link metric may be used in a computation of a path metric.

3.s12 Mesh Multicast: A frame forwarding mechanism for transporting MSDUs to a group of Mesh Points within a Mesh.

3.s13 Mesh Neighbor: Any Mesh Point that is directly connected to another Mesh Point via a Mesh Link.

3.s14 Mesh Path: A concatenated set of connected Mesh Links from a source Mesh Point to a destination Mesh Point.

3.s15 Mesh Path Selection: The process of selecting a Mesh PathsPath.

3.s16 Mesh Point (MP): Any IEEE 802.11 entity that contains an IEEE 802.11-conformant Medium Access Control (MAC) and Physical Layer (PHY) interface to the Wireless Medium (WM), that is within a Mesh, and that supports Mesh Services.

3.17 Path Metric: Aggregate multi-hop criterion used for Mesh Path Selection.

3.s18 Mesh Portal: A point at which MSDUs exit and enter a Mesh to and from other parts of a DS or to and from a non-802.11 network. A Mesh Portal can be collocated with an IEEE 802.11 portal.

3.s19 Mesh Services: The set of services defined in this standard that together with other 802.11 MAC services provide for the creation and operation of mesh networks using 802.11 PHY services.

3.s20 Mesh Topology: A graph consisting of the full set of Mesh Points and Mesh Links in a WLAN Mesh.

3.s21 Mesh NeighborUnicast: Any Mesh Point that is directly connected A frame forwarding mechanism for transporting MSDUs to another an individual Mesh Point with within a Mesh LinkMesh.

3.22 Mesh Unicast: Frame forwarding mechanism for transporting MSDUs to an individual Mesh Point within a WLAN Mesh.

3.23 Mesh Multicast: Frame forwarding mechanism for transporting MSDUs to a group of Mesh Points within a WLAN Mesh.

3.24 Mesh Broadcast: Frame forwarding mechanism for transporting MSDUs to all Mesh Points within a WLAN Mesh.

3.25 Path Metric: An aggregate multi-hop criterion used for Mesh Path Selection.

3.s26 Unified Channel Graph (UCG): A set of mesh point radio interfaces that are interconnected to each other via a common channel.

4. Abbreviations and acronyms

Insert the following new acronym in alphabetical order:

AODV	Ad-Ad hoc On-demand Distance Vector
BB	Beacon Broadcaster
E2E	End-to-End
EAPAIE	EAP Authentication information element
EAPMIE	EAP Message information element
EMSAMSA	Efficient Mesh Security Association
EMSAIEMSAIE	EMSA MSA Handshake information element
HWMP	Hybrid Wireless Mesh Protocol
KCK-KD	Key confirmation key for key distribution
KDK	Key Distribution Key
KEK-KD	Key encryption key for key distribution

LQM	Link Quality Matrix
MA	Mesh Authenticator
MA-ID	Mesh Authenticator Identifier
MANET	Mobile Ad-hoc Networks
MAP	Mesh Access Point
MDA	Mesh Deterministic Access
MDAOP	Mesh Deterministic Access Opportunity
MEKIE	Mesh encrypted key information element
MKD	Mesh Key Distributor
MKD-ID	Mesh Key Distributor Identifier
MKHSIE	Mesh key holder security information element
MKDD-ID	MKD domain Identifier
MKDDIE	MKD domain information element
MP	Mesh Point
MPP	Mesh Point collocated with a mesh Portal
MSDU	MAC Service Data Unit
OLSR	Optimized Link State Routing
PMK-MA	Mesh Authenticator PMK
PMK-MKD	Mesh Key Distributor PMK
PTK-KD	Pairwise transient key for key distribution
RA-OLSR	Radio Aware Optimized Link State Routing
RERR	Route Error
RM-AODV	Radio-Metric Ad-hoc On-demand Distance Vector
RREQRREP	Route RequestReply
RREPRREQ	Route ReplyRequest
TBC	To be confirmed
TBD	To be determined
TTL	Time to Live
UCG	Unified Channel Graph

5. General description

5.2 Components of the IEEE 802.11 architecture

Insert the following new clause after 5.2.68, renumbering figures as appropriate.

5.2.9 Wireless LAN Meshmesh

5.2.9.1 Rationale

The networks described in this document make use of layer-2 mesh path selection and forwarding (that is, a mesh network that performs routing at the link layer). Mesh networks have advantageous properties in terms of robustness, range extension and density, but also have significant potential disadvantages. In particular, power consumption and security are typical problems with such networking topologies. In addition, any implementation of a mesh network cannot assume that all devices will use this new protocol. The approach described in this document is specifically designed to address all of these problems.

EDITORIAL NOTE—Rationale moved to Introduction on page iii per CID 3488

5.2.9.2 Introduction to WLAN Meshmesh

In most wireless local area network (WLAN) deployments today, there is a clear distinction between the devices that comprise the network infrastructure and the devices that are clients that use the infrastructure to gain access to network resources. The most common WLAN infrastructure devices deployed today are access points (APs) that provide a number of services, in particular: support for power saving devices, for which it buffers traffic, authentication without mesh services, and access to the network. APs are usually directly connected to a wired network (e.g., 802.3), and simply provide wireless connectivity to client devices rather than utilizing wireless connectivity themselves. Client devices, the other hand, are typically implemented as stations (STAs) that must associate with an AP in order to gain access to the network. These simple STAs are dependent on the AP with which they are associated to communicate. The An example of the non-mesh WLAN deployment model and device classes are illustrated in Figure s1.

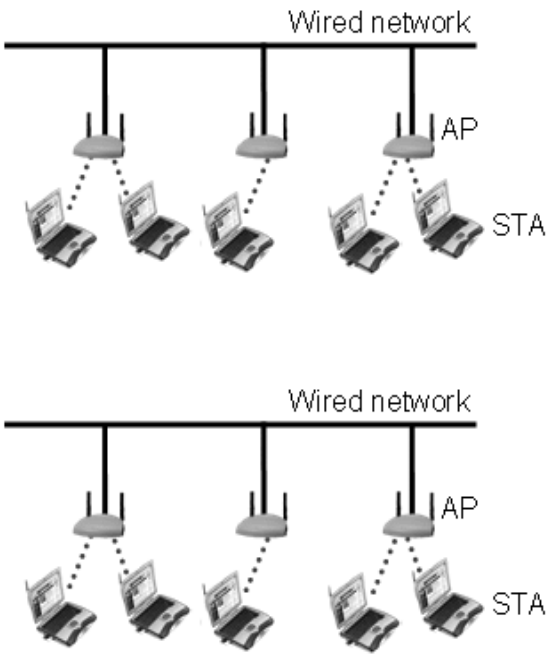

Figure s1—Non-mesh IEEE 802.11 deployment model and device classes.

There is no reason, however, that many of the Many WLAN devices under consideration can benefit from support for use in WLANs cannot support much more flexible wireless connectivity. Dedicated infrastructure class devices such as APs should Functionally, the DS of an AP can be able to establish peer-to-peer replaced with wireless links with neighboring APs to establish a mesh backhaul infrastructure, without the need for a wired network connection to each APor multi-hop paths between multiple APs. Moreover, in many cases devices Devices traditionally categorized as clients should also be able can benefit from the ability to establish peer-to-peer wireless links with neighboring clients and APs in a mesh network.

An example WLAN Mesh is illustrated in Figure s2. Mesh points (MPs) are entities that support WLAN mesh services, i.e. they participate in the formation and operation of the mesh network. Higher layer protocols and applications An MP may optionally be implemented above the MAC_SAP of a MP collocated with one or more other entities (the details e.g., AP, portal, etc.). The implementation of which are collocated entities is beyond the scope of this standard). A Mesh Point may be The configuration of an MP that is collocated with an Access Point, a configuration Point is referred to as a Mesh Access Point (MAP). Such a configuration allows a single device implementation entity to logically provide both mesh services functionalities and AP services functionalities simultaneously. STAs may associate with Mesh APs to gain access to the (mesh) network. STAs do not Only MPs participate in WLAN Mesh Services mesh functionalities such as path selection and forwarding, etc. Figure s3Figure s2 illustrates the set diagram relationship between STAs, MPs, and MAPsthis.

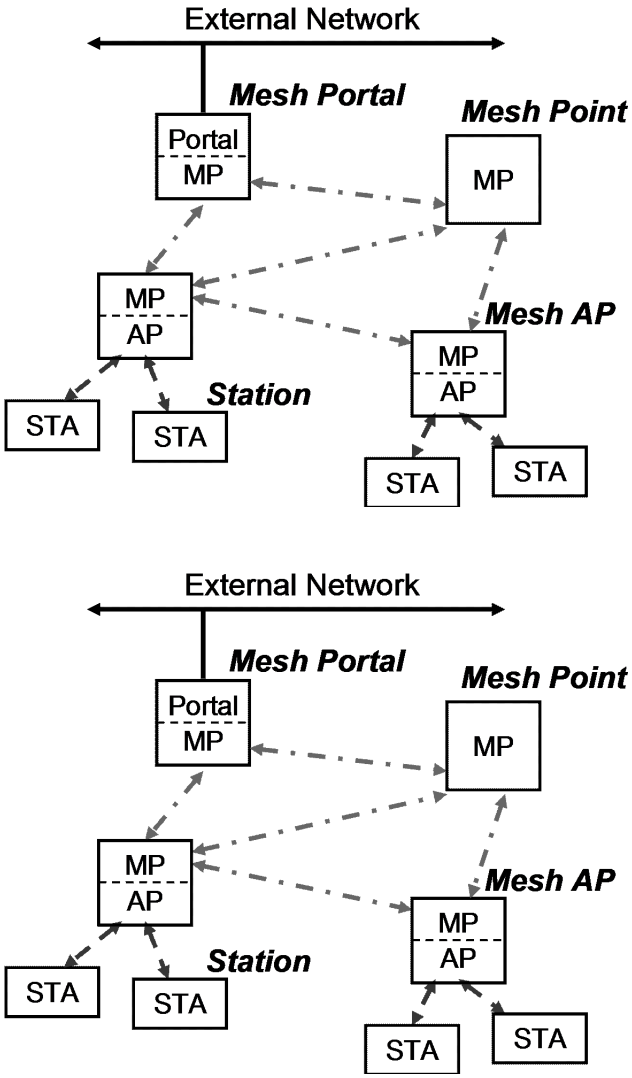

Figure s2—WLAN Mesh Containing containing MPs, MAPs, and STAs.

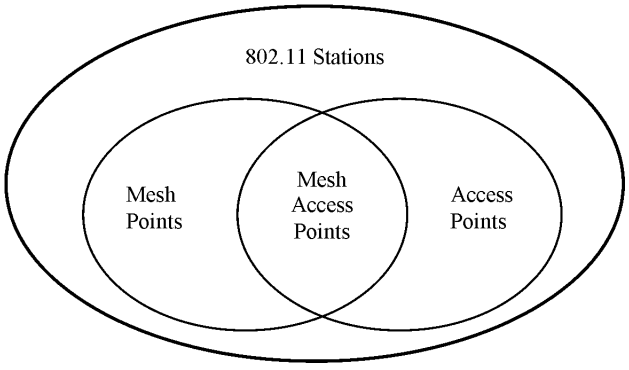

Figure s3—Set Diagram of Station, Mesh Point, and Mesh AP

Mesh points MPs may operate with different levels of functionality. Table s20 in Annex T.1 describes several examples of WLAN Mesh related entities (MPs and entities that exhibit MP functionalitiesfunctions, such as MAPs) with different levels of functionality.

5.2.9.3 WLAN Mesh Network Model

5.2.9.4 Mesh network model

A WLAN Mesh network is a layer 2 network that functions as a traditional an IEEE 802 LAN comprised of IEEE 802-style LAN.11 links and control elements to forward frames among the network members. Effectively, this means that a WLAN Mesh network appears functionally equivalent to a broadcast Ethernet from the perspective of other networks and higher layer protocols. Thus, it must appear normally appears as if all Mesh Points MPs in a WLAN Mesh are directly connected at the link layer. The Mesh protocols described in this document hide the details of this functionality from higher layer protocols by transparently providing multi-hop broadcast individually addressed and unicast group addressed data delivery at layer 2 within the mesh.mesh (see in Figure s5).

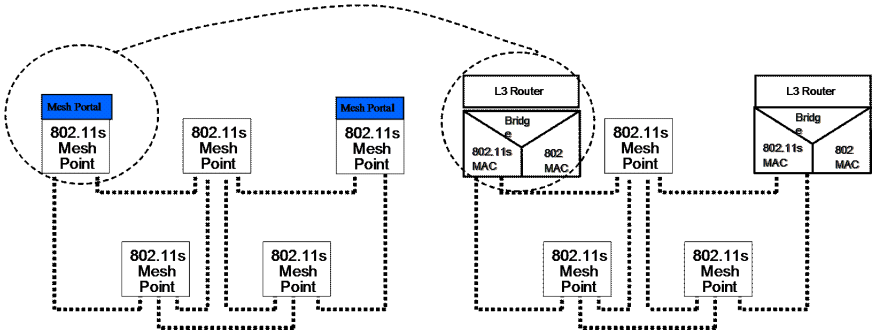

Figure s4—Reference model for WLAN Mesh interworking

As shown in Figure s4, the WLAN Mesh MAC entity appears as a single port to an 802.1 bridging relay or L3 router. Mesh portals expose the WLAN mesh behavior as an 802-style LAN segment. The mesh appears as a single loop-free broadcast LAN segment to the 802.1 bridge relay and higher layers.

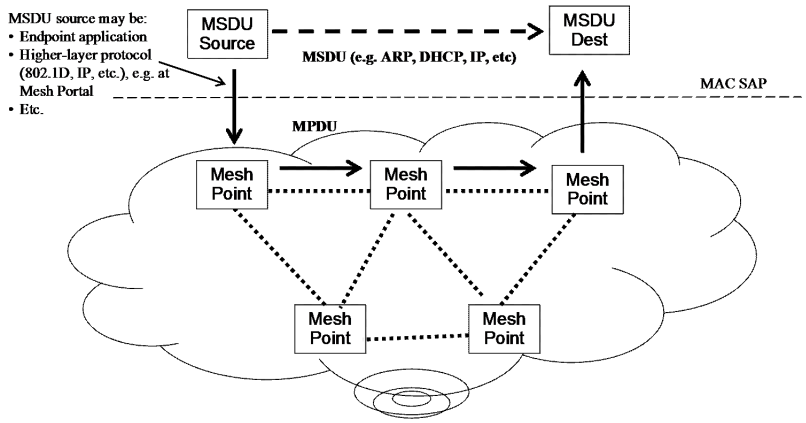

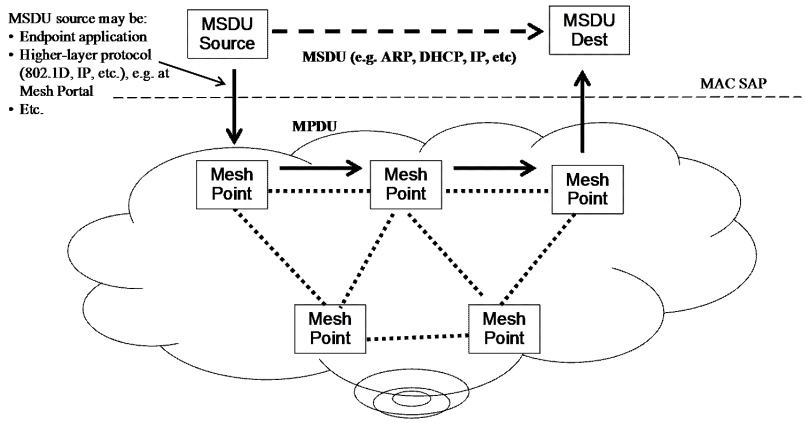

Figure s5—MAC data transport over a WLAN Mesh

As shown in Figure s5, the behavior of a WLAN Mesh is transparent to higher-layers. Internal L2 behavior of WLAN Mesh (including multi-hop message forwarding) is hidden from higher-layer protocols under the MAC-SAP.

5.2.9.5 Organization of WLAN Mesh subclauses (Informative)

The remainder of this document is organized as follows:

Functional Area	Clause
Frame Formats	7
Mesh Security	8.8
Mesh Deterministic Access (MDA)	9.14
Mesh Discovery and Peer Link Establishment	11A.1
Mesh Path Selection, Forwarding, and Interworking	11A.2, 11A.3, 11A.4, 11A.5
Intra-Mesh Congestion Control	11A.6
Mesh Beaconing and Synchronization	11A.7
Power Management in a Mesh	11A.8

5.2.9.6 Organization of mesh subclauses

Mesh functionalities are described in the subclauses shown in Table s1:

Table s1—Organization of mesh subclauses

Functional Area	Clause
Frame Formats	7
Mesh Security	8.8, 11A.2
Mesh Deterministic Access (MDA)	9.14
Mesh Discovery and Peer Link Establishment	11A.1
Mesh Path Selection, Forwarding, and Interworking	11A.3, 11A.4, 11A.5, 11A.6
Intra-Mesh Congestion Control	11A.7
Mesh Beaconing and Synchronization	11A.8
Power Management in a Mesh	11A.9

7. Frame formats

7.1 MAC frame formats

7.1.2 General frame format

Change the text in 7.1.2 and Figure 19 as shown:

The MAC frame format comprises a set of fields that occur in a fixed order in all frames. Figure 19 18 depicts the general MAC frame format. The first three fields (Frame Control, Duration/ID, and Address 1) and the last field (FCS) in Figure 19 18 constitute the minimal frame format and are present in all frames, including reserved types and subtypes. The fields Address 2, Address 3, Sequence Control, Address 4, QoS Control, Mesh Header, and Frame Body are present only in certain frame types and subtypes. When present, the Mesh Header Field field is prepended to the Frame Body and Body. The Mesh Header field is handled identically to the contents of the body with respect to MPDU processing. Each field is defined in 7.1.3. The format of each of the individual subtypes of each frame types type is defined in 7.2. The components of management frame bodies are defined in 7.3. The formats of management frames of subtype Action are defined in 7.4. The formats format of mesh management frames of subtype Mesh Action are is defined in 7.4A.

The Frame Body field is of variable size. The maximum frame body size is determined by the maximum MSDU size (2304 octets) plus any overhead from security encapsulation.

Octets: 2	**2**	**6**	**6**	**6**	**2**	**6**	**2**	**4 or 16**	**0-2314<u>2</u>**	**4**
Frame Control	Duration/ ID	Address 1	Address 2	Address 3	Sequence Control	Address 4	QoS Control	Mesh Header	Body	FCS

MAC Header

Figure 18—MAC Frame Format

Octets: 2	2	6	6	6	2	0 or 6	2	4 or 16	0-231<u>42</u>	4
Frame Control	Duration / ID	Address 1	Address 2	Address 3	Sequence Control	Address 4	QoS Control	Mesh Header	Body	FCS

7.1.3 MAC Frame Formatfields

7.1.4 Frame fields

7.1.4.1 Frame Control control field

7.1.4.1.2 Type and Subtype subtype fields

Change the contents of Table 1 as shown:

Table 1—Valid type and subtype combinations

Type value b3 b2	Type description	Subtype value b7 b6 b5 b4	Subtype description
11	Extended	0000	Mesh Data
11	Extended	0001	Mesh Management
11	~~Reserved~~Extended	~~0000~~0010-1111	Reserved

Table 2—Valid type and subtype combinations

(numeric values in Table 1 are shown in binary)

Type value b3 b2	Type description	Subtype value b7 b6 b5 b4	Subtype description
(Ed: insert unchanged table entries for completeness)			
11	Extended	0000	Mesh Data
11	Extended	0001	Mesh Data + CF-Ack
11	Mesh Management	0010	Mesh Action
11	~~Reserved~~Extended	~~0000~~0011-1111	Reserved

7.1.4.1.3 To DS and From DS fields

Change the contents of Table 3 as shown:

Table 3—To/From DS combinations in data frames

To DS and From DS values	Meaning
To DS = 1 From DS = 1	A data frame using the four-address format. ~~This standard does not define procedures for using this combination of field values.~~

7.1.4.1.8 More Data field

Insert the following text to the end of 7.1.3.1.8

Special considerations exist within a mesh. The 'more data' bit is set to 1 by MPs for individually addressed MSDU/MMPDUs sent to a neighboring MP when there are more frames to be transmitted to that MP in the current beacon interval. The 'more data' bit is set to 1 by MPs for group addressed MSDUs/MMPDUs when there are more group addressed frames to be transmitted in the current beacon interval.

Add Insert the following text to the end of Clause new clause after 7.1.3.1.85:

The 'more data' bit is set by mesh points for unicast messages sent to a neighboring mesh point operating in power save mode when there are more MSDU/MMPDUs to be transmitted to that mesh point in the current beacon interval.

The 'more data' bit is set by mesh points for broadcast/multicast MSDUs when the mesh is determined to be operating in power save scheme and there are more broadcast/multicast traffic to be transmitted in the current beacon interval.

7.1.4.5a Mesh Header field

The mesh forwarding control field is a 4- or 16 octet field which includes a time to live field for use in multi-hop forwarding to eliminate the possibility of infinite loops and a mesh end-to-end sequence number for use in controlled broadcast flooding and other services, an 8-bit mesh flags field for use in mesh-specific extensions of header processing including mesh address extension, and optionally a 12-octet mesh addressing field for allowing total of 6 addresses to be carried in mesh data frames. The Mesh Header field is present in all frames of type Extended with subtype Mesh Data.

Octets: 1	1	2	12
Mesh Flags	Time To Live (TTL)	Mesh E2E SeqNum	Mesh Addressing

Figure s6—Mesh Header field

7.1.4.5a.1 General

The mesh header field is a 4 or 16 octet field which includes:

— an 8-bit Mesh Flags field to control mesh header processing

— a time to live field for use in multi-hop forwarding to aid in limiting the effect of transitory routing loops

— a mesh sequence number to suppress duplicates in flooding algorithms and for other services

— and optionally a 12-octet mesh address extension field containing two addresses resulting in a total of 6 addresses in mesh data frames

The Mesh Header field, shown in Figure s5, is present in all frames of type Extended with subtypes Mesh Data and Mesh Management.

Octets: 1	**1**	**2**	**12**
Mesh Flags	Mesh Time To Live (TTL)	Mesh Sequence Number	Mesh Address Extension (present in some configurations)

Figure s5—Mesh Header field

7.1.4.5a.2 Mesh Flags field

The Mesh Flags field field, shown in Figure s6, is eight 8 bits in length and the flags therein are used to control mesh-specific header processing, e.g., for mesh address extension and encapsulation/tunnelingextension.

Bits: 0	1-7
Address Extension (AE)	Reserved

B0	B1	B7
Address Extension (AE)	Reserved	
Bits: 1	7	

Figure s6—Mesh Flags field

Figure s7—Mesh Flags field

The "Address Extension (AE)" flag is used to indicate the existence of — and thereby the use of — the mesh addressing address extension field for mesh address extension based on 6-address format: If AE = 1, the mesh addressing address extension field follows the mesh header Mesh Sequence Number field and provides two additional address fields for mesh address extension (see Clause 11A.3.4 for the usage of address fields in 6-address format); if AE = is 0, there is no mesh addressing field after the mesh header address extension field and the use of existing address fields, i.e., Address 1 to Address 4, is not changedpresent. Note that The "AE" flag can be set to 1 only when both "To DS" and "From DS" fields are set to 1.

The reserved field is set to zero.

7.1.4.5a.3 Mesh TTL Time to Live field

The Mesh TTL Time to Live (TTL) field is eight 8 bits in length and is used to mitigate the possibility of transient loops in a WLAN mesh network by ensuring frames that are caught in a loop are eventually discarded.

7.1.4.5a.4 Mesh E2E Sequence number Number field

The Mesh E2E Sequence number Number field is sixteen 16 bits in length and used to control broadcast flooding detect duplicate reception and to enable ordered delivery of messages in a WLAN Mesh network.

For unicast In an individually addressed data framesframe, the Mesh E2E Sequence number is used to Number field uniquely identify identifies the frame from a given Source Mesh PointMP. This field is set by the Source Mesh PointMP, kept unchanged at the intermediate Relay Mesh PointsMPs, and used by the Destination Mesh Point MP to eliminate duplicate frames or to detect out of order frames.

7.1.4.5a.5 Mesh Addressing Address Extension field

The Mesh Addressing field Address Extension field, shown in Figure s8, is sixteen bits 12 octets in length and follows the Mesh TTL Sequence Number field only when the "AE" flag of Mesh Flags field is set to 1. Mesh Addressing filed Address Extension field provides two additional address fields, Address 5 and Address 6, for mesh address extension based on 6-address format.

Octets:6	6
Address 5	Address 6

Octets: 6 6

Address 5	Address 6

Figure s8—Mesh Address Extension field

Figure s9—Mesh Addressing field

Address 5 and Address 6 are used to carry carry, over a mesh link link, the MAC addresses of IEEE 802 entities that do not support WLAN mesh services. This is useful, but are the destination and the source of an IEEE 802 data link where there are one or more MPs for example, in the middle that take either of the following actions for incoming framescases:

— Transport of frames to and from stations outside of the mesh, e.g., involving conversion of address format, i.e., from 3-address to 6-address format or vice versa at MAPs;

— Conversion of address format, i.e., . from 34-address to 6-address format or vice versa at MAPs;MPPs.

— Redirection to a different destination either at a root node (in HWMP) or at an MPP (in interworking), which involves updating the fields of Address 3 and Address 4 based on the information in Address 5 and Address 6.

Details on the usage of these optional address fields are given in Clause 11A.3.4 as a part of that describes the frame forwarding processing descriptionprocess.

7.2 Format of individual frame types

7.2.3 Management frames

Add Insert the following text to at the end of the bulleted list in the fourth paragraph of Clause 7.2.3:

The exchange of management frames shall be supported between neighboring Mesh Points. The management frame header supports two address fields, DA and SA. The value of these address fields are as follows:

DA field: Receiving MP MAC Address (with respect to one-hop transmission)

SA field: Transmitting MP MAC address (with respect to one-hop transmission)

BSSID field: This field is not used for management frames transmitted between Mesh Points (TBD: recommended value for backward compatibility with non-mesh STAs, e.g., setting field to all 0's)

d) If the STA is an MP, the field is set to the wildcard value and is not interpreted by receiving MPs.

7.2.3.1 Beacon frame format

Change the contents of the of Table 8 as shown::

Table 8—Beacon frame body

Order	Information	Notes
4	Service Set Identifier (SSID)	When dot11WLANMeshService is true but the interface on which the beacon is being sent is not configured as an Access Point, the SSID information element is set to the wildcard value.
36	Mesh ID	The Mesh ID information element is present within Beacon frames when dot11WLANMeshService is true.
37	Mesh Capability	The Mesh Capability information element is present within Beacon frames when dot11WLANMeshService is true.
38	Mesh Neighbor List	The Mesh Neighbor List information element is present within frames with the DTIM bit set when dot11WLANMesh Service is true and the MP transmits to other MPs in power save mode.
39	Mesh DTIM	The Mesh DTIM information element is present in Beacon frames generated by the MP when dot11WLANMesh Service is true and MP is supporting Transmission to MP in power save mode.
40	Mesh Portal Reachability	The Mesh Portal Reachability information element is present within Beacon frames when dot11WLANMeshService is true.
41	Beacon Timing	The Beacon Timing information element is present within Beacon frames when dot11WLANMeshService is true.
42	MDAOP Advertisements	The MDAOP Advertisements information element is present within Beacon frames when dot11WLANMeshService is true.
43	MKDDIE	The MKDDIE element is present when dot11WLANMeshService is true.

Table 8—Beacon frame body

Order	Information	Notes
4	Service Set Identifier (SSID)	When dot11WLANMeshService is true but the interface on which the beacon is being sent is not configured as an Access Point, the SSID IE shall be set to the wildcard value. [Note: the SSID is a required IE in beacon frames. To avoid having non-mesh STAs send association requests to non-MAP Mesh Points, a valid SSID should not be included in beacons sent by non-MAP Mesh Points. To avoid backward compatibility issues, rather than removing the SSID IE from MP (non-MAP) beacons the wildcard value is used.]
(Ed: insert unchanged table entries for completeness)		
26	OFDM Parameter Set	The OFDM Parameter Set information element is present within Beacon frames generated by STAs using Clause 17 PHYs.
27	Mesh ID	The Mesh ID information element shall be present within Beacon frames only when dot11WLANMeshService is true.
28	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Beacon frames only when dot11WLANMeshService is true.

29	Neighbor List	The Neighbor List information element shall be present within DTIM Beacon frames generated when dot11WLANMesh Service is true and MP is supporting Transmission to MP in power save mode.
30	DTIM	The DTIM IE shall be present in beacon frames generated by when dot11WLANMesh Service is true and MP is supporting Transmission to MP in power save mode.
31	Mesh Portal Reachability	The Mesh Portal Reachability information element shall be present within Beacon frames only when dot11WLANMeshService is true.
32	Beacon Timing	The Beacon Timing information element shall be present within Beacon frames only when dot11WLANMeshService is true.
33	MDAOP Advertisements	The MDAOP Advertisements information element shall be present within Beacon frames only when dot11WLANMeshService is true.
34	MDAOP Set Teardown	The MDAOP Set Teardown information element shall be present within Beacon frames only when dot11WLANMeshService is true.
35	MKDDIE	The MKDDIE element shall be present only when dot11WLANMeshService is true

0.0.0.1 IBSS ATIM frame format

Change the title of 7.2.3.2 as shown:

7.2.3.3 ~~IBSS~~ ATIM frame format

7.2.3.3 Disassociation frame format

Add Insert the following to before the contents last row of Table 9 as shown:

Table 9—Disassociation frame body

Order	Information	Notes
3	Peer Link Close	The Peer Link Close information element is present when dot11WLANMeshService is true.

Table 10—Association Request frame format

Order	Information	Notes
3	Peer Link Close IE	The Peer Link Close IE shall be present only when dot11WLANMeshService is true

7.2.3.4 Association Request frame format

Add Insert the following to before the contents of entry labeled "Last" in Table 10 as shown:

Table 11—Association Request frame body

Order	Information	Notes
13	Mesh ID	The Mesh ID information element is present within Association Request frames when dot11WLANMeshService is true.
14	Mesh Capability	The Mesh Capability information element is present within Association Request frames when dot11WLANMeshService is true.
15	Peer Link Open	The Peer Link Open information element is present when dot11WLANMeshService is true.
16	MKDDIE	The MKDDIE element is present when dot11WLANMeshService is true.
17	MSAIE	The MSAIE element is present when dot11WLANMeshService is true.

Table 12—Association Request frame format

Order	Information	Notes
11	Mesh ID	The Mesh ID information element shall be present within Association Request frames only when dot11WLANMeshService is true.
12	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Association Request frames only when dot11WLANMeshService is true.
13	Active Profile Announcement	The Active Profile Announcement information element shall be present within Association Request frames only when dot11WLANMeshService is true.
14	Peer Link Open IE	The Peer Link Open IE shall be present only when dot11WLANMeshService is true
15	MKDDIE	The MKDDIE element shall be present only when dot11WLANMeshService is true
16	EMSAIE	The EMSAIE element shall be present only when dot11WLANMeshService is true

7.2.3.5 Association Response frame format

Add the following to the contents of Table 11 as shown:

Table 13—Association Response frame body

Order	Information	Notes
8	Mesh ID	The Mesh ID information element shall be present within Association Resposne frames only when dot11WLANMeshService is true.

9	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Association Response frames only when dot11WLANMeshService is true.
10	Active Profile Announcement	The Active Profile Announcement information element shall be present within Association Response frames only when dot11WLANMeshService is true.
11	Peer Link Confirm IE	The Peer Link Confirm IE shall be present only when dot11WLANMeshService is true
12	MKDDIE	The MKDDIE element shall be present only when dot11WLANMeshService is true
13	EMSAIE	The EMSAIE element shall be present only when dot11WLANMeshService is true
14	RSNIE	The RSNIE element shall be present only when dot11WLANMeshService is true

Insert the following before before the entry labeled "Last" in Table 11 as shown

Table 14—Association Response frame body

Order	Information	Notes
16	Mesh ID	The Mesh ID information element is present within Association Response frames when dot11WLANMeshService is true.
17	Mesh Capability	The Mesh Capability information element is present within Association Response frames when dot11WLANMeshService is true.
18	Peer Link Confirm	The Peer Link Confirm information element is present when dot11WLANMeshService is true.
19	MKDDIE	The MKDDIE element is present when dot11WLANMeshService is true.
20	MSAIE	The MSAIE element is present when dot11WLANMeshService is true.
21	RSNIE	The RSNIE element is present when dot11WLANMeshService is true.

7.2.3.6 Reassociation Request frame format

Insert the following before the entry labeled "Last" in Table 12 as shown

Table 15—Reassociation Request frame body

Order	Information	Notes
16	Mesh ID	The Mesh ID information element is present within Reassociation Request frames when dot11WLANMeshService is true.
17	Mesh Capability	The Mesh Capability information element is present within Reassociation Request frames when dot11WLANMeshService is true.

Add the following to the contents of Table 12 as shown:

Table 16—Reassociation Request frame body

Order	Information	Notes
12	Mesh ID	The Mesh ID information element shall be present within Reassociation Request frames only when dot11WLANMeshService is true.
13	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Reassociation Request frames only when dot11WLANMeshService is true.
14	Active Profile Announcement	The Active Profile Announcement information element shall be present within Reassociation Request frames only when dot11WLANMeshService is true.

7.2.3.7 Reassociation Response frame format

Add the following to the contents of Table 13 as shown:

Table 17—Reassociation Response frame body

Order	Information	Notes
8	Mesh ID	The Mesh ID information element shall be present within Reassociation Response frames only when dot11WLANMeshService is true.
9	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Reassociation Response frames only when dot11WLANMeshService is true.
10	Active Profile Announcement	The Active Profile Announcement information element shall be present within Reassociation Response frames only when dot11WLANMeshService is true.

Insert the following before the entry labeled "Last" in Table 13 as shown.

Table 18—Reassociation Response frame body

Order	Information	Notes
17	Mesh ID	The Mesh ID information element is present within Reassociation Response frames when dot11WLANMeshService is true.
18	Mesh Capability	The Mesh Capability information element is present within Reassociation Response frames when dot11WLANMeshService is true.

7.2.3.8 Probe Request frame format

Add the following to the contents of Table 14 as shown:

Table 19—Probe Request frame body

Order	Information	Notes
6	Mesh ID	The Mesh ID information element shall be present within Probe Request frames only when dot11WLANMeshService is true.
7	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Probe Request frames only when dot11WLANMeshService is true.

Insert the following before the entry labeled "Last" in Table 14 as shown.

Table 20—Probe Request frame body

Order	Information	Notes
35	Mesh ID	The Mesh ID information element is present within Probe Request frames when dot11WLANMeshService is true.
36	Mesh Capability	The Mesh Capability information element is present within Probe Request frames when dot11WLANMeshService is true.
37	Mesh Neighbor List	The Mesh Neighbor List information element is present within frames with the DTIM bit set when dot11WLANMesh Service is true and the MP transmits to other MPs in power save mode.
38	Mesh DTIM	The Mesh DTIM information element is present in Probe Request frames generated by the MP when dot11WLANMesh Service is true and MP is supporting Transmission to MP in power save mode.
39	Mesh Portal Reachability	The Mesh Portal Reachability information element is present within Probe Request frames when dot11WLANMeshService is true.
40	Beacon Timing	The Beacon Timing information element is present within Probe Request frames when dot11WLANMeshService is true.
41	MDAOP Advertisements	The MDAOP Advertisements information element is present within Probe Request frames when dot11WLANMeshService is true.
42	MKDDIE	The MKDDIE element is present when dot11WLANMeshService is true.

7.2.3.9 Probe Response frame format

Add the following to the contents of Table 15 as shown:

Table 21—Probe Response frame body

Order	Information	Notes
25	OFDM Parameter Set	The OFDM Parameter Set information element is present within Probe Response frames generated by STAs using Clause 17 PHYs.
26	Mesh ID	The Mesh ID information element shall be present within Probe Response frames only when dot11WLANMeshService is true.
27	WLAN Mesh Capability	The WLAN Mesh Capability information element shall be present within Probe Response frames only when dot11WLANMeshService is true.
28	DTIM	The DTIM information element is only present in probes generated when dot11WLANMeshService is true and MP supports Power Save mode.
29	Mesh Portal Reachability	The Mesh Portal Reachability information element shall be present within Probe Response frames only when dot11WLANMeshService is true.
30	Beacon Timing	The Beacon Timing information element shall be present within Probe Response frames only when dot11WLANMeshService is true.
31	MKDDIE	The MKDDIE element shall be present only when dot11WLANMeshService is true

Insert the following before the entry labeled "Last" in Table 15 as shown.

Table 22—Probe Response frame body

Order	Information	Notes
26	Mesh ID	The Mesh ID information element is present within Probe Response frames when dot11WLANMeshService is true.
27	Mesh Capability	The Mesh Capability information element is present within Probe Response frames when dot11WLANMeshService is true.
28	Mesh DTIM	The Mesh DTIM information element is present in probes generated when dot11WLANMeshService is true and MP supports Power Save mode.
29	Mesh Portal Reachability	The Mesh Portal Reachability information element is present within Probe Response frames when dot11WLANMeshService is true.
30	Beacon Timing	The Beacon Timing information element is present within Probe Response frames when dot11WLANMeshService is true.
31	MKDDIE	The MKDDIE element is present when dot11WLANMeshService is true.

Insert the following new clause after 7.2.3:

7.2.4 Extended frames

7.2.4.1 General

Frames of type "11" are referred to as extended frames. The frame format for an extended frame is dependent on the Subtype field as described in the following subclauses.

7.2.4.2 Mesh Data frame format

The frame format for a mesh data frame is identical to based on the format of a data frame as defined in clause 7.2.2 with the addition of the Mesh Header fieldfield described in 7.1.4.5a.

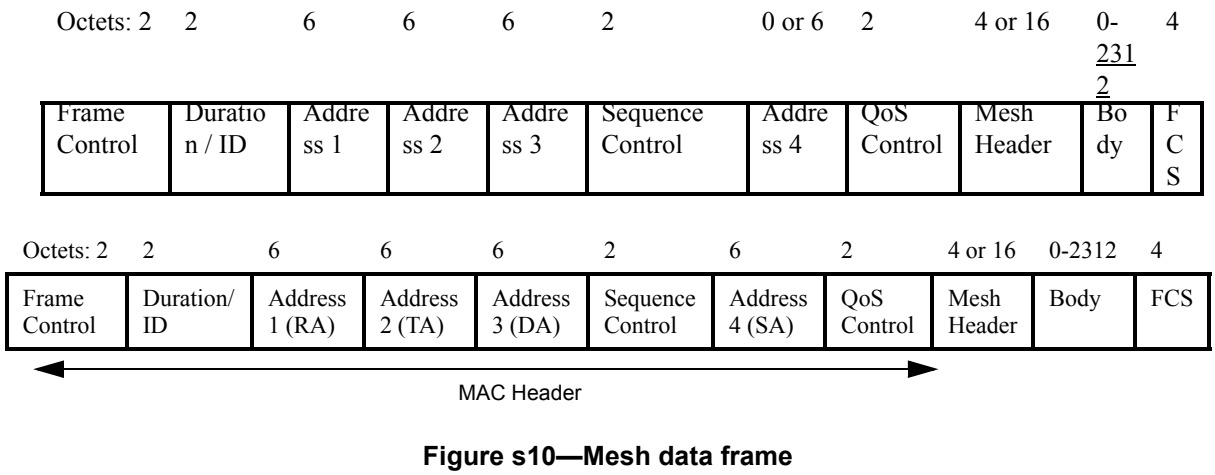

Figure s10—Mesh data frame

Figure s11—Mesh data frame

7.2.4.3 Mesh Management frame format

The frame format for a mesh management frame is independent of frame subtype and is as defined in Figure s10.

Octets: 2	2	6	6	6	2	6	4 or 16	0-23142	4
Frame Control	Duration / ID	Address 1 (RA)	Address 2 (TA)	Address 3 (DA)	Sequence Control	Address 4 (SA)	Mesh Header	Body	FCS

Figure s12—Mesh management frame

The frame format for a mesh management frame is defined in Figure s13. Mesh management frames are not

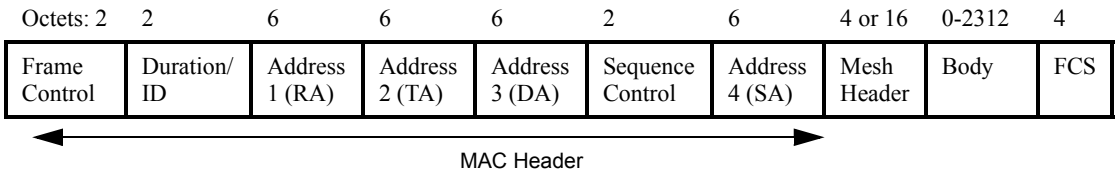

Figure s13—Mesh management frame

the same as management frames as defined in 7.2.3.

The Duration field contains a duration value as defined in 7.1.4.

The duration value calculation for the mesh management frame is based on the rules in 9.6 that determine the data rate at which the control frames in the frame exchange sequence are transmitted. If the calculated duration includes a fractional microsecond, that value is rounded up to the next higher integer. All MPs process Duration field values less than or equal to 32 767 from valid management frames to update their NAV settings as appropriate under the coordination function rules.

The address fields for mesh management frames do not vary by frame subtype.

The RA field is contains the unicast MAC address of the MP that is the immediate intended receiver of the frame or the multicast or broadcast group address of the MPs that are the immediate intended receivers of the frame.

The TA field is the address of the MP that is transmitting the frame.

The Address 3 TA field (DA) is typically contains the address of the MP that is the destination of transmitting the Mesh management frame.

The SA field is contains the address of the MP that initiated the Mesh management frame.

A An MP uses the contents of the RA field to perform address matching for receive receiving decisions. In cases where the RA field contains a group address, the TA also is validated to ensure verify that the broad-

cast or multicast group addressed frame originated from a an MP with which the receiving MP has an established link. A An MP uses the contents of the TA field to direct the acknowledgment if an acknowledgment is necessary.

The Mesh Header field is defined in 7.1.4.5a.

Detailed usage Use of the Address 3 and Mesh Header fields is specified in 11A.3.4.

The Sequence Control field is defined in 7.1.3.4.

The frame body consists of Mesh Action Data Units and a security header and trailer (if and only if the Protected Frame subfield in the Frame Control field which is set to 1)present . A Mesh Action Data Unit is a an MMPDU sent between two mesh MAC entities. The Mesh Action Data Unit contains the Mesh Action field, defined in 7.3.1.18. The Mesh Action field comprises Category and Action fields followed by the information elements defined for each Mesh Action. All fields and information elements are mandatory unless stated otherwise and appear in the specified, relative order. Destination MPs that encounter an element ID they do not recognize in the frame body of a received management frame ignore that element and continue to parse the remainder of the management frame body (if any) for additional information elements with recognizable element IDs. A MP receiving a vendor-specific IE that it does not support shall ignore the vendor-specific IE. Unused element ID codes are reserved.

The maximum size of a Mesh Action Data Unit is 2304 2312 octets.

Gaps may exist in the ordering of elements within frames. The order that remains shall be ascending.

Elements are included in ascending order.

7.3 Management frame body components

7.3.1 Fields that are not information elements

7.3.1.4 Capability Information field

Make the following changes Change to the fourth paragraph of Clause 7.3.1.4 "Capability Information field"" as shown:

APs set the ESS subfield to 1 and the IBSS subfield to 0 within transmitted Beacon, Beacon or Probe Response, Association Request, or Association Response management frames. STAs within an IBSS set the ESS subfield to 0 and the IBSS subfield to 1 in transmitted Beacon or Probe Response management frames. Non-AP MPs that are not MAPs set both the ESS subfield and IBSS subfield to 0 in transmitted Beacon, Beacon and Probe Response, Association Request, or Association Response management frames.

7.3.1.11 Action field

Change the contents of Table 24 as shown:.

Table 24—Category values

Value	Meaning	See subclause
5	Mesh	7.4.9
6~~5~~-127	Reserved	---

Table 24: Category values

Value	Meaning	See subclause
4	Mesh Management	7.4.5
5~~4~~-127		---
128-255		---

Insert the following new subclause:

Insert the following new subclause after the last subclause in 7.3.1 and renumber accordingly:

7.3.1.18 Mesh Action field

The Mesh Action field provides a mechanism for specifying mesh management actions. The format of the Mesh Action field is shown in Figure s11Figure s14.

Octets: 1	variable
Category	Mesh Action Details

Figure s14—Mesh Action field

	Category	**Mesh Action Details**
Octets:	1	variable

Figure s15—Mesh Action field

The Category field is set to one of the nonreserved values shown in Table s1Table s2. Mesh Action frames of a given category are referred to as *<category name> Mesh Action frames*.

If a STA an MP receives a unicast an individually addressed Mesh Action frame with an unrecognized Category field or some other syntactic error and the MSB of the Category field set to 0, then the STA shall return MP returns the Mesh Action frame to the source without change except that the MSB of the Category field is set to 1.

The Mesh Action Details field contains the details of the action. The field contains an action value followed by zero or more non-information element fields, and zero or more information elements. The details of the actions allowed in each category are described in the appropriate subclause referenced in Table s1Table s2.\

Table s2—Mesh Action Category values

Code	Meaning	See subclause
0	MSA	7.4A.1
1-126	Reserved	–
127	Vendor Specific	–
128-255	Error	(MSB reserved for error signalling)

Table s1—Mesh Action Category values

Code	Meaning	See subclause
0	EMSA	7.4A.1
1-126	Reserved	–
127	Vendor Specific	–
128-255	Error	–

7.3.2 Information elements

Insert the following entries in Table 26 and change the Reserved row as shown.

Table 26—Element IDs

Information element	Element ID	Length (in octets)
Mesh Capability	ABSB (see note below)	19
Active Mesh Profile Announcement	ABSB (see note below)	9
Mesh ID	ABSB (see note below)	0 to 255
Local Link state announcement	ABSB (see note below)	4
Target Transmission Rate	ABSB (see note below)	18
Offered Traffic Load	ABSB (see note below)	16
Neighborhood Congestion	ABSB (see note below)	3
Peer Link Close	ABSB (see note below)	9
Peer Link Open	ABSB (see note below)	4
Peer Link Confirm	ABSB (see note below)	8
Mesh Portal Reachability	ABSB (see note below)	1 to 255
Mesh Portal/Root Announcement	ABSB (see note below)	28 to 255
UCG Switch Announcement	ABSB (see note below)	13
Mesh Neighbor List	ABSB (see note below)	1 to 255
Mesh DTIM	ABSB (see note below)	2
Beacon Timing	ABSB (see note below)	5 to 255
MDAOP Setup Request	ABSB (see note below)	8 to 255
MDAOP Setup Reply	ABSB (see note below)	3
MDAOP Advertisements Request	ABSB (see note below)	0
MDAOP Advertisements	ABSB (see note below)	1 to 255
MDAOP Set Teardown	ABSB (see note below)	7
Connectivity Report	ABSB (see note below)	14 to 255
PANN	ABSB (see note below)	17
RANN	ABSB (see note below)	21
RREQ	ABSB (see note below)	36 to 255
RREP	ABSB (see note below)	32 to 255
RERR	ABSB (see note below)	13
RA-OLSR HELLO	ABSB (see note below)	22 to 255
RA-OLSR Topology Control (TC)	ABSB (see note below)	23 to 255

Table 26—Element IDs

Information element	Element ID	Length (in octets)
RA-OLSR Multiple Interface Declaration (MID)	ABSB (see note below)	23 to 255
RA-OLSR Local Association Base Advertisement (LABA)	ABSB (see note below)	20 to 255
RA-OLSR Local Association Base Checksum Advertisement (LABCA)	ABSB (see note below)	11 to 255
RA-OLSR Association Base Block Request (ABBR)	ABSB (see note below)	11 to 255
MKD domain	ABSB (see note below)	7
MSA Handshake	ABSB (see note below)	88 to 255
Mesh Key Holder Security	ABSB (see note below)	98
Mesh Encrypted Key	ABSB (see note below)	82 to 255
EAP Authentication	ABSB (see note below)	42
EAP Message	ABSB (see note below)	2 to 255
Reserved	ABSB (see note below)	

EDITORIAL NOTE—Assignment of values for these information elements needs to be approved by IEEE 802.11 ANA. Until that time, these values are marked as Allocate By Sponsor Ballot (ABSB). Final values will be requested from IEEE 802.11 ANA once this amendment reaches the 75% approval threshold in Sponsor Ballot.

7.3.2.1 SSID element

Change the second paragraph of 7.3.2.1 as shown:

The length of the SSID information field is between 0 and 32 octets. A 0 length information field is used within Probe Request management frames to indicate the wildcard SSID. The wildcard SSID is also used in Beacon management frames transmitted by non-AP MPs.

7.3.2.25 RSN information element

7.3.2.25.2 AKM Suites

Change Table 34 as follows:

Table 34 AKM Suite Selectors

OUI	Suite Type	Authentication type	Key management type
00-0F-AC	5	EMSA Authentication negotiated over IEEE 802.1X, or using PMKSA caching as defined in 8.4.6.2	EMSA Key Management
00-0F-AC	6	EMSA Authentication using PSK	EMSA Key Management
00-0F-AC	~~5~~7-255	Reserved	Reserved

Insert two new rows and change the existing 'Reserved' row in Table 34 as shown.

Table 34—AKM Suite Selectors

OUI	Suite Type	Authentication type	Key management type
00-0F-AC	5	MSA Authentication negotiated over IEEE 802.1X, or using PMKSA caching as defined in 8.4.6.2	MSA Key Management
00-0F-AC	6	MSA Authentication using PSK	MSA Key Management
00-0F-AC	~~5~~7-255	Reserved	Reserved

EDITORIAL NOTE—Assignment of values 5 and 6 for EMSA Key management needs to be approved by IEEE 802.11 ANA. Until that time, these values are tentative and subject to change. Final values will be requested from ANA once this amendment reaches the 75% approval threshold in Sponsor Ballot.

Insert the following new subclauses after 7.3.2.48

EDITORIAL NOTE—numbering of following subclauses based on 11r ending with 7.3.2.48

7.3.2.49 WLAN Mesh Capability element

The "WLAN Mesh Capability" element shown in Figure s16 is used to advertise WLAN Mesh services. It is contained in beacons Beacon frames transmitted by MPs, and is also contained in probe request/response messages and (re)association request/response messages.

Octets:1	1	1	4	4	2	1	1	1	4
ID	Length	Version	Active Protocol ID	Active Metric ID	Peer Capacity	Power Save capability	Synch-ronization Capability	MDA Capability	Channel Precedence

Figure s16—Mesh Capability Element

Octets: 1	1	1	4	4	2	1	1	2	4
ID	Length	Version	Active Protocol ID	Active Metric ID	Peer Capacity	Power Save capability	Synch-ronization Capability	MDA Capability	Channel Precedence

Figure s17—WLAN Mesh Capability Element

The fields contained in the element are as shown in Table s2.

Table s35—WLAN Mesh Capability Element Fields

Field	Value/description
ID	T.B.D
Length	Variable
Version	1
Active Protocol ID	Path selection protocol in use
Active Metric ID	Path selection metric in use
Peer capacity	Peer capacity value
Power Save capability	Support for power save mode and current power management mode
Synchronization Capability	Support for synchronization services and current synchronization status
MDA Capability	Support for MDA services and current status
Channel precedence	Channel precedence value

MPs may support one or more path selection protocols and path metrics. However, only one path selection protocol and one path metric may be active in a particular WLAN mesh network at any point in time.

The Element ID is set to the value given in Table 26 for this information element. The Length field is set to 18. The version is set to 1.

MPs may support one or more path selection protocols and path metrics. However, only one path selection protocol and one path metric may be active in a particular mesh network at a time. The Active Protocol ID field indicates the path selection protocol in use and is described in 7.3.2.49.1. The Active Metric ID field indicates the path selection metric in use and is described in 7.3.2.49.2.

The Peer capacity field indicates the MP's capacity for establishing additional peer links and is described in 7.3.2.49.3.

The Power Save capability field indicates support for power save mode and current power save mode and is described in 7.3.2.49.4.

The Synchronization Capability field indicates support for synchronization services and current synchronization status and is described in 7.3.2.49.5.

The MDA Capability field indicates support for MDA services and current status and is described in 7.3.2.49.6.

The channel precedence field is set to the value of the channel precedence of the unified channel graph to which the MP interface belongs.

7.3.2.49.1 Path selection protocol identifier field

The path selection protocol identifier indicates the path selection protocol for individually addressed and group addressed transmission. It has format shown in Figure s18.

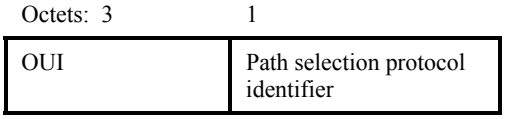

Figure s18—Path selection protocol identifier field

The path selection protocol identifier specifies the protocol which is currently used to generate routing information in this network, as defined in 11A.3. Identifier values are listed in Table s3.

Table s3—Path selection protocol identifier values

OUI	Value	Meaning
00-0F-AC	0	Hybrid Wireless Mesh Protocol (default path selection protocol)
00-0F-AC	1	Radio Aware OLSR (optional path selection protocol)
00-0F-AC	2-254	Reserved for future use
00-0F-AC	255	Null protocol
Vendor OUI	Other	Vendor specific

A Null protocol indicates the MP has no active layer 2 path selection and forwarding. An MP with Null protocol does not send or respond to path selection protocol messages.

7.3.2.49.2 Path selection metric identifier field

The path selection metric identifier is contained in the Active Metric field. This information identifies the path metric which is currently used by the active path selection protocol in the mesh network. The path selection metric identifier has the format shown in Figure s19.

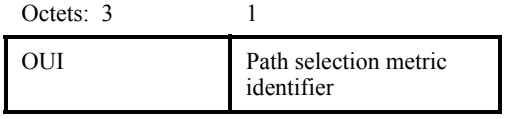

Figure s19—Path selection metric identifier field

The path selection metric identifier specifies the metric that is used when selecting routes in this network, as defined in Table s4.

Table s4—Path selection metric identifier values

OUI	Value	Meaning
00-0F-AC	0	Airtime path metric (default path metric)
00-0F-AC	1-254	Reserved for future use
00-0F-AC	255	Null metric
Vendor OUI	Other	Vendor specific

The WLAN Mesh Capability Element indicates an active null metric is used in conjunction with null path selection protocol and an active path metricsetting.

7.3.2.49.3 Peer Capacity Field

The peer capacity value is treated as a single field, with the least significant octet transmitted first. It contains four sub-fields as shown in Figure s13Figure s20.

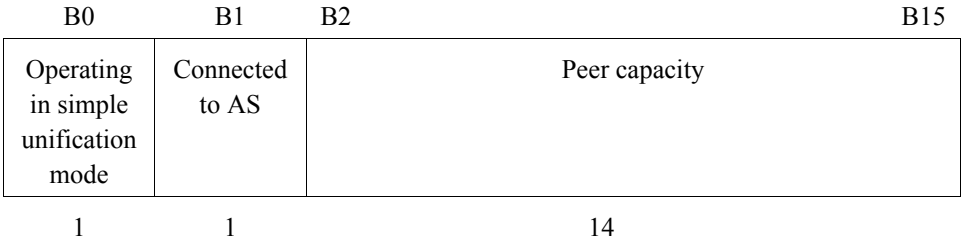

Figure s20—Peer Capacity Field

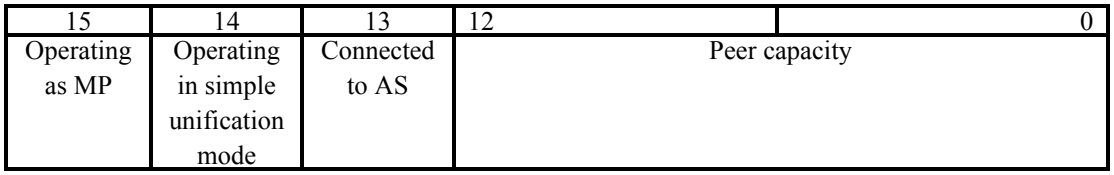

Figure s21—Peer Capacity Field

The "operating as MP" subfield is set to one if the device is currently operating as a MP, and zero otherwise. The "operating in simple unification mode (Clause 11A.1.7.4)" subfield bit is set to one if the logical radio interface PHY is currently operating in the simple unification mode, and zero otherwise. The "connected to AS" subfield bit specifies whether the mesh point MP is connected to an AS (Authentication Server), enabling the MP to support authentication and key management with IEEE 802.1X. The "peer capacity" subfield is set to the number of additional MP peers that the device can accommodate.

7.3.2.49.4 Power Save Capability field

The Power Save capability field includes 5 sub-fieldsis shown in Figure s22.

B0	B1	B2	B3	B7
Supporting Power Save Mode	Require Power Save Mode from Peer	Power Save Mode Enabled	Reserved	
Bits: 1	1	1	5	

Figure s22—Power Save Capability field

7	6	5	4-0
Supporting Power Save Mode	Require Power Save Mode from Peer	Current Power Management Mode	Reserved

Figure s23—Power Save Capability field

The "Supporting Power Save Mode" subfield bit indicates if the Mesh Point MP supports power save mode.

The "Require Power Save Mode from peer" subfield bit indicates if this mesh point MP requires peers attempting to associate or directly communicate with it to support Power Save mode.

The Current Power Management Mode" sub filed indicates the current "Power Save Management Mode of the mesh point. A Enabled" bit, when set bit to 1, indicates the mesh point MP is operating in Power Save modeMode. A cleared bit value of 0 indicates it is not operating in active modePower Save Mode.

7.3.2.49.5 Synchronization Capability field

The Synchronization Capability field includes 3 sub-fieldsfields as shown in Figure s24.

B0	B1	B2	B3	B7
Supporting Syn-chronization	Requests Syn-chronization from Peer	Synchronizing with Peer MP	Reserved	
Bits: 1	1	2	5	

Figure s24—Synchronization Capability Field

Bits: 0	1	2	3-7
Supporting Synchronization	Requests Synchronization from Peer	Synchronizing with peer MP	Reserved

Figure s25—Synchronization Capability Field

The "Supporting Synchronization" subfield is set to 1 if the Mesh Point MP supports timing synchronization with peer MPs, and 0 otherwise.

The "Requests Synchronization from Peer" subfield is set to 1 if the mesh point MP requests MP peers attempting to communicate with it to synchronize with it, and 0 otherwise.

The "Synchronizing with Peer MP" subfield is set to 1 if the non-AP MP is currently a synchronizing MP, and 0 otherwise.

7.3.2.49.6 MDA Capability field

The MDA Capability field includes 5 sub-fields:fields as shown in Figure s26.

B0	B1	B2	B3	B4	B5 B7
MDA Capable	MDA Active	MDA Active Requested in Mesh	MDA Not Allowed in Mesh	MDA EDCA Mixed Mode Enabled	Reserved
Bits: 1	1	1	1	1	3

Figure s26—MDA Capability Field

Bits: 0	1	2	3	4	5-7
MDA Capable	MDA Active	MDA Active Requested in Mesh	MDA Not Allowed in Mesh	MDA EDCA Mixed Mode Enabled	Reserved

Figure s27—Figure s16: MDA Capability Field

The "MDA Capable" subfield is set to 1 if the Mesh Point MP supports MDA servicesMDA, and 0 otherwise.

The "MDA active" subfield is set to 1 if the mesh point has MDA services is active, and 0 otherwise. When set to 1, the MP provides full MDA services as described in 9.14. When set to 0, the MP does not interpret any frames described for MDA operation and does not provides any services described in 9.14. This field is ignored and interpreted as 0, if "MDA capable" bit is set to 0.

The "MDA Active Requested in Mesh" field is set to 1 if the mesh requires that all MPs that are capable of MDA have MDA active, and to 0 otherwise. This is an informative flag and the participating MPs are not required to act on it. This field is ignored and interpreted as 0 if either of "MDA Capable" or "MDA Active" bits are set to 0. This field is a mesh wide field, and common for all MPs that belong to a single mesh.

The "MDA Not Allowed in Mesh" subfield is set to 1 if the mesh requires that all MPs participating in the mesh have MDA not activated, and 0 otherwise. If this field is set to 1, the "MDA Active" and "MDA Active Required in Mesh" fields are ignored and interpreted as 0. This field is a mesh wide field, and common for all MPs that belong to a single mesh. If this field is set to 1 in a mesh, any new MP that wishes to participate in the mesh is required to not use/invoke MDA services.

When "MDA not Allowed in Mesh" is set to 0, MPs that are capable of providing MDA services may use MDA in the Mesh for portions of their traffic.

The "MDA EDCA Mixed Mode Enabled" field is set to 1, if MDA traffic may be transmitted using EDCA access along with MDA access, and 0 otherwise. If this filed field is set to 0, any flow that is set to access the channel using MDA may only transmit during MDAOPs of the transmitter MP.

The channel precedence field is set to the value of channel precedence of the unified channel graph to which the MP interface belongs.

7.3.2.50 Path selection protocol identifier element

7.3.2.51 Active Mesh Profile Announcement element

The "Active Mesh Profile Announcement" element is used to notify the profile pair of the active path selection protocol and the active path metric to a peer MP. It is contained in association request messages transmitted by the association requesting MP. The profile pair is selected by the association requesting MP. The format of the Active Mesh Profile Announcement element is shown in Figure s28.

Octets: 1	1	1	4	4
ID	Length	Version	Active Protocol ID	Active Metric ID

Figure s28—Active Mesh Profile Announcement Element

The Element ID is set to the value given in Table 26 for this information element. The Length field is set to 9. The version field is set to 1.

The path selection protocol identifier element is contained in the Active Protocol field. This information identifies ID indicates the path selection protocol for unicast, multicast in use and broadcast transmission. Protocol profile has format shown is formatted as described in Figure s177.3.2.49.1.

Octets: 3	1
OUI	Path selection protocol identifier

Figure s29—Path selection protocol identifier element format

The path selection protocol identifier specifies the protocol which is currently used to generate routing information in this network, as defined in Clause 11A.2.

Table s5—Path selection protocol identifier Values

OUI	Value	Meaning
00-0F-AC	0	Hybrid Wireless Mesh Protocol (default path selection protocol)
00-0F-AC	1	Radio Aware OLSR (optional path selection protocol)
00-0F-AC	2-254	Reserved for future use
00-0F-AC	255	Null protocol
Vendor OUI	Other	Vender specific

A Null protocol indicates the MP has no active layer 2 path selection and forwarding. An MP with Null protocol will not send or respond to path selection protocol messages.

7.3.2.52 Path selection metric identifier element

The path selection metric identifier is contained in the Active Metric field. This information identifies the path metric which is currently used by the active path selection protocol in the WLAN mesh network. Path selection metric identifier has format shown in Figure s18.

Octets: 3	1
OUI	Path selection metric identifier

Figure s30—Path selection metric identifier format

The Active Metric ID indicates the path selection metric identifier specifies the metric that in use and is used when selecting routes in this network, formatted as defined described in Table s47.3.2.49.2.

Table s6—Path selection metric identifier Values

OUI	Value	Meaning
00-0F-AC	0	Airtime path metric (default path metric)
00-0F-AC	1-254	Reserved for future use
00-0F-AC	255	Null metric
Vendor OUI	Other	Vender specific

Null metric is used in conjunction with Null protocol setting.

7.3.2.53 Active Profile Announcement element

The "Active Profile Announcement" element is used to notify the profile pair of the active path selection protocol and the active path metric to a peer MP. It is contained in association request message transmitted by the association requesting MP. The profile pair is selected by the association requesting MPs local mechanism.

Octets: 1	1	1	4	4
ID	Length	Version	Active Protocol ID	Active Metric ID

Figure s31—Active Profile Announcement Element

The fields contained in the element are as shown in Table s5.

Table s7—Table s5: Active Profile Announcement Element Values

Field	Value/description
ID	T.B.D
Length	Variable
Version	1
Active Protocol ID	Path Selection Protocol in use
Active Metric ID	Path Selection Metric in use

7.3.2.54 Mesh ID element

The "Mesh ID" element is used to advertise the identification of a WLAN mesh networknetwork and is described in 11A.1.2. A 0 length information field may be used within Probe Request management frame to indicate the wildcard Mesh ID. The contents are shown in Figure s32

Octets: 1	1	0-32
Element ID	Length	Mesh ID

Octets: 1	1	0-32
Element ID	Length	Mesh ID

Figure s32—Mesh ID element format

Figure s33—Mesh ID element format

The Element ID is set to the value given in Table 26 for this information element.

The length of the Mesh ID is between 0 and 32 octets. A value of 0 in the length field may be used to indicate the wildcard Mesh ID.

7.3.2.55 Local Link state announcement element

A local link state announcement element is transmitted by an MP to a neighbor MP to indicate the quality of the link between them. This information may be used to ensure verify that the link quality is symmetric for all Mesh links if the path selection protocol so requires. The contents of the element are shown in Figure s34.

Octets: 1	1	2	2
ID	Length	r	e_f

Figure s34—Mesh Local Link State Announcement Element

Octets: 1	1	2	2
ID	Length	*r*	e_{pt}

Figure s35—WLAN Mesh Local Link State Announcement Element

The fields contained in the element are as shown in Table s6.

Table s8—WLAN Mesh Local Link State Announcement Element Fields

Field	Value/description
ID	TBD
Length	8
r	Transmit bit rate
e_{pt}	PER

The Element ID is set to the value given in Table 26 for this information element. The length is set to 4.

The rate field, *r*, shall be is interpreted as a 16-bit unsigned integer, with the least significant octet transmitted first, which indicates the on-air bit rate currently in use in units of 1Mbit/s.

The PERframe error rate, $e_{pf}e_f$, shall be is interpreted as a 16-bit unsigned binary fraction, with the least significant octet transmitted first, such that a value of 0xffff corresponds to a fractional value of:

$$1 - \frac{1}{2^{16}}$$

The PER frame error rate indicates an estimated packet frame error rate for a data frame containing a payload of 1000 bytes transmitted at the bit rate specified in the *r* field.

The local link state announcement element may be used in the generation of a link metric such as the airtime metric defined in 11A.5.

7.3.2.56 OFDM Parameter Target Transmission Rate Element

A new "OFDM Parameter" element is used to advertise current channel identification to the neighbor MP with OFDM PHY. It is contained in beacons transmitted by MPs, and is also contained in probe response messages.

The OFDM Parameter Set element contains information to allow channel number identification for MPs using an OFDM PHY. The information field contains a single parameter containing the dot11CurrentChannelNumber (see Clause 15.4.6.2 for values). The length of the dot11CurrentChannelNumber parameter is 1 octet. See Figure s22.

Octets: 1	1	1
ID	Length	Current Channel

Figure s36—OFDM Parameter Set Element Format

7.3.2.57 Target Transmission Rate Element

A new "The Target Transmission Rate" Rate element is used in the Flow Congestion Control Request frame by a Mesh Point an MP to indicate to its upstream neighbor the target data rate the two Mesh Points MPs should coordinate to maintain. It contains four target Data Rate fields for the four EDCA access categories and an Expiration Timer. The format of the element is shown in Figure s37.

Octets: 1	1	4	4	4	4	2
ID	Length	Target Data Rate (AC_BK)	Target Data Rate (AC_BE)	Target Data Rate (AC_VI)	Target Data Rate (AC_VO)	Expiration Timer

Figure s37—Target Transmission Rate Element Format

Octets: 1	1	4	4	4	4	2
ID	Length	Target Data Rate (AC_BK)	Target Data Rate (AC_BE)	Target Data Rate (AC_VI)	Target Data Rate (AC_VO)	Expiration Timer

Figure s38—Target Transmission Rate Element Format

The Element ID is set to the value given in Table 26 for this information element. The length is set to 18.

Target Data Rate for a given Access Category indicates the mean transmission rate in bitsbit/s from the upstream Mesh Point MP to the downstream neighboring Mesh Point MP for this AC. The upstream Mesh Point must MP does not exceed the target Data Rate; otherwise it may result in congestion at the downstream nodeMP.

Expiration timer indicates the valid period of Target Data Rates information provided in this element. Expiration Timer is represented in TU.

7.3.2.58 Offered Traffic Load Element

A new An "Offered Traffic Load" element is used in the Flow Congestion Control Response frame by a Mesh Point an MP to indicate to its downstream neighbor the incoming traffic load between itself and the downstream neighbor. It contains four Offered Traffic Load fields for the four access categories. The format of the element is shown in Figure s39.

Octets: 1	1	4	4	4	4
ID	Length	Offered Traffic Load (AC_BK)	Offered Traffic Load (AC_BE)	Offered Traffic Load (AC_VI)	Offered Traffic Load (AC_VO)

Figure s39—Offered Traffic Load Element Format

Octets: 1	1	4	4	4	4
ID	Length	Offered Traffic Load (AC_BK)	Offered Traffic Load (AC_BE)	Offered Traffic Load (AC_VI)	Offered Traffic Load (AC_VO)

Figure s40—Offered Traffic Load Element Format

The Element ID is set to the value given in Table 26 for this information element. The length is set to 16.

Offered Traffic Load for a given Access Category indicates the incoming traffic rate in bitsbit/s estimated or measured at the MAC interface. This information can be used by the downstream neighboring Mesh Point MP to better estimate the Target Data Rate.

7.3.2.59 Neighborhood Congestion Element

A new "Neighborhood Congestion" element is used by a Mesh Point an MP to indicate to its neighbors its congestion level. It contains a congestion level field and an expiration timer field. The format of the element is shown in Figure s41.

Octets: 1	1	1	2
ID	Length	Congestion Level	Expiration Timer

Figure s41—Neighborhood Congestion Element Format

Octets: 1	1	1	2
ID	Length	Congestion Level	Expiration Timer

Figure s42—Neighborhood Congestion Element Format

The Element ID is set to the value given in Table 26 for this information element. The length is set to 3.

The congestion level is TBD.

Expiration The expiration timer indicates the valid period of congestion information provided in this element. Expiration Timer element and is represented in TU.

7.3.2.60 Peer Link Close element

The Peer Link Close element is transmitted by a Mesh Point an MP Requesting to close tear down a link with a peer Mesh PointMP. This element may be transmitted in a Disassociate Disassociation frame sent from one Mesh Point MP to another. The format of the Peer Link Close element is shown in Figure s26Figure s43.

	Element ID	Length	Reason Code	Local Link ID	Peer Link ID
Octets:	1	1	1	4	4

Octets: 1	1	1	4	4
Element ID	Length	Reason Code	Local Link ID	Peer Link ID

Figure s43—Peer link close element

Figure s44—Peer link close element

The fields contained in the element are as shown in Table s7.

Table s9—Peer link close element fields

Field	Value/description
ID	TBD
Length	9
Reason Code	The Reason Code field is set to value that represents the reason to close a peer link. The reason code is specified in Table s8.

Local Link ID	Random value generated by local system in the effort to identify link instance with the peer
Peer Link ID	Random value received from the peer in the effort to identify the same link instance

The Element ID is set to the value given in Table 26 for this information element. The length is set to 9.

The Reason Code field is set to the value that represents the reason to tear down a peer link. The reason code is specified in Table s10.

The Local Link ID field contains a pseudo-random number generated by the local system in order to create a statistically unique identifier for this the link instance with the peer MP.

The Peer Link ID field contains a pseudo-random number received from the peer, via a Peer Link Open or a Peer Link Confirm frame. The pair <Local Link ID, Peer Link ID> together with both MPs' identifier (e.g., their MAC addresses) uniquely identifies this the link instance to be established between these the two MPs.

Table s10—Peer link close reason code field values

Reason Code value	Meaning
0	Cancelled
1	Close received
2	Invalid parameters
3	Exceed maximum retries
4	Timeout

Table s11—Peer link close reason code field values

Reason Code value	Meaning
0	Cancelled
1	Close received
2	Invalid parameters
3	Exceed maximum retries
4	Timeout

7.3.2.61 Peer Link Open element

The Peer Link Open element is transmitted by a Mesh Point an MP requesting to open a link with a peer Mesh PointMP. This element may be transmitted in an Association Request frame sent from one Mesh Point MP to another. The format of the Peer Link Open element is shown in Figure s27Figure s45.

Octets: 1	1	4
Element ID	Length	Local Link ID

Figure s45—Peer link open element

	Element ID	Length	Local Link ID
Octets:	1	1	4

Figure s46—Peer link open element

The fields contained in the element are as shown in Table s9.

Table s12—Peer link open element fields

Field	Value/description
ID	TBD
Length	4
Local Link ID	Random value generated by local system in the effort to identify link instance with the peer

The Element ID is set to the value given in Table 26 for this information element. The length is set to 4.

The Local Link ID field contains a pseudo-random number generated by the local system in order to create a unique identifier for this the link instance with the peer MP. It is one portion of the complete link instance identifier. The pair <Local Link ID, Peer Link ID> together with both MPs' identifier (e.g., their MAC addresses) uniquely identifies this the link instance to be established between these the two MPs.

7.3.2.62 Peer Link Confirm element

The Peer Link Confirm frame element is transmitted by a Mesh Point an MP requesting to confirm a link with a peer Mesh PointMP. This element may be transmitted in an Associate Response frame sent from one Mesh Point MP to another. The format of the Peer Link Confirm element is shown in Figure s28Figure s47.

Octets: 1	1	4	4
Element ID	Length	Local Link ID	Peer Link ID

Figure s47—Peer link confirm element

	Element ID	Length	Local Link ID	Peer Link ID
Octets:	1	1	4	4

Figure s48—Peer link confirm element

The fields contained in the element are as shown in Table s10.

Table s13—Peer link confirm element fields

Field	Value/description
ID	TBD
Length	8
Local Link ID	Random value generated by local system in the effort to identify link instance with the peer
Peer Link ID	Random value received from the peer in the effort to identify the same link instance

The Element ID is set to the value given in Table 26 for this information element. The length is set to 8.

The Local Link ID field contains a pseudo-random number generated by the local system in order to create a unique identifier for this the link instance with the peer MP.

The Peer Link ID field contains a pseudo-random number received from the peer, via a Peer Link Open or a Peer Link Confirm frame. The pair <Local Link ID, Peer Link ID> together with both MPs' identifier (e.g., their MAC addresses) uniquely identifies this the link instance to be established between these the two MPs.

EDITORIAL NOTE—Former Mesh Portal Reachability* Element*element replaced by Portal Announcement

"Mesh Portal Reachability" element is used to advertise the identification of one or more Mesh Portal to which an MP is able to communicate. This information element is included in Beacon and Probe response frame.

Octets: 1	1	1	10*n	
Element ID	Length	Number of Mesh Portals	Mesh Portal Description	

Figure s49—Mesh Portal Reachability element format

Octets: 6	4
Mesh Portal MAC address	Metric

Figure s50—Mesh Portal Description format

Table s14—Mesh Portal Reachability element fields

Field	Value/description
ID	TBD
Length	Value
Number of Mesh Portals	The number of connected Mesh portals
Mesh Portal Description	The list of Mesh Portal descriptions for Mesh Portals that are reachable from the MP. This field has a description entry for each reachable Mesh Portal consisting of the Mesh Portal MAC address and the path metric from the MP to the Mesh Portal.

7.3.2.63 Mesh Portal/Root Announcement Element

Mesh Portals periodically broadcast *Portal announcement* to the WLAN Mesh network every PORTAL_ANNOUNCEMENT_TIME interval. The *Portal announcement* messages serve three purposes to the mesh network:

e) Announcing the Mesh Portal to the WLAN Mesh nodes allows the MPs to designate the Mesh Portal as a packet forwarder after receiving the annoucement message if they wish to.

f) Announcing the Mesh Portal to other Portals helps all Portals discover the uplinks present in the network and ensures that frames with unknown addresses are forwarded on all uplinks from the mesh network.

g) Allows unconfigured Root Portals to arbitrate and choose a single Portal as Root for hybrid wireless mesh protocol (HWMP) route building when route building is enabled via a management entity.

The *Root announcement* with HWMP-Registration flag set allows the HWMP topology to form. When this flag is not set, the announcements propagate in the network as per Clause 11A.5.

Octets: 1	1	1	6	1	1	6
Element ID	Length	Flags	Mesh Portal Bridge ID	Priority	Number of Mesh Portals	**Mesh Portal Address**

4	4	4	1	6*n
Root Seq. Num	Lifetime	Root Metric	Topology Maintenance Policy	Connected Mesh Portal IDs

Figure s51—Mesh Portal/Root Announcement Element

Table s15—Mesh Portal/Root Announcement Element Fields

Field	Value/description
ID	TBD
Length	Length of the IE
Flags	Bit 0: Announcement type (0 = Portal; 1 = Root) Bit 1: HWMP Registration (0 = Disabled; 1 = Enabled) If the flag is set, registration of MPs with the root portal occurs, otherwise it does not. Bit 2 – 7: Reserved
Mesh Portal Bridge ID	Mesh Portal's unique bridge ID (802.1D configuration BPDUs contain bridge ID)
Priority	A Mesh Portal has the lowest priority value becomes the default Mesh Portal. (0~255). A value of 0 means it is configured as a Root.
Number of Mesh Portals	The number of the connected Mesh Portals.
Mesh Portal Address	Mesh Portal MAC address
Root Sequence Number	The latest sequence number received in the past by the source for any route towards the Root.
Life Time	The time in milliseconds for which nodes receiving the Mesh Portal/ Root Announcement is valid.
Metric	The cumulative metric from the Root to the node advertising the announcement.
Topology Maintenance Policy	0: Policy 1 (Suboptimal) 1: Policy 2 2: Policy 3 3: Policy 4 (Optimal)
Connected Mesh Portal IDs	The list of Mesh Portal MAC addresses which have both wired and wireless connectivity. A value of 0 indicates it is a Root announcement and not a Portal announcement.

EDITORIAL NOTE—Former Mesh Portal/Root Announcement Element superceded by PANN/RANN elements.

7.3.2.64 Unified Channel Graph UCG Switch Announcement element

The Unified Channel Graph UCG Switch Announcement element is used by an MP in a WLAN Mesh to advertise when it is changing to a new channel and the channel number and precedence value of the new

channel (See 11A.1.7.5). The format of the Unified Channel Graph UCG Switch Announcement element is shown in Figure s32Figure s52.

Octets: 1	1	1	1	4	1	6
ID	Length	Channel Switch Mode	New Channel Number	New Channel Precedence Indicator	Channel Switch Count	Source Address

Octets: 1	1	1	1	4	1	6
ID	Length	Channel Switch Mode	New Chan-nel Number	New Chan-nel Prece-dence Indicator	Channel Switch Count	Source Address

Figure s52—UCG Switch Announcement Element

Figure s53—Unified Channel Graph Switch Announcement Element

The Length field shall be set to 13.

The Element ID is set to the value given in Table 26 for this information element. The length is set to 13 octets.

The Channel Switch Mode field indicates any restrictions on transmission until a channel switch. an An MP shall set sets the Channel Switch Mode field to either 0 or 1 on transmission. A Channel Switch Mode set to 1 means that the MP to which the frame containing the element is addressed is advised to transmit no further frames on the current channel until the scheduled channel switch. A Channel Switch Mode set to 0 does not impose any requirement on the receiving STA.

The New Channel Number field shall be is set to the number of the channel to which the MP is moving.

The New Channel Precedence Indicator field shall be is set to the channel precedence value of the channel to which the MP is moving.

The Channel Switch Count field either shall be is set to the number of time units (TUs) until the MP sending the Unified Channel Graph UCG Switch Announcement element switches to the new channel or shall be is set to 0. A value of 0 indicates that the switch will may occur at any time after the frame containing the element is transmitted.

The Source Address field shall be is set to the address of the Mesh Point MP that originates the frame.

The Unified Channel Graph UCG Switch Announcement element is included in Unified Channel Graph UCG Switch Announcement frames.

7.3.2.65 Mesh Neighbor List element

The Mesh Neighbor List element is used by an MP to advertise its associated neighbor list peer neighbors and their Power Management Mode. The element contains list of current associated neighbor the MAC addresses of current peer neighbors and information about the neighbor power management modetheir Power Management Mode. The Neighbor List element MP Control field contains the connectivity reporting control information.

The format of the Mesh Neighbor List element is shown in Figure s33Figure s54.

Octets: 1	1	1	1	6	6	...	6	ceiling(n/8)
ID	Length	MP control	Neighbor Count	MAC Address of neighbor 1	MAC Address of neighbor 2	...	MAC Address of neighbor n	Neighbor operating in power save mode (bitfield)

Figure s54—Mesh Neighbor List element

Octets: 1	1	1	6	6	...	6	k
ID	Length	MP control	MAC Address of terminal 1	MAC Address of terminal 2	...	MAC Address of terminal n	Neighbor power management mode

Figure s55—Neighbor List element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 1 to 255 octets.

The format of the MP control field is shown in Figure s34Figure s56.

B1 B3	B4	B5	B6	B7
Connectivity Reporting Interval	Reserved	Designated BB	BB switch	BB power management mode
Bits: 4	1	1	1	1

Figure s56—MP Control field

Bits: 0-3	4	5	6	7
Connectivity Reporting Interval	Reserved	Designated BB	BB switch	BB power management mode

Figure s57—MP Control field

The Connectivity Reporting Interval specifies an integer value of the MESH mesh DTIM beacons Beacon frames between the Connectivity Report transmissions. Connectivity Reporting Interval set to zero indicates that connectivity reporting is not used.

The Designated BB field is used to indicate that the current beacon broadcaster is a designated beacon broadcaster that operates according to the scheme described in Clause 11A.9.3.3.

The BB switch bit field is used to indicate the change of the beacon broadcaster. If the this bit in the DTIM frame is set to 1, the next beacon shall be is sent by the MP whose MAC address is the first one in the neighbor list.

The BB power management mode field indicates the beacon broadcaster's power management mode. If the BB power management mode bit is set to 1, then the beacon broadcaster sends only Mesh DTIM frames. If the BB power management mode bit is set to 0, then the beacon broadcaster is in active mode.

The MAC addresses of the terminals neighbors are set for the MPs in the BSSthat have established links to their peers, which are listed in the Connectivity Reports received by the BB within `dot11BBConnectivityReportTimeout` mesh DTIM intervals.

The neighbor operating in power management save mode field bitfield indicates the current power management save mode of each neighbor list member. Each bit of this field indicates the power save state management mode of the corresponding neighbor list member. If a bit is set to 0, then the corresponding neighbor list member is in "active mode" and if a bit is set to 1, the corresponding neighbor list member is in "Power Save mode". For example, if the neighbor Mesh Neighbor List element contains 8 MAC addresses and the neighbor operating in power management save mode field bitfield is '00110001', then the MPs with MAC addresses in positions 3, 4, and 8 in the neighbor list are in the power save mode. The neighbor operating in power management save mode field bitfield length is always zero-padded to an integer number of octets. The bits are in the same order as the MAC addresses.

7.3.2.66 Mesh DTIM element

The Mesh DTIM element is used by an MP acting as a beacon broadcaster. The element contains information about the mesh DTIM period of the Mesh.

The format of the Mesh DTIM element is shown in Figure s35Figure s58.

Octets: 1	1	1	1
ID	Length	DTIM Count	DTIM Period

Figure s58—Mesh DTIM element

Octets: 1	1	1	1
ID	Length	DTIM Count	DTIM Period

Figure s59—DTIM element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 2.

The mesh DTIM Count field indicates how many beacons Beacon frames (including the current frame) appear before the next mesh DTIM. A mesh DTIM count of 0 indicates that the current TIM is a mesh DTIM. The mesh DTIM count field is a single octet.

The mesh DTIM Period field indicates the number of Beacon intervals between successive mesh DTIMs. If all TIMs are mesh DTIMs, the mesh DTIM Period has the value 1. The mesh DTIM Period value 0 is reserved. The mesh DTIM period field is a single octet.

MAP beacons will Beacon frames may include both a TIM and a Mesh DTIM element. The mesh DTIM Period of these IEs do not have to be identical since one will be is used for the AP service while the other will be is used for the Mesh service.

7.3.2.67 Beacon Timing element

The Beacon Timing element is used by a synchronizing MP to advertise an offset between its self TSF and the Mesh TSF, and to advertise the beacon timing information of zero or more of its neighbors. The format of the Beacon Timing element is shown in Figure s60.

Octets: 1	1	4	1	1	3	1	3	...	1	3	1	5	…	1	5
ID	Length	Self Beacon Timing	Number of Synchronizing neighbors reported	Last byte of MAC Address of Synch MP 1	Synchronized Beacon Timing MP 1	Last byte of MAC Address of Synch MP 2	Synchronized Beacon Timing MP 2	...	Last byte of MAC Address of Synch MP n	Synch Beacon Timing MP n	Last byte of MAC Address of Unsynch MP 1	Unsynchronized Beacon Timing unsynch MP 1	…	Last byte of MAC Address of Unsynch MP m	Unsynchronized Beacon Timing unsynch MP m

Figure s60—Beacon Timing element

The format of the Beacon Timing element is shown in Figure s36.

Octets: 1	1	4	1	1	3	1	3	...
ID	Length	Self Beacon Timing	Number of Synchronizing neighbors reported	Last byte of MAC Address of Synch Terminal 1	Synchronized Beacon Timing terminal 1	Last byte of MAC Address of Synch Terminal 2	Synchronized Beacon Timing terminal 2	...

1	3	1	5	…	1	5
Last byte of MAC Address of Synch Terminal n	Synch Beacon Timing terminal n	Last byte of MAC Address of Unsynch Terminal 1	Unsynchronized Beacon Timing unsynch terminal 1	…	Last byte of MAC Address of Unsynch Terminal m	Unsynchronized Beacon Timing unsynch terminal m

Figure s61—Beacon Timing element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 5 to 255 octets.

The format of the Self Beacon Timing field is shown in Figure s37Figure s62.

Octets: 3	1
Self TBTT offset	MP DTIM period

Figure s62—Self Beacon timing

Bits: 0-23	24-31
Self TBTT offset	MP DTIM period

Figure s63—Self Beacon timing

The 'Self TBTT offset' subfield of the Self Beacon timing field indicates the offset, measured in microsecondsTUs, used by the MP for its TSF time compared to the Mesh TSF. The sum of the MP TSF time stamp and the offset equals the mesh TSF time.

The 'MP DTIM period' subfield of the Self Beacon timing field indicates the MP DTIM period of the specific MP (i.e., how many Beacon intervals of the MP compose a single Mesh DTIM interval).

The 'Number of Synchronizing Neighbors Reported' field specifies the number of synchronizing neighbors whose beacon timing information is reported following this field.

The beacon timing information of synchronizing neighbors is reported in terms of the last byte of MAC address field and the 'Synchronized Beacon Timing' field which are included as pairs. The Beacon Timing element may contain zero or more 'Last byte of MAC Address of Synch TerminalMP' and the 'Synchronized Beacon Timing' field pairs.

The Last Byte of MAC Address of Synch Terminal MP field indicates the last byte of the MAC address of neighbors that have a non-zero Self TBTT offset value, whose information is reported in the value of the next 'Synchronized Beacon Timing' field. Note that the The Last Byte of MAC address need not be unique, as the relevant information is the timing information that follows it. That is the The relevant information is the set of times when successful beacons Beacon frames are being received.

The format of the Synchronized Beacon Timing field is shown in Figure s38Figure s64.

Octets: 1	1	1
TBTT offset	Time since last beacon	MP DTIM period

Figure s64—Synchronized Beacon Timing field

Bits: 0-7	8-15	16-23
TBTT offset	Time since last beacon	MP DTIM period

Figure s65—Synchronized Beacon Timing field

The 'TBTT offset' field is expressed in units of TU, and indicates the offset used by the neighboring MPs for their TSF time stamps compared to the Mesh TSF. For those MPs reporting self TBTT offsets with a higher resolution than a TU, the field is rounded to the nearest TU.

The 'Time since last beacon' field indicates the time passed measured in units of Mesh DTIM intervals since last beacon was received from the specific MP.

The 'MP DTIM period' field indicates the MP DTIM period of the specific MP (i.e., how many Beacon intervals of the MP compose a single Mesh DTIM interval).

Beacon timing information for unsynchronizing MP neighbors may also be included at the end of the Beacon timing element following the synchronizing neighbor information.

The beacon timing information of unsynchronizing neighbors is reported in terms of the 'Last byte of the MAC Address of Unsynch TerminalMP' and the 'Unsynchronized Beacon Timing' fields which are included as pairs. The Beacon Timing element may contain zero or more 'Last byte of MAC Address of Unsynch Terminal' and 'Unsynchronized Beacon Timing' field pairs.

The 'Last byte of the MAC Address of Unsynch TerminalMP' indicates the last byte of the MAC address of unsynchronizing neighbor whose beacon timing information is reported in the value of the next 'Unsynchronized Beacon Timing' field.

The format of the 'Unsynchronized Beacon Timing' field is as shown in Figure s39Figure s66.

Octets: 3	2
Last Beacon Time	MP Beacon Interval

Figure s66—Unsynchronized Beacon Timing field

Bits: 0-23	24-39
Last Beacon Time	MP Beacon Interval

Figure s67—Unsynchronized Beacon Timing field

The Last Beacon Time field is measured in microseconds and specifies the time at which the last beacon from the specified MP was received relative to the beacon time stamp of the IE information element transmitting MP, or the most recent TBTT of the IE information element transmitting MP if the Beacon Timing element is encapsulated in a response frame.

The MP Beacon Interval field specifies the beacon interval being used by the MP whose information is being reported.

7.3.2.68 MDAOP Setup Request Element

The MDAOP Setup Request IE information element is used by an MP to request the setup of a set of MDAOPs, identified by a single MDAOP Set ID, between itself (transmitter) and a receiver. This IE information element is unicast transmitted in individually addressed MDA action frames. The format of the IE information element is as shown in Figure s40Figure s68.

Octets: 1	1	1	6	1	3		3	variable		variable
Element ID	Length	MDAOP Set ID	Final Destination MAC Address	Number of Individual MDAOPs	MDAOP Info 1	…	MDAOP Info n	Periodic MDAOP info 1	…	Periodic MDAOP info m

Figure s68—MDAOP Setup Request Element

Element ID (1 byte)	Length (1 byte)	MDAOP Set ID (1 byte)	Final Destination MAC Address (6 bytes)	Number of Individual MDAOPs (1 byte)	MDAOP Info 1 (3 bytes)	…	MDAOP Info n (3 bytes)	Periodic MDAOP info 1 (variable bytes)	…	Periodic MDAOP info m (variable bytes)

Figure s69—MDA Setup Request Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 8 to 255 octets.

The MDAOP Set ID field is an eight bit unsigned number that represents the ID for the requested Set.

The Final Destination MAC Address field is an informative field to indicate the ultimate destination address to the immediate receiver. The address received in this field may be used by the immediate receiver to setup MDAOPs of its own for the next hop. The description of the specific use of this information is out of scope for this documentstandard.

The Number of Individual MDAOPs field specifies the number of MDAOP Info fields following it.

The format of the MDAOP Info field is specified in table ddFigure s70. Each MDAOP Info field specifies the duration and location timing of a specific MDAOP in the Mesh DTIM interval. The MDAOP Duration field specifies the duration of the MDAOP in multiples of 32us32 μs . The MDAOP Offset field specifies the location of the MDAOP beginning start from the beginning of a Mesh DTIM Interval. The specification is in terms of multiples of 32 us32 μs .

Octets: 1	1
MDAOP Duration	MDAOP Offset

Figure s70—MDAOP Info field

MDAOP Duration (1 byte)	MDAOP Offset (2 bytes)

Figure s71—MDAOP Info field

The format of the Periodic MDAOP Info field is specified in Figure s42Figure s72.

Periodic MDAOP duration (1 byte)	Periodicity (1 byte)	Number of offsets (4 bits)	$Offset_1$ (12 bits)	…	$Offset_N$ (12 bits)	Padding if required (4 bits)

B0 B7	B8 B15	B16 B19	B20 B31			
Periodic MDAOP duration	Periodicity	Number of off-sets	$Offset_1$	…	$Offset_N$	Padding if required
Bits: 8	8	4	12		12	4

Figure s72—Periodic MDAOP Info field

Figure s73—Periodic MDAOP Info field

Each Periodic MDAOP Info field specifies information about a set of MDAOPs, each with identical periodicity within the Mesh DTIM interval, and identical durations.

Each Periodic MDAOP Info field specifies information about a set of MDAOPs, each with identical periodicity within the Mesh DTIM interval, and identical durations. The Periodic MDAOP Duration field specifies the duration of each of the MDAOPs specified in the field in multiples of 32 us32 μs .

The Periodicity field specifies the number of times the specified MDAOPs repeat themselves equidistantly within a Mesh DTIM interval.

The Number of Offsets field specifies the number of Offset fields following it.

Each offset field specifies the position of an MDAOP beginning from the beginning of the Mesh DTIM interval. Since each of these MDAOPs repeat with a periodicity, the same offset is used from the beginning of each of the sub-intervals within the Mesh DTIM interval.

The padding field is 4 bits, and is used to round the IE information element to the nearest byte.

An example of periodicity, duration, and offset values for a periodic MDAOP Info field is shown in Figure s43Figure s74. In this particular example, the periodicity equals four, so that there are four subintervals within the mesh DTIM interval. The number of offsets equals two, so that there are two MDAOPs within each subinterval. As further illustrated in the figure, the two offset values indicate the start of these MDAOPs relative to the start of these subintervals.

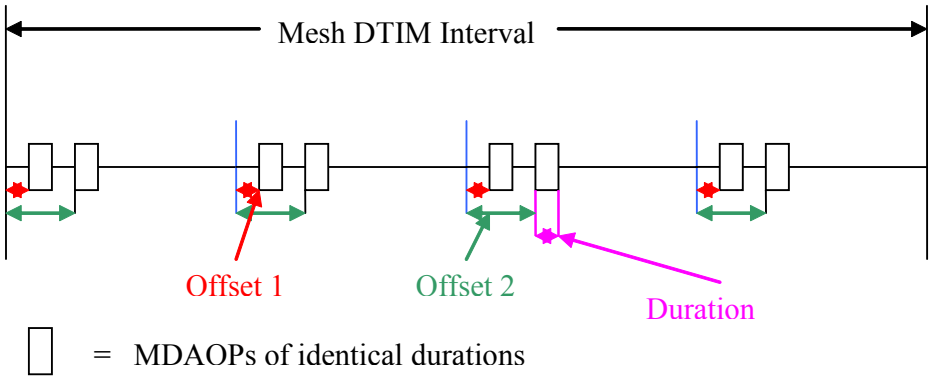

Figure s74—Values for Periodic MDAOP Info field for an example MDAOP set

7.3.2.69 MDAOP Setup Reply Element

The MDAOP Setup Reply element is used to reply to an MDAOP Setup Request. Its format is as shown in Figure s44Figure s75. This IE is unicast transmitted in MDA action frames. The reply code Element ID is 0 set to the request, and any odd number to reject the requestvalue given in Table 26 for this information element. The request code of 1 length is used set to indicate that the rejection is due to a bad choice of MDAOP locations within the Mesh DTIM interval. The rest of the code values are reserved for future use3.

Octets: 1	1	1	1	variable
Element ID	Length	MDAOP Set ID	Reply Code	Alternate suggested Request IE

Figure s75—MDAOP Setup Reply Element

Element ID (1 byte)	Length (1 byte)	MDAOP Set ID (1 byte)	Reply Code (1 byte)	Alternate suggested Request IE (variable length)

Figure s76—MDA Setup Reply Element

This information element is transmitted in individually addressed MDA action frames. The LSB set to 0 implies accept the request and 1 rejects the request. All other bits are reserved.

The Alternate Sugggested Suggested Request IE information element field includes a suggested format of the MDAOP Setup Request IE information element without the length and element ID fields, so that the setup request may be accepted. This field is an optional field and is only included/interpreted when the reply code indicates rejection.

7.3.2.70 MDAOP Advertisements Request Element

The MDAOP Advertisements Request element is used to request MDA advertisements from neighbors, one neighbor at a time. The IE information element may only be carried in an MDA action frame that is unicastindividually addressed. On the receipt of this IEinformation element, the receiver is required to broadcast its transmit group addressed MDAOP advertisements using the MDAOP advertisements IEinformation element. The format of the IE information element is shown in Figure s45Figure s77.

Octets: 1	1
Element ID	Length

Figure s77—MDAOP Advertisements Request Element

Element ID (1 byte)	Length (1 byte)

Figure s78—MDAOP Advertisements Request Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 0.

7.3.2.71 MDAOP Advertisements Element

The MDAOP Advertisements Element is used by an MP to advertise its MDA state to its neighbors. This IE information element may be carried in selected beacons Beacon frames with any chosen frequency. This IE information element may also be broadcast transmitted in an MDA action frame. The format of the IE information element is as shown in Figure s46Figure s79.

B0 B7	B8 B15	B16 B19	B20 B23		
Element ID	Length	MDA Access fraction	MDA Access fraction limit	TX-RX times report	Interfering times report
Bits: 8	8	4	4	variable	variable

Figure s79—MDAOP Advertisements Element

Element ID (1 byte)	Length (1 byte)	MDA Access fraction (0.5 byte)	MDA Access fraction limit (0.5 bytes)	TX-RX times report (variable bytes)	Interfering times report (variable bytes)

Figure s80—MDAOP Advertisements Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 1 to 255 octets.

MDA Access Fraction and MDA Access Fraction Limit fields are both 4 bit unsigned number fields. They denote a positive fraction equal to the values expressed in units of the fields times (1/16). The MDA Access Fraction field represents the current value of MDA Access Fraction at the MP rounded down (floor) to the nearest multiple of (1/16). The MDA Access Fraction Limit field represents the maximum MDA access fraction allowed at the MP. This number is always a multiple of (1/16).

The TX-RX Times Report field is a variable length field that advertises the times in the Mesh DTIM interval that are busy for the MP as a transmitter or a receiver. These times may include known and otherwise un-advertised transmission and reception times besides MDAOPs. For example, an MAP may include its HCCA times in the advertisement.

The Interference Times Report field is identical in format to the TX-RX times report field. However, through this field, an MP reports the times when one of its neighbors are is in TX or RX as reported by their (neighbors') TX-RX times report fields. This field may not include any times for which the MP is a transmitter or receiver (Such times are taken care of in the TX-RX times report). The Interfering Times reported may not be used to transmit through MDA to the reporting MP because there may be interference between any such transmission sequence and the transmission sequences already setup for the reported times. However, these reported times may possibly be used for transmissions through MDA to other MPs.

The format of the TX-RX Times and the Interference Times Report field is shown in Figure s47Figure s81. The fields involved are simlar similar to the fields involved in the MDAOP Setup Request element, and are described above. Note that while While the fields are the same as in an MDAOP setup request element, the TX-RX times and Interfering times reports can more efficiently report information compared to MDAOP setup request IEs. This is possible because MDAOPs of different MDAOP Sets may be all combined in an efficient way and reported. MDAOP Set IDs are not reported in the advertisements.

Number of Individual MDAOPs (1 byte)	MDAOP Info 1 (3 bytes)	…	MDAOP Info n (3 bytes)	Periodic MDAOP info 1 (variable bytes)	…	Periodic MDAOP info m (variable bytes)
Octets: 1	3		3	variable		variable
Number of Individual MDAOPs	MDAOP Info 1	…	MDAOP Info n	Periodic MDAOP info 1	…	Periodic MDAOP info m

Figure s81—The format of the TX-RX times report and Interfering times report fields

Figure s82—The format of the TX-RX times report and Interfering times report fields

7.3.2.72 MDAOP Set Teardown Element

The MDAOP Teardown element is as shown in Figure s48Figure s83, and is used to indicate announce the peer of teardown of an MDAOP Set. An MDAOP Set teardown IE information element may be transmitted by either the transmitter or the receiver of the MDAOP Set to tear it down. The MDAOP Set Owner field is an optional field that indicates the MAC address of the owner (transmitter) of the MDAOP Set. This field is only included if the IE information element is transmitted by the receiver in an MDAOP set, to tear it down. The MDAOP teardown element may be transmitted in beacons Beacon frames or group addressed MDA

action frames. When received in broadcast frames (e.g., beaconsBeacon frames), the MDAOP Set Owner field is ignored, even if it is present.

Octets: 1	1	1	6
Element ID	Length	MDAOP Set ID	MDAOP Set Owner

Figure s83—MDAOP Teardown element

Element ID (1 byte)	Length (1 byte)	MDAOP Set ID (1 byte)	MDAOP Set Owner (6 byte)

Figure s84—MDAOP Teardown element

7.3.2.73 Connectivity Report element

The Connectivity Report element is used by an MP to list the number of beacon broadcasters during the reporting interval and to list the neighbors that transmitted a connectivity report and their Power Management Mode. The element contains a list of neighbors neighbor MAC addresses, where the connectivity report has been received and information about the neighbor power management mode. The format of the Connectivity Report is shown in Figure s49Figure s85.

Octets: 1	1	2	6	6	…	6	6	. . .	6	K
ID	Length	Connectivity Report Control	SSID	MAC Address of Beaconing MP 1	…	MAC Address of beaconing MP n	MAC Address of Connectivity Reporting MP 1	. . .	MAC Address of Connectivity Reporting MP n	Beacon and Connectivity Reporting Power management mode

Figure s85—Connectivity Report

Octets: 1	1	2	6	6	…	6	6	...	6	K
ID	Length	Connectivity Report Control	SSID	MAC Address of Beaconing MP 1	…	MAC Address of beaconing MP n	MAC Address of Connectivity Reporting MP 1	...	MAC Address of Connectivity Reporting MP n	Beacon and Connectivity Reporting Power management mode

Figure s86—Connectivity Report

The Element ID is set to the value given in Table 26 for this information element.

The Length field shall be is set to 18+6*(number of Beacon MACs+number of terminals)+n + m) + Ceiling((n + m)/8), where n is the number of Beacon MACs+MACs and m is the number of terminals)/8)MPs. The format of the Connectivity Report Control field is shown in Figure s50Figure s87.

B0 B3	B4 B6	B7	B8 B15
Reserved	Number of Heard Beaconers	Power Management mode of Reporting MP	Number of Reported MPs
Bits: 4	3	1	8

Figure s87—Connectivity Report Control field

Bits: 0-3	4-6	7	8-15
Reserved	Amount of Heard Beacons	Power Management mode of Reporting MP	Amount of Reported MPs

Figure s88—Connectivity Report Control field

The Amount Number of Heard Beacons Beaconers is an integer number indicating the number of beacons Beacon frames received from different transmitters. Value 0 indicates that no beacon has been received correctly during the previous Connectivity Reporting Interval. Value 7 defines that 7 or more transmitters have transmitted at least one correctly received beacon during the previous Connectivity Reporting Interval.

The Power Management Mode of Reporting MP indicates the power management mode of the reporting MP. If the power management mode of the reporting MP is set to 1, then the reporting MP is required to listen only during the MESH mesh DTIM beacons Beacon frames and the following ATIM periods. If the BB power management mode bit is set to 0, then the Reporting MP is in active mode.

The Amount Number of Reported MPs is an integer number indicating the number of Connectivity Reports or Beacons Beacon frames received from different transmitters. The Amount Number of Reported MPs indicates the total number of MAC addresses, which are reported in Connectivity Report.

The SSID specifies the SSID of the BSS. This information is used to differentiate connectivity reports from multiple BSSs.

The MAC address of the Beaconing MP indicates the address, where at least one beacon has been received during the previous Connectivity Reporting Interval. The number of the MAC Addresses is controlled by the Amount Number Of Heard Beacons Beaconers field. The last received beacon in the Connectivity Reporting Interval is the first reported MAC address of the beacon transmitter. If beacons Beacon frames are received from more than 7 differnet seven different MPs during the Connectivity Reporting Interval, the addresses of the first beacon transmitters in the during the connectivity reporting interval are listed as MPs that send a Connectivity Report.

The MAC addresses of the terminals MPs are set for the MPs in the same BSS, for which the Connectivity Report is received within `dot11BBConnectivityReportTimeout`. If the Connectivity Report from the MP is not received within the `dot11BBConnectivityReportTimeout` the MAC address of the MP is not present in the Connectivity Report List Element.

The beacon and connectivity reporting power management mode field indicates the current power management mode of each Beaconing and connectivity Reporting MP list member. Each bit of this field indicates the power save state of the corresponding beaconing and connectivity reporting member. If a bit is set to 0, then the corresponding beaconing or connectivity reporting list member is in "active mode" and if a bit is set to 1, the corresponding beaconing or connectivity reporting list member is in "Power Save mode". For example, if the beaconing list contains one MAC address and the connectivity list elements contains 7 MAC addresses and the neighbor power management mode field is '00110001', then the beaconing MP is in full power and the MPs that have transmitted a connectivity report with MAC addresses in positions 2, 3, and 7 in the connectivity reporting list are in the power save mode. The beacon and connectivity reporting power management mode field length is always an integer number of octets.

7.3.2.74 PANN information element

EDITORIAL NOTE—PANN information element relocated from 11A.3

The Portal Announcement (PANN) element is used for announcing in the mesh the presence of an MP configured as a Portal MP (which has a live connection to an external network). MPs may use this information to increase the efficiency of communication with stations outside the mesh.

The format of the PANN element is shown in Figure s89.

Octets: 1	1	1	1	1	6	4	4
Element ID	Length	Flags	Hopcount	Time to Live	Originator Address	Sequence Number	Metric

Figure s89—PANN Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 17.

The Flags field is reserved for future use.

The Hop Count field indicates the number of hops from the originator to the MP transmitting the request.

The Time to Live field indicates the maximum number of hops allowed for this element.

The Originator Address is set to the MAC address of the MP that is collocated with the portal.

The Sequence Number field is set to a sequence number specific for the originator.

The Metric field indicates a cumulative metric from the originator to the MP transmitting the announcement.

Detailed usage of the PANN element is described in 11A.4.

7.3.2.75 RANN information element

EDITORIAL NOTE—RANN information element relocated from 11A.5.4.2

The RANN information element is used for announcing the presence of an MP configured as Root MP. RANN elements are sent out periodically by the Root MP.

The format of the RANN element is shown in Figure s90.

Octets: 1	1	1	1	1	6	4	4	4
Element ID	Length	Flags	Hopcount	Time to Live	Originator Address	Destination Sequence Number	Lifetime	Metric

Figure s90—RANN Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 21.

The Flags field is set as follows. Bit 0: Portal Role (0 = non-portal, 1 = portal). Bit 1 – 7: Reserved

The Hop Count field indicates the number of hops from the originator (root MP) to the MP transmitting the request.

The Time to Live field indicates the remaining number of times the RANN may be forwarded.

The Originator Address is set to the Root MP MAC address.

The Destination Sequence Number is set to a sequence number specific to the originator (root MP).

The Lifetime field is set to the time for which MPs receiving the RANN consider the forwarding information to be valid.

The Metric is set to the cumulative metric from the originator to the MP transmitting the announcement.

Detailed usage of the RANN element is described in 11A.6.4.

7.3.2.76 RREQ information element

EDITORIAL NOTE—RREQ information element relocated from 11A.5.5.2

The RREQ element is used for discovering a route to one or more destinations, building a proactive (reverse) routing tree to the root MP, and confirming a route to a destination (optional).

The format of the RANN element is shown in Figure s91.

The Element ID is set to the value given in Table 26 for this information element. The length is set to 36 to 255 octets.

Octets: 1	1	1	1	1	4	6	4	4
Element ID	Length	Flags	Hopcount	Time to Live	RREQ ID	Originator Address	Originator Sequence Number	Lifetime

4	1	6	4	...	1	6	4	
Metric	Per Destination Flags #1	Destination Address #1	Destination Seq. Num. #1	...	Per Destination Flags #N	Destination Address #N	Destination Seq. Num. #N	

Figure s91—RREQ Element

The Flags field is set as follows. Bit 0: Portal Role (0 = non-portal, 1 = portal), Bit 1: (0 = group addressed, 1 = individually addressed) (see 11A.6.2), Bit 2: Proactive RREP (0 = off, 1 = on), Bit 3 – 7: Reserved.

The Hop Count field is set to the number of hops from the originator to the MP transmitting the request.

The Time to Live field is set to the maximum number of hops allowed for this element.

The RREQ ID field is set to some unique ID for this RREQ.

The Originator Address is set to the originator MAC address.

The Originator Sequence Number is set to a sequence number specific to the originator.

The Lifetime field is set to the time for which MPs receiving the RREQ consider the forwarding information to be valid.

The Metric is set to the cumulative metric from the originator to the MP transmitting the RREQ.

The Destination Count N gives the number of Destinations (N) contained in this RREQ.

The format of the Per Destination Flags field is shown in Figure s92.

B0	B1	B2 B7
DO	RF	Reserved
Bits: 1	1	6

Figure s92—RREQ Per-Destination Flags Field Format

Per Destination Flags are set as follows.

Bit 0: DO (Destination Only): If DO=0, an intermediate MP with active forwarding information to the corresponding destination responds to the RREQ with a unicast RREP; if DO=1, only the destination can respond with a unicast RREP. The default value is 1.

Bit 1: RF (Reply-and-Forward): The RF flag controls the forwarding of RREQ at intermediate MPs. When DO=0 and the intermediate MP has active forwarding information to the corresponding destination, the RREQ is not forwarded if RF=0 and forwarded if RF=1. The default value is 1. When DO=1, the RF flag has no effect.

Bit 2-7: Reserved

The Destination Address is the MAC address of the destination MP.

The Destination Sequence Number is the latest sequence number received in the past by the originator for any route towards the destination.

Detailed usage of the RREQ element is described in 11A.6.5.

7.3.2.77 RREP information element

EDITORIAL NOTE—RREP information element relocated from 11A.5.6.2

The RREP element is used to establish a forward route to a destination and to confirm that a destination is reachable.

The format of the RREP element is shown in Figure s93.

Octets: 1	1	1	1	1	6	4	4	4
ID	Length	Mode Flags	Hopcount	Time to Live	Destina-tion Address	Destina-tion Seq.Num.	Lifetime	Metric

6	4	1	6	4		6	4
Source Address #1	Source Seq. Num.	Depen-dent MP Count N	Depen-dent MP MAC Address #1	Depen-dent MP DSN #1	...	Depen-dent MP MAC Address #N	Depen-dent MP DSN #N

Figure s93—Route Reply Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 32 to 255 octets.

The Flags field is reserved.

The Hop Count field is set to the number of hops from the route destination to the local MP.

The Time to Live field is set to the maximum number of hops allowed for this element.

The Destination Address is the MAC address [of the destination for which a route is supplied].

The Destination Sequence Number field is set to the DSN of the originator of the RREP.

The Lifetime field, if applicable, reflects the Lifetime of the RREQ this RREP responds to.

The Metric field indicates the cumulative metric from the route destination to the local MP.

The Dependent MP Count N field indicates the number of dependent MPs (N).

The Dependent MP MAC Address # indicates the MAC address of dependent MP.

The Dependent MP DSN # indicates the Destination Sequence Number associated with MAC address of dependent MP.

The detailed usage of the RREP element is described in 11A.6.6.

7.3.2.78 RERR Information Element

EDITORIAL NOTE—RERR information element relocated from 11A.5.7.2

The Route Error (RERR) element is used for announcing a broken link to all traffic sources that have an active path over this broken link.

The format of the RERR element is shown in Figure s94.

Octets: 1	1	1	1	6	4
ID	Length	Mode Flags	Num of Des-tinations	Destination Address	Destination MP Seq. Num

Figure s94—Route Error Element

The Element ID is set to the value given in Table 26 for this information element. The length is set to 13 octets.

The Mode Flags field is reserved.

The Number of Destinations indicates the number of announced destinations in RERR (destination address and destination MP sequence number).

The Destination Address indicates the detected unreachable destination MAC address.

The Destination Sequence Number indicates the sequence number of detected unreachable destination MP.

The detailed usage of the RERR element is described in 11A.6.7.

7.3.2.79 RA-OLSR Information Elements

RA-OLSR IEs are carried in RA-OLSR frames (defined in clause 7.4.9.14) and have the following common format:format shown in Figure s95.

Octets: 1	1	1	6	1	1	2	Variable
ID	Length	Vtime	Originator Address	Time To Live (TTL)	Hop Count	Message Sequence Number (MSN)	Information Element-specific fields

Figure s95—Format of RA-OLSR information elements

Octets: 1	1	1	6	1	1	2	Variable
ID	Length	Vtime	Originator Address	Time To Live (TTL)	Hop Count	Message Sequence Number (MSN)	IE-specific fields

Figure s96—Format of RA-OLSR information elements

Table s16—Fields common to all RA-OLSR information elements

Field	Value/Description
ID	TBD
Length	Length of the IE
Vtime	Validity time during which a node must consider the information contained in the IE as valid after its reception, unless a more recent update to the information is received. The validity time is represented by its mantissa (four highest bits of Vtime field) and by its exponent (four lowest bits of Vtime field). In other words: Vtime = C * (1+a/16) * 2^b [in seconds] where 'a' is the integer represented by the four highest bits of Vtime field and 'b' the integer represented by the four lowest bits of Vtime field. The recommended value of the scaling factor 'C' is specified in clause 11A.6.14.
Originator Address	The MAC address of the node that has originally generated this IE. This may be different from the sender-address in case that multiple IEs are encapsulated in one frame.
Time To Live (TTL)	The maximum number of retransmissions of an IE. Before an IE is retransmitted, the TTL must be decremented by 1. An IE with TTL equal to 0 or 1 must not be retransmitted under any circumstances.
Hop Count	The number of hops an IE has attained. Before an IE is retransmitted, the Hop Count must be incremented by 1. The Hop Count is initialized to 0 by the originator of the IE.
Message Sequence Number (MSN)	A sequence number assigned by the originator node, which ensures that each message (i.e., IE) can be uniquely identified in the network. The MSN is increased by 1 for each IE originating from the node. "Wrap-around" is handled as described in clause 11A.6.15.
IE-Specific Fields	Fields specific to each IE

The Element ID is set to the value given in Table 26 for the information element.

The Vtime field indicates the validity time (in seconds) during which an MP considers the information contained in the information element as valid after its reception, unless a more recent update to the information is received. The validity time is represented by its mantissa (four highest bits of V_{time} field) and by its exponent (four lowest bits of V_{time} field). In other words:

$$V_{time} = C \times \left(1 + \frac{a}{16}\right) \times 2$$

where 'a' is the integer represented by the four highest bits of V_{time} field and 'b' the integer represented by the four lowest bits of V_{time} field. The recommended value of the scaling factor 'C' is specified in 11A.7.14.

The Originator Address is the MAC address of the MP that has originally generated this information element. This may be different from the sender-address in case that multiple IEs are encapsulated in one frame.

The Time To Live (TTL) field indicates the maximum number of times an information element may be forwarded. Before an information element is forwarded, the TTL is decremented by 1. An information element with TTL equal to 0 or 1 is not forwarded under any circumstances.

The Hop Count field indicates the number of hops an information element has attained. Before an information element is forwarded, the Hop Count is incremented by 1. The Hop Count is initialized to 0 by the originator of the information element.

The Message Sequence Number (MSN) field indicates a sequence number assigned by the originator MP, which ensures that each message (i.e., IE) can be uniquely identified in the network. The MSN is increased by 1 for each information element originating from the MP. "Wrap-around" is handled as described in 11A.7.15.

The Information Element-Specific Fields are fields specific to each IE Information Element and are described in the following subclauses.

7.3.2.79.1 RA-OLSR HELLO Element

Octets: 1	1	1	2	6	4	…	6	4
Htime	Willingness	Link Code	Link Message Size	Neighbor Interface Address	Link Metric	…	Neighbor Interface Address	Link Metric

…	1	2	6	4	…	6	4
…	Link Code	Link Message Size	Neighbor Interface Address	Link Metric	…	Neighbor Interface Address	Link Metric

Figure s97—Fields specific to RA-OLSR HELLO element

Table s17—Fields specific to RA-OLSR HELLO element

Field	Value/Description

Htime	HELLO emission interval used by the node on this particular interface. The HELLO emission interval is represented by its mantissa (four highest bits of Vtime field) and by its exponent (four lowest bits of Vtime field). In other words: Vtime = C * (1+a/16) * 2^b [in seconds] where 'a' is the integer represented by the four highest bits of Vtime field and 'b' the integer represented by the four lowest bits of Vtime field. The recommended value of the scaling factor 'C' is specified in clause 11A.6.14.
Willingness	The willingness of an MP to carry and forward traffic for other MPs (i.e., to be selected as MPR). An MP with willingness WILL_NEVER must never be selected as MPR by any MP. An MP with willingness WILL_ALWAYS must always be selected as MPR. By default, an MP should advertise a willingness of WILL_DEFAULT (see clause 11A.6.14.6 for willingness constants).
Link Code	Link state (see Clause 11A.1.5). One additional link state: MPR_NEIGH – indicating that the neighbors have at least one symmetrical link AND have been selected as MPR by the sender.
Link Message Size	The size of each link message counted in octets and measured from the beginning of the preceding "Link Code" field until the next "Link Code" field (or to the end of the IE if there are no more link types).
Neighbor Interface Address	The MAC address of an interface of a neighbor MP.
Link Metric	An example is the Airtime cost in clause 11A.4.

The fields specific to the RA-OLSR HELLO Element are shown in Figure s98.

Octets: 1	1	1	2	6	4	…	6	4
Htime	Will-ingness	Link Code	Link Mes-sage Size	Neigh-bor Inter-face Addres	Link Metric	…	Neigh-bor Inter-face Addres	Link Metric

…	1	2	6	4	…	6	4
…	Link Code	Link Mes-sage Size	Neigh-bor Inter-face Addres	Link Metric	…	Neigh-bor Inter-face Addres	Link Metric

Figure s98—Fields specific to RA-OLSR HELLO element

The Htime field indicates the HELLO emission interval used by the MP on this particular interface. The HELLO emission interval is represented by its mantissa (four highest bits of Vtime field) and by its

exponent (four lowest bits of Vtime field). In other words:Vtime = C * (1+a/16) * 2^b [in seconds], where 'a' is the integer represented by the four highest bits of Vtime field and 'b' the integer represented by the four lowest bits of Vtime field. The recommended value of the scaling factor 'C' is specified in 11A.7.14.

The Willingness field indicates the willingness of an MP to carry and forward traffic for other MPs (i.e., to be selected as MPR). An MP with willingness WILL_NEVER is never selected as MPR by any MP. An MP with willingness WILL_ALWAYS is always selected as MPR. By default, an MP should advertise a willingness of WILL_DEFAULT (see 11A.7.14.6 for willingness constants).

The Link Code field indicates the link state (see 11A.1.5). One additional link state: MPR_NEIGH – indicating that the neighbors have at least one symmetrical link AND have been selected as MPR by the sender.

The Link Message Size field indicates the size of each link message counted in octets and measured from the beginning of the preceding "Link Code" field until the next "Link Code" field (or to the end of the information element if there are no more link types).

The Neighbor Interface Address indicates the MAC address of an interface of a neighbor MP.

The Link Metric indicates the metric of the link. An example is the Airtime cost in 11A.5.

7.3.2.79.2 RA-OLSR Topology Control (TC) Element

Octets: 2	6	4	…	6	4
Advertised Neighbor Sequence Number (ANSN)	Advertised Neighbor Main Addresses	Link Metric	…	Advertised Neighbor Main Addresses	Link Metric

Figure s99—Fields specific to RA-OLSR TC element

Table s18—Fields specific to RA-OLSR TC element

Field	Value/Description
ANSN	A sequence number associated with the advertised neighbor set. Every time an MP detects a change in its advertised neighbor set, it increments this sequence number ("Wrap-around" is handled as described in clause 11A.6.15). This number is sent in this ANSN field of the TC IE to keep track of the most recent information. When an MP receives a TC IE, it can decide on the basis of this advertised ANSN, whether or not the received information about the advertised neighbors of the originator MP is more recent than what it already has.

Advertised Neighbor Main Address	The main address of a neighbor MP. All main addresses of the advertised neighbors of the originator MP are put in the TC IE. If the resulting IE cannot fit into one frame (due to maximum allowed frame size as imposed by the network), more TC IEs will be generated for any advertised neighbor addresses that have not been transmitted and carried in separate frames until the entire advertised neighbor set has been sent. Extra main addresses of neighbor MPs may be included, if redundancy is desired. Advertisement Neighbor Main Address pairs with its link metric. If an advertised neighbor is reachable through more than one link, the link with the best quality (smallest cost value) is selected and advertised.
Link Metric	An example is the Airtime cost in clause 11A.4.

The fields specific to the RA-OLSR Topology Control Element are shown in Figure s100.

Octets: 2	6	4	…	6	4
Advertised Neighbor Sequence Number (ANSN)	Advertised Neighbor Main Addresses	Link Metric	…	Advertised Neighbor Main Addresses	Link Metric

Figure s100—Fields specific to RA-OLSR TC element

The ANSN field indicates a sequence number associated with the advertised neighbor set. Every time an MP detects a change in its advertised neighbor set, it increments this sequence number ("Wrap-around" is handled as described in 11A.7.15). This number is sent in this ANSN field of the TC information element to keep track of the most recent information. When an MP receives a TC information element, it can decide on the basis of this advertised ANSN, whether or not the received information about the advertised neighbors of the originator MP is more recent than what it already has.

The Advertised Neighbor Main Address field indicates the main address of a neighbor MP. All main addresses of the advertised neighbors of the originator MP are put in the TC information element. If the resulting information element cannot fit into one frame (due to maximum allowed frame size as imposed by the network), more TC IEs are generated for any advertised neighbor addresses that have not been transmitted and carried in separate frames until the entire advertised neighbor set has been sent. Extra main addresses of neighbor MPs may be included, if redundancy is desired.

Advertisement Neighbor Main Address pairs with its link metric. If an advertised neighbor is reachable through more than one link, the link with the best quality (smallest cost value) is selected and advertised.

The Link Metric field indicates the metric of the link. An example is the Airtime cost in 11A.5.

7.3.2.79.3 RA-OLSR Multiple Interface Declaration (MID) Element

Octets: 6	6	…
RA-OLSR Interface Address	RA-OLSR Interface Address	…

Figure s101—Fields specific to RA-OLSR MID element

Table s19—Fields specific to RA-OLSR MID element

Field	Value/Description
RA-OLSR Interface Address	The address of an RA-OLSR interface of the MP, excluding the MPs main address (which is already indicated in the Originator Address field). All interface addresses other than the main address of the originator MP are put in the MID IE. If the resulting IE cannot fit into one frame (due to maximum allowed frame size as imposed by the network), more MID IEs will be generated for any interface addresses that have not been transmitted and carried in separate frames until the entire interface addresses have been sent.

The fields specific to the RA-OLSR Multiple Interface Declaration element are shown in Figure s102.

Octets: 6	6	…	6
RA-OLSR Interface Address 1	RA-OLSR Interface Address 2	…	RA-OLSR Interface Address n

Figure s102—Fields specific to RA-OLSR MID element

The RA-OLSR Interface Address indicates the address of an RA-OLSR interface of the MP, excluding the MPs main address (which is already indicated in the Originator Address field). All interface addresses other than the main address of the originator MP are put in the MID information element. If the resulting information element cannot fit into one frame (due to maximum allowed frame size as imposed by the network), more MID IEs are generated for any interface addresses that have not been transmitted and carried in separate frames until the entire interface addresses have been sent.

7.3.2.79.4 RA-OLSR Local Association Base Advertisement (LABA) Element

Octets: 1	1	6	1	…	6	1	…
Block Index (1)	Block Message Size	STA Address	STA Sequence Number	…	STA Address	STA Sequence Number	…

1	1	6	1	…	6	1
Block Index (N)	Block Message Size	STA Address	STA Sequence Number	…	STA Address	STA Sequence Number

Figure s103—Fields specific to RA-OLSR LABA element

Table s20—Fields specific to RA-OLSR LABA element

Field	Value/Description
Block Index	The index of a block that stores a list of STA addresses with their sequence numbers

Block Message Size	The size of a block counted in octets and measured from the beginning of the preceding "Block Index" field until the next "Block Index" field (or to the end of the IE if there are no more blocks).
STA Address	The MAC address of the associated STA. Note that a station address does not include the "Group MAC address bit" in the 48-bit MAC address as described for Local Association Base (LAB) and Global Association Base (GAB) in clause 11A.6.4.
STA Sequence Number	The sequence number in the association management frame sent by the STA

The fields specific to the RA-OLSR Local Association Base Advertisement Element are shown in Figure s104.

Octets: 1	1	6	1	…	6	1	…
Block 1 Index	Block 1 Message Size	Block 1 STA Address 1	Block 1 STA Sequence Number 1	…	Block 1 STA Address x	Block 1 STA Sequence Number x	…

1	1	6	1	…	6	1
Block n Index	Block n Message Size	Block n STA Address 1	Block n STA Sequence Number 1	…	Block n STA Address y	Block n STA Sequence Number y

Figure s104—Fields specific to RA-OLSR LABA element

The Block Index field indicates the index of a block that stores a list of STA addresses with their sequence numbers.

The Block Message Size field indicates the size of a block counted in octets and measured from the beginning of the preceding "Block Index" field until the next "Block Index" field (or to the end of the information element if there are no more blocks).

The STA Address field indicates the MAC address of the associated STA. A station address does not include the "Group MAC address bit" in the 48-bit MAC address as described for Local Association Base (LAB) and Global Association Base (GAB) in 11A.7.4.

The STA Sequence Number field indicates the sequence number in the association management frame sent by the STA.

7.3.2.79.5 RA-OLSR Local Association Base Checksum Advertisement (LABCA) Element

Octets: 1	16*	…	1	16
Block Index (1)	Block Checksum	…	Block Index (N)	Block Checksum

* This is based on the assumption that MD5 is used in checksum calculation.

Figure s105—Fields specific to RA-OLSR LABCA element

Table s21—Fields specific to RA-OLSR LABCA element

Field	Value/Description
Block Index	The index of a block that stores a list of STA addresses with their sequence numbers
Block Checksum	The checksum of a block (see clause 11A.6.13.6.5 for details of the checksum calculation).

The fields specific to the RA-OLSR Local Association Base Checksum Advertisement Element are shown in Figure s106.

Octets: 1	16	…	1	16
Block Index (1)	Block Check-sum	…	Block Index (N)	Block Check-sum

Figure s106—Fields specific to RA-OLSR LABCA element

The Block Index field indicates the index of a block that stores a list of STA addresses with their sequence numbers

The Block Checksum field indicates the checksum of a block (see 11A.7.13.6.5 for details of the checksum calculation). This field length is assumes MD5 is uded in checksum calculation.

7.3.2.79.6 RA-OLSR Association Base Block Request (ABBR) Element

Octets: 1	…	1
Block Index	…	Block Index

Figure s107—Fields specific to RA-OLSR ABBR element

Table s22—Fields specific to RA-OLSR ABBR element

Field	Value/Description
Block Index	The index of a block that has been detected to be inconsistent and is requested for readvertisement by the MP generating this IE.

The fields specific to the RA-OLSR Association Base Block Request Element are shown in Figure s108.

Octets: 1	…	1
Block Index	…	Block Index

Figure s108—Fields specific to RA-OLSR ABBR element

The Block Index field indicates the index of a block that has been detected to be inconsistent and is requested for readvertisement by the MP generating this information element.

7.3.2.80 MKD domain Domain information element [MKDDIE]

The MKD domain information element contains the MKD domain Identifier. A mesh authenticator uses the MKD domain information element to advertise its status as a an MA, and to advertise that it is included in the group of MAs that constitute a an MKD domain. The format for this information element is given in Figure s58Figure s109.

Octets: 1	1	6	1
Element ID	Length	MKD domain ID	Mesh Security Configuration

Figure s109—MKD domain information element

	Element ID	Length	MKD domain ID	Mesh Security Configuration
Octets:	1	1	6	1

Figure s110—MKD domain information element

The Element ID of this element shall be TBD. The Length field shall be set to 7.

The Element ID is set to the value given in Table 26 for this information element. The Length field is set to 7.

The MKD domain Identifier is a 6-octet value, following the ordering conventions from 7.1.1.

The Mesh Security Configuration field is one octet and is defined in Figure s59Figure s111.

B0	B1	B2	B3 B7
Mesh Authenticator	Connected to MKD	Default Role Negotiation	Reserved
Bits: 1	1	1	5

Figure s111—Mesh Security Configuration field

	Mesh Authenticator	Connected to MKD	Default Role Negotiation	Reserved
Bit:	B0	B1	B2	B3-B7

Figure s112—Mesh Security Configuration field

The Mesh Authenticator bit is set to one to indicate that a an MP is configured as a mesh authenticator in the MKD domain identified in this IEinformation element, and that the MP may act in the IEEE 802.1X Authenticator role during an EMSA MSA handshake.

The Connected to MKD bit is set to one to indicate that the MP has a valid route to the MKD and a current security association with the MKD. The Connected to MKD bit shall is not be set to one if the Mesh Authenticator bit is set to zero.

The interpretation of the Mesh Authenticator and Connected to MKD bits is described in Table s20Table s23.

Table s23—Meaning of Mesh Security Configuration bits

Mesh Authenticator	Connected to MKD	Meaning
0	0	The device is not configured to act as a mesh authenticator.
0	1	Invalid
1	0	The device is configured to act as a mesh authenticator but does not have a connection to the MKD. In this case the device may successfully act as an IEEE 802.1X authenticator, for example, if it possesses a cached key for the supplicant MP.
1	1	The device is configured to act as a mesh authenticator and has a connection to the MKD.

Table s24—Meaning of Mesh Security Configuration bits

Mesh Authenticator	Connected to MKD	Meaning
0	0	The device is not configured to act as a mesh authenticator.
0	1	Invalid
1	0	The device is configured to act as a mesh authenticator but does not have a connection to the MKD. The device may successfully act as an IEEE 802.1X authenticator, for example, if it possesses a cached key for the supplicant MP.
1	1	The device is configured to act as a mesh authenticator and has a connection to the MKD.

The Default Role Negotiation bit is set to one by an MP if it is using the default method to select IEEE 802.1X Authenticator and Supplicant roles, and shall is set this bit to 0 otherwise. When set to 0, the Authenticator/Supplicant selection method to use is specified by a mechanism outside the scope of this standard.

7.3.2.81 EMSA MSA Handshake element [EMSAIEMSAIE]

The EMSA MSA handshake information element includes information needed to perform the authentication sequence during an EMSA MSA handshake. This information element is depicted shownshown in Figure s60Figure s113.

	Element ID	Length	ANonce	SNonce	MA-ID	Optional Parameters	MIC Control	MIC
Octets:	1	1	32	32	6	variable	2	16

Octets: 1	1	32	32	6	variable	2	16
Element ID	Length	ANonce	SNonce	MA-ID	Optional Parameters	MIC Con-trol	MIC

Figure s113—MSA Handshake information element

Figure s114—EMSA Handshake information element

The Element ID of is set to the value given in Table 26 for this element shall be TBDinformation element. The Length field for this information element indicates the number of octets in the information field (fields following the Element ID and Length fields).

The ANonce field contains a nonce pseudo-random value chosen by the MA. It is encoded following the conventions from 7.1.1.

The SNonce field contains a nonce pseudo-random value chosen by the supplicant MP. It is encoded following the conventions from 7.1.1.

The MA-ID field contains the MA's identity, which is used by the supplicant MP for deriving the PMK-MA. It is encoded following the conventions from 7.1.1.

The format of the optional parameters is shown in Figure s61Figure s115.

Octets: 1	1	variable
Sub-element ID	Length	Data

Figure s115—Optional parameters field

	Sub-element ID	Length	Data
Octets:	1	1	variable

Figure s116—Optional parameters field

The Sub-element ID is one of the values from Table s21Table s25.

Table s25—Sub-element IDs

Value	Contents of data field	Length
0	Reserved	
1	MKD-ID	6
2	GTK	variable
3	EAP Transport List	variable
4-255	Reserved	

Table s26—Sub-element IDs

Value	Contents of data field	Length
0	Reserved	
1	MKD-ID	6
2	GTK	variable
3	EAP Transport List	variable
4-255	Reserved	

MKD-ID indicates the MKD that the supplicant MP may contact to initiate the mesh key holder security handshake.

GTK contains a KDE containing the Group Key; it is encrypted. The KDE is defined in Figures 143 and 144 of 8.5.2.

EAP Transport List contains a series of transport type selectors that indicate the EAP transport mechanism. A transport type selector has the format shown in Figure s62Figure s117.

Octets: 3	1
OUI	Transport Type

Figure s117—Transport type selector format

	OUI	Transport Type
Octets:	3	1

Figure s118—Transport selector format

The order of the organizationally unique identifier (OUI) field shall follow follows the ordering convention for MAC addresses from 7.1.1. The transport selectors types defined by this amendment standard are provided in Table s22Table s27.

Table s27—Transport types

OUI	Transport Type	Meaning
00-0F-AC	0	EAP Transport mechanism as defined in 11A.2.5.
00-0F-AC	1-255	Reserved
Vendor OUI	Any	Vendor specific
Other	Any	Reserved

Table s28—Table s22: Transport selectors

OUI	Transport Type	Meaning
00-0F-AC	0	EAP Transport mechanism as defined in8.8.3.3.3.
00-0F-AC	1-255	Reserved
Vendor OUI	Any	Vendor specific
Other	Any	Reserved

The transport selector type 00-0F-AC:0 shall be is the default transport type selector value.

The MIC Control field is two octets and is defined in Figure s63Figure s119.

B0 B3	B4 B7	B8 B15
MIC Algo-rithm	Reserved	Information element count
Bits: 4	4	8

Figure s119—MIC Control field

	MIC Algorithm	Reserved	Information element count
Bit:	B0-B3	B4-B7	B8-B15

Figure s120—MIC Control field

The MIC algorithm is one of the values from Table s23Table s29.

Table s29—MIC Algorithms

Value	Algorithm description
0	Reserved
1	HMAC-SHA1-128
2-15	Reserved

Table s30—MIC Algorithms

Value	Algorithm description
0	Reserved
1	HMAC-SHA1-128
2-15	Reserved

The Information Element count of the MIC Control field contains the number of information elements that are included in the MIC calculation. A value of zero indicates no MIC is present.

The MIC field contains a Message Integrity Check, calculated using the algorithm selected by the MIC algorithm field of the MIC Control field.

7.3.2.82 Mesh Key Holder Security element [MKHSIE]

The Mesh key holder security information element includes information needed to perform the authentication sequence during a mesh key holder security handshake. This information element is depicted shown in Figure s64Figure s121.

Octets: 1	1	32	32	6	6	4	2	16
Element ID	Length	MA-Nonce	MKD-Nonce	MA-ID	MKD-ID	Transport Type Selector	MIC Control	MIC

Figure s121—Mesh key holder security information element

	Element ID	Length	MA-Nonce	MKD-Nonce	MA-ID	MKD-ID	Transport Type Selector	MIC Control	MIC
Octets:	1	1	32	32	6	6	4	2	16

Figure s122—Mesh key holder security information element

The Element ID field shall be TBD. The Length field shall contain the value 98.

The Element ID is set to the value given in Table 26 for this information element. The Length field is set to 98.

The MA-Nonce field contains a nonce pseudo-random value chosen by the MA. It is encoded following the conventions from 7.1.1.

The MKD-Nonce field contains a nonce pseudo-random value chosen by the MKD. It is encoded following the conventions from 7.1.1.

The MA-ID field contains the MAC address of the MA. It is encoded following the conventions from 7.1.1.

The MKD-ID field contains the MAC address of the MKD. It is encoded following the conventions from 7.1.1.

The Transport Type Selector field contains a single transport type selector that indicates the supported transport types. The transport type selector format is given in Figure s62Figure s117. The transport selectors types defined by this amendment standard are provided in Table s22Table s27.

The MIC Control field is two octets and is defined in Figure s63Figure s119.

The MIC algorithm is one of the values from Table s23Table s29.

The Information Element count of the MIC Control field contains the number of information elements that are included in the MIC calculation. A value of zero indicates no MIC is present.

The MIC field contains a Message Integrity Check, calculated using the algorithm selected by the MIC algorithm field of the MIC Control field.

7.3.2.83 Mesh Encrypted Key element [MEKIE]

The Mesh Encrypted Key information element is used to request, deliver, or confirm delivery of a PMK-MA. It is also used to request deletion of a PMK-MA. This information element is depicted shown in Figure s65Figure s123.

Octets: 1	1	8	6	16	32	2	variable	2	16
Element ID	Length	Replay Counter	SPA	PMK-MKD Name	ANonce	Encrypted Contents Length	Encrypted Contents	MIC Control	MIC

Figure s123—Mesh Encrypted Key information element

	Element ID	Length	Replay Counter	SPA	PMK-MKD Name	ANonce	Encrypted Contents Length	Encrypted Contents	MIC Control	MIC
Octets:	1	1	8	6	16	32	2	variable	2	16

Figure s124—Mesh Encrypted Key information element

The Element ID field shall be TBDis set to the value given in Table 26 for this information element. The Length field shall contain is set to the number of octets in the information field (fields following the Element ID and Length fields).

The Replay Counter field contains a sequence number, represented as an unsigned binary number, used to detect replayed frames.

The SPA field contains the MAC address of the supplicant MP that, during its Initial EMSA MSA Authentication, created the PMK-MA that is being requested, delivered, confirmed, or deleted.

The PMK-MKDName field contains the identifier of the PMK-MKD that was used to derive the PMK-MA that is being requested, delivered, confirmed, or deleted.

The ANonce field contains a pseudo-random nonce selected by the MKD and used in the derivation of the PMK-MKD identifier provided in the PMK-MKDName field.

The Encrypted Contents Length field indicates the number of octets contained in the Encrypted Contents field.

The Encrypted Contents field contains an PMK-MA and related key context information. All information in the field shall be is encrypted using an algorithm selected using the value of the MIC algorithm subfield.

The MIC Control field is two octets and is defined in Figure s63Figure s119.

The MIC algorithm is one of the values from Table s23Table s29.

The Information Element count of the MIC Control field contains the number of information elements that are included in the MIC calculation. A value of zero indicates no MIC is present.

The MIC field contains a Message Integrity Check, calculated using the algorithm selected by the MIC algorithm field of the MIC Control field.

7.3.2.84 EAP Authentication element [EAPAIE]

The EAP Authentication information element is depicted shown in Figure s66Figure s125.

Octets: 1	1	1	16	6	1	2	16
Element ID	Length	EAP Message Type	Message Token	SPA	Message Fragments	MIC Control	MIC

Figure s125—EAP Authentication information element

	Element ID	Length	EAP Message Type	Message Token	SPA	Message Fragments	MIC Control	MIC
Octets:	1	1	1	16	6	1	2	16

Figure s126—EAP Authentication information element

The Element ID field shall be TBD. The Length field shall contain the value 42.

The Element ID is set to the value given in Table 26 for this information element. The Length field is set to 42.

The EAP Message Type field identifies the type of EAP Encapsulation message, and is set to a value described in Table s24Table s31.

Table s31—EAP Message Type values

Value	Message Type
0	Reserved
1	Request
2	Response – Accept
3	Response – Reject
4-10	Reserved
11	Response
12-255	Reserved

Table s32—EAP Message Type values

Value	Message Type
0	Reserved
1	Request
2	Response – Accept
3	Response – Reject
4-10	Reserved
11	Response
12-255	Reserved

The Message Token field contains a random nonce transaction identifier in messages of type request. In messages of type response, response-accept, and response-reject, the Message Token field contains the value of the Message Token field in the request message to which the response message corresponds.

The SPA field contains the MAC address of the supplicant MP that is performing EAP authentication.

The Message Fragments field is one octet and contains an unsigned binary integer indicating the number of EAP Message elements that follow this EAPAIE, containing fragments of an EAP message. The Message Fragments field may contain the value zero if no EAP Message elements follow.

The MIC Control field is two octets and is defined in Figure s63Figure s119.

The MIC algorithm is one of the values from Table s23Table s29.

The Information Element count of the MIC Control field contains the number of information elements that are included in the MIC calculation. A value of zero indicates no MIC is present.

The MIC field contains a Message Integrity Check, calculated using the algorithm selected by the MIC algorithm field of the MIC Control field.

7.3.2.85 EAP Message element [EAPMIE]

The EAP Message element is defined in Figure s67Figure s127.

Octets: 1	1	1	variable
Element ID	Length	Fragment Control	EAP Message/Fragment

Figure s127—EAP Message information element

	Element ID	Length	Fragment Control	EAP Message Fragment
Octets:	1	1	1	variable

Figure s128—EAP Message information element

The Element ID field shall be TBDis set to the value given in Table 26 for this information element. The Length field shall contain is set to the number of octets in the information field (fields following the Element ID and Length fields).

The Fragment Control field is one octet and contains an unsigned binary integer indicating the number of each fragment of an EAP message. The fragment number is set to 0 in the first or only fragment of an EAP message and is incremented by one for each successive fragment of that EAP message.

The EAP Message Message/Fragment field contains an EAP packet, or fragment thereof, with format as defined in IETF RFC 3748.

7.4 Action frame format details

EDITORIAL NOTE—11ma ends with 7.4.5, 11k adds 7.4.6, 11r adds 7.4.7, 11n adds 7.4.8.

Insert the following new clause after 7.4.8:

7.4.9 Mesh management action frame details (3-addr action frames)

Action frame formats for mesh management are defined in this sectionsubclause. An Action field, in the octet field immediately after the Category field, differentiates the frame format. The Action field values associated with each frame format are defined in Table s25Table s33.

Table s33—Mesh Management Action field values (3-addr action frames)

Action field value	Description	Application
0	Local Link state announcement	Neighbor discovery
1	Route Request	HWMP routing
2	Route Reply	HWMP routing

Table s33—Mesh Management Action field values (3-addr action frames)

3	Route Error	HWMP routing
4	Route Reply Ack	HWMP routing
5	Congestion Control Request	Congestion Control
6	Congestion Control Response	Congestion Control
7	Neighborhood Congestion Announcement	Congestion Control
8	Mesh Deterministic Access (MDA)	MDA
9	Beacon Timing Request	Beaconing and Synchronization
10	Beacon Timing Response	Beaconing and Synchronization
11	Non-mesh Action Encapsulation	Action
12	Connectivity Report	Connectivity reporting
13	Radio Aware Optimized Link State Routing (RA-OLSR)	RA-OLSR
14-220	Reserved	
221	Vendor Specific Mesh Management	Vendor Specific
222-255	Reserved	

Table s34—Mesh Management Action field values (3-addr action frames)

Action field value	Description	Application
0	Local Link state announcement	Neighbor discovery
1	Peer Link Disconnect	Neighbor discovery
2	Route Request	HWMP routing
3	Route Reply	HWMP routing
4	Route Error	HWMP routing
5	Route Reply Ack	HWMP routing
6	Congestion Control Request	Congestion Control
7	Congestion Control Response	Congestion Control
8	Neighborhood Congestion Announcement	Congestion Control
9	Mesh Deterministic Access (MDA)	MDA
10	Beacon Timing Request	Beaconing and Synchronization
11	Beacon Timing Response	Beaconing and Synchronization
12	Non-mesh Action Encapsulation	Action
13	Connectivity Report	Connectivity reporting
13	Radio Aware Optimized Link State Routing (RA-OLSR)	RA-OLSR
15-254	Reserved	
255	Vender Specific Mesh Management	Vender Specific

7.4.9.1 Local Link State Announcement frame format

The Local Link State Announcement frame format uses the Action frame body format and is transmitted by an MP to a neighbor MP in WLAN mesh to advertise metric information. This frame is transmitted in a unicast an individually addressed manner. The frame format is shown in Figure s68Figure s129.

Octets: 1	1	Variable
Category	Action	Local Link State Announcement Element

Octets: 1	1	Variable
Category	Action	Local Link State Announcement Element

Figure s129—Local Link State Announcement frame format

Figure s130—Local Link State Announcement frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 0 (representing Local Link state announcement)

7.4.9.2 Peer Link Disconnect frame format

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

7.4.9.3 Route Request frame format

The Route Request frame format uses the Action frame body format and is transmitted by a source MP to discover the path to the destination MP. This frame is typically transmitted in a broadcast mannergroup addressed frames. The intermediate MP rebroadcasts forwards this frame. The frame format is shown in Figure s69Figure s131.

Octets: 1	1	Variable
Category	Action	Route Request Element

Figure s131—Route Request frame format

Octets: 1	1	Variable
Category	Action	Route Request Element

Figure s132—Route Request frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 2 (representing Route Request)

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Route Request elements field shall be is set as described in Clause 11A.6.5.

7.4.9.4 Route Reply frame format

The Route Reply frame format uses the Action frame body format and is transmitted by a destination MP to a source MP in WLAN the mesh to determine the path between the source and destination MP. This frame is typically transmitted in a unicast an individually addressed manner. The frame format is shown in Figure s70Figure s133.

Octets: 1	1	Variable
Category	Action	Route Reply Element

Figure s133—Route Reply frame format

Octets: 1	1	Variable
Category	Action	Route Reply Element

Figure s134—Route Replay frame format

The Category field shall be is set to 5 (representing mesh management)the value in Table 24 for category Mesh Management.

The Action field shall be set to 3 (representing Route Request)

The Action field is set to the value in Table s33 for this action frame type.

The Route Reply elements field shall be is set as described in Clause 11A.6.6.

7.4.9.5 Route Error frame format

The Route Request Error frame format uses the Action frame body format and is transmitted by the MP detected the link failure on a certain path to the precursor MP. This frame is typically transmitted in a unicast an individually addressed manner. The frame format is shown in Figure s71Figure s135.

Octets: 1	1	Variable
Category	Action	Route Error Element

Octets: 1	1	Variable
Category	Action	Route Error Element

Figure s135—Route Error frame format

Figure s136—Route Error frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 4 (representing Route Error)

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Route Error elements field shall be is set as described in Clause 11A.6.7.

7.4.9.6 Route Reply Ack frame format

The Route Reply Ack frame format uses the Action frame body format and is transmitted by a source MP to a destination MP in response to a Route Reply frame. This frame is typically transmitted in a unicast an individually addressed manner. The frame format is shown in Figure s72Figure s137.

Octets: 1	1	Variable
Category	Action	Route Reply Ack Element

Figure s137—Route Reply Ack frame format

Octets: 1	1	Variable
Category	Action	Route Reply Ack Element

Figure s138—Route Reply Ack frame format

The Category field shall be is set to 5 (representing mesh management)the value in Table 24 for category Mesh Management.

The Action field shall be set to 5 (representing Route Reply Ack)

The Action field is set to the value in Table s33 for this action frame type.

The Route Reply Ack elements field shall be is set as described in Clause 11A.6.7.

7.4.9.7 Congestion Control Request frame format

The Congestion Control Request frame format uses the Action frame body format and is sent by an MP to its upstream neighboring MP to indicate the target data rate it desires, or the peak data rate it wants its neighbor not to exceed so as to avoid or control congestion. The Congestion Control Request frame can also be used to inform its upstream neighbors about the current traffic load in case of congestion. This frame is transmitted in a unicast an individually addressed manner. The frame format is shown in Figure s73Figure s139.

Octets: 1	1	20
Category	Action	Target Transmission Rate Element

Figure s139—Congestion Control Request frame format

Octets: 1	1	20
Category	Action	Target Transmission Rate Element

Figure s140—Congestion Control Request frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 6 (representing Congestion Control Request)

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Target Transmission Rate element field shall be is set following the guidelines described in Clause 11A.8.4.

7.4.9.8 Congestion Control Response frame format

The Congestion Control Response frame format uses the Action frame body format and is sent by an MP as a response to the "Congestion Control Request" from its downstream MP. This frame is transmitted in a uni-cast an individually addressed manner. The frame format is shown in Figure s74Figure s141.

Octets: 1	1	18
Category	Action	Offered Traffic Load Element

Figure s141—Congestion Control Response frame format

Octets: 1	1	18
Category	Action	Offered Traffic Load Element

Figure s142—Congestion Control Response frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 7 (representing Congestion Control Response).

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Offered Traffic Load element field shall be is set following the guidelines described in Clause 11A.8.

7.4.9.9 Neighborhood Congestion Announcement frame format

The Neighborhood Congestion Announcement frame format uses the Action frame body format and is sent by an MP that is congested due to the high channel load in the neighborhood. This frame is transmitted in a broadcast mannerusing group addressing. The frame format is shown in Figure s75Figure s143.

Octets: 1	1	5
Category	Action	Neighborhood Congestion Element

Figure s143—Neighbor Congestion Announcement frame format

Octets: 1	1	5
Category	Action	Neighborhood Congestion Element

Figure s144—Neighbor Congestion Announcement frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 8 (representing Neighborhood Congestion Announcement).

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Neighborhood Congestion Announcement element field shall be is set following the guidelines described in Clause 11A.8.3.

7.4.9.10 Mesh Deterministic Access frame format

The Mesh Deterministic Access frame format uses the Action frame body format and is transmitted by an MDA-active MP to one (in unicast individually addressed manner) or more neighbor MDA-active MPs (in broadcast mannergroup addressed frames) in WLAN mesh, depending on the MDA elements that it carries. This frame is always transmitted with immediate Ack policy. The frame format is shown in Figure s76Figure s145.

Octets: 1	1	5
Category	Action	One or more MDA Elements

Figure s145—Mesh Deterministic Access frame format

Octets: 1	1	5
Category	Action	One or more MDA Elements

Figure s146—Mesh Deterministic Access frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 9 (representing Mesh Deterministic Access).

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The MDA elements are described from Clause 7.3.2.68 to Clause 7.3.2.72.

7.4.9.11 Beacon Timing Request frame format

The Beacon Timing Request frame format uses the Action frame body format and is used to request a peer MAC its neighbors' beacon timing informationinformation from neighbors. The frame format is shown in Figure s77Figure s147.

Octets: 1	1
Category	Action

Figure s147—Beacon Timing Request frame format

Octets: 1	1
Category	Action

Figure s148—Beacon Timing Request frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 10 (Beacon Timing Request).

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

7.4.9.12 Beacon Timing Response frame format

The Beacon Timing Response frame format uses the Action frame body format and is used to respond to a Beacon Timing Request frame with neighbors' beacon timing information. The frame format is shown in Figure s78Figure s149.

Octets: 1	1	8	Variable
Category	Action	Most Recent TBTT	Beaconing Timing Element

Figure s149—Beacon Timing Response frame format

Octets: 1	1	8	Variable
Category	Action	Most Recent TBTT Time	Beaconing Timing Element

Figure s150—Beacon Timing Response frame format

The Category field shall be set to 5 (representing mesh management).

The Action Category field is set to 11 (Beacon Timing Response)the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Most Recent TBTT Time information is an eight-octet field that reflects the most recent TBTT time of the transmitting MP, so that the beacon timing IE information element reflects information as if it was transmitted in a beacon at that TBTT.

The Beacon Timing elements field shall be is set as described in Clause 7.3.2.67.

7.4.9.13 Non-mesh Action Encapsulation frame format

Numerous information elements for action frames are already defined in IEEE 802.11. Some of these frames (for example, DFS, TPC, measurement) are also applicable in multi-hop mesh network topologies. The "Non-Mesh Action Encapsulation" " frame format is designed to transfer these non-mesh information elements through multi-hop WLAN mesh network topologies and gives a framework to encapsulate them. The Non-Mesh Action Encapsulation frame uses the Action frame body format and is transmitted by a source MP to a destination MP. The frame format is shown in Figure s79Figure s151.

Octets: 1	1	6	6	Variable
Category	Action	Source MP Address	Destination MP Address	Encapsulated Action Frame Body

Figure s151—Non-Mesh Action Encapsulation frame format

Octets: 1	1	6	6	Variable
Category	Action	Source MP Address	Destination MP Address	Encapsulated Action Frame

Figure s152—Non-Mesh Action Encapsulation frame format

The Category field shall be set to 5 (representing mesh management).

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field shall be is set to 12 (representing non-mesh the value in Table s33 for this action encapsulation)frame type.

The Source MP Address shall be is set to the value of Source MP MAC address.

The Destination MP Address shall be is set to the value of Destination MP MAC address.

The encapsulated action frame body constructed by a Source MP shall contain contains the original action frame. The format of this field obeys the original frame format defined in the IEEE 802.11 standard. This original format generally consists of three types of field (Category, Action and Element(s)) or four types of

field (Category, Action, Dialog Token and Element(s)). Upon reception of a non-mesh action encapsulation frame, the destination MP shall perform performs the requested action as defined in the encapsulated frame. Similarly, if the action requires a response message, the destination MP shall generate generates such response and send it back with a non-mesh action encapsulation frame.

Example of Non-Mesh Action Encapsulation frame to carry Measurement Request action frame is shown in Table s26Table s35.

Table s35—Example of Non-Mesh Action Encapsulation frame (Measurement Request)

Field		Value/description
Category		5 (mesh management)
Action		10 (non-mesh action encapsulation)
Source MP Address		Source MP MAC address
Destination MP Address		Destination MP MAC address.
Encapsulated Action Frame Body	Category	0 (spectrum management)
	Action	0 (Measurement Request)
	Dialog Token	1
	Element(s)	Measurement Request Elements

Table s36—Example of Non-Mesh Action Encapsulation frame (Measurement Request)

Field		Value/description
Category		5 (mesh management)
Action		10 (non-mesh action encapsulation)
Source MP Address		Source MP MAC address
Destination MP Address		Destination MP MAC address.
Encapsulated Action Frame	Category	0 (spectrum management)
	Action	0 (Measurement Request)
	Dialog Token	1
	Element(s)	Measurement Request Elements

7.4.9.14 RA-OLSR frame format

The RA-OLSR frame format uses the Action frame body format and is transmitted by an MP to one (in unicast individually addressed manner) or more neighbor MPs (in broadcast mannergroup addressed frames) in a WLAN mesh where the RA-OLSR protocol is in use as its active path selection protocol. The frame format is shown in Figure s80Figure s153.

Octets: 1	1	Variable
Category	Action	One or more RA-OLSR elements

Figure s153—RA-OLSR frame format

Octets: 1	1	Variable
Category	Action	One or more RA-OLSR elements

Figure s154—RA-OLSR frame format

The Category field shall be set to 4 (representing mesh management).

The Action field shall be set to 13 (representing RA-OLSR).

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The RA-OLSR elements are described in clause 7.3.2.79.

7.4.9.15 Connectivity Report frame

The connectivity Report frame is a group addressed management frame. It is used to list the number of beacon broadcasters during the reporting interval and to list the neighbors that transmitted a connectivity report and the Power Management Mode of each neighbor. The element contains list of neighbor MAC addresses, where the connectivity report has been received and information about the neighbor power management mode. The format of the Connectivity Report is shown in Figure s155.

Octets: 1	1	Variable
Category	Action	Connectivity Report element

Figure s155—Connectivity Report frame format

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The Connectivity Report Element is defined in 7.3.2.73.

7.4.9.16 Vender Vendor Specific Mesh Management frame format

The Vender Vendor Specific Mesh Management frame format uses the Action frame body format and is used to carry information not defined in this standard within a single defined frame format, so that reserved Action IDs are not usurped for non-standard purposes and so that interoperability is more easily achieved in the presence of non-standard management frames. The frame format is shown in Figure s81Figure s156.

Octets: 1	1	3	Variable
Category	Action	OUI	Vendor specific content

Figure s156—Vendor Specific Mesh Management frame format

Octets: 1	1	3	Variable
Category	Action	OUI	Vendor specific content

Figure s157—Vender Specific Mesh Management frame format

The Category field shall be set to 5 (representing mesh management).

The Action field shall be set to 255 (representing vendor specific mesh management)

The Category field is set to the value in Table 24 for category Mesh Management.

The Action field is set to the value in Table s33 for this action frame type.

The OUI field shall be is a public OUI assigned by the IEEE.

The Vender Vendor specific contents shall be is set to the suitable value defined by each vendervendor's rule.

7.4.9.17 Connectivity Report frame

The connectivity Report frame is an broadcasted management frame. It is used to list the number of beacon broadcasters during the reporting interval and to list the neighbors that transmitted a connectivity report and their Power Management Mode. The element contains list of neighbors MAC addresses, where the connectivity report has been received and information about the neighbor power management mode. The format of the Connectivity Report is shown in Figure s82.

Octets: 1	1	Variable
Category	Action	Connectivity Report element

Figure s158—Connectivity Report frame format

The Category field shall be set to 5

The Action field shall be set to 13

The Connectivity Report Element is defined in 7.3.2.73.

Insert the following new subclause after 7.4:

7.4A Mesh Action (4-addr action frames)

This subclause describes the Mesh Action frame formats, including the Mesh Action Details field, allowed in each of the mesh action categories defined in Table s1 Table s2 in 7.3.1.18.

7.4A.1 EMSA MSA mesh action details

Seven Mesh Action frame formats are defined for EMSAMSA. An Action Value field, in the octet field immediately after the Category field, differentiates the five formats. The Action Value field values associated with each frame format are defined in Table s27Table s37.

Table s37—MSA Action field values

Action Field Value	Description
0	Mesh key holder security establishment
1	PMK-MA delivery push
2	PMK-MA confirm
3	PMK-MA request
4	PMK-MA delivery pull
5	PMK-MA delete
6	Mesh EAP encapsulation
7-255	Reserved

Table s38—EMSA Action field values

Action Field Value	Description
0	Mesh key holder security establishment
1	PMK-MA delivery push
2	PMK-MA confirm
3	PMK-MA request
4	PMK-MA delivery pull
5	PMK-MA delete
6	Mesh EAP encapsulation
7-255	Reserved

7.4A.1.1 Mesh key holder security establishment frame format

The Mesh key holder security establishment frame uses the Mesh Action frame body format and is transmitted by a mesh key holder to perform the mesh key holder security handshake. The format of the mesh key holder security establishment frame body is shown in Figure s83Figure s159.

Octets: 1	1	variable	9	100
Category	Action Value	Mesh ID IE	MKD-DIE	MKHSIE

Figure s159—Mesh key holder security establishment frame body format

	Category	Action Value	Mesh ID IE	MKDDIE	MKHSIE
Octets:	1	1	variable	9	100

Figure s160—Mesh key holder security establishment frame body format

The Category field shall be is set to 0 (representing EMSAMSA).

The Action Value field shall be is set to 0 (representing a Mesh key holder security establishment frame).

The Mesh ID IE shall be information element is set as described in 7.3.2.51.

The MKDDIE shall be is set as described in 7.3.2.80.

The MKHSIE shall be is set as described in 7.3.2.82.

7.4A.1.2 PMK-MA delivery push frame format

The PMK-MA delivery push frame uses the Mesh Action frame body format and is transmitted by a an MKD in the mesh key transport push protocol. The format of the PMK-MA delivery push frame body is shown in Figure s84Figure s161.

Octets: 1	1	9	variable
Category	Action Value	MKD-DIE	MEKIE

Figure s161—PMK-MA delivery push frame body format

	Category	Action Value	MKDDIE	MEKIE
Octets:	1	1	9	variable

Figure s162—PMK-MA delivery push frame body format

The Category field shall be is set to 0 (representing EMSAMSA).

The Action Value field shall be is set to 1 (representing a PMK-MA delivery push frame).

The MKDDIE shall be is set as described in 7.3.2.80.

The MEKIE shall be is set as described in 7.3.2.83.

7.4A.1.3 PMK-MA confirm frame format

The PMK-MA confirm frame uses the Mesh Action frame body format and is transmitted by a an MA in the mesh key transport push or the mesh key delete protocol. The format of the PMK-MA confirm frame body is shown in Figure s85Figure s163.

Octets: 1	1	9	variable
Category	Action Value	MKD-DIE	MEKIE

Figure s163—PMK-MA confirm frame body format

	Category	Action Value	MKDDIE	MEKIE
Octets:	1	1	9	variable

Figure s164—PMK-MA confirm frame body format

The Category field shall be is set to 0 (representing EMSAMSA).

The Action Value field shall be is set to 2 (representing a PMK-MA confirm frame).

The MKDDIE shall be is set as described in 7.3.2.80.

The MEKIE shall be is set as described in 7.3.2.83.

7.4A.1.4 PMK-MA request frame format

The PMK-MA request frame uses the Mesh Action frame body format and is transmitted by a an MA in the mesh key transport pull protocol. The format of the PMK-MA request frame body is shown in Figure s86Figure s165.

Octets: 1	1	9	variable
Category	Action Value	MKD-DIE	MEKIE

Figure s165—PMK-MA request frame body format

	Category	Action Value	MKDDIE	MEKIE
Octets:	1	1	9	variable

Figure s166—PMK-MA request frame body format

The Category field shall be is set to 0 (representing EMSAMSA).

The Action Value field shall be is set to 3 (representing a PMK-MA request frame).

The MKDDIE shall be is set as described in 7.3.2.80.

The MEKIE shall be is set as described in 7.3.2.83.

7.4A.1.5 PMK-MA delivery pull frame format

The PMK-MA delivery pull frame uses the Mesh Action frame body format and is transmitted by a an MKD in the mesh key transport pull protocol. The format of the PMK-MA delivery pull frame body is shown in Figure s87Figure s167.

Octets: 1	1	9	variable
Category	Action Value	MKD-DIE	MEKIE

Figure s167—PMK-MA delivery pull frame body format

	Category	Action Value	MKDDIE	MEKIE
Octets:	1	1	9	variable

Figure s168—PMK-MA delivery pull frame body format

The Category field shall be is set to 0 (representing EMSAMSA).

The Action Value field shall be is set to 4 (representing a PMK-MA delivery pull frame).

The MKDDIE shall be is set as described in 7.3.2.80.

The MEKIE shall be is set as described in 7.3.2.83.

7.4A.1.6 PMK-MA delete frame format

The PMK-MA delete frame uses the Mesh Action frame body format and is transmitted by a an MKD in the mesh key delete protocol. The format of the PMK-MA delete frame body is shown in Figure s88Figure s169.

Octets: 1	1	9	variable
Category	Action Value	MKD-DIE	MEKIE

Figure s169—PMK-MA delete frame body format

	Category	Action Value	MKDDIE	MEKIE
Octets:	1	1	9	variable

Figure s170—PMK-MA delete frame body format

The Category field shall be is set to 0 (representing EMSAMSA).

The Action Value field shall be is set to 5 (representing a PMK-MA delete frame).

The MKDDIE shall be is set as described in 7.3.2.80.

The MEKIE shall be is set as described in 7.3.2.83.

7.4A.1.7 Mesh EAP encapsulation frame format

The Mesh EAP encapsulation frame uses the Mesh Action frame body format and is transmitted by a mesh key holder in the mesh EAP message transport protocol. The frame body of the Mesh EAP encapsulation frame contains the information shown in Table s28Table s39.

Table s39—Mesh EAP encapsulation frame body

Order	Information
0	Category
1	Action Value
3	EAPAIE
4–n	EAPMIE (optional)

Table s40—Mesh EAP encapsulation frame body

Order	Information
0	Category
1	Action Value
3	EAPAIE
4–n	EAPMIE (optional)

The Category field is one octet and shall be is set to 0 (representing EMSAMSA).

The Action Value field is one octet and shall be is set to 6 (representing a Mesh EAP encapsulation frame).

The EAPAIE shall be is set as described in 7.3.2.84.

There may be zero or more EAPMIEs in the Mesh EAP encapsulation frame. If present, EAPMIEs shall be is set as described in 7.3.2.85.

8. Security

8.5 Keys and key distribution

8.5.2 EAPOL-Key frames

8.5.2.1 EAPOL-Key frame notation

Change the text as follows:

Lifetime	is the key lifetime KDE used for sending the expiry timeout value for SMK used during PeerKey Handshake for STA-to-STA SMK key identification. The lifetime KDE is also used during EMSA MSA Authentication in a mesh to express the timeout value of the PMK-MA.

Add Insert the following new clause after Clause 8.7:

EDITORIAL NOTE—protection for management frames defined in the base standard is outside the scope of this clause. This will change if 802.11w is finished prior to the completion of 802.11s. In that case, this clause may be updated to protect management frames from the base standard as well.

8.8 Mesh Link Security

8.8.1 Overview of EMSA

Efficient mesh security association (EMSA) services are used to permit efficient establishment of link security between two MPs in a wireless mesh network, and support both centralized and distributed authentication schemes. EMSA services are provided through the use of a mesh key hierarchy, a hierarchy of derived keys that is established through the use of a PSK or when a MP performs IEEE 802.1X authentication.

The operation of EMSA relies on mesh key holders, which are functions that are implemented at MPs within the wireless mesh network. Two types of mesh key holders are defined: mesh authenticators (MAs) and mesh key distributors (MKDs). A single MP may implement both types of key holders, or the MKD and the MA key holders may be implemented on different MPs.

EMSA requires information to be exchanged during a MP's initial security association with a MA, and is referred to as "Initial EMSA Authentication." Subsequent security associations to other MAs within the same MKD domain (and the same WLAN mesh, as identified by the Mesh ID) may utilize the mesh key hierarchy that is established during Initial EMSA Authentication.

If a MP implements a MA key holder but does not implement a MKD key holder,, EMSA provides mechanisms for secure communications between mesh key holders. The "Mesh Key Holder Security Association" provides the mechanism for establishing a security association between a MA and MKD. Secure mesh key transport protocols and an optional EAP message transport protocol are defined.

8.8.1.1 Mesh Key Holders

Mesh key holders, MAs and MKDs, manage the mesh key hierarchy by performing key derivation and secure

key distribution. A mesh key distributor (MKD) domain is defined by the presence of a single MKD. Within the MKD domain, several MAs may exist, each implemented at a MP, and each MA maintains both a route to and a security association with the MKD. The MKD derives keys to create a mesh key hierarchy, and distributes derived keys to MAs.

A MP implementing the MA key holder function may play the IEEE 802.1X Authenticator role during an EMSA exchange (e.g., Initial EMSA Authentication) as determined according to the procedures in 8.8.1.3. The MA receives derived keys from the MKD, and derives additional keys for use in securing a link with a supplicant MP.

The design of EMSA assumes that the AS and MKD have a trustworthy channel between them that can be used to exchange cryptographic keys without exposure to intermediate parties. The IEEE 802.1X AS never exposes the MSK to any party except the MKD implementing the NAS Client functionality of the IEEE 802.1X Authenticator with which the supplicant is communicating. The communication between AS and MKD is outside the scope of this standard.

8.8.1.2 Discovery & EMSA Capability Advertisement

The support of EMSA is advertised by MPs in Beacon and Probe Response frames through the inclusion of the MKDDIE. Moreover, when a MP wants to utilize EMSA to authenticate with other MPs, it shall advertise its security policy by inserting an RSN information element into its Beacons and Probe Responses.

The MKDDIE shall be included in Beacon and Probe Response frames to advertise support for EMSA and to advertise the MKD domain identifier (MKDD-ID) and the Mesh Security Configuration field. The value of MKDD-ID that is advertised by the MP is the value received from the MKD during the mesh key holder security handshake (as specified in 8.8.10.3.1.2), or the value of dot11MeshKeyDistributorDomainID if the MP implements the MKD function. If the MP has not yet received the MKDD-ID value, it shall set the MKD domain ID field in the MKDDIE to zero, and shall set the Mesh Authenticator and Connected to MKD bits of the Mesh Security Configuration field to zero.

The Mesh Security Configuration field in the MKD domain information element shall be set as follows:

— Bit 0 (Mesh Authenticator): The MP shall set this bit to 1 if the MP is configured to play the IEEE 802.1X Authenticator role during an EMSA handshake. The selection of the IEEE 802.1X Authenticator and Supplicant roles is described in 8.8.1.3.

— Bit 1 (Connected to MKD): The MP shall set this bit to 0 if bit 0 (Mesh Authenticator) is set to 0. Otherwise, the MP shall set this bit to 1 if the MP has a security association with the MKD and has a valid route to the MKD. If the MA and MKD are both implemented at the MP and bit 0 is set to 1, the MP shall set this bit to 1.

— Bit 2 (Default Role Negotiation): The MP shall set this bit to 1 if it uses the mesh default role determination scheme specified in 8.8.1.3. The MP shall set this bit to 0 if it uses some other role determination scheme, such as a proprietary scheme. The specification of other schemes is outside the scope of this standard.

A MKD may support one or more EAP transport mechanisms. A MA advertises the mechanisms supported by the MKD with which it has a security association during the Initial EMSA Authentication (using the EAP Transport List optional parameter in the EMSAIE).

8.8.1.3 Role Determination

When EMSA is used, roles must be selected prior to link establishment and policy selection. In this case, the two MPs shall determine the IEEE 802.1X Authenticator and Supplicant roles through the use of the Mesh Authenticator (Bit 0) and Connected to MKD (Bit 1) bits of the Mesh Security Configuration field (in the MKDDIE) as follows:

— If one of the MPs has set Bit 0 to 1 and the other to 0, then the MP that set Bit 0 to 1 shall assume the IEEE 802.1X Authenticator role, and the MP that set Bit 0 to 0 shall assume the IEEE 802.1X Supplicant role.

— If both MPs set Bit 0 to 1, then

  • If both or neither MPs have set Bit 1 to 1, then the MP with the higher MAC address shall assume the IEEE 802.1X Authenticator role, and the MP with the lower MAC address shall assume the IEEE 802.1X Supplicant role.
  • If one of the MPs has set Bit 1 to 1 and the other to 0, then the MP that set Bit 1 to 1 shall assume the IEEE 802.1X Authenticator role, and the MP that set Bit 1 to 0 shall assume the IEEE 802.1X Supplicant role.

— If both MPs set Bit 0 to 0, then the MP with the higher MAC address shall assume the IEEE 802.1X Authenticator role, and the MP with the lower MAC address shall assume the IEEE 802.1X Supplicant role.

Note that when both MPs set Bit 1 to 0, it is possible for the secure association to fail because one of the parties lacks credentials in its local database to authenticate and/or authorize the other.

8.8.1.4 Policy Selection

An MP may initiate the link establishment mechanism defined in 11A.1.5. This mechanism leverages Association Request and Association Response frames to exchange Peer Link Open and Peer Link Confirm information elements.

If an IEEE 802.1X-based authentication is used, the MP playing the role of the IEEE 802.1X Supplicant shall include an RSN information element in the Association Request specified by this mechanism. In the RSN information element, the Supplicant MP shall specify one pairwise ciphersuite and one authenticated key management suite.

In a mesh, all STAs must utilize the same group ciphersuite. Therefore, a Supplicant MP shall not send an Association Request frame, and shall reject Association Request frames from the Authenticator MP (with Status Code 41), if the group ciphersuite advertised by the Authenticator MP does not match its own.

The Authenticator MP shall reject the Association Request frame from the Supplicant MP if either the pairwise cipher suite (with Status Code 42) or authenticated key management suite (with Status Code 43) selected by the Supplicant is not included in the corresponding lists of pairwise ciphersuites and authenticated key management suites specified in its own Beacons and Probe Responses. The Authenticator MP may also reject the Supplicant MP's Association Request frame for other reasons unrelated to security. The Authenticator MP may accept the Association Request frame if the Supplicant selected pairwise and authenticated key management suites from among those specified by the Authenticator in its Beacons and Probe Responses.

If an IEEE 802.1X-based authentication is used, the Supplicant MP shall additionally include an MKDDIE in the Association Request frame. The Authenticator MP shall reject the Association Request frame from the Supplicant MP if the MKDD-ID included in the MKDDIE does not match the value advertised by the Authenticator MP in its beacons and probe responses.

Selection of the EAP Transport mechanism to be used between a MP and MKD is performed during the mesh key holder security handshake described in 8.8.10.3.1.2. The MP shall decline to establish a mesh key holder security association with the MKD if the EAP transport mechanisms supported by the MP and MKD do not overlap.

8.8.1.5 Initial EMSA Authentication

Pre-RSNA authentication shall not be supported for mesh link establishment.

The Initial EMSA Authentication mechanism permits an MP to enable the use of the mesh key hierarchy when establishing security for subsequent links.

If the link establishment mechanism specified in 11A.1.5 succeeds in creating a link, and if it selects IEEE 802.1X authentication, then the Authenticator MP shall initiate the authentication. Clause 8.4.4 specifies the authentication procedure used when IEEE 802.1X is selected. If pre-shared keys (PSKs) are selected instead, then the PMK is derived from the PSK.

If authentication succeeds from the Authenticator MP's perspective, then it shall initiate a 4-Way Handshake, as specified in 8.8.10.1. After the 4-Way Handshake completes, either MP may initiate a Group Key Handshake (Clause 8.5.4) at any time during the link's lifetime, to update the GTK.

The Initial EMSA Authentication sequence is depicted in Figure s89, with procedures specified in 8.8.10.1.

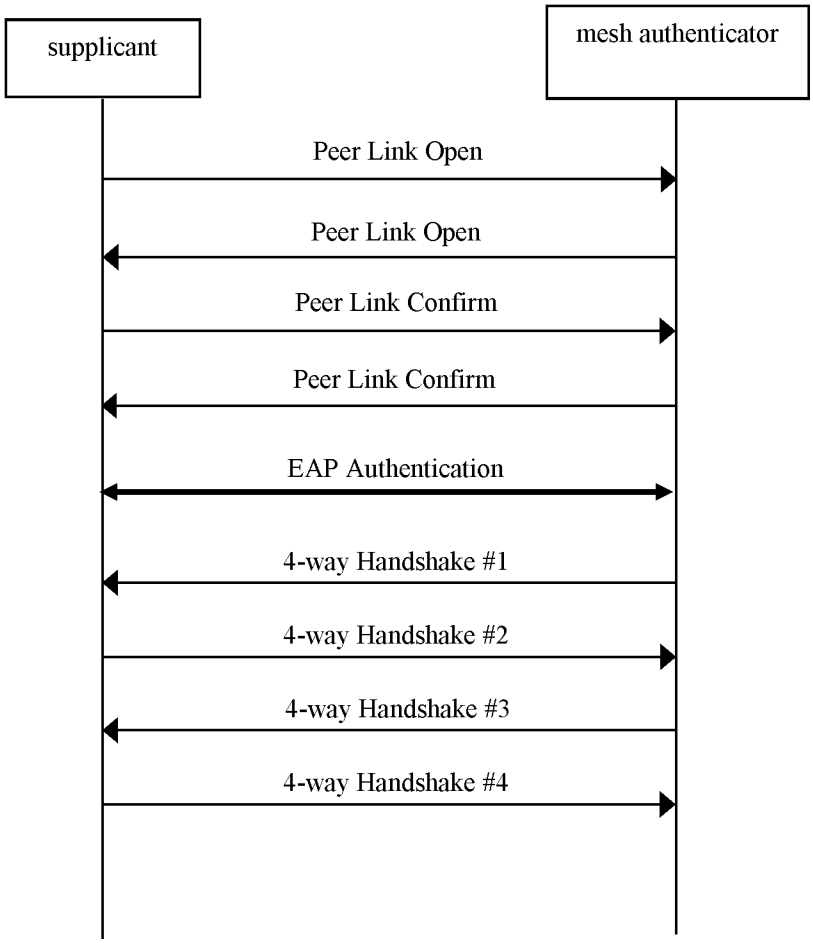

Figure s89—Initial EMSA Authentication

EDITORIAL NOTE—Contents of 8.8 moved to 11A.10, with the exception of 8.8.2 (per CID 2924).

8.8.1.6 Subsequent EMSA Authentication

Pre-RSNA authentication shall not be supported for mesh link establishment.

The Subsequent EMSA Authentication mechanism permits an MP to establish security for subsequent links with other MPs in the mesh once the mesh key hierarchy has been established.

An example Subsequent EMSA Authentication sequence is depicted in Figure s90 with procedures specified in 8.8.10.2.

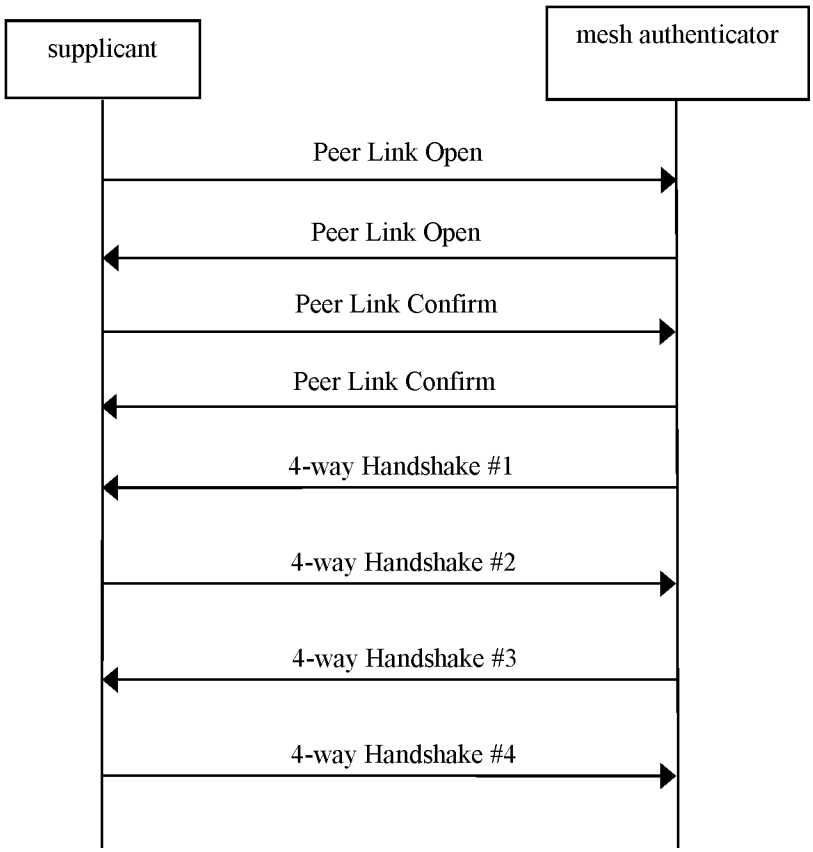

Figure s90—Subsequent EMSA Authentication

8.8.1.7 Mesh Key Holder Security Association

The mesh key holder security association establishes a security association between an MP and a MKD, permitting the MP to begin operating as a MA. The MP may initiate the mesh key holder security handshake after it has completed Initial EMSA Authentication. The mesh key holder security handshake is shown in Figure s91, with procedures specified in 8.8.10.3.1.

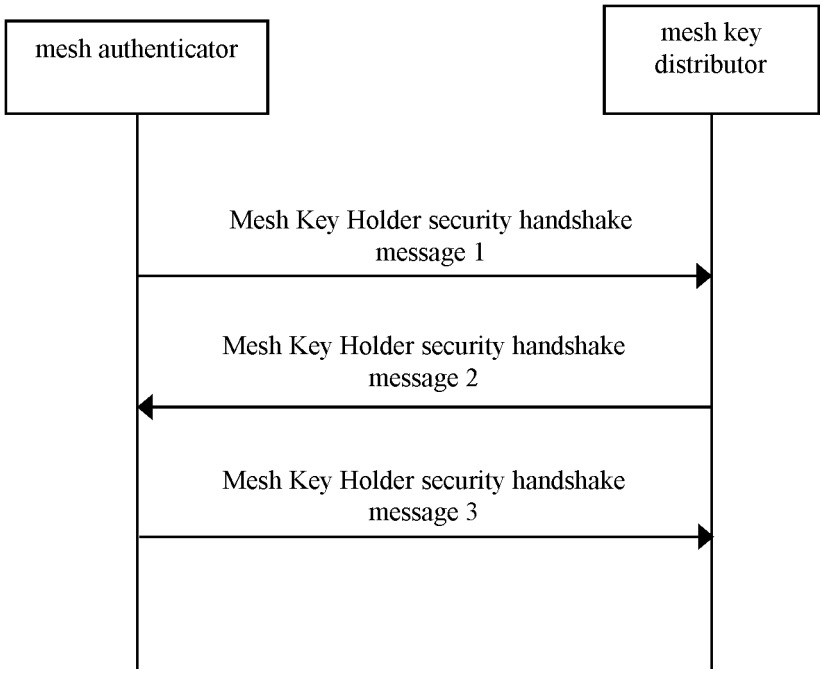

Figure s91—Mesh key holder security handshake

8.8.1.8 Mesh key and EAP message transport protocols

The mesh key transport protocol comprises three mechanisms for performing key delivery and key management within a mesh key hierarchy.

The delivery pull protocol is initiated by the MA to request delivery of a PMK-MA, is shown in Figure s92, and is specified in 8.8.10.3.2.1.

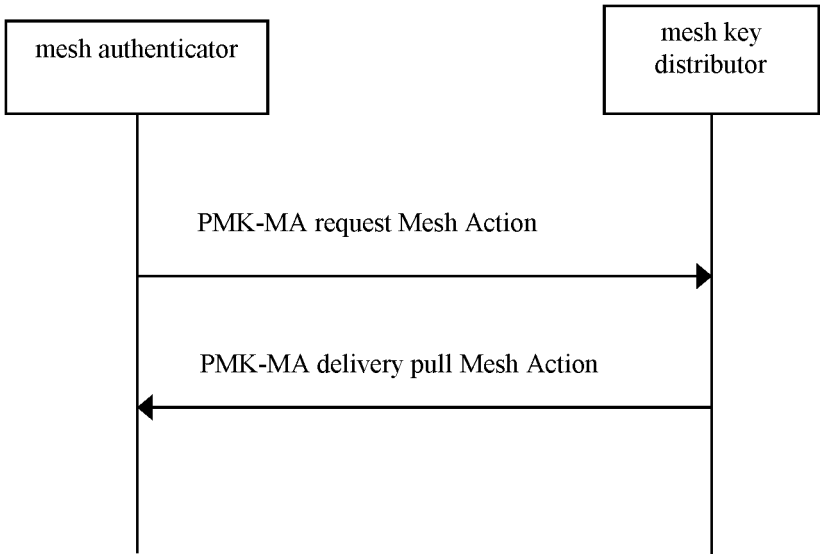

Figure s92—Mesh key transport delivery pull protocol

The delivery push protocol is initiated by the MKD to deliver a PMK-MA, is shown in Figure s93, and is specified in 8.8.10.3.2.2.

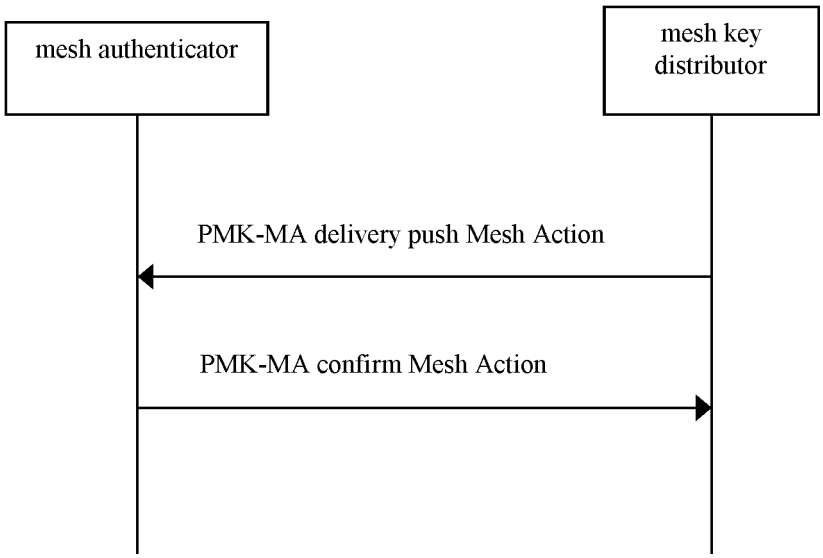

Figure s93—Mesh key transport delivery push protocol

The key delete protocol is initiated by the MKD to request revocation of a PMK-MA, is shown in Figure s94, and is specified in8.8.10.3.2.3.

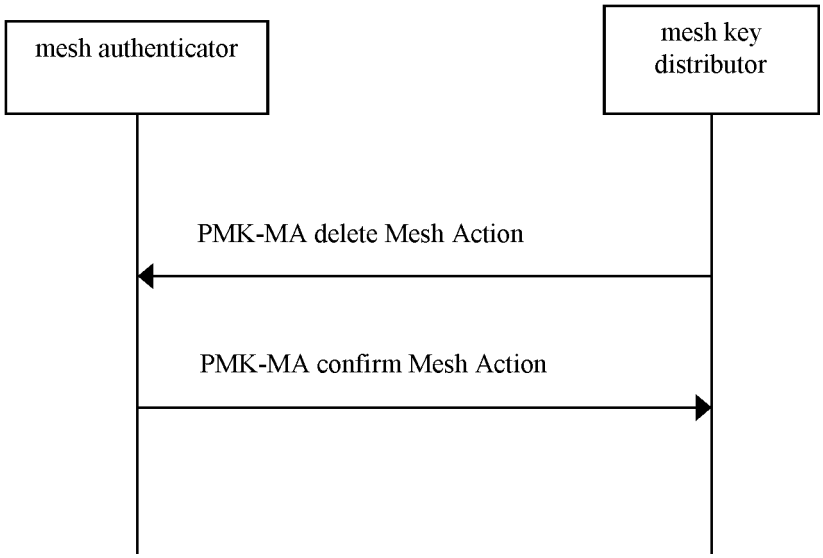

Figure s94—Mesh key delete protocol

The optional EAP message transport protocol may be initiated by the MA to facilitate EAP authentication with the supplicant during a supplicant MP's Initial EMSA Authentication. The protocol permits an EAP message received from the supplicant to be transported from MA to MKD, and permits EAP messages received from the authentication server to be transported from MKD to MA.

A single request/response EAP message transport frame exchange is depicted in Figure s95. The authentication of a supplicant will typically require several such exchanges. The optional protocol is specified in8.8.10.3.3.

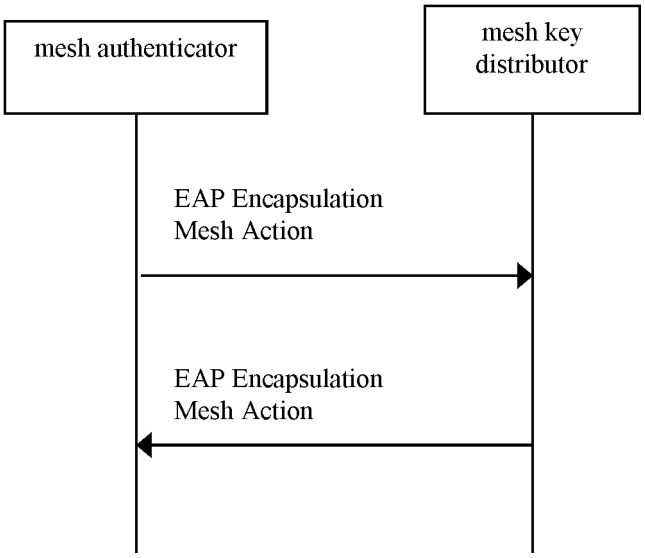

Figure s95—EAP message transport protocol (single exchange)

8.8.1.9 Secure Link Operation

In the case when the 4-Way Handshake completes successfully, then both the IEEE 802.1X Authenticator and Supplicant shall open their respective controlled ports, to permit data traffic to be exchanged using the selected ciphersuites.

When key management completes, each MP further uses the session key to protect the contents of mesh action data units using the agreed upon ciphersuites. Each MP permits MPDUs protected by the session key and group key using the agreed upon ciphersuites, and discards received MPDUs that are unprotected.

8.8 Key Distribution distribution for EMSAMSA

This subclause describes the mesh key hierarchy and its supporting architecture. The mesh key hierarchy permits a an MP to create secure associations with peer MPs without the need to perform an IEEE 802.1X authentication each time. The mesh key hierarchy can be used with either IEEE 802.1X authentication or PSKPSK authentication. It is assumed by this standard that the PSK is specific to a single MP and a single MKD.

8.8.1 Overview

A key hierarchy consisting of two branches is introduced for use within a mesh. A link security branch consists of three levels, supporting distribution of keys between mesh key holders to permit the use of the mesh key hierarchy between a supplicant MP and a an MA. A key distribution branch provides keys to secure the transport and management of keys between mesh key holders.

As depicted shown in Figure s93, the mesh key distributor generates the first level key for both branches from either the PSK or from the MSK resulting (per IETF RFC 3748) from a successful IEEE 802.1X Authentication between the AS and the supplicant MP. The second level keys in both branches are generated by the MKD as well.

In the link security branch, the first level key (PMK-MKD) is derived by the MKD from either the PSK or MSK. The second level keys (PMK-MA keys) are generated by the MKD as well. The PMK-MA keys are delivered from the MKD to the MAs using a secure protocol, as described in 8.8.10.3.211A.2.4. The PMK-MA keys are used for PTK generation.

In the key distribution branch, the first level key (KDK) is derived by the MKD from either the PSK or MSK. The second level key (PTK-KD) is generated by the MKD as well, during the mechanism described in 8.8.10.3.1.211A.2.3.2.

Upon a successful authentication between a supplicant MP and the MKD, the supplicant MP and the MKD shall delete the prior PMK-MKD, KDK, and PTK-KD keys and all PMK-MA keys which were created between the supplicant MP and the same MKD domain. Upon receiving a new PMK-MA key for a supplicant MP, a an MA shall delete the prior PMK-MA key and all PTKs derived from the prior PMK-MA key.

The lifetime of all keys derived from the PSK or MSK are bound to the lifetime of the PSK or MSK. For example, the IEEE 802.1X AS may communicate the MSK key lifetime with the MSK. If such an attribute is provided, the lifetimes of the PMK-MKD and KDK shall be not more than the lifetime of the MSK. If the MSK lifetime attribute is not provided, or for PSK, the key lifetime shall be the value of the MIB variable dot11MeshTopLevelKeyLifetime.

The lifetime of the PTK and PMK-MA shall be the same as that of the PMK-MKD and the lifetime of the

PTK-KD shall be the same as that of the KDK, as calculated above. When the key lifetime expires, each key holder shall delete their respective derived keys.

The mesh key hierarchy derives its keys using a the Key Derivation Function (KDF) as defined in 8.8.3 with separate labels to further distinguish derivations.

The mesh key hierarchy is depicted shown in Figure s92.

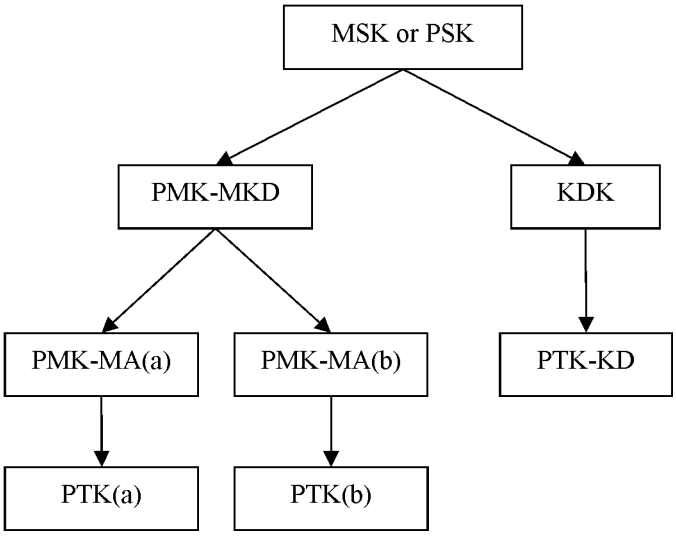

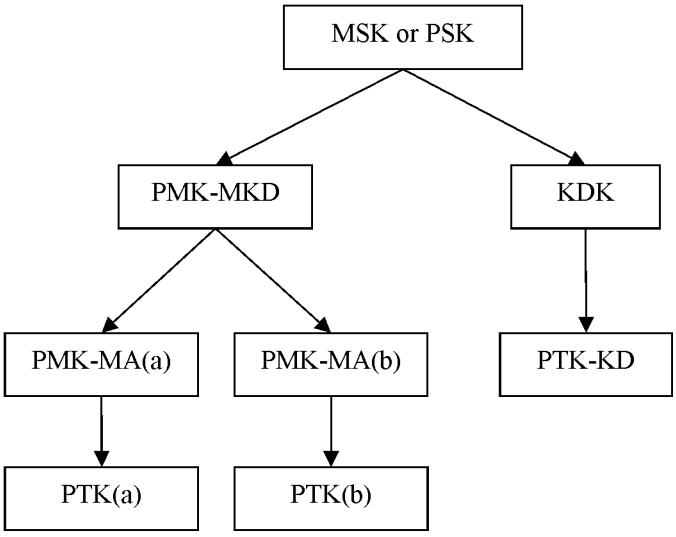

Figure s92—Mesh key hierarchy

The operations performed by mesh key holders and the movement of keys within the mesh key hierarchy are shown in Figure s93.

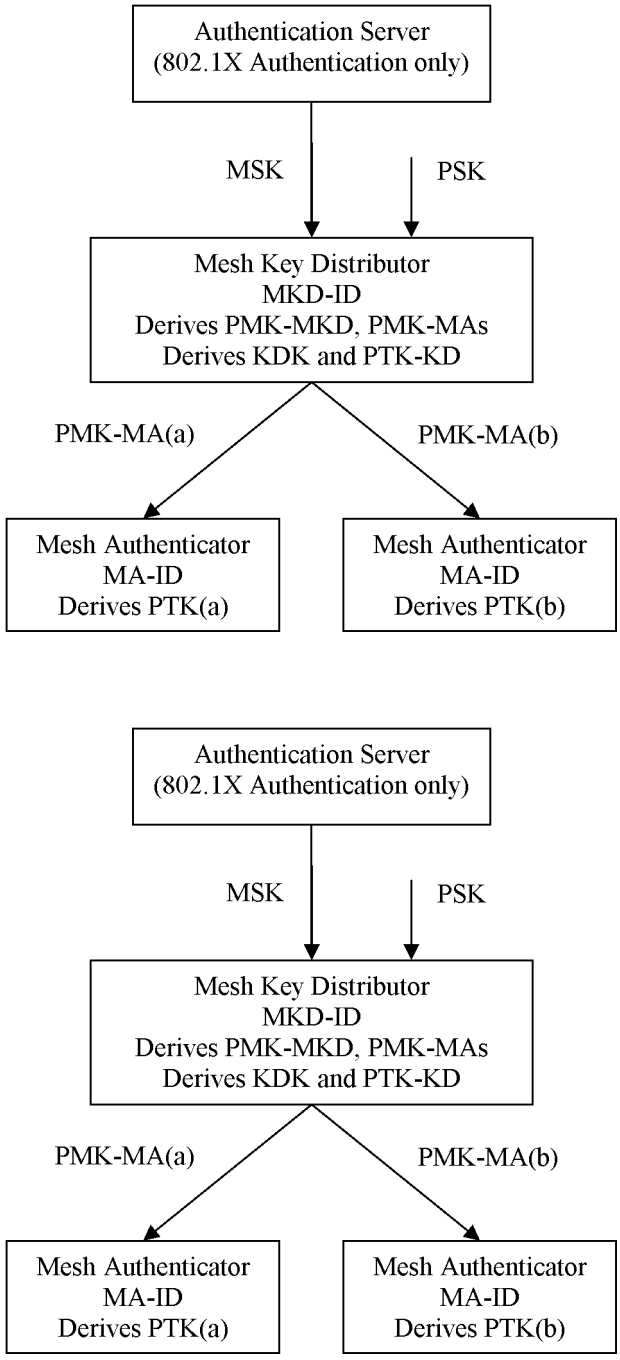

Figure s93—Mesh key holders

The construction of the key hierarchy ensures that compromise of keying material within the link security branch is isolated to only that portion, or sub-branch, of the hierarchy. For example, a mesh authenticator only has knowledge to decrypt those sessions protected by the PTK derived from its PMK-MA.

In some key management systems, PMK-MKD key may be deleted by the MKD after PMK-MA keys have been derived. Such an operation lends itself to the good security practice of protecting the key hierarchy in cases where the PMK-MKD is no longer needed. In such cases, the key management system only needs to

maintain information about the PMK-MA keys. Such a removal of the PMK-MKD key does not indicate the invalidity of the key hierarchy.

8.8.2 Key Hierarchyhierarchy

The mesh key hierarchy consists of two branches whose keys are derived using a the KDF described in 8.8.3.

The first branch, the link security branch, consists of three levels and results in a PTK for use in securing a link.

— PMK-MKD – the The first level of the link security branch, this key is derived as a function of the MSK or PSK and the Mesh ID. It is stored by the supplicant MP and the PMK-MKD key holder, namely the MKD. This key is mutually derived by the supplicant MP and the MKD. There is only a single PMK-MKD derived between the supplicant MP and the MKD domain.

— PMK-MA – the The second level of the link security branch, this key is mutually derived by the supplicant MP and the MKD. It is delivered by the MKD to a an MA to permit completion of an EMSA MSA handshake between the supplicant MP and the MA.

— PTK – the The third level of the link security branch that defines the IEEE 802.11 and IEEE 802.1X protection keys. The PTK is mutually derived by the supplicant and the PMK-MA key holder, namely the MA.

The PTK is used as defined by 8.5 for secure link operation.

The second branch, the key distribution branch, consists of two levels and results in a PTK-KD for use in allowing a an MP to become a an MA, and in securing communications between a an MA and the MKD.

— KDK – the The first level of the key distribution branch, this key is derived as a function of the MSK or PSK and the Mesh ID and stored by the supplicant MP and the MKD. This key is mutually derived by the supplicant MP and the MKD. There is only a single KDK derived between the supplicant MP and the MKD.

— PTK-KD – the The second level of the key distribution branch that defines protection keys for communication between MA and MKD. The PTK-KD is mutually derived by the supplicant MP (when it becomes a an MA) and the MKD.

8.8.3 Key Derivation Functionderivation function

The key derivation function for the mesh key hierarchy, KDF, is a variant of the PRF function defined in 8.5.1.1, and defined as follows:

Output = KDF-Length (K, label, Context) where
Input: *K*, a 256 bit key derivation key
label, a string identifying the purpose of the keys derived using this KDF
Context, a bit string that provides context to identify the derived key
Length, the length of the derived key in bits
Output: a *Length*-bit derived key

result = ""
iterations = (*Length*+255)/256
do i = 1 **to** *iterations*
result = *result* || HMAC-SHA256(*K*, *i* || *label* || 0x00 || *Context* || *Length*)
od
return first *Length* bits of *result,* and securely delete all unused bits

In this algorithm, *i* and *Length* are encoded as 16-bit unsigned integers, represented using the bit ordering conventions of 7.1.1.

8.8.4 PMK-MKD

The top level key of the mesh key hierarchy link security branch, PMK-MKD binds the SPA, MKD domain Identifieridentifier, and Mesh ID with the keying material resulting from the negotiated AKM. The PMK-MKD is the top level 256-bit keying material used to derive the next level keys (PMK-MAs):

PMK-MKD = KDF-256(XXKey, "MKD Key Derivation", MeshIDlength || MeshID || MKDD-ID || 0x00 || SPA)

where

- KDF-256 is the KDF function as defined in 8.8.3 used to generate a key of length 256 bits.
- If the AKM negotiated is 00-0F-AC:5, then XXKey shall be the second 256 bits of the MSK (MSK being derived from the IEEE 802.1X authentication), i.e., XXKey = L(MSK, 256, 256). If the AKM negotiated is 00-0F-AC:6, then XXKey shall be the PSK.
- "MKD Key Derivation" is 0x4D4B44204B6579204465726976617469 6F6E.
- MeshIDLength is a single octet whose value is the number of octets in the Mesh ID.
- Mesh ID is the mesh identifier, a variable length sequence of octets, as it appears in the Beacons Beacon frames and Probe ResponsesResponse frames.
- MKDD-ID is the 6-octet MKD domain identifier field from the MKD domain information element that was used during Initial EMSA MSA Authentication.
- SPA is the supplicant MP's MAC address.

The PMK-MKD is referenced and named as follows:

PMK-MKDName = Truncate-128(SHA-256("MKD Key Name" || MeshIDlength || MeshID || MKDD-ID || 0x00 || SPA || ANonce))

where

- "MKD Key Name" is 0x4D4B44204B6579204E616D65.
- ANonce is an unpredictable 256-bit pseudo-random value generated by the PMK-MKD holder (MKD), delivered along with PMK-MA to the MA, and provided by the MA to the supplicant MP during Initial EMSA MSA Authentication.
- Truncate-128(-) returns the first 128 bits of its argument, and securely destroys the remainder.

8.8.5 PMK-MA

The second level key of the mesh key hierarchy link security branch, PMK-MA, is a 256-bit key used to derive the PTK. The PMK-MA binds the SPA, MKD, and MA:

PMK-MA = KDF-256(PMK-MKD, "MA Key Derivation", PMK-MKDName || MA-ID || 0x00 || SPA)

where

- KDF-256 is the KDF function as defined in 8.8.3 used to generate a key of length 256 bits.
- PMK-MKD is the key defined in 8.8.4.

— "MA Key Derivation" is 0x4D41204B6579204465726976617469 6F6E.
— PMK-MKDName is defined in 8.8.4.
— MA-ID is the identifier of the holder of PMK-MA (MA).
— SPA is the supplicant MP's MAC address.

The PMK-MA is referenced and named as follows:

PMK-MAName = Truncate-128(SHA-256("MA Key Name" || PMK-MKDName || MA-ID || 0x00 || SPA))

where

— "MA Key Name" is 0x4D41204B6579204E616D65.

8.8.6 PTK

The third level key of the mesh key hierarchy link security branch is the PTK. This key is mutually derived by the Supplicant MP and the MA with the key length being a function of the negotiated cipher suites as defined by Table 60 in 8.5.2.

The PTK derivation is as follows:

PTK = KDF-PTKLen(PMK-MA, "Mesh PTK Key derivation", SNonce || ANonce || SPA || MA-ID || PMK-MAName)

where

— KDF-PTKLen is the KDF function as defined in 8.8.3 used to generate a PTK of length PTKLen.
— PMK-MA is the key that is shared between the Supplicant MP and the MA
— "Mesh PTK Key derivation" is 0x4D6573682050544B204B6579206465726976617469 6F6E.
— SNonce is a 256 bit pseudo-random bit string contributed by the Supplicant MP
— ANonce is a 256 bit pseudo-random string contributed by the MKD or MA
— SPA is the Supplicant MP's MAC address
— MA-ID is the MAC address of the MA.
— PMK-MAName is defined in 8.8.5
— PTKlen is the total number of bits to derive, e.g., number of bits of the PTK. The length is dependent on the negotiated cipher suites as defined by Table 60 in 8.5.2.

Each PTK has three associated component keys, KCK, KEK, and TK, derived as follows:

The KCK shall be computed as the first 128 bits (bits 0-127) of the PTK:

KCK = L(PTK, 0, 128)

where L(-) is defined in 8.5.1.

The KCK is used to provide data origin authenticity between a supplicant MP and the MA, as defined in 8.8.1011A.2.2.

The KEK shall be computed as bits 128-255 of the PTK:

KEK = L(PTK, 128, 128)

The KEK is used to provide data confidentiality between a supplicant MP and the MA, as defined in 8.8.1011A.2.2.

Temporal keys (TK) shall be computed as bits 256-383 (for CCMP) or bits 256-511 (for TKIP) of the PTK:

TK = L(PTK, 256, 128), or
TK = L(PTK, 256, 256)

The temporal key is configured into the Supplicant MP through the use of the MLME-SETKEYS.request primitive. The MP uses the temporal key with the pairwise cipher suite; interpretation of this value is cipher-suite specific.

The PTK is referenced and named as follows:

PTKName = Truncate-128(SHA-256(PMK-MAName || "Mesh PTK Name" || SNonce || ANonce || MA-ID || SPA))

where

— "Mesh PTK Name" is 0x4D6573682050544B204E616D65.

8.8.7 KDK

The first level key of the key distribution branch, KDK binds the MA-ID (the MAC address of the MP establishing the KDK to become a an MA), MKD domain identifier, and Mesh ID with the keying material resulting from the negotiated AKM. The KDK is used to derive the PTK-KD.

KDK = KDF-256(XXKey, "Mesh Key Distribution Key", MeshIDLength || MeshID || MKDD-ID || 0x00 || MA-ID)

where

— KDF-256 is the KDF function as defined in 8.8.3 used to generate a key of length 256 bits.

— If the AKM negotiated is 00-0F-AC:5, then XXKey shall be the second 256 bits of the MSK (MSK being derived from the IEEE 802.1X authentication), i.e., XXKey = L(MSK, 256, 256). If the AKM negotiated is 00-0F-AC:6, then XXKey shall be the PSK.

— "Mesh Key Distribution Key" is 0x4D657368204B65792044697374726962757469 6F6E204B6579.

— MeshIDLength is a single octet whose value is the number of octets in the Mesh ID.

— Mesh ID is the mesh identifier, a variable length sequence of octets, as it appears in the Beacons Beacon frames and Probe ResponsesResponse frames.

— MKDD-ID is the 6-octet MKD domain identifier field from the MKD domain information element that was used during Initial EMSA MSA Authentication.

— MA-ID is the MAC address of the MP establishing a security association with the MKD in order to become configured as a an MA.

The KDK is referenced and named as follows:

KDKName = Truncate-128(SHA-256("KDK Name" || MeshIDLength || MeshID || MKDD-ID || 0x00 || MA-ID || ANonce))

where

— "KDK Name" is 0x4B444B204E616D65.

— Truncate-128(-) returns the first 128 bits of its argument, and securely destroys the remainder.

— ANonce is identical to the value used to calculate PMK-MKDName, as described in 8.8.4.

8.8.8 PTK-KD

The second level key of the key distribution branch, PTK-KD, is a 256-bit key that is mutually derived by a an MA and a an MKD. The PTK-KD is derived:

PTK-KD = KDF-256(KDK, "Mesh PTK-KD Key", MA-Nonce || MKD-Nonce || MA-ID || MKD-ID)

where

— KDK is the key defined in 8.8.7.
— "Mesh PTK-KD Key" is 0x4D6573682050544B2D4B44204B6579.
— MA-Nonce is a 256-bit pseudo-random string contributed by the MA.
— MKD-Nonce is a 256-bit pseudo-random string contributed by the MKD.
— MA-ID is the MAC address of the MA.
— MKD-ID is the MAC address of the MKD.

The PTK-KD has two associated component keys, the Key confirmation key for key distribution (KCK-KD) and the Key encryption key for key distribution (KEK-KD), derived as follows:

The KCK-KD shall be computed as the first 128 bits (bits 0-127) of the PTK-KD:

KCK-KD = L(PTK-KD, 0, 128)

where L(-) is defined in 8.5.1.

The KCK-KD is used to provide data origin authenticity in messages exchanged between MA and MKD, as defined in 8.8.10.211A.2.2.4.

The KEK-KD shall be computed as bits 128-255 of the PTK-KD:

KEK-KD = L(PTK-KD, 128, 128)

The KEK-KD is used to provide data confidentiality in messages exchanged between MA and MKD, as defined in 8.8.10.211A.2.2.4.

The PTK-KD is referenced and named as follows:

PTK-KDName = Truncate-128(SHA-256(KDKName || "PTK-KD Name" || MA-Nonce || MKD-Nonce || MA-ID || MKD-ID))

where

— "PTK-KD Name" is 0x50544B2D4B44204E616D65.

8.8.9 Mesh key holders

8.8.9.1 Key holder requirements

The MKD and MA are responsible for the derivation of keys in the mesh key hierarchy. For EMSAMSA, the functions of the IEEE 802.1X Authenticator are distributed between the MKD and MA. Each mesh key holder shall have an identity that is communicated to the supplicant MP and other key holders which is bound

into the key derivation. Each identity shall be mapped to a physical entity where it resides.

The MKD shall meet the following requirements.

— The MKD shall be co-resident with the NAS client functionality of the IEEE 802.1X Authenticator.

— The MKD domain identifier (MKDD-ID) uniquely identifies a an MKD (i.e., there is a one-to-one mapping between a an MKD domain and a an MKD). A An MKD's MKDD-ID shall be set to the value of dot11MeshKeyDistributorDomainID. The MKDD-ID is bound into the derivation of the first level keys (PMK-MKD and KDK).

— The mesh key distributor identifier (MKD-ID) shall be set to the MAC address of the physical entity that stores the MKD. The MKD-ID is used in the generation of the PTK-KD.

— When the PMK-MKD lifetime expires, the MKD shall delete the PMK-MKD SA and should revoke all PMK-MAs derived from the PMK-MKD. A mechanism for PMK-MA revocation is provided in 8.8.10.3.2.311A.2.4.3.

The MA shall meet the following requirements.

— The mesh authenticator identity (MA-ID) shall be set to the MAC address of the physical entity that stores the PMK-MA and uses it to generate the PTK. That same MAC address shall be used to advertise the MA identity to mesh points MPs and to the MKD.

— The MA shall provide the IEEE 802.1X Authenticator function to derive and distribute the GTK to connected MPs.

— When the PMK-MA lifetime expires, the MA shall delete the PMK-MA SA and shall revoke all PTKs derived from the PMK-MA using the MLME-DELETEKEYS primitive.

8.8.9.2 PMK-MA Distribution distribution within a an MKD domain

A An MKD domain is identified by the MKD domain identifier (MKDD-ID). A An MKD domain contains a single MKD, and the MKD uses the MKDD-ID to identify itself. A An MKD domain comprises at least one MA, which has established a security association with the MKD.

An MP creates its mesh key hierarchy during the Initial EMSA MSA Authentication, utilizing information forwarded from the MKD by the MA. During the Initial EMSA MSA Authentication, the MKD derives the PMK-MKD from the MSK acquired during IEEE 802.1X authentication, when the negotiated AKM is 00-0F-AC:5, or from the PSK, when the negotiated AKM is 00-0F-AC:6.

Additionally, the MKD is responsible for deriving a PMK-MA for each MA within the MKD domain. The MKD is responsible for transmitting the derived PMK-MA keys securely to those key holders, along with the PMK-MAName, the key lifetime, and key context information associated with that PMK-MA. The MKD shall also securely transmit the ANonce used in the calculation of PMK-MKDName to the MA for use in the Initial EMSA MSA Authentication mechanism.

The secure transmission of keys and key information from MKD to MA shall be through the use of the mesh key transport protocol described in 8.8.10.211A.2.2.4.

Each MA shall derive the PTK mutually with the supplicant MP.

8.8.10 EMSA Establishment Procedure

EMSA defines the following procedures for use within a mesh:

— Initial EMSA Authentication (8.8.3.1) is used by a MP to authenticate and establish the mesh key hierarchy that may be used when securing future links.

— Subsequent EMSA Authentication (8.8.3.2) is used by a MP to securely establish links with peer MPs after it has established a mesh key hierarchy using Initial EMSA Authentication.

— EMSA Key Holder Communication comprises three related mechanisms:

- The procedure for establishing communications and security between a MA and a MKD is the mesh key holder security association (8.8.3.3.1).
- The mesh key transport protocol (8.8.3.3.2) describes the mechanisms for key delivery and key management within the mesh key hierarchy.
- The optional mesh EAP message transport protocol (8.8.3.3.3) describes a mechanism for transporting EAP messages between MKD and MA to facilitate authentication of a supplicant MP.

8.8.10.1 Initial EMSA Authentication Mechanism

During its first authentication in a mesh, a MP establishes the mesh key hierarchy to be used when securing future links. This is referred to as the Initial EMSA Authentication Mechanism, and contains communication exchanged between an MP and a MA with which it is associating.

In this sequence, a MP issues an association request frame containing a Peer Link Open IE and an indication (the MKDDIE) that it wishes to establish the mesh key hierarchy. The MP receives an association response frame containing a Peer Link Confirm IE and information required for the MP to perform key derivations and establish link security. If required, 802.1X authentication occurs next, followed by an EMSA 4-way handshake.

The supplicant MP in the Initial EMSA Authentication mechanism sends an association request frame to the MA. The association request frame shall contain:

— Peer Link Open IE, which shall be set according to 11A.1.5

— MKDDIE, configured exactly as advertised by the supplicant MP in its beacons and probe responses.

— RSNIE, which shall be set according to the policy in 8.4.3 and 8.8.1.4, and with the PMKID list field empty.

The association response frame from the MA shall contain:

— Peer Link Confirm IE, which shall be set according to 11A.1.5

— MKDDIE, configured exactly as advertised by the MA in its beacons and probe responses.

— EMSAIE, where

- MA-ID is set to the MAC address of the MA
- The Optional Parameters field includes the MKD-ID, which contains the identifier of the MKD with which the MA has a security association.
- The Optional Parameters field includes the EAP Transport List, which contains the list of transport types supported by the MKD with which the MA has a security association.
- All other fields are set to zero.

— RSNIE, configured exactly as advertised by the MA in its beacons and probe responses, with the PMKID list field empty.

After successful peer link establishment, the supplicant MP and the MA proceed with IEEE 802.1X authentication, if required. The IEEE 802.1X exchange is sent between the supplicant MP and the MA using EAPOL messages carried in IEEE 802.11 data frames. The MA initiates the IEEE 802.1X exchange with the supplicant MP and may transport the 802.1X exchange to the MKD using the optional mesh EAP message transport protocol, as specified in 8.8.10.3.3.

Upon successful completion of the IEEE 802.1X authentication, the MKD receives the MSK and authorization attributes associated with it and with the supplicant MP. If a mesh key hierarchy already exists for this supplicant, the MKD shall delete the old PMK-MKD SA and PMK-MA SAs. It then calculates the PMK-

MKD and PMK-MKDName. The PMK-MKD SA includes:

- MKDD-ID
- PMK-MKD
- PMK-MKDName
- SPA, and
- authorization information including PMK-MKD lifetime.

The MKD then generates a PMK-MA for the MA. The PMK-MA SA includes:

- PMK-MA,
- PMK-MA lifetime,
- PMK-MAName,
- MA-ID,
- PMK-MKDName, and
- SPA

The MKD then delivers the PMK-MA to the MA using the mesh key distribution protocol defined in 8.8.10.3.2. Once the PMK-MA is delivered, the MA and supplicant MP then perform an EMSA 4-way handshake. The EAPOL-Key frame notation is defined in 8.5.2.2.

MA -> Supplicant: Data(EAPOL-Key(0, 0, 1, 0, P, 0, 0, ANonce, 0, DataKD_M1)) where DataKD_M1 = 0.

Supplicant -> MA: Data(EAPOL-Key(0, 1, 0, 0, P, 0, 0, SNonce, MIC, DataKD_M2)) where DataKD_M2 = (RSNIE, MKDDIE, GTK KDE).

MA -> Supplicant: Data(EAPOL-Key(1, 1, 1, 1, P, 0, 0, ANonce, MIC, DataKD_M3)) where DataKD_M3 = (RSNIE, MKDDIE, EMSAIE, GTK KDE, Lifetime KDE).

Supplicant -> MA: Data(EAPOL-Key(1, 1, 0, 0, P, 0, 0, MIC, DataKD_M4)) where DataKD_M4 = 0.

The message sequence is similar to that of 8.5.3. The contents of each message shall be as described in 8.5.3, except as follows:

- Message 1: ANonce is the value received by the MA from the MKD during PMK-MA delivery. The Key Data field is empty.
- Message 2: The RSNIE shall contain only the PMK-MAName in the PMKID list field. The remaining fields of the RSNIE and the MKDDIE shall be the same as that provided in the association request frame sent by the supplicant MP. The GTK KDE shall contain the GTK of the supplicant MP. The Key Data field shall be encrypted.
- Message 3: The RSNIE shall contain only the PMK-MAName, as calculated by the MA, in the PMKID list field. The remaining fields of the RSNIE, as well as the MKDDIE and EMSAIE shall be the same as that provided in the association response frame sent by the MA. The Lifetime KDE shall contain the lifetime of the PMK-MA, expressed in seconds.

The PTK shall be calculated by the supplicant and the MA according to the procedures given in 8.8.6.

Following a successful EMSA 4-way handshake, the IEEE 802.1X controlled port shall be opened on both the supplicant and MA. Subsequent EAPOL-Key frames shall use the Key Replay Counter to ensure they are not replayed.

8.8.10.2 Subsequent EMSA Authentication Mechanism

The subsequent EMSA authentication mechanism is used by a MP after it has established a mesh key hierar-

chy using the initial EMSA authentication mechanism. The subsequent EMSA authentication mechanism may only be performed between MPs that advertise the same MKD domain identifiers in the MKDDIE in beacons and probe responses.

The subsequent EMSA authentication mechanism follows the procedure specified in 8.8.10.1, Initial EMSA authentication mechanism, with the following modifications.

The RSNIE sent by the supplicant MP in the association request frame (peer link open message) shall contain a single entry in the PMKID list field. The value of the entry is PMK-MKDName, identifying the PMK-MKD the supplicant MP created during its initial EMSA authentication.

The RSNIE sent by the authenticator MP in the association response frame (peer link confirm message) shall contain the PMKID list entry sent by the supplicant MP in the peer link open message.

After successful peer link establishment, the MA shall calculate the PMK-MAName using the PMK-MKD-Name sent by the supplicant MP in the peer link open message. If the MA does not have the key identified by PMK-MAName, it may attempt to retrieve that key using the mesh key transport protocol according to 8.8.10.3.2. If the MA is unable to retrieve the PMK-MA, it shall initiate IEEE 802.1X authentication (if required) to establish a mesh key hierarchy for the supplicant MP, and continue with the Initial EMSA Authentication mechanism as specified in 8.8.10.1.

Upon obtaining the specified PMK-MA, the MA omits IEEE 802.1X authentication, and initiates the EMSA 4-way handshake, as specified in 8.8.10.1. However, the ANonce value in message 1 is a random nonce created by the MA, and is not the value received during PMK-MA delivery.

Following a successful EMSA 4-way handshake, the IEEE 802.1X controlled port shall be opened on both the supplicant and MA. Subsequent EAPOL-Key frames shall use the Key Replay Counter to ensure they are not replayed.

8.8.10.3 EMSA Key Holder Communication

In order to support the mesh key hierarchy, mesh key holders must communicate securely to provide the following services to mesh points:

— transporting EAP authentication traffic between key holders to permit a supplicant mesh point to perform 802.1X authentication, and

— securely delivering derived keys to facilitate the use of a derived key hierarchy.

A mesh point invokes the mesh key holder security handshake to establish a security association with a mesh key distributor (MKD). The mechanism permits the mesh point to operate as a mesh authenticator (MA). Subsequently, the MA advertises, in beacons and probe responses, its capability to authenticate mesh points using the mesh key hierarchy.

8.8.10.3.1 Mesh key holder security association

A security association is established between a MA and MKD to provide secure communications between these two key holders within a mesh. The mesh key holder security association is used to provide message integrity and data origin authenticity in all messages passed between MA and MKD after the security association is established. Further, it provides encryption of derived keys and key context during key delivery protocols. Establishing the mesh key holder security association begins with discovery of the MKD, followed by a handshake initiated by the MA. The result of the security association is the pairwise transient key for key derivation (PTK-KD), used to provide the security services between MA and MKD.

8.8.10.3.1.1 Mesh key distributor discovery

Prior to initiating the mesh key holder security handshake described in clause 8.8.10.3.1.2, a MA must obtain the address of its MKD. If the MA is not also a MKD, it may obtain the MKD-ID address of its MKD from the EMSAIE conveyed in the Association Response frame received during its initial EMSA security association.

8.8.10.3.1.2 Mesh key holder security handshake

The mesh key holder security handshake may commence after a mesh point has completed its Initial EMSA Authentication. The mechanism permits the MP to establish a security association with the MKD that derived its PMK-MKD during Initial EMSA Authentication.

The MP initiates the exchange by constructing mesh key holder security handshake message 1, and sending the message to the MKD identified by the MKD-ID received in the Association Response frame during the MP's Initial EMSA Authentication. The MP selects an EAP Transport mechanism from among those listed in the EMSAIE received in the Association Response frame during the MP's Initial EMSA Authentication. The MP shall decline to establish a mesh key holder security association with the MKD if the EAP transport mechanisms supported by the MP and MKD do not overlap. The contents of handshake message 1 are given in

Upon receiving handshake message 1, the MKD chooses MKD-Nonce, a value chosen randomly, and computes the PTK-KD using the MA-Nonce received in handshake message 1 and MKD-Nonce, as specified in 8.8.8. The MKD verifies that it supports the selected EAP Transport mechanism; if not, the handshake fails. The MKD sends handshake message 2, with contents as given in 8.8.10.3.1.2.2. Upon receiving handshake message 2, the MP computes the PTK-KD, and sends handshake message 3, with contents as given in 8.8.10.3.1.2.3.

After completing the handshake, the MP sets both the "Mesh Authenticator and "Connected to MKD" bits to 1 in the MKDDIE in its beacons and probe responses to advertise that it is configured as a mesh authenticator that is connected to the MKD. The MKDDIE shall contain the MKDD-ID that is received from the MKD in mesh key holder security handshake message 2. A MA must maintain a route to the MKD. If the route is lost and cannot be repaired, the MA shall set the "Connected to MKD" bit to 0 in the MKDDIE. In such a case, the "Mesh Authenticator" bit may be set to 1 to advertise the ability to act in the IEEE 802.1X Authenticator role using, for example, cached keys. After the route is re-established, the MP may again set the "Connected to MKD" bit to 1.

The MA and the MKD will maintain separate key replay counters for sending messages providing mesh key transport that are protected using the PTK-KD. Immediately upon deriving the PTK-KD, both the MKD and MA shall reset their replay counters to zero.

8.8.10.3.1.2.1 Mesh key holder security handshake message 1

Mesh key holder security handshake message 1 is a mesh key holder security establishment EMSA mesh action frame with the following contents:

The MAC address of the MKD shall be asserted in the DA field of the message header.

The MAC address of the MP shall be asserted in the SA field of the message header.

The Mesh ID IE shall contain the Mesh ID that the MP advertises in its beacons and probe responses.

The MKDDIE shall contain the value of MKDD-ID that was contained in the MKDDIE received in the Association Response frame during the MP's Initial EMSA Authentication. The Mesh Security Configuration

field shall be set to zero.

The MKHSIE shall be set as follows:

— MA-Nonce shall be set a value chosen randomly by the MP, following the recommendations of 8.5.8.

MA-ID Each MA shall be set to derive the MAC address of PTK mutually with the supplicant MP.

— MKD-ID shall be set to the MAC address of the MKD.

— The Transport Type Selector field shall contain a single transport selector (with format as given in Figure s62). The specified transport type shall be from among those listed in the EMSAIE received in the Association Response frame during the MP's Initial EMSA Authentication.

— All other fields shall be set to zero.

8.8.10.3.1.2.2 Mesh key holder security handshake message 2

Mesh key holder security handshake message 2 is a mesh key holder security establishment EMSA mesh action frame with the following contents:

The MAC address of the MP shall be asserted in the DA field of the message header.

The MAC address of the MKD shall be asserted in the SA field of the message header.

The Mesh ID IE shall contain the Mesh ID as configured in dot11MeshID.

The MKDDIE shall contain the MKDD-ID as configured in dot11MeshKeyDistributorDomainID. The Mesh Security Configuration field shall be set to zero.

The MKHSIE shall be set as follows:

— MA-Nonce, MA-ID, and MKD-ID shall be set to the values contained in handshake message 1.

— MKD-Nonce shall be set to a value chosen randomly by the MKD, following the recommendations of 8.5.8.

— The Transport Type Selector field shall be set to the value contained in handshake message 1.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 3, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MP MAC address
- MKD MAC address
- Handshake sequence number (1 octet), set to the value 2.
- Contents of the Mesh ID IE, from the element ID to the end of the Mesh ID IE.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MKHSIE, from element ID through MIC Control fields, and omitting the MIC field.

8.8.10.3.1.2.3 Mesh key holder security handshake message 3

Mesh key holder security handshake message 3 is a mesh key holder security establishment EMSA mesh action frame with the following contents:

The MAC address of the MKD shall be asserted in the DA field of the message header.

The MAC address of the MP shall be asserted in the SA field of the message header.

The Mesh ID IE shall contain the Mesh ID IE received in handshake message 2.

The MKDDIE shall contain the MKDDIE received in handshake message 2.

The MKHSIE shall be set as follows:

— MA-Nonce, MKD-Nonce, MA-ID, MKD-ID, and Transport Type Selector shall be set to the values contained in handshake message 2.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 3, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MP MAC address
- MKD MAC address
- Handshake sequence number (1 octet), set to the value 3.
- Contents of the Mesh ID IE, from the element ID to the end of the Mesh ID IE.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MKHSIE, from element ID through MIC Control fields, and omitting the MIC field.

8.8.10.3.2 Mesh key transport protocol

The mesh key transport protocol describes the method by which the MKD securely transmits a derived PMK-MA to a MA, along with key context and additional related information. An additional management protocol permits the MKD to request the MA delete a key that has previously been delivered.

Three protocols are defined for mesh key delivery and management, each consisting of 2 messages. The pull protocol is initiated by the MA by sending a request message, followed by the MKD delivering the PMK-MA. The push protocol is initiated by the MKD delivering (unsolicited) the PMK-MA, followed by the MA sending a confirmation message. Finally, the key delete protocol is initiated by the MKD by sending a message requesting key deletion to the MA, followed by the MA sending a confirmation message.

The MA and MKD maintain separate key replay counters for use in these three protocols. In the pull protocol, the MA's key replay counter is used to protect the first message, which the MA sends. In both the push protocol and the key delete protocol, the MKD's key replay counter is used to protect the first message, which the MKD sends.

In each protocol, prior to sending the first message, the sender shall increment the value of its replay counter. Upon receiving the first message, the recipient shall verify that the replay counter value contained in the first message is a value not yet used by the sender in a first message. If the replay counter value has been previously used, the message shall be discarded. Thus, MA and MKD must each maintain the state of two replay counters: the counter used to generate a value for first messages that it sends, and a counter used to detect replay in first messages that it receives.

Further, the second message of each protocol shall contain a replay counter value that equals the value in the first message of the protocol, to permit matching messages within a protocol instance.

8.8.10.3.2.1 Mesh key transport pull protocol

The key transport pull protocol is a two-message exchange consisting of a PMK-MA request message sent to

the MKD, followed by a key delivery sent to the MA. Both messages contain a MIC for integrity protection, and the PMK-MA being delivered is encrypted.

Mesh key transport pull message 1 is a PMK-MA request EMSA mesh action frame. The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header. Prior to constructing the message, the value of the MA's replay counter associated with the PTK-KD shall be incremented by 1.

The MKDDIE shall be configured exactly as advertised by the MA in its beacons and probe responses.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of the MA's replay counter.

— SPA shall be set to the MAC address of the MP that, during its Initial EMSA Authentication, generated the mesh key hierarchy that includes the PMK-MA being requested

— PMK-MKDName shall be set to the identifier of the key from which the PMK-MA being requested was derived.

— ANonce shall be set to zero.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Action Value field of the PMK-MA request EMSA mesh action frame, which contains the value 3.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 1, the MKD shall verify the MIC, and shall verify that the Replay counter field contains a value not previously used with the PTK-KD in a first message sent by the MA. If verified, the MKD may attempt to derive the PMK-MA for use between the MP identified by SPA and the MA that sent message 1, using the key identified by PMK-MKDName. Subsequently, the MKD constructs and sends message 2.

Mesh key transport pull message 2 is a PMK-MA delivery pull EMSA mesh action frame. The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header.

The MKDDIE shall contain the MKDDIE received in message 1.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of replay counter in message 1.

— SPA shall be set to the value contained in message 1.

— PMK-MKDName shall be set to the value contained in message 1 if an encrypted PMK-MA is included in the Encrypted Contents field. If the Encrypted Contents field is omitted, then PMK-MKDName shall be set to zero.

— ANonce shall be set to the random value that was selected by the MKD for derivation of the PMK-MKDName that was indicated in message 1. If the PMK-MKDName field is set to zero, then the ANonce shall be set to zero.

— Encrypted Contents Length field shall be set to the length in octets of the Encrypted Contents field, or shall be set to zero if the Encrypted Contents field is omitted.

— Encrypted Contents shall be set as follows:

  • If the MKD does not have a PMK-MA to send to the MA (e.g., it was unable to derive the key), the Encrypted Contents field shall be omitted.
  • If the MKD is sending an PMK-MA to the MA, then the Encrypted Contents field shall contain the concatenation: key_data = {PMK-MA || PMK-MAName || Lifetime KDE}.
  • Lifetime KDE is defined in Figures 144 and 149. The KDE contains a 4-octet value containing the number of seconds remaining in the lifetime of the PMK-MA.
  • If the MIC algorithm is 1 (HMAC-SHA1-128), then the concatenation key_data shall be encrypted using NIST AES Key Wrap algorithm, with the KEK-KD, as defined in RFC 3394, prior to being inserted in the Encrypted Contents field.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address
  • MKD MAC address
  • Action Value field of the PMK-MA delivery pull EMSA mesh action frame, which contains the value 4.
  • Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
  • Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 2, the MA shall verify the MIC, and shall verify that the Replay counter field contains the value as in message 1.

8.8.10.3.2.2 Mesh key transport push protocol

The key transport push protocol is a two-message exchange consisting of a PMK-MA delivery message sent to the MA, followed by a confirmation message sent in reply. Both messages contain a MIC for integrity protection, and the PMK-MA being delivered is encrypted.

Mesh key transport push message 1 is a PMK-MA delivery push EMSA mesh action frame. The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header. Prior to constructing the message, the value of the MKD's replay counter associated with the PTK-KD shall be incremented by 1.

The MKDDIE shall contain the MKDD-ID as configured in dot11MeshKeyDistributorDomainID.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of the MKD's replay counter.

— SPA shall be set to the MAC address of the MP that, during its Initial EMSA Authentication, generated the mesh key hierarchy that includes the PMK-MA being delivered

— PMK-MKDName shall be set to the identifier of the key from which the PMK-MA being delivered was derived.

— ANonce shall be set to the random value that was selected by the MKD for derivation of the PMK-MKDName indicated in this message

— Encrypted Contents Length field shall be set to the length in octets of the Encrypted Contents field.

— Encrypted Contents field shall contain the concatenation: key_data = {PMK-MA || PMK-MAName || Lifetime KDE}

  • Lifetime KDE is defined in Figures 144 and 149. The KDE contains a 4-octet value containing the number of seconds remaining in the lifetime of the PMK-MA.
  • If the MIC algorithm is 1 (HMAC-SHA1-128), then the concatenation key_data shall be encrypted using NIST AES Key Wrap algorithm, with the KEK-KD, as defined in RFC 3394, prior to being inserted in the Encrypted Contents field.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address
  • MKD MAC address
  • Action Value field of the PMK-MA delivery push EMSA mesh action frame, which contains the value 1.
  • Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
  • Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 1, the MA shall verify the MIC, and shall verify that the replay counter field contains a value not previously used with the PTK-KD in a first message sent by the MKD. If verified, the MA shall send a confirmation message to the MKD.

Mesh key transport push message 2 is a PMK-MA confirm EMSA mesh action frame. The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header.

The MKDDIE shall contain the MKDDIE received in message 1.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of replay counter in message 1.

— SPA, PMK-MKDName, and ANonce shall be set to the values contained in message 1.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address
  • MKD MAC address
  • Action Value field of the PMK-MA confirm EMSA mesh action frame, which contains the value 2.
  • Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
  • Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 2, the MKD shall verify the MIC, and shall verify that the Replay counter field contains the value as in message 1.

8.8.10.3.2.3 Mesh key delete protocol

The MKD may initiate the mesh key delete protocol in order to request that a previously-delivered PMK-MA be revoked. Revocation of the PMK-MA implies that the PMK-MA shall be deleted and all keys derived from the PMK-MA shall be deleted.

The key delete protocol is a two-message exchange consisting of a PMK-MA delete message sent to the MA, followed by a confirmation message sent in reply. Both messages contain a MIC for integrity protection.

Mesh key delete message 1 is a PMK-MA delete EMSA mesh action frame. The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header. Prior to constructing the message, the value of the MKD's replay counter associated with the PTK-KD shall be incremented by 1.

The MKDDIE shall contain the MKDD-ID as configured in dot11MeshKeyDistributorDomainID.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of the MKD's replay counter.

— SPA shall be set to the MAC address of the MP that, during its Initial EMSA Authentication, generated the mesh key hierarchy that includes the PMK-MA that shall be deleted.

— PMK-MKDName shall be set to the identifier of the key from which the PMK-MA that shall be deleted was derived.

— ANonce shall be set to zero.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Action Value field of the PMK-MA delete EMSA mesh action frame, which contains the value 5.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 1, the MA shall verify the MIC, and shall verify that the replay counter field contains a value not previously used with the PTK-KD in a first message sent by the MKD. If verified, the MA shall compute the value of PMK-MAName using the PMK-MKDName and SPA included in message 1. The MA shall revoke the PMK-MA named by PMK-MAName, and shall send a confirmation message to the MKD.

Mesh key delete message 2 is a PMK-MA confirm EMSA mesh action frame. The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header.

The MKDDIE shall contain the MKDDIE received in message 1.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of replay counter in message 1.

— SPA, PMK-MKDName, and ANonce shall be set to the values contained in message 1.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Action Value field of the PMK-MA confirm EMSA mesh action frame, which contains the value 2.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 2, the MKD shall verify the MIC, and shall verify that the Replay counter field contains the value as in message 1.

8.8.10.3.3 Mesh EAP message transport protocol (optional)

This optional protocol describes how the MA may initiate and perform EAP authentication with the supplicant during the supplicant MP's Initial EMSA Authentication. The use of this protocol is selected during the mesh key holder security handshake defined in 8.8.10.3.1.2 and is described by transport selector 00-0F-AC:0. When the transport selector specifies any other value, the mechanism for EAP Transport is outside the scope of this standard.

EAP, as described in RFC 3748, is a "lock-step protocol," with request messages sent from the supplicant always receiving a response from the AS. The mesh authentication message transport protocol permits transport of these request and response messages through the mesh, between the MA and the MKD.

The MA initiates 802.1X authentication with the supplicant by sending a first EAP message to the supplicant. If the MA is configured with the appropriate first EAP message to send, then the MA does so. Otherwise, the MA may request the first EAP message from the AS, using the EAP-Start indication described below. When the MA receives an EAP message from the supplicant, the MA sends an EAP Encapsulation EMSA mesh action frame to the MKD that contains the received EAP message. When the MKD has an EAP message, received from the AS and destined for the supplicant, it sends an EAP Encapsulation EMSA mesh action frame to the MA containing the EAP message.

The final EAP Encapsulation EMSA mesh action frame of a sequence will be sent by the MKD, and is given a special type to provide information to the MA. If the EAP authentication of the supplicant provided an "accept" indication to the MKD, then the MKD sends the final message with type "accept" to indicate to the MA that the supplicant should be granted access. Alternatively, if EAP authentication failed, the MKD sends the final message with type "reject" to the MA. Upon reception of an EAP Encapsulation EMSA mesh action frame of type "reject," the MA shall terminate the peer link with the supplicant.

When an EAP message is included in a EAP Encapsulation EMSA mesh action frame, it is encapsulated within one or several EAP Message IEs. The maximum length EAP message that may be transported is 2231 octets. If the EAP message has length greater than 254 octets, fragmentation is required. In such a case, the EAP message shall be separated into fragments. Each fragment shall be of length 254 octets except the last or only fragment. The maximum number of fragments shall be 9. The fragments are included in one or several EAPMIEs, which each contain a Fragment Control field value indicating the sequence of the fragment,

beginning with the value zero. Upon reception, the contents of the EAPMIE EAP Message Fragment fields are concatenated according to the order indicated in the Fragment Control fields to reconstruct an IETF 3748 EAP message.

The EAP-Start indication is sent from MA to MKD by constructing an EAP Encapsulation request message that contains only a single EAP Authentication IE and no EAP Message IEs.

8.8.10.3.3.1 EAP Encapsulation request message

An EAP Encapsulation mesh action message with EAP message type request is sent from MA to MKD, either to transport an EAP message from the supplicant, or to request the AS to initiate EAP authentication ("EAP-Start").

EAP Encapsulation request message is an EAP Encapsulation EMSA mesh action frame. The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header. The contents of the EAP Authentication IE are as follows:

— EAP Message Type shall be set to 1 to indicate "request".

— Message Token shall be set to a unique nonce value chosen by the MA.

— SPA shall be set to the MAC address of the supplicant mesh point that is participating in EAP authentication.

— Message Fragments shall indicate the number of EAP Message IEs that are included in this EAP Encapsulation request message.

  • If the MA is sending an "EAP-Start" notification, the Message Fragments field shall be set to zero, and no EAP Message IEs are included in the EAP Encapsulation request message.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 1, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address
  • MKD MAC address
  • Contents of the EAP Authentication IE, from element ID through MIC Control fields, and omitting the MIC field.
  • If present, contents of all EAPMIEs, from the element ID field of the first EAPMIE, through the EAP Message Fragment field of the last EAPMIE. The EAPMIEs shall be ordered with increasing Fragment Control field values.

Zero or more EAP Message IEs may be present. If present, the contents of each EAP Message IE are as follows:

— Fragment Control contains the number of the fragment contained in the EAP Message Fragment field.

— EAP Message Fragment contains an EAP message with format as defined in IETF RFC 3748, or portion thereof. The maximum size of the EAP message portion is 254 octets.

Upon receiving a request message, the MKD shall verify the MIC, and store the Message Token for use in constructing the response message.

8.8.10.3.3.2 EAP Encapsulation response message

An EAP Encapsulation mesh action message with EAP message type response, accept, or reject is sent from MKD to MA, to transport an EAP message from the AS, and, in the final response message of a sequence,

provide an indication of the success of the EAP authentication.

EAP Encapsulation response message is an EAP Encapsulation EMSA mesh action frame. The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header. The contents of the EAP Encapsulation IE are as follows:

— EAP Message Type shall be set as follows:

  • If this is the final message of the sequence, and the EAP authentication of the supplicant resulted in an "accept" indication, EAP Message Type shall be set to 2, to indicate "accept."

  • If this is the final message of the sequence, and the EAP authentication of the supplicant resulted in a "reject" indication, EAP Message Type shall be set to 3, to indicate "reject."

  • Otherwise, EAP Message Type shall be set to 11, to indicate "response."

— Message Token shall be set to the value contained in the request message to which this response corresponds.

— SPA shall be set to the value contained in the request message to which this response corresponds.

— Message Fragments shall indicate the number of EAP Message IEs that are included in this EAP Encapsulation request message.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.64.

— The Information element count field of the MIC control field shall be set to 1, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address

  • MKD MAC address

  • Contents of the EAP Authentication IE, from element ID through MIC Control fields, and omitting the MIC field.

  • Contents of all EAPMIEs, from the element ID field of the first EAPMIE, through the EAP Message Fragment field of the last EAPMIE. The EAPMIEs shall be ordered with increasing Fragment Control field values.

One or more EAP Message IEs shall be present. The contents of each EAP Message IE are as follows:

— Fragment Control contains the number of the fragment contained in the EAP Message Fragment field.

— EAP Message Fragment contains an EAP message with format as defined in IETF RFC 3748, or portion thereof. The maximum size of the EAP message portion is 254 octets.

Upon receiving a response message, the MA shall verify the MIC, and verify that the Message Token received in the message matches the value sent in the most recent request message. If the final response message receive has type "reject," the MA shall terminate the peer link with the supplicant.

9. MAC sublayer functional description

Add the following after Clause 9.9.1.6:

9.9.1.7 Recommendations for use of EDCA in Mesh Points (Informative)

9.9.1.7.1 General

EDCA is used as the basis for the WLAN Mesh Media Access Mechanism. A set of recommendations on how to optimize EDCA for Mesh Point without changing the basic Media Access Mechanism is presented here.

9.9.1.7.2 Forwarding and BSS Traffic Interaction

Since Mesh Point is only a logical entity, it is possible that it is physically collocated with an Access Point within one single device: a device that has both AP and MP functionalities collocated in it is called a Mesh AP. It is also possible to have a Mesh Point implemented on a device that also acts as an Application End Point, that is, in addition to participating in the mesh and forwarding frames on behalf of other Mesh Points, it also generate its own application traffic. In both of these cases, one single device has to forward a mixture of mesh traffic (with 4-address frame formats) and BSS traffic (with 3-address frame formats). How these two different kinds of traffic are being handled within a single device can have a profound impact on overall network performance. For example, the forwarding traffic tends to traverse the network through multi-hop path and hence may have already consumed significant amount of network resources before reaching a certain Mesh Point. Dropping such traffic basically renders all such resource as being wasted. The frames originated from a local BSS and destined to the mesh have only traversed one hop, and so dropping such traffic only has local impact. It is also possible that an aggressive STA with heavy traffic backlog in the BSS can potentially starve the neighboring mesh points in the network. Such traffic prioritization may have different implication from the point of view of fairness and the prioritization policy may very well depend on the mesh network deployment scenario and the business model used. There are many implementation choices as how to best support such traffic prioritization within a single device like MAP. For example, one may choose to employ multiple radios to separate the BSS traffic and mesh forwarding traffic into different radios operating at different channels. Such choice is entirely implementation matter, outside the scope of this specification, but it is highly recommended that consideration is taken into account to separate BSS traffic and mesh forwarding traffic as much as possible and regulate the interaction between the traffic when complete separation is not possible. For example, it is recommended in Clause 11A.7 that a BSS traffic rate control mechanism be used in conjunction with intra-mesh congestion control to ensure the overall network performance.

EDITORIAL NOTE—Recommendation for use of EDCA in MPs moved from 9.9.1.7 to Annex T per CID 3100

Add Insert the following new clause after Clause 9.13:

9.14 MDA (Optional)

Mesh Deterministic Access (MDA) is an optional access method that allows supporting MPs that support MDA to access the channel at selected times with lower contention than would otherwise in selected timesbe possible. This specification standard does not require all MPs to use MDA. MDA can be used by a subset of MPs in a WLAN Mesh. However, MDA connections can only be setup among MDA-supporting MPs. The performance of MDA may be impacted by devices that do not respect MDA reservations. MDA sets up time periods in mesh neighborhoods when a number of MDA-supporting MPs that may potentially interfere with each others others' transmissions or receptions are set to not initiate any transmission sequences. For each such time period, supporting MPs that set up the state for the use of these time periods are allowed to

access the channel using MDA access parameters (CWMin, CWMax, and AIFSN). In order to use the MDA method for access, an MP must shall be a synchronizing MP. The MDA method is described in detail below.

9.14.1 MDA opportunity (MDAOP)

An MDAOP is a period of time within every Mesh DTIM interval that is set up between a transmitter and a receiver such that the following are satisfied. Once an MDAOP is setup,

— The transmitter that owns the MDAOP uses CSMA/CA and backoff to obtain a TXOP as described in 11e 9.9.1 and using parameters MDACWmin, MDACWmax, and MDAIFSN. The ranges of values allowed for MDACWMin, MDACWMax, and MDAIFSN parameters are identical to that those allowed for EDCA.

— Once the setup of an MDAOP is advertised, all MPs supporting MDA that hear these advertisements except the transmitter that set up the MDAOP are required to not initiate any new transmission during the TXOP initiated in the MDAOP. This can be done by setting their NAVs for the duration of the MDAOP at the beginning of the MDAOP, or by using enhanced carrier sensing (ECS) that achieves the same result.

9.14.2 MDAOP Setssets

A set of MDAOPs may be setup for unicast individually addressed transmissions from a transmitter to a receiver by the transmitter. Such a set is identified by a unique ID called the MDAOP Set ID. The MDAOP Set ID has to be unique for a transmitter, so that the MDAOP set ID and the transmitter (or set owner) MAC address uniquely identifies an MDAOP set in the mesh. The MDAOP set ID is a handle that allows operation such as setup and teardown to be conducted together for the entire set of MDAOPs in an MDAOP set. An example use of MDAOP set concept is to establish an MDAOP set for a single QoS flow.

MDAOP set ID is an 8 8-bit unsigned number. The special value of MDAOP set ID, when all bits are set to 1 1, is reserved to mean all MDAOPs.

9.14.3 MDA TXOP

Any TXOP that is obtained by an MP by using MDA parameters in an MDAOP is called an MDA TXOP. An MDA TXOP is required to end within the MDAOP in which it was obtained. Thus, an MDA TXOP ends latest either by MDA TXOP limit time after it began or by MDAOP end time, whichever is earlier.

9.14.4 Neighborhood MDAOP Times times at an MP

At a neighborhood centered at an MP, all the TX-RX times reported by its neighbors (in their MDAOP advertisements) form a set of MDAOPs that are already being used in the neighborhood. No new MDAOPs may be set up by the MP during these times. These times are referred to as Neighborhood MDAOP times for the MP. In effect, Neighborhood MDAOP Times at an MP include all MDAOPs for which the MP and its neighbors are either the transmitters or receivers.

9.14.5 Neighbor MDAOP Interfering Times interfering times for an MP

The Interfering times reported by an MP in its MDAOP advertisements are times that may not be used for a new MDAOP with that specific MP. While the MP itself is not the transmitter or receiver in an MDAOP during these times, one of its neighbor's neighbors is. Any new MDAOPs to the MP during these times may experience interference. However, new MDAOPs may be setup with another MP during these interfering times. Thus, for every neighbor, there is a set of times that are interfering. These times are referred to as Neighbor MDAOP interfering times for that neighbor.

9.14.6 MDA Access Fraction access fraction (MAF)

The MDA access fraction at an MP is the ratio of the total duration of its 'Neighborhood MDAOP Times' (see definition above) in a Mesh DTIM interval to the duration of the Mesh DTIM interval. This parameter may be used to limit the use of MDA in a neighborhood centered at an MP to a certain fraction of the total channel time. The maximum value for MAF that is allowed at an MP is specified by the dot11MAFlimit parameter.

The dot11MAFlimit is copied in into the MDA Access Fraction Limit field of the MDAOP Advertisements IEinformation element. Before attempting to set up an MDAOP Set with a neighbor, an MP is required to make sure ensure that the new MDAOP set does not cause the MAF of any of its neighbors to exceed their MAF Limit. An MDAOP setup request may be refused by the intended receiver if the MAF limit of any of its own neighbors is exceeded due to the new setup.

9.14.7 Action Frames frames for MDAOPs setup, teardown, and MDAOP advertisements

The IEs that are used for MDA may be carried in action frames. The format of such actions frames that carry the IEs and the MDA frame is described in clause 7.4.9.9.

9.14.8 MDAOP Setup Proceduresetup procedure

The setup of an MDAOP set is initiated by the intended transmitter, and is accepted/rejected by the intended receiver. Once accepted, the transmitter is referred to as the owner of the MDAOP. The setup procedure for an MDAOP set is as follows:

a) The MP that intends to be the transmitter in a new MDAOP set builds a map of Neighborhood MDAOP times in the Mesh DTIM interval after hearing Advertisements from all of its neighbors that have MDA active. If no advertisement was heard from a neighbor in the last dot11MDAdvertPeriodMax, the MP may request the neighbor for MDAOP Advertisement.

b) The intended transmitter MP also considers the Neighbor MDAOP Interfering Times of the intended receiver.

c) Based on traffic characteristiccharacteristics, it then chooses MDAOP locations and durations in the Mesh DTIM interval that do not overlap with either its Neighborhood MDAOP Times or the Neighbor MDAOP Interfering Times of the intended receiver. It also avoids using times that are known to it as being used by itself or one of its neighbors for other activities such as beacons beacon transmissions.

d) It then verifies that the new MDAOP Set will not cause the MAF limit to be crossed for any of its neighbors. If MAF limit would be crossed for any of its neighbors, due to the new MDAOP Set, it suspends the setup process.

e) If the MAF limits at all neighbors are respected despite the new MDAOP set, it transmits an MDAOP Setup request IE information element to the intended receiver with chosen MDAOP locations and durations.

f) The receiver of the MDAOP Setup Request IE information element checks to see if the MDAOP times have any overlap with its Neighborhood MDAOP Times. The receiver also checks if the new MDAOP Set will cause the MAF limit to be crossed for any of its neighbors. The MDAOP Setup Reply IE information element is used to reply to a setup request.

g) The receiver rejects the setup request if there are any overlaps of the requested MDAOP set with its Neighborhood MDAOP Times, or other times that it knows are set to be used by itself or its neighbors for activities such as beacon transmissions. It may suggest alternate times by including the optional field Alternate Suggested Request IE information element in the MDAOP Setup Reply element.

h) The receiver also rejects the setup request if the MAF limit of itself or any of its neighbors will be exceeded due to the new setup.

i) If suitable, the receiver accepts the setup.

j) After successful setup, both the MDAOP owner (the transmitter) and the receiver advertise the MDAOP Set times in the TX-RX Times Report field of the MDAOP Advertisement IEinformation element.

9.14.9 MDAOP Advertisements advertisements

Every MP that has MDA active is required to advertise TX-RX and Interfering times using the MDAOP Advertisements IEinformation element, at least once in dot11MDAdvertPeriodMax. These advertisements are always transmitted in broadcast group addressed frames; either in beacons Beacon frames or MDA action frames. The advertised times include:

a) TX-RX times report:

1) All MDAOP times for which the MP is the transmitter or the receiver.

2) All other times that it knows are busy/reserved such that it is either the transmitter or the receiver. A non exhaustive list includes expected HCCA times for an MAP and self or neighbor's expected beacon times.

b) Interfering times report:

1) All TX-RX times reported by the MP's neighbors so that the MP is neither the transmitter nor the receiver during those times.

9.14.10 MDAOP Set Teardownset teardown

An MDAOP set is successfully torn down, down once both the transmitter and the receiver stop advertising the set in their TX-RX times. Either the transmitter or the receiver may indicate a teardown by transmitting the MDAOP Set Teardown IE information element to the other communicating end (transmitter or the receiver). The teardown is assumed successful once the ACK is received, or maximum retry attempts are exceeded.

The transmitter assumes a successful teardown and stops using or advertising (in TX-RX times report) an MDAOP set if any of the following happens:

a) Its MDAOP Set Teardown IE information element is successfully Acked.

b) The maximum retries for the teardown IE information element it is transmitting are exceeded.

c) The receiver's advertisement does not include the MDAOP set

d) The receiver is inactive unreachable for greater than dot11MDAOPtimeout time

The receiver assumes a successful teardown and stops advertising an MDAOP set if any of the following happens:

a) Its MDAOP Set Teardown IE information element is successfully Acked.

b) The maximum retries for the teardown IE information element it is transmitting are exceeded.

c) The transmitter's advertisement does not include the MDAOP set.

d) The transmitter is unreachable inactive for greater than dot11MDAOPtimeout time

The interfering times are directly derived from neighbors' TX-RX times report. The interfering times report reflects the latest TX-RX times reports from the neighbors.

9.14.11 Access during MDAOPs

MPs that have MDA active maintain the Neighborhood MDAOP Times state of MDAOPs when either they or their neighbors are transmitters or receivers. The access behavior for such MPs during the Neighborhood MDAOP Times is described as below.

a) Access by MDAOP owners:

If the MP is the owner of the MDAOP, it attempts to access the channel using CSMA/CA and backoff using MDACWmax, MDACWmin, and MDAIFSN parameters. If the MP successfully captures an MDA TXOP, before the end of its MDAOP, it may transmit until the end of the MDAOP or until a time less than MDA TXOP limit from the beginning of the MDA TXOP, whichever is earlier. The retransmission rules for access in an MDAOP are the same as that of EDCA. Specifically, if there is loss inferred during the MDA TXOP, retransmissions require that a new TXOP be obtained using the MDA access rules in the MDAOP. No MDA TXOPs may cross MDAOP boundaries.

If the MP intends to end the TXOP with enough time before the end of the MDAOP, it is responsible for relinquishing the remaining MDAOP time by using any of the methods that reduce NAV as defined in 802.11NAV.

The MP shall not use MDA access parameters to access the channel outside of MDAOPs owned by it.

b) Access by non-owners of MDAOP:

At the beginning of an MDAOP that is not owned by the MP (it is not the transmitter) but is part of the Neighborhood MDAOP Times, the MP sets its NAV to the end of the MDAOP. Instead of setting the NAV, it can also use other means to not initiate any new transmission sequence during the MDAOP. The NAV setting may be reduced to a shorter time on the receipt of either a QoS+CF-Poll frame with a duration of 0 or a CF-end frame. MPs may then attempt channel access even before the end of MDAOP through access mechanisms other than MDA.

10. Layer management

10.3 MLME SAP interface

EDITORIAL NOTE—11ma ended with 10.3.29, 11k added 30-33, 11r added 34-35, 11y added 36, next is 37

Add Insert the following new clause after Clause 10.3.2936:

10.3.37 PassivePeerLinkOpen

The following primitives describe how a mesh entity passively starts a peer link establishment process.

10.3.37.1 MLME-PassivePeerLinkOpen.request

10.3.37.1.1 Function

This primitive requests that the mesh entity start the link establishment protocol passively.

10.3.37.1.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-PassivePeerLinkOpen.request(
RetryTimeout,
OpenTimeout,
CancelTimeout,
MaxReqs,
CapabilityInformation,
ListenInterval
)

Name	Type	Valid range	Description
RetryTimeout	Integer	1	Specifies the time limit (in TU) the mesh entity will wait waits for a confirmation from the neighbor mesh entity before retrying trying to send another peer link open request.
OpenTimeout	Integer	1	Specifies the time limit (in TU) the mesh entity will wait waits for the peer link open request from the neighbour neighbor mesh entity after receiving the corresponding confirm, before declaring the failure of the link establish-mentestablishment failure.
CancelTimeout	Integer	1	Specifies a time limit (in TU) after which the link will be is cancelled completely after the mesh entity sets the link status as holding.
MaxReqs	Integer	1	Specifies the maximal number of retries the mesh entity will issue issues to the neighbor mesh entity before declaring link establishment failure.
CapabilityInformation	As defined in frame format	As defined in frame format	Specifies the requested operational capabili-ties to the neighbor mesh entity.
ListenInterval	Integer	0	Specifies the number of beacon intervals that can pass before the mesh entity awakens and listens for the next beacon.

Additional parameters needed to perform PassivePeerLinkOpen procedure are not included in the primitive parameter list since the MLME already has that data (maintained as internal state).

10.3.37.1.3 When generated

This primitive is generated when the mesh entity wishes to establish a link with a neighbor mesh entity, but does not specify a particular neighbor.

10.3.37.1.4 Effect of receipt

This primitive initiates a peer link establishment procedure. The MLME subsequently issues an MLME-PassivePeerLinkOpen.confirm that reflects the results.

10.3.37.2 MLME-PassivePeerLinkOpen.confirm

10.3.37.2.1 Function

This primitive reports the results of a passive open attempt.

10.3.37.2.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-PassivePeerLinkOpen.confirm(
ResultCode,
Local Link ID
)

Name	Type	Valid range	Description
ResultCode	Enumeration	SUCCESS, INVALID_PARAMETERS, TIMEOUT, FAILED_OUT_OF_MEMORY	Indicates the result of the MLME-PassivePeerLinkOpen.request.
Local Link ID	Integer	1—2^{32}-1	Specifies the pseudo-random number generated by the local mesh entity in the an effort of identifying the link instance about to be established with a neighbor mesh entity

10.3.37.2.3 When generated

This primitive is generated as a result of an MLME-PassivePeerLinkOpen.request.

10.3.37.2.4 Effect of receipt

The SME is notified of the results of the passive open procedure.

10.3.38 ActivePeerLinkOpen

The following primitives describe how a mesh entity actively starts a peer link establishment procedure with a specified peer MAC entity that is within a mesh entity.

10.3.38.1 MLME-ActivePeerLinkOpen.request

10.3.38.1.1 Function

This primitive requests that the mesh entity start the link establishment procedure actively with a specified peer MAC entity that is within a mesh entity.

10.3.38.1.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-ActivePeerLinkOpen.request(
PeerAddress,
RetryTimeout,
OpenTimeout,
CancelTimeout,
MaxReqs,
CapabilityInformation,
ListenInterval
)

Name	Type	Valid range	Description
PeerAddress	MACAddress	Any valid individual MAC address	Specifies the address of the peer MAC entity with which to perform the link establishment process.
RetryTimeout	Integer	1	Specifies the time limit (in TU) the mesh entity will wait waits for a confirmation from the neighbor mesh entity before retrying trying to send another peer link open request.
OpenTimeout	Integer	1	Specifies the time limit (in TU) the mesh entity will wait waits for the peer link open request from the neighbour neighbor mesh entity after receiving the corresponding confirm, before declaring the failure of the link establishmentestablishment failure.
CancelTimeout	Integer	1	Specifies a time limit (in TU) after which the link will be is cancelled completely after the mesh entity sets the link status as holding.
MaxReqs	Integer	1	Specifies the maximum number of retries the mesh entity will issue issues to the neighbor mesh entity before declaring link establishment failure.
CapabilityInformation	As defined in frame format	As defined in frame format	Specifies the requested operational capabilities to the neighbor mesh entity.
ListenInterval	Integer	0	Specifies the number of beacon intervals that can pass before the mesh entity awakens and listens for the next beacon.

Additional parameters needed to perform active open procedure are not included in the primitive parameter list since the MLME already has that data (maintained as internal state).

10.3.38.1.3 When generated

This primitive is generated when the mesh entity wishes to establish a link with a neighbor mesh entity.

10.3.38.1.4 Effect of receipt

This primitive initiates a peer link establishment procedure. The MLME subsequently issues an MLME-ActivePeerLinkOpen.confirm that reflects the results.

10.3.38.2 MLME-ActivePeerLinkOpen.confirm

10.3.38.2.1 Function

This primitive reports the results of an active open attempt.

10.3.38.2.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-ActivePeerLinkOpen.confirm(
PeerAddress,
ResultCode,
Local Link ID
)

Name	Type	Valid range	Description
PeerAddress	MACAddress	Any valid individual MAC address	Specifies the address of the peer MAC entity with which to perform the link establishment process.
ResultCode	Enumeration	SUCCESS, DUPLICATED, INVALID_PARAMETERS, TIMEOUT, FAILED_OUT_OF_MEMORY	Indicates the result of the MLME-ActivePeerLinkOpen.request.
Local Link ID	Integer	1—2^{32}-1	Specifies the pseudo-random number generated by the local mesh entity in the an effort of identifying the link instance about to be established with a neighbor mesh entity

10.3.38.2.3 When generated

This primitive is generated as a result of an MLME-ActivePeerLinkOpen.request.

10.3.38.2.4 Effect of receipt

The SME is notified of the results of the active open procedure.

10.3.39 SignalPeerLinkStatus

The following primitives report the link status to the mesh entity as the result of peer link establishment, at the end of peer link establishment procedure.

10.3.39.1 MLME-SignalPeerLinkStatus.indication

10.3.39.1.1 Function

This primitive indicates that the mesh entity has finishes finished the link establishment procedure with a specified peer mesh entity and reports the status of the link.

10.3.39.1.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-SignalPeerLinkStatus.indication(
PeerAddress,
Local Link ID,
Peer Link ID,
Status
)

Name	Type	Valid range	Description
PeerAddress	MACAddress	Any valid individual MAC address	Specifies the address of the peer MAC entity with for which to perform the link establishment processstatus is needed.
Local Link ID	Integer	1—2^{32}-1	Specifies the pseudo-random number generated by the local mesh entity to identify this the link instance
Peer Link ID	Integer	0—2^{32}-1	Specifies the pseudo-random number generated by the peer mesh entity and received by the local mesh entity that in order to identify this the link instance. The value "0" indicates the Peer Link ID is unknown.
Status	Enumeration	SUCCESS, FAIURFAILURE-CANCELLED, FAILURE-CLOSE, FAILURE-INVALID-PARAMETERS, FAILURE-MAX-REQS, FAILURE-TIMEOUT	Indicates the result of the peer link establishment procedure

10.3.39.1.3 When generated

This primitive is generated when the mesh entity finishes the link establishment procedure.

10.3.39.1.4 Effect of receipt

This primitive enables the mesh entity to handle the link status and to end a peer link establishment procedure.

10.3.40 CancelPeerLink

This mechanism supports the process of cancelling the link with a specified peer mesh entity.

10.3.40.1 MLME-CancelPeerLink.request

10.3.40.1.1 Function

This primitive requests the link with a specified peer mesh entity be cancelled.

10.3.40.1.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-CancelPeerLink.request(
PeerAddress,
Local Link ID,
Peer Link ID
)

Name	Type	Valid range	Description
PeerAddress	MACAddress	Any valid individual MAC address	Specifies the address of the peer MAC entity with which to perform the link establishment process.
Local Link ID	Integer	1—2^{32}-1	Specifies the pseudo-random number generated by the local mesh entity to identify this the link instance
Peer Link ID	Integer	0—2^{32}-1	Specifies the pseudo-random number generated by the peer mesh entity and received by the local mesh entity that in order to identify this the link instance. The value "0" indicates the Peer Link ID is unknown.

10.3.40.1.3 When generated

This primitive is generated by the SME to cancel a link instance with a specified peer mesh entity.

10.3.40.1.4 Effect of receipt

This primitive sets the mesh entity to get ready to close the peer link with the specified peer mesh entity. The MLME subsequently issues a MLME-CancelPeerLink.confirm to reflect the results.

10.3.40.2 MLME-CancelPeerLink.confirm

10.3.40.2.1 Function

This primitive reports the result of cancel link request.

10.3.40.2.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-CancelPeerLink.confirm(
PeerAddress,
Local Link ID,
Peer Link ID,
Result codes
)

Name	Type	Valid range	Description
PeerAddress	MACAddress	Any valid individual MAC address	Specifies the address of the peer MAC entity with which to perform the link establishment process.
Local Link ID	Integer	1—2^{32}-1	Specifies the pseudo-random number generated by the local mesh entity to identify this the link instance
Peer Link ID	Integer	0—2^{32}-1	Specifies the pseudo-random number generated by the peer mesh entity and received by the local mesh entity that in order to identify this the link instance. The value "0" indicates the Peer Link ID is unknown.
Result codes	Enumeration	SUCCESS FAILURE-NOT-FOUND	Indicate the result of cancel link request

Either or both of Local Link ID and Peer Link ID fields can be **null** if the SME doesn't know the values yet.

10.3.40.2.3 When generated

This primitive is generated by the MLME as the result of an MLME-CancelPeerLink.request.

10.3.40.2.4 Effect of receipt

The SME is notified of the results of the cancel link procedure.

10.3.41 Mesh Layer Management layer management (Informative)

The mesh MAC defines a set of protocol independent layer entities, service access points (SAPs) and management objects which support implementation of path selection and forwarding protocols. The layer entities are responsible for:

a) Re-transmission Forwarding and filtering of unicast individually addressed and multicast/broadcast group addressed MPDUs

b) Maintenance of the information required to make re-transmission forwarding and filtering decisions

c) Management of the above

The re-transmission forwarding of the MPDUs is handled by the mesh MAC. The maintenance of the information needed to make re-transmission forwarding and filtering decisions is handled by the MAC Layer Management Entity (MLME) and the Station Management Entity (SME). The management of these is handled by the SME.

The re-transmission forwarding and filtering processing are time critical functions that are in the data path. The maintenance of the information needed to make re-transmission forwarding and filtering decisions is not time critical and only impacts the performance of the path selection protocol.

Figure s94 illustrates the functional architecture provided by the Extensible Routing Framework entities and the principles of operation.

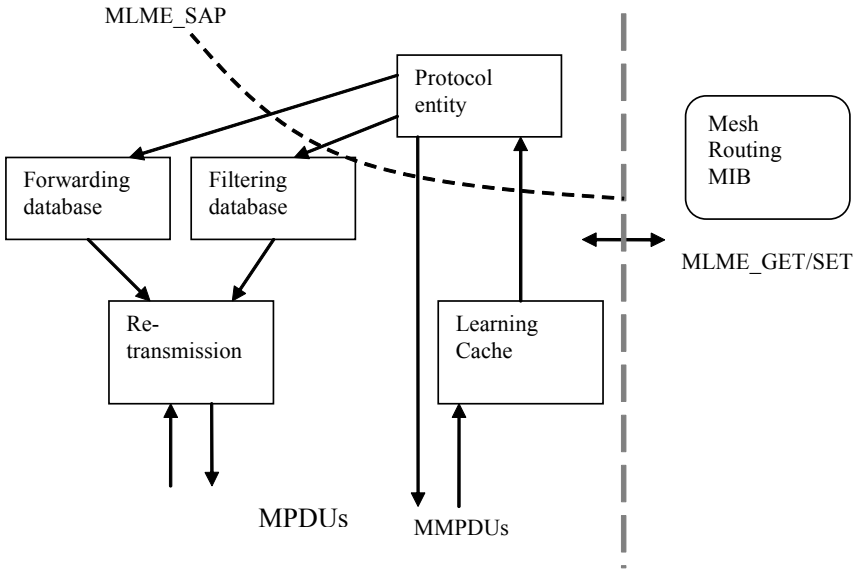

Figure s94—Extensible Routing Framework system architecture

10.3.41.1 Principles of Operationoperation

The MPDUs are received and re-transmitted by the MAC Re-transmission forwarding entity. The Protocol entity is responsible for implementing the routing algorithm. It uses the internal MLME_SAP interface to send MMPDUs for route discovery and maintenance. The received MMPDUs are cached in the Learning cache. The Protocol entity retrieves data from the cache. Its internal algorithm is used to compute routes and maintain network topology. The Protocol entity adds and removes entries in the Forwarding and Filtering databases based on the internal computations. The Mesh Routing MIB is used to manage the entities that comprise the Extensible Routing Framework.

10.3.41.2 Inter-Layer Managementlayer management

Figure 11 in IEEE Std. 802.11 Figure s95 needs to be modified to include illustrates the following entities supporting routing, as illustrated in Figure s95:

d) Re-transmit process (forwarding)
e) Filtering database
f) Forwarding database
g) Learning cache
h) Protocol entity

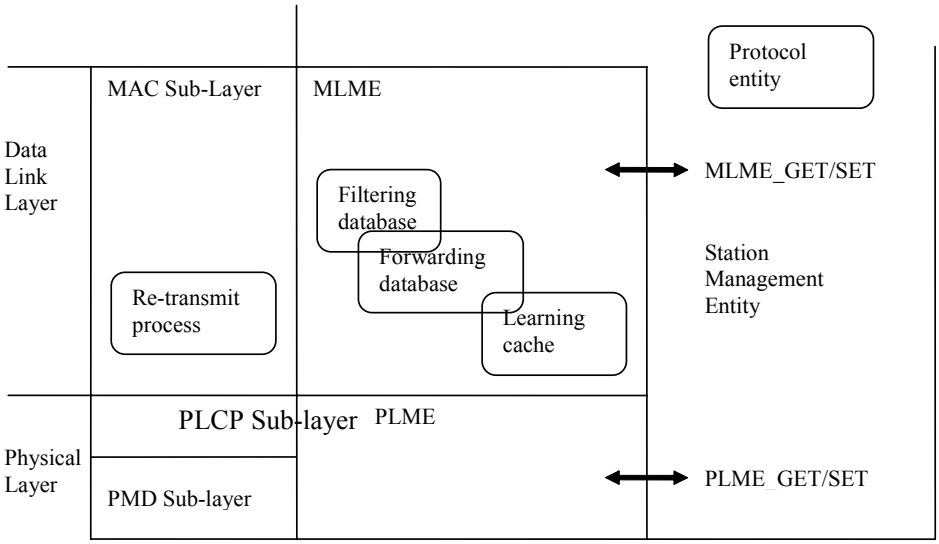

Figure s95—Inter-layer management entities and their relationship and service access points (SAPs) used for internal communication.

10.3.41.3 Re-transmit Process

When an IEEE 802.11 MPDU is received, its header is examined to determine the destination. If the destination is local, it is forwarded to the upper layer protocol. If its destination is not local, it is re-transmitted. The re-transmit process uses the information in the forwarding table to determine the address of the next hop destination.

10.3.41.4 Filtering database

The filtering decisions are based on the destination and source addresses and any group addressing policy.

10.3.41.5 Forwarding database

The forwarding database is used for making the re-transmission forwarding decision for unicast individually addressed MPDUs.

10.3.41.6 Learning cache

The learning cache keeps track of the protocol specific topology metrics collected from the MMPDUs received. This is a simple cache of the received protocol MMPDUs. It is maintained by the MLME and can be periodically queried by the protocol entity.

10.3.41.7 Protocol entity

The protocol entity implements the specific path discovery protocol logic, sends MMPDUs, queries the learning database, and updates the forwarding and filtering databases.

10.3.41.8 Service Primitives

The SME uses the MLME SAP interface to gain access to the entities needed to control the data path processing, send MMPDUs, retrieve status based on the received MMPDUs. These service primitives are defined in addition to the standard IEEE 802.11 GET and SET primitives which operate on the Mesh Routing MIB objects. The exact implementation details of these entities are beyond the scope of this standard.

10.3.41.8.1 MLME-SendMeshMgmt.request

This primitive is used to send an MMPDU.

MLME-SendMeshMgmt.request(

Interval,

Count,

Action,

OUI OUI,

Contents

)

Name	Type	Valid Range	Description
Interval	Integer		Time interval for sending MMPDUs
Count	Enumeration		Number of times to sent send the MMPDU
Action	Enumeration		Management action
OUI	Octet String		IEEE OUI
Contents	Octet String		Protocol dependent information

10.3.41.8.2 MLME-SendMeshMgmt.confirm

This service primitive indicates the result of the request.

10.3.41.8.3 MLME-RecvMeshMgmt.request

This service primitive is used to retrieve one or more received MMPDUs.

MLME-RecvMeshMgmt.request (

Count,
Action,
OUI,
Vendor specific contents

)

Name	Type	Valid Range	Description
Count	Enumeration		Number of items to retrieve
Action	Enumeration		Management action
OUI	Octet String		IEEE OUI
Contents	Octet String		Protocol dependent information

10.3.41.8.4 MLME-RecvMeshMgmt.confirm

This service primitive indicates the result of the request.

10.3.41.8.5 MLME-PathAdd.request

This service primitive is used to add a path to the forwarding table.

```
MLME-PathAdd.request(
                        OUI,
                        Id,
                        Dest,
                        Gw,
                        Metrics
                    )
```

Name	Type	Valid Range	Description
OUI	OctetString		IEEE OUI
Id	OctetString		Protocol Identifier
Dest	MACAddress		Destination MAC address matched with the contents of DA field
Gw	MACAddress		Next hop MAC address matched with the contents of the RA field
Metrics	OctetString		Re-transmission forwarding metrics used to rank alternate routes. May be matched with the user priority bits for QoS – type routing.

10.3.41.8.6 MLME-PathAdd.confirm

This service primitive indicates the result of the request.

10.3.41.8.7 MLME-PathRemove.request

This service primitive is used to remove a path from the forwarding table.

```
MLME-pathRemove.request(
                        OUI,
```

Id,
Dest,
Gw,
Metrics
)

Name	Type	Valid Range	Description
OUI	OctetString		IEEE OUI
Id	OctetString		Protocol Identifier
Dest	MACAddress		Destination MAC address matched with the contents of DA field
Gw	MACAddress		Next hop MAC address matched with the contents of the RA field
Metrics	OctetString		Re-transmission forwarding metrics used to rank alternate routes. May be matched with the user priority bits for QoS – type routing.

10.3.41.8.8 MLME-PathRemove.confirm

This service primitive indicates the result of the request.

11. MLME

11.9 DFS procedures

11.9.7 Selecting and advertising a new channel

Add Insert the following new clause after Clause 11.109.7.2:

11.9.7.3 Selecting and advertising a new channel in a WLAN Meshmesh

If a mesh point an MP detects the need to switch the channel of a logical radio interface PHY (e.g., due to regulatory requirement for radar avoidance), the mesh point must MP shall inform neighboring mesh points MPs to which an active association exists a peer link has been established on the logical radio interface PHY of the need to channel switch. The mesh point MP may request other mesh points MPs or STAs to scan and provide measurements and/or run other algorithms to make a choice on which candidate channel to switch to (the detailed algorithm for choosing an appropriate channel is beyond the scope of this specificationstandard).

Once the mesh point MP identifies the candidate channel to switch its logical radio interface PHY to, it creates a new candidate channel precedence indicator value by adding a pseudo-random number (in the range TBD) to the current channel precedence value. The mesh point MP then executes the UCG switch procedure described in Clause 11A.1.7.5.

Insert the following new clause after 11:

11A. WLAN Mesh Networkingnetworking

11A.1 Mesh Discovery discovery and Peer Link Establishmentpeer link establishment

11A.1.1 General

Mesh discovery and peer link establishment require that mesh points MPs have sufficient information about themselves and potential neighbors. This process requires detection of potential mesh neighbors through beacons Beacons or through active scanning using Mesh Probe Requests. Mesh peer link establishment is a continuous process that entails monitoring of neighbors so as to detect and react to changes in connectivity.

11A.1.2 Use of Mesh Identifiermesh identifier

The Mesh Identifier is used as a shorthand for an established mesh network with known properties, created by a known administrative authority. The Mesh ID may be installed in mesh capable devices by a variety of means, all outside the scope of this documentstandard. In addition, the Mesh ID may be advertised – possibly together with other data relevant for the location and identification of a mesh network – by a variety of means, all outside the scope of this standard. In the simplest case, the Mesh ID is set by the user, e.g., "Mike's Mesh".

NOTE--Conceptually, the Mesh ID is similar in purpose to an SSID, which is used to allow simple STAs to identify candidate APs with which to associate. Given that SSIDs are used in STA implementations for AP discovery, to enable MP-to-MP discovery in a WLAN mesh while avoiding confusing non-mesh STAs, a new mesh-specific identifier is specified rather than reusing the existing overloaded SSID identifier. To avoid having STAs send association requests to non-MAP Mesh PointsMPs, a valid SSID should not be included in beacons Beacon frames sent by non-MAP Mesh PointsMPs. To avoid backward compatibility issues, rather than removing the SSID IE information element from MP (non-MAP) beacons Beacon frames, the wildcard value is used.

11A.1.3 Profiles for Extensibilityextensibility

A device must shall support at least one profile. A profile consists of:

1. A Mesh ID
2. A path selection protocol identifier
3. A path selection metric identifier

The path selection protocol and path selection metrics in use may be different for different profiles.

11A.1.4 Neighbor Discoverydiscovery

The purpose of this procedure neighbor discovery is to discover neighbor MP devices and their properties.

A configured MP, by definition, has at least one active Mesh profile.

A An MP that is not a member of any WLAN Mesh performs passive or active scanning to discover neighboring MPs. In case of passive scanning, a device shall be considered a neighbor MP if and only if all of the following conditions are met (a similar mechanism with probe response can be used for active scanning):

1. A beacon is received from that device

2. The received beacon contains a Mesh ID that matches the Mesh ID of at least one of the profiles on the MP
3. The received beacon contains a WLAN Mesh capability element (see clause 7.3.2.49) which contains
    a. A supported version number
    b. An MP active indication
    c. A path selection protocol identifier and metric identifier matching the selected profile

A neighbor MP shall also be considered a candidate peer if and only if, in addition:

4. The beacon contains an WLAN Mesh capability element (see clause 7.3.2.49) which contains a nonzero peer link available value

The MP attempts to discover all neighbor and candidate peer devices, and maintains the neighbor MP information (see clause 0T.9.1) indicating the MAC address of each device, the most recently observed link state parameters, the received channel number and with state equal to *neighbor* or *candidate peer* as determined by the rules in this sectionstate.

When mesh point MP devices are discovered, advertising a Mesh ID for which the device has a profile, the path selection protocol and metric should be checked for a match with the profile. If there is no match, the newly discovered device should be ignored.

If an MP is unable to detect any neighbor MPs, it may adopt a Mesh ID from one of its profiles, and proceed to the active state. This may occur, for example, when the MP is the first device to power on (or multiple MPs power on simultaneously). Any peer MP links will be are established later as part of the continuous mesh formation procedures.

11A.1.5 Mesh Peer Link Establishmentpeer link establishment

11A.1.5.1 Overview

The purpose of this procedure mesh peer link establishment is to establish at least one, and in many cases several, peer links with one or more peer MPMPs. MPs shall not transmit data frames or management frames other than the ones used for link establishment until the peer link has been established successfully. The MP shall drop such frames if it considers the link is not established.

A An MP must shall be able to establish at least one mesh link with a peer MP, and may be able to establish many such links simultaneously. It is possible that there are more candidate peer MPs than the device is capable of maintaining peer links with simultaneously. In this case, the The MP must shall select which MPs to establish peer links with based on some measure of signal quality or other information received from candidate neighbor MPs.

The MP shall start the peer link establishment protocol in either of the following two cases. In each case, the MP is issued a command by the IEEE 802.11 SME. In the first case, the MP has not reached the maximum number of neighbors and is willing to accept new connections. In thise case, the The command MLME-PassivePeerLinkOpen(().request causes the MP to start a link establishment protocol state machine instance and listen to incoming connection requests. In the second case, the MP agrees with the profile carried in a beacon or a probe response from a peer MP and it has not reached the maximum number of neighbors. The command MLME-ActivePeerLinkOpen(peerId).request is issued that causes the MP to initiate a link establishment protocol state machine instance with the peer MP identified as peerId and actively send connection request requests to this peer MP.

The MP shall end the protocol state machine instance in three cases. In the first case, the MP encounters failure during the peer link establishment procedure. As the a result, the MP shall send a connection close request. In the second case, the MP receives a connection close request from the peer MP. In the third case, the connection is closed because the MP receives a peer link cancel signal. The IEEE 802.11 SME shall issue a MLME-CancelPeerLink().request command. This signal can be triggered by some internal event. For instance, the MP discovers the data transmission failure. The actual specification of the internal events is out of the scope of this specificationstandard.

The MP uses three types of management frames to handle the execution of execute the protocol:

— Association Request frame including a Peer Link Open IE information element (referred to as a Peer Link Open message)

— Association Response frame including a Peer Link Confirm IE information element (referred to as a Peer Link Confirm message)

— Disassociate Disassociation frame including a Peer Link Close IE information element (referred to as a Peer Link Close message)

Peer Link Open message requests that the establishment of a link be established between the Peer Link Open sender and the receiver. The Peer Link Confirm message responds to the Peer Link Open message. The Peer Link Close message tries to close the connection between two MPs. The protocol design is based on the rule:

— A link is established when both MPs have sent and received both Peer Link Open and Peer Link Confirm messages

If one of the MPs, say for example MP A, decides to send the Peer Link Open message first, the receiving MP, say for example MP B, will respond responds by sending the Peer Link Open and Peer Link Confirm messages. Then MP A will send sends back a Peer Link Confirm message once upon receiving the Peer Link Open and Peer Link Confirm messages from MP B. In Once the end, once the MP B receives the Peer Link Confirm message from MP A, the peer link is considered estlished established between MP A and MP B.

In the case where both MP A and MP B initiate a Peer Link Open message simultaneously, both MPs shall send back a Peer Link Confirm message. Once both MPs receive the corresponding Peer Link Confirm, the link is considered established.

The protocol has a retransmission forwarding process to make sure ensure the robustness of the protocol. After the MP sends a Peer Link Open message, a timer, RetryTimer, is set. If the MP has not received any Peer Link Confirm message when the timer expires, the MP shall resend the Peer Link Open message. If the MP has resent MAX-REQS Peer Link Open messages without receiving a corresponding Peer Link Confirm, the MP declares the failure of link establishment and shall send a Peer Link Close message to close the link. If the Peer Link Confirm message is received and processed before the incoming Peer Link Open message, another timer, OpenTimer, is set to guard the time for waiting for the Peer Link Open message. If the timer expires without receiving the corresponding Peer Link Open message, the MP declares the failure of link establishment and shall send a Peer Link Close message to close the link.

When the message processing encounters any error or the MP makes the local decision to close the link, the MP shall send a Peer Link Close message to close the link. The MP shall wait for some time for the close to close the link complete. This graceful grace time period is governed by a timer—CancelTimer. Before the CancelTimer expires, the MP shall continue to respond to any additional Peer Link Open message from the peer MP by sending back the Peer Link Close message. When the CancelTimer expires, the MP eventually closes the link by releasing the local resource allocated for the link instance. This mechanism allows the protocol to recover from failure quickly.

The link establishment management messages carry the identifiers to specify the link instance between two peer MPs. Each peer MP generates a pseudo-random number to contribute to the link instance identifier. The link instance is identified by <myId, peerId, LocalLinkID, PeerLinkID>, where myId is the MAC address of

the local MP and the peerId is the MAC address of the peer MP. Link establishment management messages carry these identifiers to bind the communication with a particular instance of the link. As we'll see later, the The finite state automaton that implements the protocol has an instance as well. One automaton handles one link instance. The receiving of a Peer Link Open message with a different identifier shall cause the mesh entity to generate a new finite state automaton.

Once a peer link is successfully established, the MP with the higher MAC address is referred to as the superordinate and the MP with the lower MAC address is referred to as the subordinate for this peer link. Note that the The terms superordinate and subordinate are arbitrary and only used to uniquely distinguish each peer MP when referring to a particular peer link.

11A.1.5.2 Processing Peer Link Establishment Messagespeer link establishment messages

When receiving a peer link establishment message, the MP shall process the message and report the result of the process. Either the MP accepts the information carried in the message and reports "success" or it denies the message and reports "failure".

The MP shall reject the Peer Link Close message if it carries a mismatched instance identifier. The received instance identifier is considered a mismatch in three cases.

— In case one, the MP does not have a record of instance identifier for the peer MP yet.

— In case two, the message carries the pseudo-random value in the Local Link ID field, but it does not match the PeerLinkID recorded locally. In case three, the value in the Peer Link ID field (if it is not null) does not match the local record of LocalLinkID.

— The MP shall reject the Peer Link Close message if it carries a mismatched instance identifier. The received instance identifier is considered a mismatch in three cases. In case one, the MP does not have a record of instance identifier for the peer MP yet. In case two, the message carries the random value in the Local Link ID field, but it does not match the *PeerLinkID* recorded locally. In case three, the value in the Peer Link ID field (if it is not **null**) does not match the local record of *LocalLinkID*. In the rest of the cases, the instance identifier is considered a match, and the MP shall accept the message. Doing so, the MP shall record the link ID from the peer MP.

The MP shall reject the Peer Link Open message if it carries a mismatched instance identifier or the configuration parameters in the received message are not acceptable by the MP. The instance identifier carried in the Peer Link Open contains only the link ID provided by the peer. It is considered a mismatch only when the MP has the local record of the PeerLinkID and it does not match the received value in the Local Link ID field. Besides the instance identifier, the Peer Link Open message contains configuration parameters. The MP shall reject the message if these parameters do not match the local configuration and policy. In other cases, the MP shall accept the message and record the received configurations that are useful for operations once the link has been established. The mesh entity shall also store the received value in the Local Link ID field as PeerLinkID if the mesh entity does not have a record yet.

The MP shall reject the Peer Link Confirm message if it carries a mismatched instance identifier or the configuration parameters in the received message are not acceptable by the MP. It is considered a mismatch of the instance identifier if the value in the Peer Link ID field of the received message does not match the locally recorded LocalLinkID, or the value in the Local Link ID field of the received message does not match the locally recorded PeerLinkID (if the MP has a local record of PeerLinkID). Besides the instance identifier, the Peer Link Confirm message contains configuration parameters as well. The MP shall reject the message if the received parameters do not match the local configuration and policy, or if they are not consistent with the parameters received earlier in the Peer Link Open message.

11A.1.5.3 Finite State Automatonstate automaton

This clause defines the finite state automaton that specifies the peer link establishment protocol. The finite state automaton has seven states. The terminate states are IDLE state and ESTABLISHED state. That is, in At the end of the protocol, either the link is closed (IDLE state) or the link is established successfullysuccessfully (ESTABLISHED state).

11A.1.5.3.1 States

State 0 represents the IDLE state where the local system refuses any attempt to establish a connection from the remote system

State 1 represents the LISTEN state where the local system is passively listening to incoming Peer Link Open message to establish a link from a peer

State 2 represents the OPEN_SENT state where the local system has actively sent a Peer Link Open message and is waiting for the incoming Peer Link Open and Peer Link Confirm messages

State 3 represents the CONFIRM_RCVD state where the local system has received a Peer Link Confirm message, but has not received a Peer Link Open message. This means that the local system hasn't sent out the corresponding Peer Link Confirm message either.

State 4 represents the CONFIRM_SENT state where the local system has sent a Peer Link Confirm message upon receiving a Peer Link Open message. But it hasn't received a Peer Link ConfirmConfirm message.

State 5 represents the ESTABLISHED state where the local system has sent and received both the Peer Link Open and Peer Link Confirm messages.

State 6 represents the HOLDING state where the local system is closing the connection. If the receiving The system will ignore any received peer link establishment message, the system is going to ignore it, message except for the Peer Link Open message.

11A.1.5.3.2 Events and Actionsactions

The following table Table s15 summarizes the events and actions in the finite state automaton.

Table s15—Peer link establishment events and actions

— Events			
• CancelPeerLink (CNCL)	• CloseReceived (CLR)	• Timeout (RetryTimer) (TOR)	
• ActivePeerLinkOpen (ACT)	• OpenReceived (OPR)	• Timeout (OpenTimer) (TOO)	
• PassivePeerLinkOpen (PAS)	• ConfirmReceived (CNR)	• Timeout (CancelTimer) (TOC)	
— Actions			
• Send-Open (SOP)	• Send-Confirm (SCN)	• Send-Close (SCL)	

— Events • CancelPeerLink (CNCL) • ActivePeerLinkOpen (ACT) • PassivePeerLinkOpen (PAS)	• • CloseReceived (CLR) • OpenReceived (OPR) • ConfirmReceived (CNR)	• • Timeout (Retry-Timer) (TOR) • Timeout (Open-Timer) (TOO) • Timeout (CancelT-imer) (TOC)
— Actions • Send-Open (SOP)	• • Send-Confirm (SCN)	• • Send-Close (SCL)

The automaton uses three kinds of events: local commands, events corresponding to protocol messages, and "internal" events.

The local commands are:

— MLME-CancelPeerLink(peerId, LocalLinkID, PeerLinkID).request. The MLME-CancelPeer-Link(peerId, LocalLinkID, PeerLinkID).request event represents a local decision to end a link with the peer whose MAC address is peerId. The link to be cancelled is identified as <myId, peerId, LocalLinkID, PeerLinkID>. The value for PeerLinkID is null if it is unknown.

— MLME-PassivePeerLinkOpen().request. The MLME-PassivePeerLinkOpen.request event represents a local decision to start listen to connection request.

— MLME-ActivePeerLinkOpen(peerId).request. The MLME-ActivePeerLinkOpen(peerId).request event represents a local decision to for open a link with the peer whose MAC address is peerId.

The events corresponding to protocol messages are:

— CloseReceived(peerId, myId, PeerLinkID, LocalLinkID). This event represents the reception of a Peer Link Close message from the peer to close link instance named by <myId, peerId, LocalL-inkID, PeerLinkID>.

— OpenReceived(peerId, myId, PeerLinkID). This event represents the reception of a Peer Link Open message from the peer named by peerId. This message was addressed by myId, and conveys the link instance identifier PeerLinkID generated by the peer MP, which must shall be non-null.

— ConfirmReceived(peerId, myId, PeerLinkID, LocalLinkID). This event represents the reception of a Peer Link Confirm message from the peer named by peerId. This message was addressed by myId, and conveys the link instance identifiers PeerLinkID and LocalLinkID, both of which must shall be non-null.

The internal events are:

— Timeout(item). This event represents a timeout identified locally by item. There are three types of timers. The RetryTimer controls the procedure of resending the Peer Link Open message. The RetryCounter is used to record the number of Peer Link Open messages that have been resent to the peer MP. The OpenTimer controls the procedure of receiving the Peer Link Open message after processing the Peer Link Confirm message. The CancelTimer controls the procedure of staying in the HOLDING state and eventually transitioning back to IDLE state.

11A.1.5.3.3 State transitions

The following table Table s16 gives the sketch table of state transitions. Columns are states and rows are events. The contents in each table entry represent the action and the result of state transition, encoded as "action/state". The follow list Table s15 contains abbreviations of important events and actions (related to messages over the air).

Table s16—Peer link establishment finite state machine state transitions

States Events	0 (IDLE)	1 (LISTEN)	2 (OPEN_SENT)	3 (CONFIRM_ RCVD)	4 (CONFIRM_S ENT)	5 (ESTABLISH ED)	6 (HOLDIN G)
CNCL	- / 0	- / 0	SCL / 6	SCL / 6	SCL / 6	SCL / 6	- / 6
PAS	- / 1	- / 1	- / 2	- / 3	- / 4	- / 5	- / 6
ACT	SOP / 2	SOP/2	- / 2	- / 3	- / 4	- / 5	- / 6
CLR	- / 0	- / 1	- / 6 or - / 2	- / 6 or - /3	- / 6 or - / 4	- / 6 or - /5	- / 6
OPR	- / 0	SOP, SCN / 4 or SCL / 1	SCN / 4 or SCL/6	SCN / 5 or - / 3	SCN / 4 or - / 4	SCN / 5 or - / 5	SCL / 6 or - / 6
CNR	- / 0	- / 1	- / 3 or SCL/6	- / 3	- / 5 or SCL / 6	- / 5	SCL / 6 or - / 6
TOR	- / 0	- / 1	SOP / 2 or SCL / 6	- / 3	SOP / 4 or SCL / 6	- / 5	- / 6
TOO	- / 0	- / 1	- / 2	SCL / 6	- / 4	- /5	- / 6
TOC	- / 0	- / 1	- / 2	- / 3	- / 4	- / 5	SCL / 0

States Events	0	1	2	3	4	5	6
CNCL	- / 0	- / 0	SCL / 6	SCL / 6	SCL / 6	SCL / 6	- / 6
PAS	- / 1	- / 1	- / 2	- / 3	- / 4	- / 5	- / 6
ACT	SOP / 2	SOP/2	- / 2	- / 3	- / 4	- / 5	- / 6
CLR	- / 0	- / 1	- / 6 or - / 2	- / 6 or - /3	- / 6 or - / 4	- / 6 or - /5	- / 6
OPR	- / 0	SOP, SCN / 4 or SCL / 1	SCN / 4 or SCL/6	SCN / 5 or - / 3	SCN / 4 or - / 4	SCN / 5 or - / 5	SCL / 6 or - / 6
CNR	- / 0	- / 1	- / 3 or SCL/6	- / 3	- / 5 or SCL / 6	- / 5	SCL / 6 or - / 6
TOR	- / 0	- / 1	SOP / 2 or SCL / 6	- / 3	SOP / 4 or SCL / 6	- / 5	- / 6
TOO	- / 0	- / 1	- / 2	SCL / 6	- / 4	- /5	- / 6
TOC	- / 0	- / 1	- / 2	- / 3	- / 4	- / 5	SCL / 0

The following subclauses specify the detailed state transitions.

11A.1.5.3.4 IDLE state (0)

In the IDLE state the local MP shall not respond to incoming messages from a remote any other MP. It only responds to the internal commands.

When MLME-PassivePeerLinkOpen().request occurs, the MP allocates necessary resource for a new peer link, generates a new link identifier, LocalLinkID, and start to listen to connection request. The finite state automaton transitions to LISTEN state.

When MLME-ActivePeerLinkOpen(peerId).request occurs, the local MP actively tries to establish a peer link with the peer MP, identified as peerId. The local MP allocates necessary resource for the peer link, generates a new link identifier, LocalLinkID, and sends out a Peer Link Open message to the peer MP. The RetryTimer is set. The finite state automaton transitions to OPEN_SENT state.

The timers shall be cleared if a timeout event occurs.

All other events are ignored in this state.

11A.1.5.3.5 LISTEN state (1)

In the LISTEN state, the local MP listens to the incoming connection request.

The MLME-CancelPeerLink(null**,** LocalLinkID**,** null).request event means that the local MP is no longer willing to accept more peer link connections. The allocated resource shall be cleared and the finite state automaton transitions to IDLE state.

PassivePeerLinkOpen() command shall be ignored.

If MLME-ActivePeerLinkOpen(peerId).request event occurs, the local MP actively tries to establish a peer link with the peer MP, identified by peerId, by sending a Peer Link Open message, which contains the LocalLinkID. The RetryTimer is set. The finite state automaton transitions to OPEN_SENT state.

If CloseReceived event or ConfirmReceived event occur, the MP shall verify the received instance identifier. If the number matches LocalLinkID, it means that the peer MP is able to predict my the local MP's pseudo-random number. The local MP shall raise an exception to the higher layer to indicate that the local pseudo-random number generator is broken. If the identifiers do not match, both of these events are silently ignored.

If OpenReceived event occurs, the MP shall process the incoming Peer Link Open message. (11A.1.5.2 defines the procedure of processing Peer Link Open message). If the process succeeds, the MP shall send the corresponding Peer Link Open message and Peer Link Confirm message. The RetryTimer is set and the finite state automaton transitions to CONFIRM_SENT state. If the process fails, the MP shall send the Peer Link Close message to the peer MP to notify the failure and close the link. The Peer Link Close message shall carry the Peer Link ID value but not the LocalLinkID value. The CancelTimer is set, the failure reason code is recorded as "FAILURE-INVALID-PARAMETERS", and the finite state automaton transitions to HOLDING state.

The timers shall be cleared if a timeout event occurs.

11A.1.5.3.6 OPEN_SENT state (2)

In the OPEN_SENT state, the local MP has sent out a Peer Link Open message to the peer MP and is waiting for a response. In this state, the RetryTimer is set.

If MLME-CancelPeerLink(peerId, LocalLinkID, null).request occurs, the connection attempt to connect to the peer peerId shall abortbe aborted. A Peer Link Close message shall be sent to indicate a closing link. The CancelTimer is set, the failure reason code is recorded as "FAILURE-CANCELLED", and the finite state automaton transitions to HOLDING state.

The PassivePeerLinkOpen().request event and ActivePeerLinkOpen(peerId).request event shall be ignored.

If CloseReceived event occurs, the MP shall process the incoming Peer Link Close message. (11A.1.5.2 defines the procedure of processing a Peer Link Close message). If the process succeeds, the RetryTimer shall be cleared, the CancelTimer shall be set, the failure reason code is recorded as "FAILURE-CLOSE", and the finite state automaton transitions to HOLDING state. If the process fails, the Peer Link Close message is silently ignored.

If OpenReceived event occurs, the MP shall clear the RetryTimer and process the incoming Peer Link Open message. If the process fails, it means the local MP denies the connection request. A Peer Link Close message shall be sent to close the connection. The CancelTimer is set, the failure reason code is recorded as "FAILURE-MACMAX-REQS" and the finite state automaton transitions to HOLDING state. If the process succeeds, a Peer Link Confirm message shall be sent, the retryTimer is set, and the finite state automaton transitions to CONFIRM_SENT state.

If ConfirmReceived event occurs, the MP shall clear the RetryTimer and process the incoming Peer Link Confirm message. (11A.1.5.2 defines the procedure of processing a Peer Link Confirm message). If the process fails, it means the local MP denies the response from the peer MP. In this case, the The local MP shall close the link by sending a Peer Link Close to the peer MP to indicate the failure and the reason. The CancelTimer is set, the failure reason code is recorded as "FAILURE-INVALID-PARAMETERS" and the finite state automaton transitions to HOLDING state. If the process succeeds, the OpenTimer shall be set and the finite state automaton transitions to CONFIRM_RCVD state.

If Timeoute(RetryTimer) event occurs, the MP shall verify the RetryCounter. If the RetryCounter exceeds the MAX-REQS limit, the MP shall send a Peer Link Close message to the peer MP, set the CancelTimer, record the failure reason code as "FAILURE-MACMAX-REQS", and transition to the HOLDING state. If the An MP that has not reached the MAX-REQS limit, it limit shall resend the Peer Link Open message to the peer MP. The RetryCounter is incremented by 1. The backoff algorithm is used to compute the next retry timeout value. The RetryTimer is set to the updated timeout value. No state transition occurs.

If Timeout(OpenTimer) event occurs, clear the OpenTimer, and stays in OPEN_SENT state.

If Timeout(CancelTimer) event occurs, clear the CancelTimer, and stays in OPEN_SENT state.

11A.1.5.3.7 CONFIRM_RCVD state (3)

In the CONFIRM_RCVD state, the MP has received a Peer Link Confirm message and is waiting for a Peer Link Open message.

If CancelPeerLink(peerId, LocalLinkID, PeerLinkID).request event occurs, the connection attempt to connect to the peer peerId shall abortbe aborted. A Peer Link Close message shall be sent to indicate a closing link. The Peer Link Close message shall contain the instance identifiers from both ends. The CancelTimer is set, the failure reason code is recorded as "FAILURE-CANCELLED", and the finite state automaton transitions to HOLDING state.

The PassivePeerLinkOpen().request event and ActivePeerLinkOpen(peerId).request event shall be ignored in this state.

If CloseReceived event occurs, the MP shall process the incoming Peer Link Close message. (11A.1.5.2 defines the procedure of processing a Peer Link Close message). If the process succeeds, the RetryTimer shall be cleared, the CancelTimer shall be set, the recoreded recorded failure reason code is set as "FAILURE-CLOSE", and the finite state automaton transitions to HOLDING state. If the process fails, the Peer Link Close message is silently ignored.

If OpenReceived event occurs, the MP shall process the incoming Peer Link Open message. If the process succeeds, the OpenTimer shall be cleared, a Peer Link Confirm message shall be sent to the peer MP, the finite state automaton transitions to ESTABLISHED state, and the status code "SUCCESS" is reported to the higher layer. The peer link has been successfully established. If the process fails, the MP shall ignore the received message and continues to wait for the Peer Link Open message from the peer MP.

If ConfirmReceived event occurs, the MP shall ignore the incoming message.

If Timeout(RetryTimer) event occurs, clear the RetryTimer, and stays in CONFIRM_RCVD state.

If Timeout(OpenTimer) event occurs, the MP declares the failure of the link establishment and shall send a Peer Link Close message to close the link. The CancelTimer is set, the recorded failure reason code is set as "FAILURE-TIMEOUT", and the finite state machine transitions to HOLDING state.

If Timeout(CancelTimer) event occurs, clear the CancelTimer, and stays in CONFIRM_RCVD state.

11A.1.5.3.8 CONFIRM_SENT state (4)

In the CONFIRM_SENT state, the MP has received a Peer Link Open message and sent the corresponding Peer Link Confirm message. The incoming Peer Link Confirm is expected.

If CancelPeerLink(peerId, LocalLinkID, PeerLinkID) command is received, the connection attempt to connect to the peer peerId shall abortbe aborted. A Peer Link Close message shall be sent to indicate a closing link. The Peer Link Close message shall contain both ends' pseudo-random numbers for this session. The CancelTimer is set, the failure reason code is recorded as "FAILURE-CANCELLED", and the finite state automaton transitions to HOLDING state.

The PassivePeerLinkOpen() event and ActivePeerLinkOpen event shall be ignored.

If CloseReceived event occurs, the MP shall process the incoming Peer Link Close message. (11A.1.5.2 defines the procedure of processing a Peer Link Close message). If the process succeeds, the RetryTimer shall be cleared, the cancelTimer shall be set, the failure reason code is recorded as "FAILURE-CLOSE", and the finite state automaton transitions to HOLDING state. If the process fails, the Peer Link Close message is silently ignored.

If OpenReceived event occurs, the MP shall process the incoming Peer Link Open message. If the process succeeds, a Peer Link Confirm message shall be sent to the peer MP and the finite state automaton shall stay in the CONFIRM_SENT state. If the process fails, the MP shall ignore the received message and continue to wait for the Peer Link Confirm message from the peer MP.

If ConfirmReceived event occurs, the MP shall process the Peer Link Confirm message and clear the RetryTimer. If the process succeeds, the finite state automaton transitions to ESTABLISHED state, and the status code "SUCCESS" is reported to the higher layer. If the process fails, the MP shall clear the RetryTimer, send the Peer Link Close message to notify the failure, set the CancelTimer, the failure reason code is recorded as "FAILURE-INVALID-PARAMETERS", and the finite state automaton transitions to HOLDING state.

If Timeoute(RetryTimer) event occurs, the MP shall verify the RetryCounter. If the RetryCounter exceeds the MAX-REQS limit, the MP shall send a Peer Link Close message to the peer MP, set the CancelTimer, the failure reason code is recorded as "FAILURE-MAX-REQS", and transition to the HOLDING state. If the An MP that has not reached the MAX-REQS limit, it limit shall resend the Peer Link Open message to the peer MP. The RetryCounter is incremented by 1. The backoff algorithm is used to compute the next retry timeout value. The RetryTimer is set to the updated timeout value. No state transition occurs.

If Timeout(OpenTimer) event occurs, the MP clears the OpenTimer, and stays in CONFIRM_SENT state.

If Timeout(CancelTimer) event occurs, the MP clears the CancelTimer, and stays in CONFIRM_SENT state.

11A.1.5.3.9 ESTABLISHED state (5)

In the ESTABLISHED state, the MP has successfully established a peer link with the peer MP. The MPs can start the data transmission and routing functions.

If CancelPeerLink(peerId, LocalLinkID, PeerLinkID).request event occurs, the connection attempt to connect to the peer peerId shall abortbe aborted. A Peer Link Close message shall be sent to indicate a closing link. The Peer Link Close message shall contain the complete instance identifier from both parties. The CancelTimer is set, the failure reason code is recorded as "FAILURE-CANCELLED", and the finite state automaton transitions to HOLDING state.

The PassivePeerLinkOpen().request event and ActivePeerLinkOpen(peerId).request event shall be ignored.

If CloseReceived event occurs, the MP shall process the incoming Peer Link Close message. (11A.1.5.2 defines the procedure of processing a Peer Link Close message). If the process succeeds, the RetryTimer shall be cleared, the CancelTimer shall be set, the failure reason code is recorded as "FAILURE-CLOSE", and the finite state automaton transitions to HOLDING state. If the process fails, the Peer Link Close message is silently ignored.

If OpenReceied OpenReceived event occurs, the MP shall process the Peer Link Open message. If the process succeeds, a Peer Link Confirm message shall be sent to the peer MP and the finite state automaton shall stay in the ESTABLISHED state. If the process fails, the MP shall ignore the received message and continue to stay in the ESTABLISHED state.

The ConfirmReceived event shall be ignored in this state.

The timeout events shall be ignored in this state. The corresponding timer shall be cleared.

11A.1.5.3.10 HOLDING state (6)

If either an OpenReceived or ConfirmReceived event occurs, the MP shall process the message. If the process fails, the incoming message is ignored. If the process succeedsprocessed successfully, the MP shall send the Peer Link Close message to notify that the MP is closing the link.

If Timeout(CancelTimer) event occurs, the MP shall send the final Peer Link Close message to the peer MP, signal the failure reason code of the link establishment protocol to the higher layer, release all resources for this session, and the finite state automaton transitions to IDLE state.

All other events are ignored in this state.

11A.1.6 Link Quality Measurementquality measurement

Once a peer link has been established with a neighboring MP, it is necessary to establish a measure of the quality of the link. This information is required to allow the path selection algorithm/metric to function properly.

Before a link quality measurement has been established, the link is marked as being "Down". Path selection frames received from a node an MP whose link is "Down" should be discarded. Once a link quality measurement has been established, the link is marked as being "Up". At this point, the MP is able to fully participate in the selection of paths (and by extension, it is able to fully participate in frame forwarding).

Since an MP may use any path selection metric, it is necessary for the MP to measure a link quality value that is relevant to the path selection metric. For example, clause 7.3.2.52 describes a Local Link State Announcement element comprise comprised of a transmit bit rate and a packet frame error rate. 11A.5.1 describes a procedure for using this Local Link State Announcement to ensure that the link quality measurement is symmetrical (i.e. the same values are used at each end of a link).

11A.1.7 Mesh Network Channel Selection

11A.1.7.1 Overview of Single-Channel and Multi-Channel Operation in a WLAN Mesh

11A.1.8 Mesh network channel selection

11A.1.8.1 RF Channel Interfaces and Unified Channel Graphs

In its simplest form, a WLAN Mesh operates only on one channel. For multi-channel operation, devices either need multiple radios PHYs or channel switching capability. Devices with more than one radio interface PHY tune each radio interface PHY to a different channel. Optionally, devices with switching capability can dynamically switch to any of the available channels for a short period. An overview of the resulting multi-channel operation is provided here.

11A.1.8.1.1 RF Channel Interfaces and Unified Channel Graphs

A WLAN Mesh network topology may include mesh points MPs with one or more radio interfaces PHY and may utilize one or more channels for communication between mesh pointsMPs. When channel switching is not supported, each radio interface PHY on a mesh point an MP operates on one channel at a time, but the channel may change during the lifetime of the mesh network according to DFS requirements. The specific channel selection scheme used in a WLAN Mesh network may vary with different topology and application requirements. Figure s96 illustrates three example mesh point MP channel allocation schemes. (a) illustrates a simple deployment case with single interface mesh points MPs using a single channel throughout the mesh network. This specification standard includes a protocol to enable a set of mesh point radio interfaces MP PHYs to coalesce to a common channel for communication to enable this type of simplesimple mesh network. (b) and (c) illustrate two advanced channel allocation schemes in which one or more mesh points MPs have more than one radio interface PHY and more than one channel is used across the mesh network. Flexibility is supported to allow implementation of many different possible advanced channel allocation schemes to meet special application requirements.

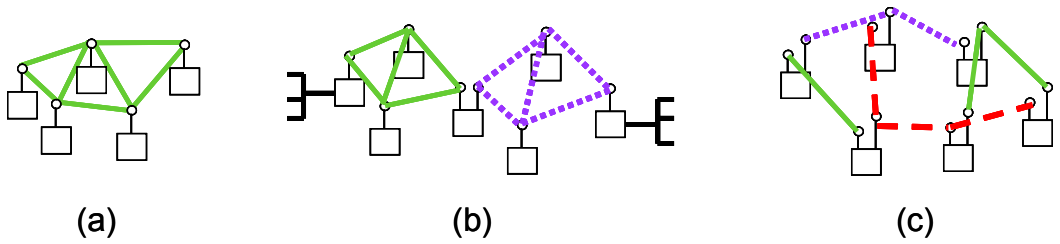

Figure s96—Example channel configurations in a WLAN Meshmesh.

Note that in In each of the example topologies, two or more mesh point radio interfaces MP PHYs are connected to each other using a common channel. A set of mesh point radio interfaces MP PHYs that are interconnected to each other via a common channel is referred to as a unified channel graph (UCG). The same device may belong to different UCGs. As illustrated in Figure s97, a simple, single-channel mesh network has only one UCG, while more sophisticated topologies may include multiple UCGs. A framework is provided for coordinated switching of the channel used within a UCG when it is necessary for channels to change in an operating mesh network, e.g., due to regulatory DFS requirements. Each UCG in a WLAN Mesh shares a common channel precedence value which may be used to coalesce (see 11A.1.8.5) or switch the channel in the UCG (see 11A.1.8.6).

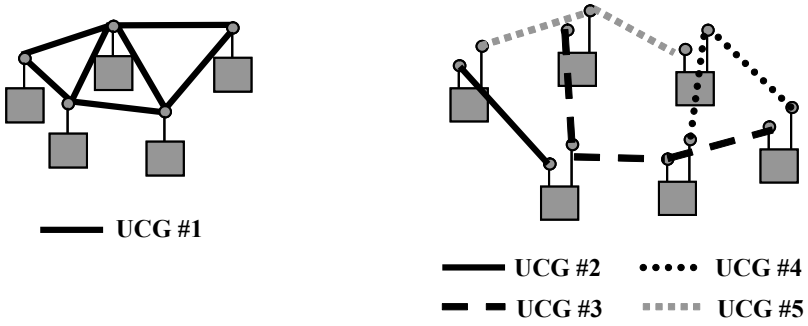

Figure s97—Example unified channel graphs in a WLAN Meshmesh.

11A.1.8.2 Single and Multiple Radio Devicesmultiple PHY devices

A device may have one or more radiosPHYs. A device that can operate as both an 802.11a a 5GHz device and an 802a 2.11b 4GHz device, but not both at the same time, shall be considered to have one radioPHY, even if it physically contains two radiosPHYs. A multiple radio PHY device need not operate each radio PHY in a different band; it is permissible to have a device containing more than one 802.11a radios5GHz PHY, for example.

A multiple radio PHY device shall use a different MAC address for each interface and will be is treated as multiple MPs connected by a non-blocking interconnect.

11A.1.8.3 Channel Selection Modes for Mesh Point Logical Radio Interfaces

11A.1.8.4 Channel selection modes for MP PHYs

A Mesh Point An MP shall specify the channel selection mode of each logical radio interface PHY as either simple unification mode or advanced mode, and advertise this mode using the simple channel unification mode flag in the mesh capability element included in beacon Beacon and probe response Probe Response frames.

A mesh point logical radio interface An MP PHY that is in simple unification mode shall select a channel in a controlled way such that it enables the formation of a unified channel graph that becomes merged and hence fully connected. The mesh point logical radio interface MP PHY shall establish links with neighbors that match the Mesh ID and Mesh Profile and select its channel based on the highest channel precedence value.

A mesh point logical radio interface An MP PHY that is in advanced mode may set its channel based on advanced management rules (beyond the scope of this specificationstandard). In this mode, the mesh point MP logical interface shall establish links with neighboring mesh point MP logical interfaces that match the Mesh ID and Mesh Profile that are on the same channel as the logical interface according to the advanced channel setting rules.

11A.1.8.5 Simple Channel Unification Protocolchannel unification protocol

This protocol is executed on mesh point logical radio interfaces MP PHYs that are configured in simple unification mode.

A An MP logical radio interface PHY that is configured in simple channel unification mode shall periodically perform passive or active scanning to discover neighboring MPs. If an MP is unable to detect any neighbor MPs, it may adopt a Mesh ID from one of its profiles, select a channel for operation, and select an initial channel precedence value. The initial channel precedence value shall be initialized to the number of microseconds since the boot time of the mesh point MP plus a pseudo-random value.

In the event that a mesh point logical radio interface an MP's PHY that is configured in simple channel unification mode discovers a disjoint mesh, that is, the list of candidate peer Mesh Points MPs spans more than one channel, it the MP shall select the channel that is indicated by the candidate peer Mesh Point MP that has the numerically highest channel precedence indicator to be the unification channel.

If the identified unification channel is different than from the current operating channel of the mesh point logical radio interfaceMP PHY, the mesh point MP shall execute the channel graph switch protocol described in 11A.1.8.6.

11A.1.8.6 Channel Graph Switch Protocolgraph switch protocol

This section subclause describes the procedure used for a mesh point an MP to initiate switching of a unified channel graph to a new channel, with a new channel precedence indicator. Due to the possibility of more than one mesh point MP of a unified channel graph executing this procedure the channel graph switch protocol concurrently, this procedure protocol includes a mechanism to resolve such possible conflicts by introducing a UCG switch wait timer that assures adequate time for the decision process of this procedureprotocol.

The mesh point MP that determines the need to switch the channel of its UCG first chooses a UCG switch wait time of TBD time units (TUs). The mesh point MP sets a local timer with this wait time and then sends a UCG switch announcement Switch Announcement frame to each peer mesh point MP to which an active association exists a peer link has been established in the unified channel graph, copying the value of the new candidate channel and new candidate channel precedence indicator and setting the channel switch count to TBD TUs.

If a mesh point an MP receives a UCG switch Switch frame with a channel precedence value larger than the current channel precedence value of the logical interface on which the frame was received, the mesh point MP shall set a UCG switch wait timer equal to the channel switch count value of the frame and then sends a UCGswitch UCG Switch announcement frame to each peer mesh point MP to which an active association exists a peer link has been established on the logical radio interfacePHY, copying the values from the received UCG switch Switch announcement frame.

Note that it It is possible that more than one mesh point MP in the unified channel graph may independently detect the need to switch channels and send separate UCG switch Switch announcement frames. If a mesh point an MP receives more than one UCG switch Switch announcement frame, it only acts upon the frame if the channel precedence value is larger than the channel precedence value of a previously received UCG switch Switch announcement frame. In case a newly received UCG announcement frame has the same channel precedence value as a previously received frame, the new frame is acted upon only if the source address is smaller than the source address from the previously received frame. If the mesh point MP acts upon the newly received UCG switch Unified Channel Switch frame, it updates its candidate channel and candidate channel precedence indicator, sets its UCG switch Switch wait timer to the channel switch count value of the frame and then sends a UCG switch Switch announcement frame to each peer mesh point MP to which an active association exists a peer link has been established on the logical radio interfacePHY, copying the values from the received UCG switch Switch announcement frame.

If a UCG switch wait timer has been set on a mesh pointan MP, the mesh point MP shall not originate a new UCG switch Switch announcement frame during the duration of the UCG switch wait timer. When the UCG switch wait timer expires on a mesh point an MP the mesh point MP switches its radio interface PHY to the candidate channel and updates its channel precedence indicator to the candidate channel precedence indicator.

EDITORIAL NOTE—Mesh Link Security Clause moved from 8.8 to 11A.2 per CID 2924 (with exception of 8.8.2)

11A.2 Mesh link security

11A.2.1 Overview of MSA

Mesh security association (MSA) services are used to permit establishment of link security between two MPs in a wireless mesh network, and support both centralized and distributed authentication schemes. MSA services are provided through the use of a mesh key hierarchy, a hierarchy of derived keys that is established through the use of a PSK or when an MP performs IEEE 802.1X authentication.

The operation of MSA relies on mesh key holders, which are functions that are implemented at MPs within the wireless mesh network. Two types of mesh key holders are defined: mesh authenticators (MAs) and mesh key distributors (MKDs). A single MP may implement both MKD and MA key holders, or an MA or neither.

MSA requires information to be exchanged during an MP's initial security association with an MA, and is referred to as "Initial MSA Authentication." Subsequent security associations to other MAs within the same MKD domain (and the same mesh, as identified by the Mesh ID) may utilize the mesh key hierarchy that is established during Initial MSA Authentication.

MSA provides mechanisms for secure communications between mesh key holders. The "Mesh Key Holder Security Handshake" provides the mechanism for establishing a security association between an MA and MKD. Secure mesh key transport protocols and an optional EAP message transport protocol are defined.

11A.2.1.1 Mesh key holders

Mesh key holders, MAs and MKDs, manage the mesh key hierarchy by performing key derivation and secure key distribution. A mesh key distributor (MKD) domain is defined by the presence of a single MKD. Within the MKD domain, several MAs may exist, each implemented at an MP, and each MA maintains both a route to and a security association with the MKD. The MKD derives keys to create a mesh key hierarchy, and distributes derived keys to MAs.

An MP implementing the MA key holder function may play the IEEE 802.1X Authenticator role during an MSA exchange (e.g., Initial MSA Authentication) as determined according to the procedures in 11A.2.1.3. The MA receives derived keys from the MKD, and derives additional keys for use in securing a link with a supplicant MP.

The design of MSA assumes that the AS and MKD have a trustworthy channel between them that can be used to exchange cryptographic keys without exposure to intermediate parties. The IEEE 802.1X AS never exposes the MSK to any party except the MKD implementing the NAS Client functionality of the IEEE 802.1X Authenticator with which the supplicant is communicating. The communication between AS and MKD is outside the scope of this standard.

11A.2.1.2 Discovery & MSA capability advertisement

The support of MSA is advertised by MPs in Beacon and Probe Response frames through the inclusion of the MKDDIE. Moreover, an MP that wants to utilize MSA to authenticate with other MPs shall advertise its security policy by inserting an RSN information element into its Beacon frames and Probe Response frames.

The MKDDIE shall be included in Beacon and Probe Response frames to advertise support for MSA and to advertise the MKD domain identifier (MKDD-ID) and the Mesh Security Configuration. The value of MKDD-ID that is advertised by the MP is the value received from the MKD during the mesh key holder security handshake (as specified in 11A.2.3.2), or the value of dot11MeshKeyDistributorDomainID if the MP implements the MKD function. An MP that has not yet received the MKDD-ID value shall set the MKD domain ID field in the MKDDIE to zero and shall set the Mesh Authenticator and Connected to MKD bits of the Mesh Security Configuration field to zero.

The Mesh Security Configuration field in the MKD domain information element shall be set as follows:

— Bit 0 (Mesh Authenticator): The MP shall set this bit to 1 if the MP is configured to play the IEEE 802.1X Authenticator role during an MSA handshake. The selection of the IEEE 802.1X Authenticator and Supplicant roles is described in 11A.2.1.3.

— Bit 1 (Connected to MKD): The MP shall set this bit to 0 if bit 0 (Mesh Authenticator) is set to 0. Otherwise, the MP shall set this bit to 1 if the MP has a security association with the MKD and has a valid route to the MKD. If the MA and MKD are both implemented at the MP and bit 0 is set to 1, the MP shall set this bit to 1.

— Bit 2 (Default Role Negotiation): The MP shall set this bit to 1 if it uses the mesh default role determination scheme specified in 11A.2.1.3. The MP shall set this bit to 0 if it uses some other role determination scheme, such as a proprietary scheme. The specification of other schemes is outside the scope of this standard.

An MKD may support one or more EAP transport mechanisms. An MA advertises the mechanisms supported by the MKD with which it has a security association during the Initial MSA Authentication (using the EAP Transport List optional parameter in the MSAIE).

11A.2.1.3 Role determination

When MSA is used, roles shall be selected prior to link establishment and policy selection. The two MPs shall determine the IEEE 802.1X Authenticator and Supplicant roles through the use of the Mesh Authenticator (Bit 0) and Connected to MKD (Bit 1) bits of the Mesh Security Configuration field (in the MKDDIE) as follows:

— If one of the MPs has set Bit 0 to 1 and the other MP has set Bit 0 to 0, then the MP that set Bit 0 to 1 shall assume the IEEE 802.1X Authenticator role, and the MP that set Bit 0 to 0 shall assume the IEEE 802.1X Supplicant role.

— If both MPs set Bit 0 to 1, then

- If both or neither MPs have set Bit 1 to 1, then the MP with the higher MAC address shall assume the IEEE 802.1X Authenticator role, and the MP with the lower MAC address shall assume the IEEE 802.1X Supplicant role.
- If one of the MPs has set Bit 1 to 1 and the other to 0, then the MP that set Bit 1 to 1 shall assume the IEEE 802.1X Authenticator role, and the MP that set Bit 1 to 0 shall assume the IEEE 802.1X Supplicant role.

— If both MPs set Bit 0 to 0, then the MP with the higher MAC address shall assume the IEEE 802.1X Authenticator role, and the MP with the lower MAC address shall assume the IEEE 802.1X Supplicant role.

When both MPs set Bit 1 to 0, it is possible for the secure association to fail because one of the parties lacks credentials in its local database to authenticate and/or authorize the other.

11A.2.1.4 Policy selection

An MP may initiate the link establishment mechanism defined in 11A.1.5. This mechanism leverages Association Request and Association Response frames to exchange Peer Link Open and Peer Link Confirm information elements.

If an IEEE 802.1X-based authentication is used, the MP playing the role of the IEEE 802.1X Supplicant shall include an RSN information element in the Association Request. In the RSN information element, the Supplicant MP shall specify one pairwise ciphersuite and one authenticated key management suite.

In a mesh, all STAs shall utilize the same group ciphersuite. Therefore, a Supplicant MP shall not send an Association Request frame, and shall reject Association Request frames from the Authenticator MP (with Status Code 41), if the group ciphersuite advertised by the Authenticator MP does not match its own.

The Authenticator MP shall reject the Association Request frame from the Supplicant MP if either the pairwise cipher suite (with Status Code 42) or authenticated key management suite (with Status Code 43) selected by the Supplicant is not included in the corresponding lists of pairwise ciphersuites and authenticated key management suites specified in its own Beacon frames and Probe Response frames. The Authenticator MP may also reject the Supplicant MP's Association Request frame for other reasons unrelated to security. The Authenticator MP may accept the Association Request frame if the Supplicant selected pairwise and authenticated key management suites from among those specified by the Authenticator in its Beacon frames and Probe Response frames.

If an IEEE 802.1X-based authentication is used, the Supplicant MP shall additionally include an MKDDIE in the Association Request frame. The Authenticator MP shall reject the Association Request frame from the Supplicant MP if the MKDD-ID included in the MKDDIE does not match the value advertised by the Authenticator MP in its Beacon frames and Probe Response frames.

Selection of the EAP Transport mechanism to be used between an MP and MKD is performed during the mesh key holder security handshake described in 11A.2.3.2. The MP shall decline to establish a mesh key holder security association with the MKD if the EAP transport mechanisms supported by the MP and MKD do not overlap.

11A.2.1.5 Initial MSA authentication

Pre-RSNA authentication shall not be supported for mesh link establishment.

The Initial MSA Authentication mechanism permits an MP to enable the use of the mesh key hierarchy when establishing security for subsequent links.

If the link establishment mechanism specified in 11A.1.5 succeeds in creating a link, and if it selects IEEE 802.1X authentication, then the Authenticator MP shall initiate the authentication. 8.4.4 specifies the authentication procedure used when IEEE 802.1X is selected. If pre-shared keys (PSKs) are selected instead, then the PMK is derived from the PSK.

If authentication succeeds from the Authenticator MP's perspective, then the Authenticator MP shall initiate a 4-Way Handshake, as specified in 11A.2.2.2. After the 4-Way Handshake completes, either MP may initiate a Group Key Handshake (8.5.4) at any time during the link's lifetime, to update the GTK.

The Initial MSA Authentication sequence is shown in Figure s98, with procedures specified in 11A.2.2.2.

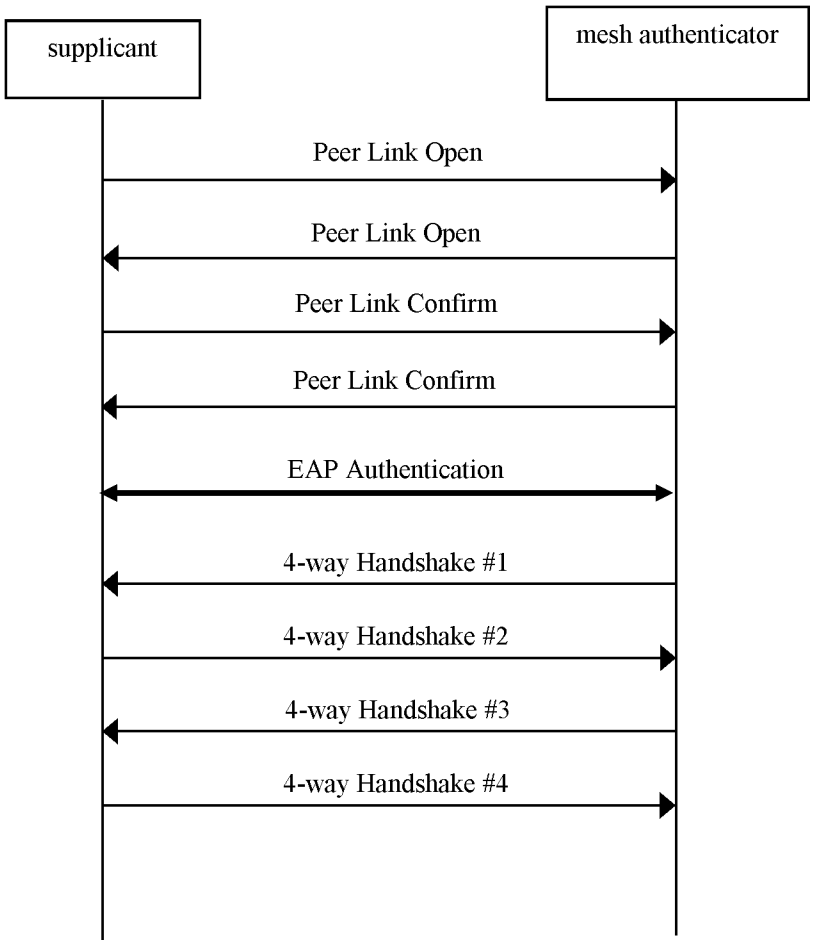

Figure s98—Initial MSA Authentication

11A.2.1.6 Subsequent MSA authentication

Pre-RSNA authentication shall not be supported for mesh link establishment.

The Subsequent MSA Authentication mechanism permits an MP to establish security for subsequent links with other MPs in the mesh once the mesh key hierarchy has been established.

An example Subsequent MSA Authentication sequence is shown in Figure s99 with procedures specified in 11A.2.2.3.

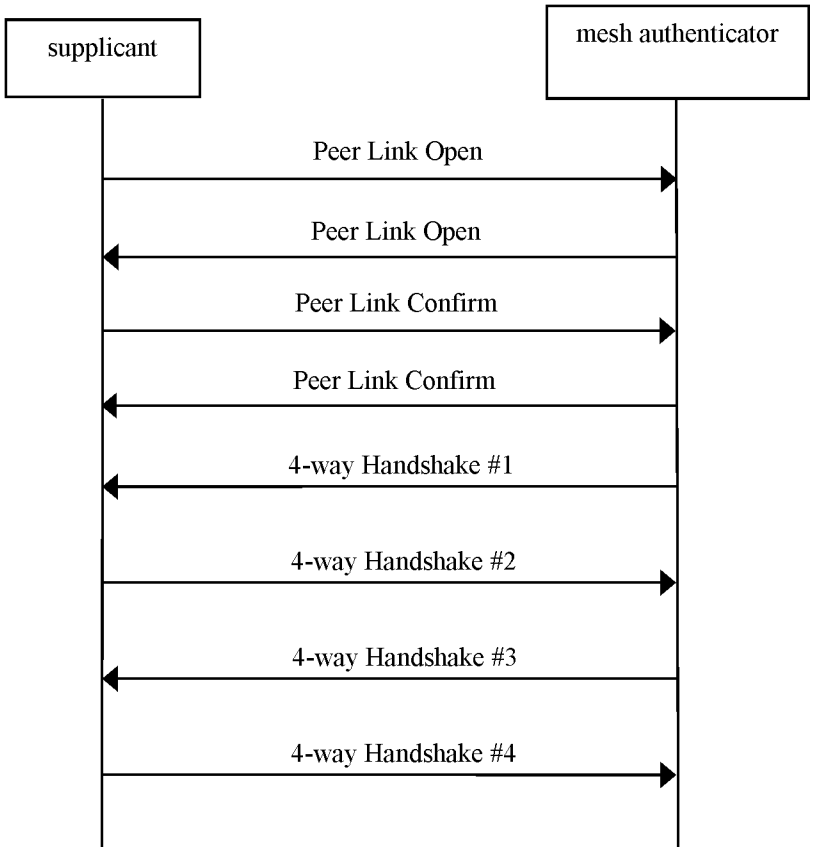

Figure s99—Subsequent MSA Authentication

11A.2.1.7 Mesh key holder security association

The mesh key holder security association establishes a security association between an MP and an MKD, permitting the MP to begin operating as an MA. The MP may initiate the mesh key holder security handshake after it has completed Initial MSA Authentication. The mesh key holder security handshake is shown in Figure s100, with procedures specified in 11A.2.3.

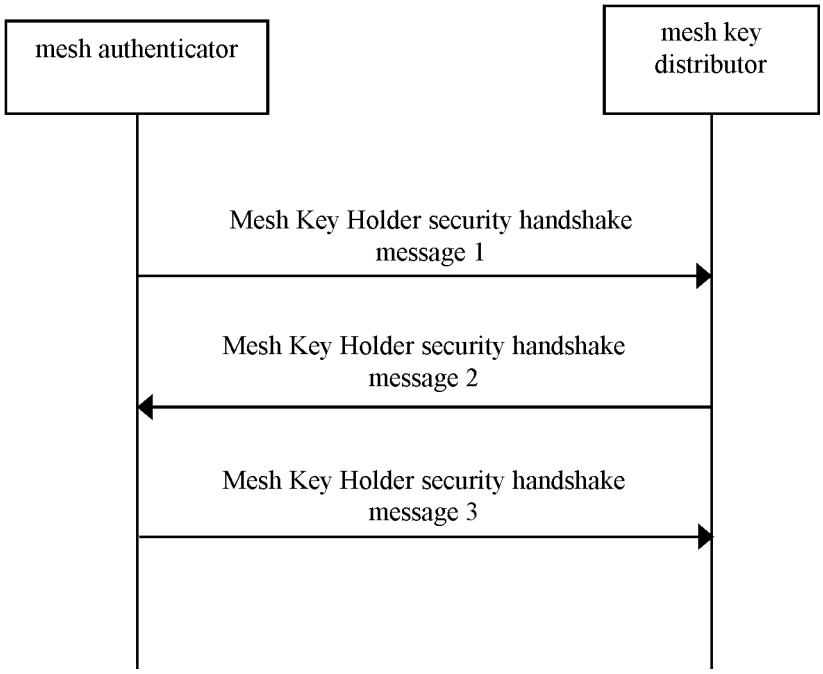

Figure s100—Mesh key holder security handshake

11A.2.1.8 Mesh key and EAP message transport protocols

The mesh key transport protocol comprises three mechanisms for performing key delivery and key management within a mesh key hierarchy.

The delivery pull protocol is initiated by the MA to request delivery of a PMK-MA, is shown in Figure s101, and is specified in 11A.2.4.1.

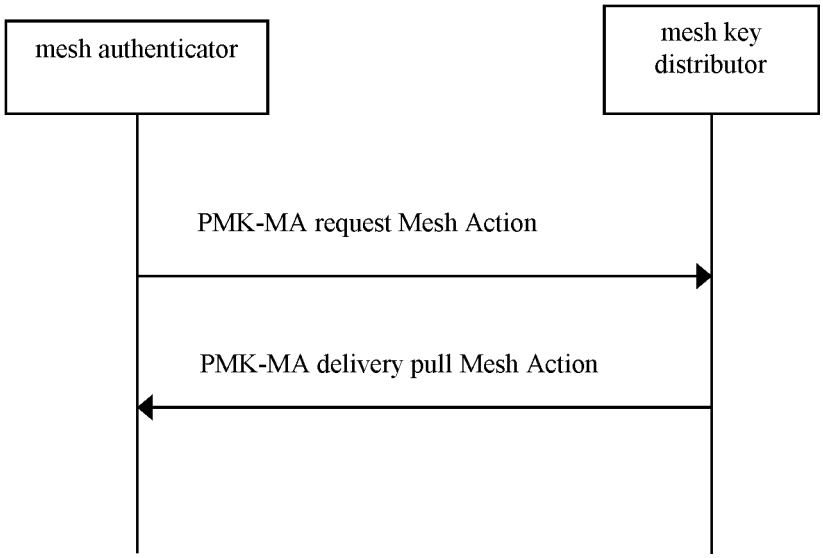

Figure s101—Mesh key transport delivery pull protocol

The delivery push protocol is initiated by the MKD to deliver a PMK-MA, is shown in Figure s102, and is specified in 11A.2.4.2.

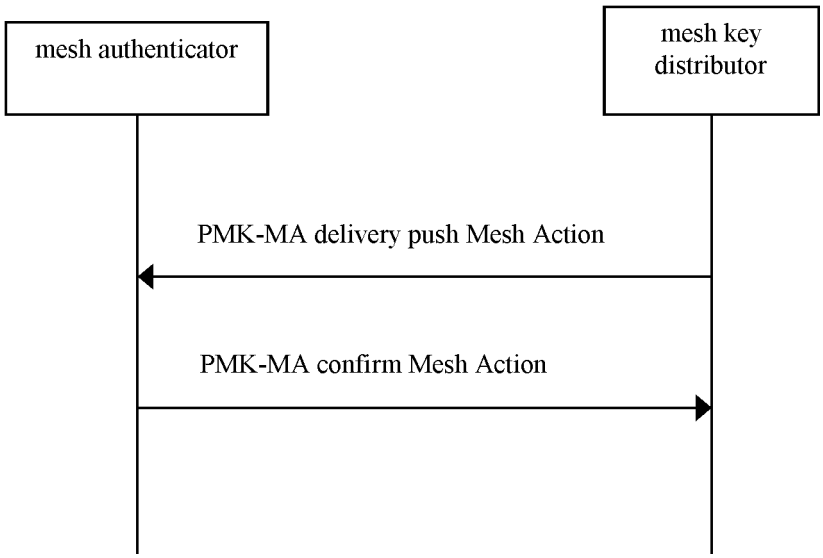

Figure s102—Mesh key transport delivery push protocol

The key delete protocol is initiated by the MKD to request revocation of a PMK-MA, is shown in Figure s103, and is specified in11A.2.4.3.

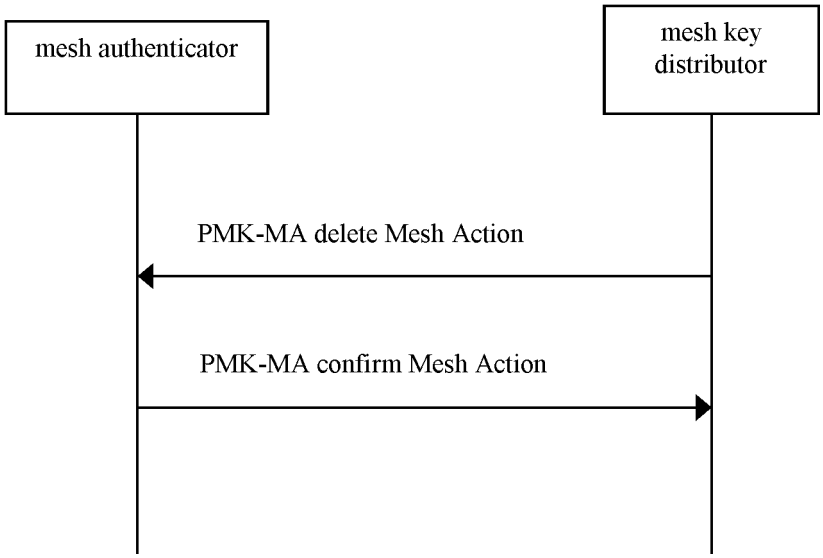

Figure s103—Mesh key delete protocol

The optional EAP message transport protocol may be initiated by the MA to facilitate EAP authentication with the supplicant during a supplicant MP's Initial MSA Authentication. The protocol permits an EAP message received from the supplicant to be transported from MA to MKD, and permits EAP messages received from the authentication server to be transported from MKD to MA.

A single request/response EAP message transport frame exchange is shown in Figure s104. The authentication of a supplicant typically requires several such exchanges. The optional protocol is specified in 11A.2.5.

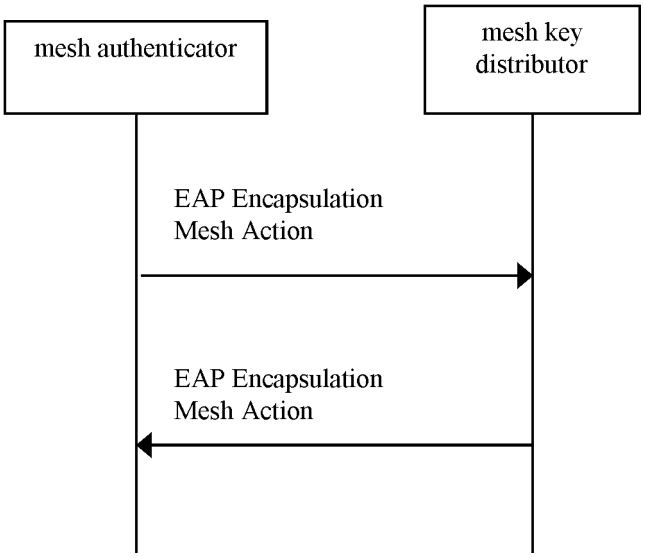

Figure s104—EAP message transport protocol (single exchange)

11A.2.1.9 Secure link operation

In the case when the 4-Way Handshake completes successfully, then both the IEEE 802.1X Authenticator and Supplicant shall open their respective controlled ports, to permit data traffic to be exchanged using the selected ciphersuites.

When key management completes, each MP further uses the session key to protect the contents of mesh action data units using the agreed upon ciphersuites. Each MP permits MPDUs protected by the session key and group key using the agreed upon ciphersuites, and discards received MPDUs that are unprotected.

11A.2.2 MSA establishment procedure

11A.2.2.1 General

MSA defines the following procedures for use within a mesh:

— Initial MSA Authentication (8.8.3.1) is used by an MP to authenticate and establish the mesh key hierarchy that may be used when securing future links.

— Subsequent MSA Authentication (8.8.3.2) is used by an MP to securely establish links with peer MPs after it has established a mesh key hierarchy using Initial MSA Authentication.

— MSA Key Holder Communication comprises three related mechanisms:

- The procedure for establishing communications and a security association between an MA and an MKD is the mesh key holder security handshake (8.8.3.3.1).
- The mesh key transport protocol (8.8.3.3.2) describes the mechanisms for key delivery and key management within the mesh key hierarchy.
- The optional mesh EAP message transport protocol (8.8.3.3.3) describes a mechanism for transporting EAP messages between MKD and MA to facilitate authentication of a supplicant MP.

11A.2.2.2 Initial MSA authentication mechanism

During its first authentication in a mesh, an MP establishes the mesh key hierarchy to be used when securing future links. This is referred to as the Initial MSA Authentication Mechanism, and contains communication exchanged between an MP and an MA with which it is associating.

In this sequence, an MP issues an association request frame containing a Peer Link Open information element and an indication (the MKDDIE) that it wishes to establish the mesh key hierarchy. The MP receives an association response frame containing a Peer Link Confirm information element and information required for the MP to perform key derivations and establish link security. If required, IEEE 802.1X authentication occurs next, followed by an MSA 4-way handshake.

The supplicant MP in the Initial MSA Authentication mechanism sends an association request frame to the MA. The association request frame shall contain:

— Peer Link Open information element, which shall be set according to 11A.1.5

— MKDDIE, configured exactly as advertised by the supplicant MP in its Beacon frames and Probe Response frames.

— RSNIE, which shall be set according to the policy in 8.4.3 and 11A.2.1.4, and with the PMKID list field empty.

The association response frame from the MA shall contain:

— Peer Link Confirm information element, which shall be set according to 11A.1.5

— MKDDIE, configured exactly as advertised by the MA in its Beacon frames and Probe Response frames.

— MSAIE, where
  - MA-ID is set to the MAC address of the MA
  - The Optional Parameters field includes the MKD-ID, which contains the identifier of the MKD with which the MA has a security association.
  - The Optional Parameters field includes the EAP Transport List, which contains the list of transport types supported by the MKD with which the MA has a security association.
  - All other fields are set to zero.
— RSNIE, configured exactly as advertised by the MA in its Beacon frames and Probe Response frames, with the PMKID list field empty.

After successful peer link establishment, the supplicant MP and the MA proceed with IEEE 802.1X authentication, if required. The IEEE 802.1X exchange is sent between the supplicant MP and the MA using EAPOL messages carried in IEEE 802.11 data frames. The MA initiates the IEEE 802.1X exchange with the supplicant MP and may transport the IEEE 802.1X exchange to the MKD using the optional mesh EAP message transport protocol, as specified in 11A.2.5.

Upon successful completion of the IEEE 802.1X authentication, the MKD receives the MSK and authorization attributes associated with it and with the supplicant MP. If a mesh key hierarchy already exists for this supplicant, the MKD shall delete the old PMK-MKD SA and PMK-MA SAs. It then calculates the PMK-MKD and PMK-MKDName. The PMK-MKD SA includes:

— MKDD-ID
— PMK-MKD
— PMK-MKDName
— SPA, and
— authorization information including PMK-MKD lifetime.

The MKD then generates a PMK-MA for the MA. The PMK-MA SA includes:

— PMK-MA,
— PMK-MA lifetime,
— PMK-MAName,
— MA-ID,
— PMK-MKDName, and
— SPA

The MKD then delivers the PMK-MA to the MA using the mesh key distribution protocol defined in 11A.2.4. Once the PMK-MA is delivered, the MA and supplicant MP then perform an MSA 4-way handshake. The EAPOL-Key frame notation is defined in 8.5.2.2.

MA -> Supplicant: Data(EAPOL-Key(0, 0, 1, 0, P, 0, 0, ANonce, 0, DataKD_M1)) where DataKD_M1 = 0.

Supplicant -> MA: Data(EAPOL-Key(0, 1, 0, 0, P, 0, 0, SNonce, MIC, DataKD_M2)) where DataKD_M2 = (RSNIE, MKDDIE, GTK KDE).

MA -> Supplicant: Data(EAPOL-Key(1, 1, 1, 1, P, 0, 0, ANonce, MIC, DataKD_M3)) where DataKD_M3 = (RSNIE, MKDDIE, MSAIE, GTK KDE, Lifetime KDE).

Supplicant -> MA: Data(EAPOL-Key(1, 1, 0, 0, P, 0, 0, MIC, DataKD_M4)) where DataKD_M4 = 0.

The message sequence is similar to that of 8.5.3. The contents of each message shall be as described in 8.5.3, except as follows:

— Message 1: ANonce is the value received by the MA from the MKD during PMK-MA delivery. The Key Data field is empty.

— Message 2: The RSNIE shall contain only the PMK-MAName in the PMKID list field. The remaining fields of the RSNIE and the MKDDIE shall be the same as that provided in the association request frame sent by the supplicant MP. The GTK KDE shall contain the GTK of the supplicant MP. The Key Data field shall be encrypted.

— Message 3: The RSNIE shall contain only the PMK-MAName, as calculated by the MA, in the PMKID list field. The remaining fields of the RSNIE, as well as the MKDDIE and MSAIE shall be the same as that provided in the association response frame sent by the MA. The Lifetime KDE shall contain the lifetime of the PMK-MA, expressed in seconds.

The PTK shall be calculated by the supplicant and the MA according to the procedures given in 8.8.6.

Following a successful MSA 4-way handshake, the IEEE 802.1X controlled port shall be opened on both the supplicant and MA. Subsequent EAPOL-Key frames shall use the Key Replay Counter to ensure they are not replayed.

11A.2.2.3 Subsequent MSA authentication mechanism

The subsequent MSA authentication mechanism is used by an MP after it has established a mesh key hierarchy using the initial MSA authentication mechanism. The subsequent MSA authentication mechanism may only be performed between MPs that advertise the same MKD domain identifiers in the MKDDIE in Beacon frames and Probe Response frames.

The subsequent MSA authentication mechanism follows the procedure specified in 11A.2.2.2, Initial MSA authentication mechanism, with the following modifications.

The RSNIE sent by the supplicant MP in the association request frame (peer link open message) shall contain a single entry in the PMKID list field. The value of the entry is PMK-MKDName, identifying the PMK-MKD the supplicant MP created during its initial MSA authentication.

The RSNIE sent by the authenticator MP in the association response frame (peer link confirm message) shall contain the PMKID list entry sent by the supplicant MP in the peer link open message.

After successful peer link establishment, the MA shall calculate the PMK-MAName using the PMK-MKD-Name sent by the supplicant MP in the peer link open message. If the MA does not have the key identified by PMK-MAName, it may attempt to retrieve that key using the mesh key transport protocol according to 11A.2.4. An MA that is unable to retrieve the PMK-MA shall initiate IEEE 802.1X authentication (if required) to establish a mesh key hierarchy for the supplicant MP, and continue with the Initial MSA Authentication mechanism as specified in 11A.2.2.2.

Upon obtaining the specified PMK-MA, the MA omits IEEE 802.1X authentication, and initiates the MSA 4-way handshake, as specified in 11A.2.2.2. However, the ANonce value in message 1 is a pseudo-random nonce created by the MA, and is not the value received during PMK-MA delivery.

Following a successful MSA 4-way handshake, the IEEE 802.1X controlled port shall be opened on both the supplicant and MA. Subsequent EAPOL-Key frames shall use the Key Replay Counter to ensure they are not replayed.

11A.2.2.4 MSA key holder communication

In order to support the mesh key hierarchy, mesh key holders shall communicate securely to provide the following services to MPs:

— transporting EAP traffic between key holders to permit a supplicant MP to perform 802.1X authentication, and

— securely delivering derived keys to facilitate the use of a derived key hierarchy.

An MP invokes the mesh key holder security handshake to establish a security association with a mesh key distributor (MKD). The mechanism permits the MP to operate as a mesh authenticator (MA). Subsequently, the MA advertises, in Beacon frames and Probe Response frames, its capability to authenticate MPs using the mesh key hierarchy.

11A.2.3 Mesh key holder security association

A security association is established between an MA and MKD to provide secure communications between these two key holders within a mesh. The mesh key holder security association is used to provide message integrity and data origin authenticity in all messages passed between MA and MKD after the security association is established. Further, it provides encryption of derived keys and key context during key delivery protocols. Establishing the mesh key holder security association begins with discovery of the MKD, followed by a handshake initiated by the MA. The result of the security association is the pairwise transient key for key derivation (PTK-KD), used to provide the security services between MA and MKD.

11A.2.3.1 Mesh key distributor discovery

Prior to initiating the mesh key holder security handshake described in 11A.2.3.2, an MA shall obtain the address of its MKD. If the MA is not also an MKD, it may obtain the MKD-ID address of its MKD from the MSAIE conveyed in the Association Response frame received during its initial MSA security association.

11A.2.3.2 Mesh key holder security handshake

The mesh key holder security handshake may commence after an MP has completed its Initial MSA Authentication. The mechanism permits the MP to establish a security association with the MKD that derived its PMK-MKD during Initial MSA Authentication.

The MP initiates the exchange by constructing mesh key holder security handshake message 1, and sending the message to the MKD identified by the MKD-ID received in the Association Response frame during the MP's Initial MSA Authentication. The MP selects an EAP Transport mechanism from among those listed in the MSAIE received in the Association Response frame during the MP's Initial MSA Authentication. The MP shall decline to establish a mesh key holder security association with the MKD if the EAP transport mechanisms supported by the MP and MKD do not overlap. The contents of handshake message 1 are given in

Upon receiving handshake message 1, the MKD chooses MKD-Nonce, a value chosen randomly, and computes the PTK-KD using the MA-Nonce received in handshake message 1 and MKD-Nonce, as specified in 8.8.8. The MKD verifies that it supports the selected EAP Transport mechanism; if not, the handshake fails. The MKD sends handshake message 2, with contents as given in 11A.2.3.2.2. Upon receiving handshake message 2, the MP computes the PTK-KD, and sends handshake message 3, with contents as given in 11A.2.3.2.3.

After completing the handshake, the MP sets both the "Mesh Authenticator" and "Connected to MKD" bits to 1 in the MKDDIE in its Beacon frames and Probe Response frames to advertise that it is configured as a mesh authenticator that is connected to the MKD. The MKDDIE shall contain the MKDD-ID that is received from the MKD in mesh key holder security handshake message 2.

An MA shall maintain a route to the MKD. If the route is lost and cannot be repaired, the MA shall set the "Connected to MKD" bit to 0 in the MKDDIE. In such a case, the "Mesh Authenticator" bit may be set to 1 to advertise the ability to act in the IEEE 802.1X Authenticator role using, for example, cached keys. After

the route is re-established, the MP may again set the "Connected to MKD" bit to 1.

The MA and the MKD maintain separate key replay counters for sending messages providing mesh key transport that are protected using the PTK-KD. Immediately upon deriving the PTK-KD, both the MKD and MA shall reset their replay counters to zero.

11A.2.3.2.1 Mesh key holder security handshake message 1

Mesh key holder security handshake message 1 is a mesh key holder security establishment MSA mesh action frame (see 7.4A.1.1) with the following contents:

The MAC address of the MKD shall be asserted in the DA field of the message header.

The MAC address of the MP shall be asserted in the SA field of the message header.

The Mesh ID information element shall contain the Mesh ID that the MP advertises in its Beacon frames and Probe Response frames.

The MKDDIE shall contain the value of MKDD-ID that was contained in the MKDDIE received in the Association Response frame during the MP's Initial MSA Authentication. The Mesh Security Configuration field shall be set to zero.

The MKHSIE shall be set as follows:

- MA-Nonce shall be set to a value chosen randomly by the MP, following the recommendations of 8.5.7.
- MA-ID shall be set to the MAC address of the MP.
- MKD-ID shall be set to the MAC address of the MKD.
- The Transport Type Selector field shall contain a single transport type selector (with format as given in Figure s63). The specified transport type shall be from among those listed in the MSAIE received in the Association Response frame during the MP's Initial MSA Authentication.
- All other fields shall be set to zero.

11A.2.3.2.2 Mesh key holder security handshake message 2

Mesh key holder security handshake message 2 is a mesh key holder security establishment MSA mesh action frame with the following contents:

The MAC address of the MP shall be asserted in the DA field of the message header.

The MAC address of the MKD shall be asserted in the SA field of the message header.

The Mesh ID information element shall contain the Mesh ID as configured in dot11MeshID.

The MKDDIE shall contain the MKDD-ID as configured in dot11MeshKeyDistributorDomainID. The Mesh Security Configuration field shall be set to zero.

The MKHSIE shall be set as follows:

- MA-Nonce, MA-ID, and MKD-ID shall be set to the values contained in handshake message 1.
- MKD-Nonce shall be set to a value chosen randomly by the MKD, following the recommendations of 8.5.7.
- The Transport Type Selector field shall be set to the value contained in handshake message 1.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 3, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MP MAC address
- MKD MAC address
- Handshake sequence number (1 octet), set to the value 2.
- Contents of the Mesh ID information element, from the element ID to the end of the Mesh ID information element.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MKHSIE, from element ID through MIC Control fields, and omitting the MIC field.

11A.2.3.2.3 Mesh key holder security handshake message 3

Mesh key holder security handshake message 3 is a mesh key holder security establishment MSA mesh action frame with the following contents:

The MAC address of the MKD shall be asserted in the DA field of the message header.

The MAC address of the MP shall be asserted in the SA field of the message header.

The Mesh ID information element shall contain the Mesh ID information element received in handshake message 2.

The MKDDIE shall contain the MKDDIE received in handshake message 2.

The MKHSIE shall be set as follows:

— MA-Nonce, MKD-Nonce, MA-ID, MKD-ID, and Transport Type Selector shall be set to the values contained in handshake message 2.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 3, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MP MAC address
- MKD MAC address
- Handshake sequence number (1 octet), set to the value 3.
- Contents of the Mesh ID information element, from the element ID to the end of the Mesh ID information element.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MKHSIE, from element ID through MIC Control fields, and omitting the MIC field.

11A.2.4 Mesh key transport protocol

The mesh key transport protocol describes the method by which the MKD securely transmits a derived PMK-MA to an MA, along with key context and additional related information. An additional management protocol permits the MKD to request that the MA delete a key that has previously been delivered.

Three protocols are defined for mesh key delivery and management, each consisting of 2 messages. The pull

protocol is initiated by the MA by sending a request message, followed by the MKD delivering the PMK-MA. The push protocol is initiated by the MKD delivering (unsolicited) the PMK-MA, followed by the MA sending a confirmation message. Finally, the key delete protocol is initiated by the MKD by sending a message requesting key deletion to the MA, followed by the MA sending a confirmation message.

The MA and MKD maintain separate key replay counters for use in these three protocols. In the pull protocol, the MA's key replay counter is used to protect the first message, which the MA sends. In both the push protocol and the key delete protocol, the MKD's key replay counter is used to protect the first message, which the MKD sends.

In each protocol, prior to sending the first message, the sender shall increment the value of its replay counter. Upon receiving the first message, the recipient shall verify that the replay counter value contained in the first message is a value not yet used by the sender in a first message. If the replay counter value has been previously used, the message shall be discarded. MA and MKD shall each maintain the state of two replay counters: the counter used to generate a value for first messages that it sends, and a counter used to detect replay in first messages that it receives.

Further, the second message of each protocol shall contain a replay counter value that equals the value in the first message of the protocol, to permit matching messages within a state machine instance.

11A.2.4.1 Mesh key transport pull protocol

The key transport pull protocol is a two-message exchange consisting of a PMK-MA request message sent to the MKD, followed by a key delivery sent to the MA. Both messages contain a MIC for integrity protection, and the PMK-MA being delivered is encrypted.

Mesh key transport pull message 1 is a PMK-MA request MSA mesh action frame (see 7.4A.1.4). The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header. Prior to constructing the message, the value of the MA's replay counter associated with the PTK-KD shall be incremented by 1.

The MKDDIE shall be configured exactly as advertised by the MA in its Beacon frames and Probe Response frames.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of the MA's replay counter.

— SPA shall be set to the MAC address of the MP that, during its Initial MSA Authentication, generated the mesh key hierarchy that includes the PMK-MA being requested

— PMK-MKDName shall be set to the identifier of the key from which the PMK-MA being requested was derived.

— ANonce shall be set to zero.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address
  • MKD MAC address
  • Action Value field of the PMK-MA request MSA mesh action frame, which contains the value shown for "PMK-MA request" in Table s13.

- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 1, the MKD shall verify the MIC, and shall verify that the Replay counter field contains a value not previously used with the PTK-KD in a first message sent by the MA. If verified, the MKD may attempt to derive the PMK-MA for use between the MP identified by SPA and the MA that sent message 1, using the key identified by PMK-MKDName. Subsequently, the MKD constructs and sends message 2.

Mesh key transport pull message 2 is a PMK-MA delivery pull MSA mesh action frame (see 7.4A.1.5). The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header.

The MKDDIE shall contain the MKDDIE received in message 1.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of replay counter in message 1.

— SPA shall be set to the value contained in message 1.

— PMK-MKDName shall be set to the value contained in message 1 if an encrypted PMK-MA is included in the Encrypted Contents field. If the Encrypted Contents field is omitted, then PMK-MKDName shall be set to zero.

— ANonce shall be set to the pseudo-random value that was selected by the MKD for derivation of the PMK-MKDName that was indicated in message 1. If the PMK-MKDName field is set to zero, then the ANonce shall be set to zero.

— Encrypted Contents Length field shall be set to the length in octets of the Encrypted Contents field, or shall be set to zero if the Encrypted Contents field is omitted.

— Encrypted Contents shall be set as follows:

- If the MKD does not have a PMK-MA to send to the MA (e.g., it was unable to derive the key), the Encrypted Contents field shall be omitted.
- If the MKD is sending an PMK-MA to the MA, then the Encrypted Contents field shall contain the concatenation: key_data = {PMK-MA || PMK-MAName || Lifetime KDE}.
- Lifetime KDE is defined in Figures 143 and 149. The KDE contains a 4-octet value containing the number of seconds remaining in the lifetime of the PMK-MA.
- If the MIC algorithm is 1 (HMAC-SHA1-128), then the concatenation key_data shall be encrypted using NIST AES Key Wrap algorithm, with the KEK-KD, as defined in RFC 3394, prior to being inserted in the Encrypted Contents field.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Action Value field of the PMK-MA delivery pull MSA mesh action frame, which contains the value shown for “PMK-MA delivery pull” in Table s13.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 2, the MA shall verify the MIC, and shall verify that the Replay counter field contains the value as in message 1.

11A.2.4.2 Mesh key transport push protocol

The key transport push protocol is a two-message exchange consisting of a PMK-MA delivery message sent to the MA, followed by a confirmation message sent in reply. Both messages contain a MIC for integrity protection, and the PMK-MA being delivered is encrypted.

Mesh key transport push message 1 is a PMK-MA delivery push MSA mesh action frame (see 7.4A.1.2). The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header. Prior to constructing the message, the value of the MKD's replay counter associated with the PTK-KD shall be incremented by 1.

The MKDDIE shall contain the MKDD-ID as configured in dot11MeshKeyDistributorDomainID.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of the MKD's replay counter.

— SPA shall be set to the MAC address of the MP that, during its Initial MSA Authentication, generated the mesh key hierarchy that includes the PMK-MA being delivered

— PMK-MKDName shall be set to the identifier of the key from which the PMK-MA being delivered was derived.

— ANonce shall be set to the pseudo-random value that was selected by the MKD for derivation of the PMK-MKDName indicated in this message

— Encrypted Contents Length field shall be set to the length in octets of the Encrypted Contents field.

— Encrypted Contents field shall contain the concatenation: key_data = {PMK-MA || PMK-MAName || Lifetime KDE}

  • Lifetime KDE is defined in Figures 144 and 149. The KDE contains a 4-octet value containing the number of seconds remaining in the lifetime of the PMK-MA.
  • If the MIC algorithm is 1 (HMAC-SHA1-128), then the concatenation key_data shall be encrypted using NIST AES Key Wrap algorithm, with the KEK-KD, as defined in RFC 3394, prior to being inserted in the Encrypted Contents field.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

  • MA MAC address
  • MKD MAC address
  • Action Value field of the PMK-MA delivery push MSA mesh action frame, which contains the value shown for "PMK-MA delivery push" in Table s13.
  • Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
  • Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 1, the MA shall verify the MIC, and shall verify that the replay counter field contains a value not previously used with the PTK-KD in a first message sent by the MKD. If verified, the MA shall send a confirmation message to the MKD.

Mesh key transport push message 2 is a PMK-MA confirm MSA mesh action frame (see 7.4A.1.3). The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header.

The MKDDIE shall contain the MKDDIE received in message 1.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of replay counter in message 1.

— SPA, PMK-MKDName, and ANonce shall be set to the values contained in message 1.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Action Value field of the PMK-MA confirm MSA mesh action frame, which contains the value shown for "PMK-MA confirm" in Table s13.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 2, the MKD shall verify the MIC, and shall verify that the Replay counter field contains the value as in message 1.

11A.2.4.3 Mesh key delete protocol

The MKD may initiate the mesh key delete protocol in order to request that a previously-delivered PMK-MA be revoked. Revocation of the PMK-MA implies that the PMK-MA shall be deleted and all keys derived from the PMK-MA shall be deleted.

The key delete protocol is a two-message exchange consisting of a PMK-MA delete message sent to the MA, followed by a confirmation message sent in reply. Both messages contain a MIC for integrity protection.

Mesh key delete message 1 is a PMK-MA delete MSA mesh action frame (see 7.4A.1.6). The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header. Prior to constructing the message, the value of the MKD's replay counter associated with the PTK-KD shall be incremented by 1.

The MKDDIE shall contain the MKDD-ID as configured in dot11MeshKeyDistributorDomainID.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of the MKD's replay counter.

— SPA shall be set to the MAC address of the MP that, during its Initial MSA Authentication, generated the mesh key hierarchy that includes the PMK-MA that shall be deleted.

— PMK-MKDName shall be set to the identifier of the key from which the PMK-MA that shall be deleted was derived.

— ANonce shall be set to zero.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address

- MKD MAC address
- Action Value field of the PMK-MA delete MSA mesh action frame, which contains the value shown for "PMK-MA delete" in Table s13.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 1, the MA shall verify the MIC, and shall verify that the replay counter field contains a value not previously used with the PTK-KD in a first message sent by the MKD. If verified, the MA shall compute the value of PMK-MAName using the PMK-MKDName and SPA included in message 1. The MA shall revoke the PMK-MA named by PMK-MAName, and shall send a confirmation message to the MKD.

Mesh key delete message 2 is a PMK-MA confirm MSA mesh action frame (see 7.4A.1.3). The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header.

The MKDDIE shall contain the MKDDIE received in message 1.

The contents of the MEKIE are as follows:

— Replay counter shall be set to the value of replay counter in message 1.

— SPA, PMK-MKDName, and ANonce shall be set to the values contained in message 1.

— Encrypted Contents Length field shall be set to 0. The Encrypted Contents field shall be omitted.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 2, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Action Value field of the PMK-MA confirm MSA mesh action frame, which contains the value shown for "PMK-MA confirm" in Table s13.
- Contents of the MKDDIE, from the element ID to the end of the MKDDIE.
- Contents of the MEKIE, from element ID through MIC Control fields, and omitting the MIC field.

Upon receiving message 2, the MKD shall verify the MIC, and shall verify that the Replay counter field contains the value as in message 1.

11A.2.5 Mesh EAP message transport protocol (optional)

This optional protocol describes how the MA may initiate and perform authentication via EAP with the supplicant during the supplicant MP's Initial MSA Authentication. The use of this protocol is selected during the mesh key holder security handshake defined in 11A.2.3.2 and is described by transport type selector 00-0F-AC:0. When the transport type selector specifies any other value, the mechanism for EAP Transport is outside the scope of this standard.

EAP, as described in IETF RFC 3748, is a "lock-step protocol," with request messages sent from the supplicant always receiving a response from the AS. The mesh authentication message transport protocol permits transport of these request and response messages through the mesh, between the MA and the MKD.

The MA initiates IEEE 802.1X authentication with the supplicant by sending a first EAP message to the supplicant. If the MA is configured with the appropriate first EAP message to send, then the MA does so. Oth-

erwise, the MA may request the first EAP message from the AS, using the EAP-Start indication described below. When the MA receives an EAP message from the supplicant, the MA sends an EAP Encapsulation MSA mesh action frame to the MKD that contains the received EAP message. When the MKD has an EAP message, received from the AS and destined for the supplicant, it sends an EAP Encapsulation MSA mesh action frame to the MA containing the EAP message.

The final EAP Encapsulation MSA mesh action frame of a sequence is sent by the MKD, and is given a special type to provide information to the MA. If the EAP authentication of the supplicant provided an "accept" indication to the MKD, then the MKD sends the final message with type "accept" to indicate to the MA that the supplicant should be granted access. Alternatively, if EAP failed, the MKD sends the final message with type "reject" to the MA. Upon reception of an EAP Encapsulation MSA mesh action frame of type "reject," the MA shall terminate the peer link with the supplicant.

When an EAP message is included in a EAP Encapsulation MSA mesh action frame, it is encapsulated within one or several EAP Message IEs. The maximum length EAP message that may be transported is 2231 octets. If the EAP message has length greater than 254 octets, fragmentation is required. In such a case, the EAP message shall be separated into fragments. Each fragment shall be of length 254 octets except the last or only fragment. The maximum number of fragments shall be 9. The fragments are included in one or several EAPMIEs, which each contain a Fragment Control field value indicating the sequence of the fragment, beginning with the value zero. Upon reception, the contents of the EAPMIE EAP Message/Fragment fields are concatenated according to the order indicated in the Fragment Control fields to reconstruct an IETF RFC 3748 EAP message.

The EAP-Start indication is sent from MA to MKD by constructing an EAP Encapsulation request message that contains only a single EAP Authentication information element and no EAP Message IEs.

11A.2.5.1 EAP encapsulation request message

An EAP Encapsulation mesh action message with EAP message type request is sent from MA to MKD, either to transport an EAP message from the supplicant, or to request the AS to initiate EAP ("EAP-Start").

EAP Encapsulation request message is an EAP Encapsulation MSA mesh action frame (see 7.4A.1.7). The MAC address of the MKD shall be asserted in the DA field of the message header, and the MAC address of the MA shall be asserted in the SA field of the message header. The contents of the EAP Authentication information element are as follows:

- EAP Message Type shall be set to 1 to indicate "request".
- Message Token shall be set to a unique nonce value chosen by the MA.
- SPA shall be set to the MAC address of the supplicant MP that is participating in EAP.
- Message Fragments shall indicate the number of EAP Message IEs that are included in this EAP Encapsulation request message.
  - If the MA is sending an "EAP-Start" notification, the Message Fragments field shall be set to zero, and no EAP Message IEs are included in the EAP Encapsulation request message.
- The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.
- The Information element count field of the MIC control field shall be set to 1, the number of information elements in this frame.
- The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:
  - MA MAC address
  - MKD MAC address
  - Contents of the EAP Authentication information element, from element ID through MIC Control fields, and omitting the MIC field.

- If present, contents of all EAPMIEs, from the element ID field of the first EAPMIE, through the EAP Message/Fragment field of the last EAPMIE. The EAPMIEs shall be ordered with increasing Fragment Control field values.

Zero or more EAP Message IEs may be present. If present, the contents of each EAP Message information element are as follows:

— Fragment Control contains the number of the fragment contained in the EAP Message/Fragment field.

— EAP Message/Fragment contains an EAP message with format as defined in IETF RFC 3748, or portion thereof. The maximum size of the EAP message portion is 254 octets.

Upon receiving a request message, the MKD shall verify the MIC, and store the Message Token for use in constructing the response message.

11A.2.5.2 EAP encapsulation response message

An EAP Encapsulation mesh action message with EAP message type response, accept, or reject is sent from MKD to MA, to transport an EAP message from the AS, and, in the final response message of a sequence, provide an indication of the success of EAP.

EAP Encapsulation response message is an EAP Encapsulation MSA mesh action frame (see 7.4A.1.7). The MAC address of the MA shall be asserted in the DA field of the message header, and the MAC address of the MKD shall be asserted in the SA field of the message header. The contents of the EAP Encapsulation information element are as follows:

— EAP Message Type shall be set as follows:

- If this is the final message of the sequence, and the EAP authentication of the supplicant resulted in an “accept” indication, EAP Message Type shall be set to 2, to indicate “accept.”
- If this is the final message of the sequence, and the EAP authentication of the supplicant resulted in a “reject” indication, EAP Message Type shall be set to 3, to indicate “reject.”
- Otherwise, EAP Message Type shall be set to 11, to indicate “response.”

— Message Token shall be set to the value contained in the request message to which this response corresponds.

— SPA shall be set to the value contained in the request message to which this response corresponds.

— Message Fragments shall indicate the number of EAP Message IEs that are included in this EAP Encapsulation request message.

— The MIC algorithm of the MIC control field shall be set to one of the values given in 7.3.2.77.

— The Information element count field of the MIC control field shall be set to 1, the number of information elements in this frame.

— The MIC shall be calculated using the KCK-KD, by the algorithm selected by the MIC algorithm subfield, on the concatenation in the following order, of:

- MA MAC address
- MKD MAC address
- Contents of the EAP Authentication information element, from element ID through MIC Control fields, and omitting the MIC field.
- Contents of all EAPMIEs, from the element ID field of the first EAPMIE, through the EAP Message/Fragment field of the last EAPMIE. The EAPMIEs shall be ordered with increasing Fragment Control field values.

One or more EAP Message IEs shall be present. The contents of each EAP Message information element are as follows:

— Fragment Control contains the number of the fragment contained in the EAP Message/Fragment field.

— EAP Message/Fragment contains an EAP message with format as defined in IETF RFC 3748, or portion thereof. The maximum size of the EAP message portion is 254 octets.

Upon receiving a response message, the MA shall verify the MIC, and verify that the Message Token received in the message matches the value sent in the most recent request message. If the final response message receive has type "reject," the MA shall terminate the peer link with the supplicant.

11A.3 Mesh Path Selection path selection and Forwarding Frameworkforwarding framework

11A.3.1 Overview

The terms "mesh path selection" and "mesh forwarding" are used to describe selection of single-hop or multi-hop paths and forwarding of data frames across these paths between mesh points MPs at the link layer. Data messages use the 802.11 standard four address format with some additional mesh specific information. Client STA nodes associate STAs Associate/Reassociate with one of the Mesh AP devices as normal, since the Mesh AP devices are logically a collection of APs that are part of the same ESS.

Path selection messages are also transported at the link layer, using 802.11 management frames (MMPDUs). Mesh path selection services consist of baseline management messages for neighbor discovery, local link state measurement and maintenance, and identification of an active path selection protocol. Each WLAN Mesh uses a single method to determine paths through the Mesh, although a single device may be capable of supporting several.

11A.3.2 Extensible Path Selection Framework

11A.3.3 Extensible path selection framework

This specification standard includes an extensible framework to enable flexible implementation of path selection protocols and metrics within the mesh framework. The specification standard includes a default mandatory protocol and metric for all implementations, to ensure baseline interoperability between devices from different vendors. However, the specification standard also allows any vendor to implement any protocol and/or metric in the mesh framework to meet special application needs. A mesh point An MP may include multiple protocol implementations (e.g., including the default protocol, optional protocols, future standard protocols, etc.), but only one protocol will shall be active on in a particular link mesh at a time. Different WLAN Meshes may have different active path selection protocols, but a particular mesh will shall have one active protocol at a time.

As described in 11A.1.3 and 11A.1.4, a mesh point an MP uses the WLAN Mesh Capability Information Element to discovery discover which protocol and metric an established WLAN Mesh is using, allowing the mesh point MP to identify if and how it should participate in the mesh. Note that this specification This standard does not force an existing WLAN Mesh that is using a protocol other than the default protocol to switch to the "least common denominator" default protocol when a new mesh point MP requests associationpeer link establishment. While it is possible, in principle, to implement such behavior, an algorithm to coordinate such reconfiguration is beyond the scope of this specificationstandard.

11A.3.4 Path Selection Metrics selection metrics and Protocolsprotocols

As described above, the The mesh extensibility framework allows, in principle, a WLAN Mesh to be implemented with any path selection metric(s) and/or any path selection protocol(s). 11A.5 defines a default radio-aware path selection metric to enable baseline interoperability. This document standard describes two

example path selection protocols that can be implemented in the extensible mesh framework. The first protocol described in 11A.6 is the default path selection protocol that must shall be implemented on all mesh devices to ensure interoperability. The second protocol, described in 11A.7, is an additional optional protocoloptional. Note that the The extensible path selection framework allows these or any other path selection protocols to be implemented in the mesh framework.

11A.3.5 Forwarding of Mesh Data Frames mesh data frames and Mesh Management Framesmesh management frames

11A.3.5.1 General

In a WLAN mesh network, a data frame is forwarded based on its four addresses in the MAC header at each intermediate MP. However, the existing four address fields are not enough to carry the address information for both true source/destination entities and MAPs/MPPs a serving as their proxies when an IEEE 802 entity that does not support WLAN mesh services — e.g., legacy STAs associated with MAPs or external non-IEEE 802.11 stations whose proxies are MPPs — is either a source or a destination of a data linkframe. Also, additional address information is needed when there exist forwarding MPPs that redirect incoming frames to other MPPs or final destinations.

In this regard, as shown in Clause 7.1, the frames of type "Extended with subtype Mesh Data" can use two optional address fields — i.e., "Address 5" and "Address 6" — to carry such additional address information, the existence of which is indicated by the AE flag in the "Mesh Header" field. The use of these optional address fields in the forwarding processing process is described in 11A.3.5.5.

11A.3.5.2 MSDU Ordering

11A.3.5.3 MSDU ordering and duplicate detection

In a WLAN Mesh network, path selection and forwarding operations are implemented as layer-2 mechanisms. When data frames are forwarded in such a multi-hop mesh network, multipath routing (either due to load balancing or dynamic route changes) can easily result in arrival of out-of-order and duplicate frames at the Destination Mesh Pointdestination MP. The probability of having out-of-order and duplicate frames increases as the rate of topology changes, load level variations, and/or wireless channel fluctuations increases. Note that the The Sequence Control field in 802.11 Data Frame data frame headers is meant to be used on a hop-by-hop basis to detect duplicates or missing frames at each hop and is changed by each intermediate Mesh PointMP. Hence, it cannot be used to detect out-of-order or duplicate frame delivery in an end-to-end fashion.

A new "The Mesh E2E Sequence Number" Number in the Mesh Header field is added to uniquely identify the data frames sent from a given Source Mesh PointMP. By the pair of Source MP Address and Mesh E2E Sequence Number, the Destination Mesh Point MP is able to detect out-of-order and duplicate frames. Duplicate frames must shall be discarded, while out-of-order frames must shall be buffered temporarily before they can be re-ordered and delivered to LLC. The goal here is to manage the buffer in a way that strives to deliver all MAC frames in the correct order. To avoid excessive delay due to such buffering, a timer may be used locally by the Mesh Point MP so that it does not wait indefinitely. When the local timer expires, the Destination Mesh Point MP gives up waiting for the missing frames, delivers the queued frames and considers the missing frames as dropped. Note that such Such an end-to-end ordered delivery of unicast individually addressed data frames is only guaranteed between the Source Mesh Point MP and the Destination Mesh PointMP, and it is possible that frames may arrive at an intermediate Mesh Point MP out of order. However, there shall be no reordering of unicast individually addressed MSDUs received at the MAC service interface of any Mesh Point MP with the same traffic identifier value and the same Source Mesh PointMP.

11A.3.5.4 Unicast Forwarding

11A.3.5.5 individually addressed frame forwarding

11A.3.5.5.1 General

In this subclause, an overview of unicast individually addressed frame forwarding based on the mesh address extension is given for the cases of MP-to-MP, STA-to-STA, and outside-mesh-to-outside-mesh communications. Hybrid cases, like MP-to-STA communications, are not described here, but their operations can be easily derived from the given ones; in fact, they are special cases (i.e., a subset) of STA-to-STA communications.

Regarding the details of mesh address extension involved with forwarding MPPs, refer to Clauses subclauses specific to path selection protocols and interworking support (e.g., 11A.6 and 11A.7).

11A.3.5.5.2 At Source source MPs

For simple MP-to-MP communications where there is no forwarding MPP involved, a source MP shall use four-address frames (with AE flag set to 0) where address fields from Address 1 to Address 4 have their usual meanings — i.e. are Address 1 for RA, Address 2 for TA, Address 3 for DA, and Address 4 for SA, respectively.

In case of STA-to-STA communications, STAs associated with MAPs shall use three-address frames to send their unicast individually addressed data frames to MAPs as in the BSS. If a unicast an individually addressed data frame received is from an associated and authenticated STA, and the destination STA address corresponds to an address in the forwarding table[1],, a an MAP shall reformat the frame as a six-address frame (with AE flag set to 1) where the addresses of source and destination STAs are carried by Address 6 and Address 5 fields respectively, and transmit it to the appropriate peer MP[2] listed in the forwarding table as the next hop address for that destination address.

In the general case of outside-mesh to outside-mesh communications through a mesh network, the outside mesh parts use their respective frame formats which at least contain an IEEE 802 source address (SSA) and an IEEE 802 destination address (DDA). The first MP, i.e., the MPP connecting the outside mesh part and the mesh network, puts the 802 Destination and 802 Source MAC addresses of D and S of incoming data frames from outside the outside-mesh into Address 5 and Address 6 of the corresponding, outgoing mesh data frames. The MPP also puts into Address 3 and Address 4 the corresponding mesh destination for D DA derived from its ingress list and its own MAC address, respectively.

The TTL field in the mesh header field for both the all preceding cases shall be set to 255/255 (representing the initial value).

11A.3.5.5.3 At Intermediate intermediate and destination MPs

On receipt of a four address unicast individually addressed data frame, the MP deciphers it and checks for authenticity. If it is not from an authentic source, the frame shall be silently discarded. If the "untrusted" bit is set in the QoS control field, and transport of untrusted traffic has not been enabled, the frame shall be discarded silently.

---

[1]Actual implementation of routing/forwarding tables for associated STAs at MAPs is specific to a path selection protocol in use and therefore beyond the scope of the current standard.

[2] A peer MP is a member of the Mesh network with which a secure link has been established.

The MP then checks to see whether the mesh destination address in "Address 3" field is known; if it is an unknown address, the MP may either silently discard the frame or trigger a routing route discovery procedure depending on the path selection protocol that is currently active in the mesh.

If the mesh destination address does not match the MP's own address, but does match one of the known MAC addresses in the forwarding table, the TTL field in the QoS control field is decremented. If zero has been reached, reached the frame is discarded. Otherwise, the frame is queued for transmission to the next-hop MP as determined from the Mesh forwarding table.

If the mesh destination address matches the MP's own address, then the MP checks the AE flag in the "Mesh Header" field and takes the following actions based on its value:

— If the AE flag is set to 0, which means the current MP is a the final destination of the current frame (i.e., MP-to-MP communications), the MP processes and sends it to an upper layer;

— If the AE flag is set to 1, which means the current MP is a an MAP associated with STAs (i.e., STA-to-STA communications), the MP checks to see whether the destination STA address in "Address 5" field is known. If the destination address corresponds to a STA associated with this MAP, the frame is translated to the three address format and queued for transmission to the final destination.

11A.3.5.6 Broadcast Forwarding frame forwarding

On receipt of a four address data frame with Address 1 set to the all-1s broadcast group address, the MP deciphers it and checks for authenticity. If it is not from a peer MP, the frame shall be silently discarded.

The tuple of SA and Mesh E2E sequence number Sequence Number from the frame header may be used as a unique message signature for tracking messages. The MP checks whether the message has previously been forwarded by this nodeMP. If this is the case, the frame shall be discarded. Otherwise, the signature for this message is retained for later use.

The MP then decrements the TTL field in the mesh control field. If the TTL value has reached zero, the message is discarded. Otherwise, the frame is queued for transmission as a four address frame to all neighboring MPs that are associated and authenticated to the MP.

If the node MP is a Mesh AP, it also creates a three-address broadcast group addressed frame with the same body contents as the received frame and transmits it to the STAs associated with it.

11A.3.5.7 Multicast Forwarding forwarding of Fourfour-Address Framesaddress frames

On receipt of a four address multicast group addressed data frame, the same process used for broadcast group addressed frame forwarding in clause 11A.3.5.6 is used for the multicast data frame.

The MP may implement multicast filtering technology to reduce multicast traffic flooding in the WLAN mesh network. This may be achieved, for example, by using the GARP Multicast Registration Protocol (GMRP) defined in IEEE802IEEE 802.1D. This filtering technology is beyond the scope of this specificationstandard.

Support for special multicast capabilities is an implementation choice and requires invoking the extensibility feature of this (draft standard).

11A.4 Interworking Frameworkframework

11A.4.1 Overview of Interworking interworking in a WLAN Meshmesh

This clause describes the behavior of a Mesh Portal (MPP) to enable bridging of a layer-2 WLAN Mesh to other IEEE 802 LAN segments in a manner compatible with IEEE 802.1D.

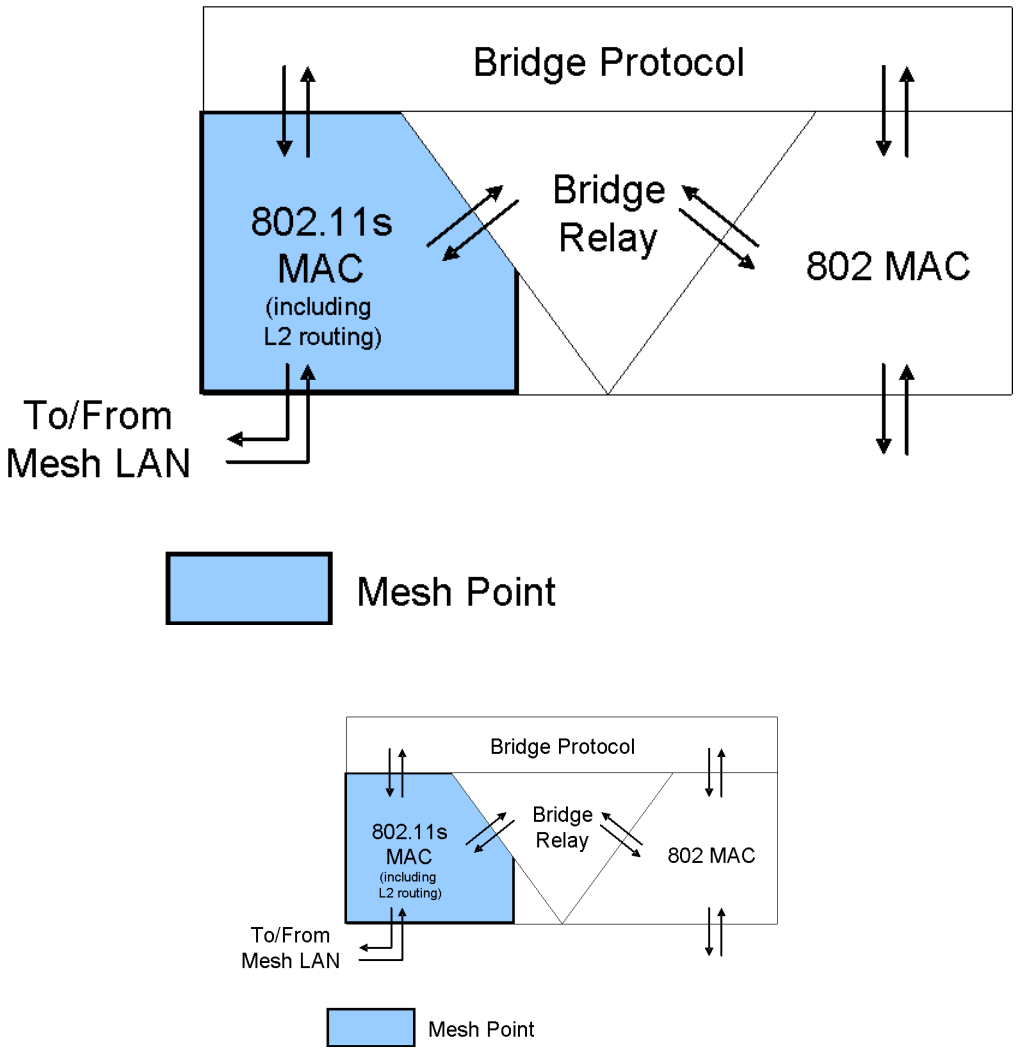

Figure s105—The logical architecture of a Mesh Portal (MPP).

The interaction between WLAN Mesh route selection and layer-2 bridging occurs at the MPP. The MPP consists of an MP (i.e. it has the behavior of a Mesh Pointan MP) collocated with bridging functionality, e.g., based on IEEE 802.1D. The bridging functionality defines the external (relative to the mesh network) behavior of the MPP and is outside the scope of this standard. In general, a mesh network with one or more portals, will behave behaves as a wired LAN segment with one or more bridges.

MPPs use the MPP Announcement Protocol described below to inform other MPs of their presence. Enabling the Announcement Protocol in a given MPP is outside the scope of this documentstandard.

11A.4.2 MPP Announcement Protocolannouncement protocol

This subclause describes the function, generation and processing of the Portal Announcement IEinformation element. This IE information element may be sent in a management frame, possibly in combination with other IEs.

11A.4.2.1 Function

The Portal Announcement IE (PANN) element, described in 7.3.2.69, is used for announcing to announce in the mesh the presence of a an MP configured as a Portal MP an MPP (which has a live connection to an external network). MPs may use this information to increase the efficiency of communication with stations outside the mesh.

An MP which receives a Portal Announcement (PANN) message PANN element shall retransmit forward the PANN as described below but may ignore its contentbelow.

11A.4.2.2 PANN information element

Figure s106—PANN Element

Octets: 1	1	1	1	1	6	4	4
Element ID	Length	Flags	**Hopcount**	**Time to Live**	**Originator Address**	**Sequence Number**	**Metric**

Table s29—PANN Element Fields

Field	Value/description
ID	TBD
Length	Length of the IE
Flags	Not used
Hop Count	The number of hops from the Originator to the MP transmitting the request
Time to Live	Maximum number of hops allowed for this IE
Originator Address	Portal MAC address
Sequence Number	A sequence number specific for the originator.
Metric	The cumulative metric from the Originator to the MP transmitting the announcement.

A routing protocol may provide a means to identify a portal -- in this case the Portal Announcement element need not be sent.

11A.4.2.3 Conditions for generating and sending a PANN

An MP will send out shall transmit a PANN in the following cases:

Case A: Original transmission

All of the following applies:

- The MP is configured as a Portal MP an MPP
- At every PORTAL_ANNOUNCEMENT_INTERVAL

Content:

Field	Value
ID	PANN information element ID
Length	As required
Flags	Not used
Hop Count	0
Time to Live	Maximum number of hops allowed for this IE
Originator Address	Own MAC address
Sequence Number	A sequence number specific for the originator.
Metric	0

Field	Value
ID	PANN IE ID
Length	As required
Flags	Not used
Hop Count	0
Time to Live	Maximum number of hops allowed for this IE
Originator Address	Own MAC address
Sequence Number	A sequence number specific for the originator.
Metric	0

Case B: Re-transmissionForwarding

All of the following applies:

- The MP has received a PANN
- PORTAL_PROPAGATION_DELAY has expired
- The decremented TTL of the PANN is equal to or greater than 1

Content:

Content:

Field	Value
ID	As received
Length	As received
PANN Flags	As received
Hop Count	As received + 1
Time to Live	As received – 1
Originator Address	As received
Sequence Number	As received
Metric	As received + own metric towards the transmitting MP

11A.4.2.4 PANN processing

Received PANN IEs are subject to certain acceptance criteria. Processing depends on the contents of the PANN and the information available at the receiving MP.

11A.4.2.4.1 Acceptance criteria

The PANN will be is discarded if the DSN Sequence Number of the PANN is lower then the DSN Sequence Number of the PANN previously received from this portal.

11A.4.2.4.2 Effect of receipt

The following applies only to PANN messages elements that were not discarded during the check described in 11A.4.2.4.1 above.

a) The receiving MP shall initiate a PANN_PROPAGATION_DELAY

b) The receiving MP shall transmit a PANN as defined in 11A.4.2.3, Case B

11A.4.3 MP behavior

When an MP receives PANN messages elements sent by MPPs in the mesh, it processes these messages and records the MAC addresses and route path metrics to all active MPPs in the mesh.

When an MP has a data message to send, it first follows the data forwarding procedures defined in 11A.3.5. If the MP is not able to determine an intra-mesh route to the destination MAC address, the MP shall assume that the destination is outside the mesh and it shall forward the message to all active MPPs in the mesh (see MPP Egress message handling below11A.4.4.1).

11A.4.4 MPP behavior

MPPs may participate in transparent layer-2 bridging, allowing users to build networks that include a WLAN Mesh in combination with other layer-2 networks. The bridging functions of an MPP are outside the scope of this standard.

MPPs learn the addresses of the MPs and of devices attached to these MPs through

a) the usual bridge learning process

b) the receipt of routing messages

11A.4.4.1 Egress message handling

A packet sent by a MP frame received at the MPP in the WLAN Mesh has the following possible final destinations:

a) A node An MP or STA attached device to an MAP that the MPP knows is inside the Mesh

In this case the The MPP forwards the packet frame to the destination MP.

b) A node An MP that the MPP knows is outside the Mesh

In this case the The MPP forwards the packet frame on the external network.

c) A node An MP unknown to the MPP

In this case the The MPP forwards the packet frame on the mesh and on the external network.

11A.4.4.2 Ingress message handling

A packet frame received by an MPP from an external network (i.e., not from the mesh) has two possible destinations:

a) A An MP or attached station that the MPP knows is inside the Mesh

In this case the The MPP forwards the packet frame to the destination MP.

b) A node An MP unknown to the MPP

In this case the The MPP has two options:

1) To attempt to establish a route to the destination

2) To broadcast the packet frame within the mesh

The criteria for making this choice are outside the scope of this standard.

11A.4.5 Operational Considerations (informative)

11A.4.5.1 Formation and Maintenance of the IEEE 802.1D Spanning Tree

No special action is required to support formation of the 802.1D spanning tree. Spanning tree control messages are typically delivered to bridges in multicast packets. These packets are data packets from the point of view of the WLAN Mesh.

11A.4.5.2 Node Mobility

— Node mobility in a bridged network can be within or between physical LANs. Four cases can occur:

— *Mobility of a node within the mesh.* This kind of mobility is handled through the mesh path selection mechanisms.

— *A node may move from one LAN outside the Mesh to another LAN outside the Mesh.* In this case, the MPPs through which the node can be reached by nodes in the mesh may change. This case occurs in typical bridged networks and can be handled through bridge learning and timing out of old bridge table entries.

— *A node may move from inside the Mesh to outside the Mesh.* When an on-demand routing protocol is used, the movement is detected through the route maintenance mechanisms of the protocol, which triggers route repair procedures. When a proactive routing protocol is used, node failure and information on the new whereabouts of an node are disseminated during triggered and periodic route update rounds.

— *A node may move from outside the Mesh to inside the Mesh.* See 11A.4.3 above.

11A.4.5.3 VLAN support in a WLAN Mesh

A WLAN mesh behaves as a LAN segment which delivers data frames between MPs within the mesh and MPPs which connect the mesh to an external LAN segment, as described above. In order to be conformant with external IEEE 802 networks, a WLAN mesh is required to be compatible with the IEEE 802 architecture, including IEEE802.1D, IEEE802.1Q and IEEE802.1F.

For example, when VLAN tagging as defined in IEEE802.1Q is used, a WLAN mesh network is required to carry VLAN tag information between an MP and an MPP.

A VLAN tag consists of two fields: TPID (The Tag Protocol Identifier) and TCI (The Tag Control Information). TPID is two octets in length and used to represent MAC protocol type. The TCI is two octets in length and contains user_priority, CFI and VID (VLAN Identifier) fields.

IEEE802.1Q defines two header formats: Ethernet-encoded header and SNAP-encoded header. The Ethernet-encoded header format requires a change to the IEEE 802.11 MAC header to adopt a VLAN tag field, whereas the SNAP-encoded header does not require any revisions of the IEEE 802.11 MAC header.

EDITORIAL NOTE—Operational considerations for interworking moved to Annex T per CID 3218

11A.5 Airtime Link Metric Computation Procedureslink metric computation procedures

In order to compute the unicast forwarding table for individually addressed frames from the cached link state information generated by each nodeMP, the MP must shall first calculate the link cost for each pairwise link in the Mesh. This section subclause defines a default link metric that may be used by a path selection protocol to identify an efficient radio-aware path. Note that the The extensibility framework allows this metric to be overridden by any routing metric as specified in the active profile.

The cost function for establishment of the radio-aware paths is based on airtime cost. Airtime cost reflects the amount of channel resources consumed by transmitting the frame over a particular link. This measure is approximate and designed for ease of implementation and interoperability.

The airtime cost for each link is calculated as:

$$c_a = \left[O_{ca} + O_p + \frac{B_t}{r} \right] \frac{1}{1 - e_{pt}}$$

Where O_{ca}, O_p and B_t are constants listed in Table s30, and the input parameters r and e_{pt} e_f are the bit rate in Mbit/s and the frame error rate for the test frame size B_t respectively. The rate r represents the rate at which the mesh point MP would transmit a frame of standard size (B_t) based on current conditions and its estimation is dependent on local implementation of rate adaptation; the frame error rate e_{pt} e_f is the probability that when a frame of standard size (B_t) is transmitted at the current transmission bit rate (r), the frame is corrupted due to transmission error, and its estimation is a local implementation choice. Packet Frame drops due to exceeding TTL should not be included in this estimate as they are not correlated with link performance. The parameters r and e_{pt} e_f are determined during the Local Link State Discovery phase, described in clause 11A.5.1.

Figure s108 shows an example topology with airtime path metric.

Table s31—Airtime Cost Constants

Table s30—Airtime cost constants

Parameter	Value (802.11a)	Value (802.11b)	Description
O_{ca}	75μs	335μs	Channel access overhead
O_p	110μs	364μs	Protocol overhead
B_t	8224	8224	Number of bits in test frame

Parameter	Value (802.11a)	Value (802.11b)	Description
O_{ca}	75μs	335μs	Channel access overhead
O_p	110μs	364μs	Protocol overhead
B_t	8224	8224	Number of bits in test frame

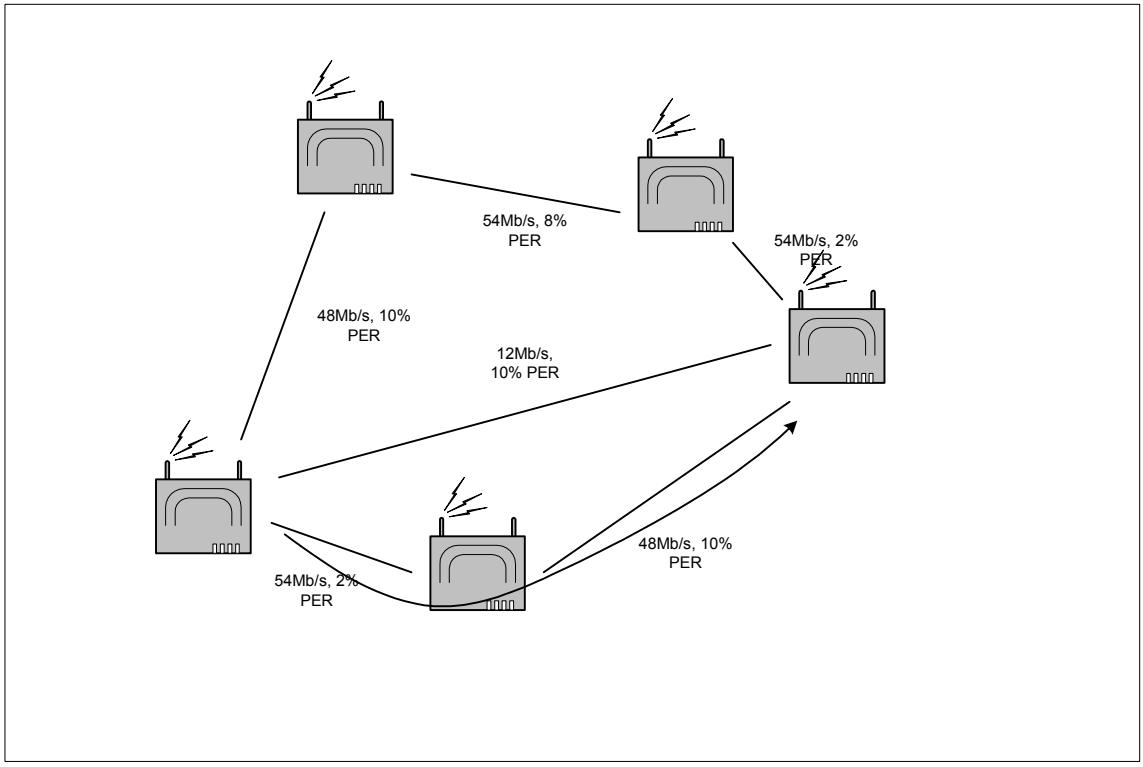

Figure s107—Example Unicast Cost Function based on Airtime Link Metrics

Figure s108—Example cost function for individually addressed frames based on airtime link metrics

11A.5.1 Local Link State Discoverylink state discovery

The purpose of the local link state discovery procedure is to populate the r and e_{pt} e_f fields used by the default airtime metric for each peer MP in the neighbor table.

By definition, a peer link is considered "Down" when there is no value assigned to the *r* and e_{pt} e_f fields. As soon as an initial local link state discovery is completed, and the values are assigned, the link is considered "Up" and remains so until a disassociation event. Since all such links created during initialization are "Down", the state transitions from subordinate, link down, to subordinate, link up, on completion of each link discovery.

The procedures for local link state discovery and maintenance are described in 11A.5.2.

As soon as the first peer link is in an "Up" state, the MP may start to receive frames attempting to establish paths.

11A.5.2 Local Link State Maintenance Procedureslink state maintenance procedures

A pairwise link consists of one node MP designated as superordinate and one designated as subordinate. These labels are illustrative and represent no hierarchical relationship. If it is necessary for the measured link state to be symmetric, the following procedure may be followed.

The superordinate node MP determines the link quality. It may use any method it chooses. The link quality is determined by the following two parameters:

r current bit rate in use, that is, the modulation mode

$e_{pt}e_f$ packet frame error rate at the current bit rate for a data frame with a 1000 1000-byte payload

A superordinate node MP may make this determination for a link in the superordinate, down, state and at future intervals at its option. On making such a determination, it may include the information in a local link state announcement frame and transmit it to the subordinate nodeMP. On successful transmission of the frame, it may update the values in its neighbor MP table with the new values, changing the state from superordinate, down to superordinate, up if this is an initial assessment.

A subordinate node MP shall update the values in its neighbor MP table whenever a local link state announcement message is received.

11A.6 Hybrid Wireless Mesh Protocol (HWMP): Default default path selection protocol for interoperability

11A.6.1 Overview

11A.6.1.1 Rules shared by all routing modes

11A.6.1.2 General

The Hybrid Wireless Mesh Protocol (HWMP) is a mesh routing protocol that combines the flexibility of on-demand routing with proactive topology tree extensions. The combination of reactive and proactive elements of HWMP enables optimal and efficient path selection in a wide variety of mesh networks (with or without infrastructure).

HWMP uses a common set of protocol primitives, generation and processing rules taken from inspired by Ad Hoc On Demand Distance Vector (AODV) routing protocol [IETF RFC 3561] adapted for Layer-2 address-based routing and link metric awareness. AODV forms the basis for finding on-demand routes within a mesh network while additional primitives are used HWMP is completely specified herein and does not require reference to proactively set up a distance-vector tree rooted at a single root MP. The root role that enables building of topology tree is a configurable option of an MP. AODV.

HWMP supports two modes of operation depending on the configuration. These modes are:

On demand mode: this mode allows MPs to communicate using peer-to-peer routes. The mode is used in situations where there is no root MP configured. It is also used in certain circumstances if there is a root configuredMPconfigured.

Proactive tree building mode: this can be performed by using either the RREQ or RANN mechanism.

These modes are not exclusive: on demand and proactive modes may be used concurrently.

All HWMP modes of operation utilize common processing rules and primitives. HWMP control messages information elements are the Route Request (RREQ), Route Reply (RREP), Route Error (RERR) and Root Announcement (RANN). The metric cost of the links determines which routes HWMP builds. In order to propagate the metric information between MPs, a *metric* field is used in the RREQ, RREP and RANN messages.

Routing in HWMP uses a sequence number mechanism to maintain loop-free connectivity at all times. Each MP maintains its own sequence number, which is propagated to other MPs in the HWMP control messagesinformation elements.

11A.6.1.3 On demand routing mode

If a source MP needs to find a route using the on demand routing mode, it broadcasts a RREQ with the destination MP specified in the destination list and the metric field initialized to 0.

When a an MP receives a RREQ it creates a route to the source or updates its current route if the RREQ contains a greater sequence number, or the sequence number is the same as the current route and the RREQ offers a better metric than the current route. If a new route is created or an existing route modified, the RREQ is also forwarded (re-broadcast). Each MP may receive multiple copies of the same RREQ that originated in the source, each RREQ traversing a unique path from the source to the MP.

Whenever a an MP forwards a RREQ, the metric field in the RREQ will be is updated to reflect the cumulative metric of the route to the RREQ's source. After creating or updating a route to the source, the destination MP sends a unicast RREP back to the source.

Intermediate MPs generate RREPs only if the "Destination Only (DO)" flag is not set, e.g. DO = 0, for corresponding destinations, provided that they have routes to these destinations. If the DO flag is set to 1, which is default, only the destination MP can generate a RREP. If an intermediate MP receives a RREQ with the DO flag set to 0 for a destination and this intermediate MP already has a valid route to the destination, it issues a unicast RREP to the source. Furthermore, if the "Reply and Forward (RF)" flag is set to 1 for the destination, this intermediate MP will forward forwards the RREQ with the DO flag for the destination set to 1 (the reason to set the DO flag to 1 is to suppress any RREP messages from the subsequent intermediate MPs). Otherwise, there is no RREQ forwarding at intermediate MP. The purpose of the "Destination Only" and "Reply and Forward" mechanisms is to enable a an MP to quickly establish a route using the RREP generated by the intermediate MP and send data frames with a low route discovery delay and buffer requirement, while allowing that the route with the best route path metric will be is chosen (or validated) after the reverse route establishment procedure has been completed. The source sets the DO flag to 0 and RF flag to 1 for a destination in the RREQ only when it does not have a valid route and wants to discover a new route to this destination. As described below, the DO flag in the maintenance RREQ is always set to 1.

Intermediate MPs create a route to the destination on receiving the RREP, and also forward the RREP toward the source. When the source receives the RREP, it creates a route to the destination. If the destination receives further RREQs with a better metric, then the destination updates its route to the source to the new route and also sends a fresh RREP to the source along the updated route. Thus a A bidirectional, best metric end-to-end route is established between the source and destination.

Note that the RREQ processing when the "Destination Only" flag is set to 0 but the "Reply and Forward" flag is set to 1, as specified in this standard, is different from that of the original AODV specification. Also, note that in HWMP, the RREQ processing at intermediate MPs is controlled per destination.

In HWMP the RREQ processing at intermediate MPs is controlled per destination.

11A.6.1.4 Proactive tree building mode

There are two mechanisms for proactively disseminating routing information for reaching the root MP. The first method uses a *proactive* Route Request (RREQ) message and is intended to create routes between the root MP and all MPs in the network proactively. The second method uses a Root Annoucement Announcement (RANN) message and is intended to distribute route information for reaching the root MP but the actual routes to the root MP can be built on-demand.

An MP configured as root MP would send either proactive RREQ or RANN messages periodically.

11A.6.1.4.1 Proactive RREQ mechanism

The RREQ tree building process begins with a proactive *Route Request* message sent by the root MP, with the destination address set to all ones (broadcast address), the DO flag set to 1 and the RF flag set to 1. The RREQ contains the distance metric (set to 0 by the rootroot MP) and a sequence number. The proactive *RREQ* is sent periodically by the rootroot MP, with increasing sequence numbers.

Any MP hearing a proactive RREQ creates or updates its forwarding information to the root MP, updates the metric and hop count of the RREQ, records the metric and hop count to the rootroot MP, and then transmits the updated RREQ. Thus, information Information about the presence of and distance to available rootroot MP(s) is disseminated to all MPs in the network.

Each MP may receive multiple copies of a proactive RREQ, each traversing a unique path from the root MP to the MP. A An MP updates its current route to the root MP if and only if the RREQ contains a greater sequence number, or the sequence number is the same as the current route and the RREQ offers a better metric than the current route to the rootroot MP. The processing of the proactive RREQ is the same as in the on-demand mode described in 11A.6.1.3.

If the proactive RREQ is sent with the "Proactive RREP" bit set to 0, the recipient MP may send a gratuitous RREP if required (for example, if the MP has data to send to the root MP and requires establishing a bidirectional route with the rootroot MP). If the RREQ is sent with a "Proactive RREP" bit set to 1, the recipient MP shall send a gratuitous RREP. The gratuitous RREP establishes the route from the root MP to the MP. When the route from an MP to a root MP changes, and the root MP RREQ was sent with a "Proactive RREP" bit set to 1, it the recipient MP shall send a gratuitous RREP to the root MP containing the addresses of the MPs which have established a route to the root MP through the current MP.

11A.6.1.4.2 Proactive RANN mechanism

The root MP periodically floods a RANN message into the network. The information contained in the RANN is used to disseminate route path metrics to the rootroot MP.

Upon reception of a RANN, each MP that has to create or refresh a route to the root will send MP sends a unicast RREQ to the root MP via the MP from which it received the RANN.

The unicast RREQ will follow follows the same processing rules defined in the on demand mode (11A.6.1.3)mode.

The root MP sends a RREP in response to each RREQ. The unicast RREQ creates the reverse route from the root MP to the originating MP, while the RREP creates the forward route from the MP to the rootroot MP.

When the route from an MP to a root MP changes, it may send a RREP with the addresses of the MPs which have established a route to the root MP through the current MP.

11A.6.2 Definitions

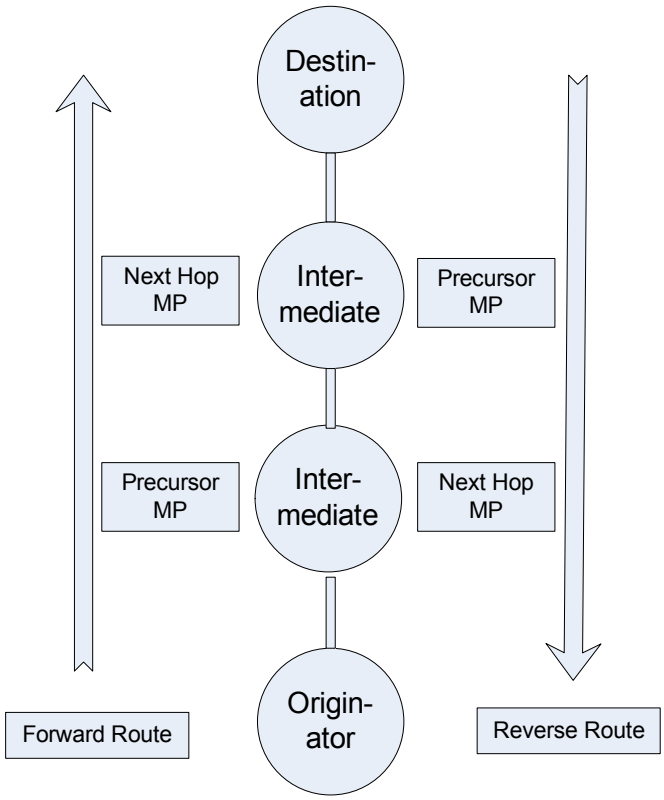

Figure s109—Illustration of definitions

Originator:

Originator: The originator is the MP that initiates an HWMP routing message that is flooded through the whole wireless mesh network (e.g. RREQ). The transmission of the corresponding information element by the originator is called “original transmission”.

Forwarding Information

Forwarding Information: The forwarding information maintained by an MP consists of at least a destination, the Destination Sequence Number (DSN), the next hop to the destination, the route path metric and the lifetime of this forwarding information.. Stored forwarding information can be active or inactive (invalidated). The latter means that the forwarding information is still known for future reference but not used for forwarding.

Forward Route:

Forward Route: The stored forwarding information of the path from the originator to the destinaton destination of a route discovery. Set up by route replies.

Reverse Route:

Reverse Route: The stored forwarding information of the path from the destination to the originator of a route discovery. Set up by route requests.

Unreachable destination:

Unreachable destination: A destination is considered unreachable in two cases:

— An MP does not have active forwarding information for the destination and the MP did not receive a RREP in response to its route request for this destination within the set waiting time.

— The destination cannot be reached over the established path, because it contains a broken link.

Destination Sequence Number: An integer sequence number associated with an MP, which is used to distinguish newer from older routing information. See also section XXX11A.6.3.3 (Destination Sequence Number).

Time-to-Live: An integer number that is used to limit the number of hops of frame frames in the WLAN mesh network. See also 11A.6.3.4.

Unicast/Broadcast RREQ: A “unicast RREQ” is a RREQ IE information element contained in a management frame that is unicast transmitted in an individually addressed frame to the next hop toward the destination of the RREQ. A “broadcast RREQ” is an a RREQ IE information element contained in a management frame that is broadcast.

Local MP: The mesh point which is the reference MP for the corresponding descriptions.

Dependent MP: A neighboring MP which has established a route to the root MP through the local MP.

11A.6.3 General rules for processing HWMP information elements

This subclause describes the rules for the processing of the following components of the HWMP messages and IEs:

— Destination Sequence Number

— TTL

— Metric

Note: It is assumed that the receiving MP will only receive receives HWMP messages from MPs with which it has established a secure link. Therefore, all HWMP messages received are presumed to have originated in the same mesh network that the receiving MP belongs to.

11A.6.3.1 Re-transmission

11A.6.3.2 Forwarding

Many HWMP IEs are intended to be flooded across a mesh network and therefore are retransmitted forwarded by MPs that receive them. Each retransmission forwarding is subject to certain rules or limitations as explained in the following subclauses. Some parameters in the HWMP IEs will be is updated before retransmissionforwarding, e.g. route path metric.

11A.6.3.3 Destination Sequence Number (DSN)

The DSN is set by the originator of a RANN, RREQ or RREP information element.

The value of the DSN is specific to each originator; it is incremented with each new instance of one of the aforementioned IE information element types. In case the DSN reaches its maximum value, the incremented DSN becomes the lowest DSN value.

In general, when an MP receives an information element with a DSN that is less than the last received DSN for that originator, it will discard discards the received IEinformation element. If they are the same, the outcome (IE discarded or not) depends on the type of the information element and some additional conditions. These cases are noted in the applicable IE information element descriptions.

When the A destination increments MP shall increment its sequence number, it shall do so number by treating the sequence number value as if it were an unsigned number. To accomplish sequence number rollover, if the sequence number has already been assigned to be the largest possible number representable as a 32-bit unsigned integer (i.e., 4294967295), then when it is incremented it will then have has a value of zero (0). On the other hand, if the sequence number currently has the value 2147483647, which is the largest possible positive integer if 2's complement arithmetic is in use with 32-bit integers, the next value will be is 2147483648, which is the most negative possible integer in the same numbering system. The representation of negative numbers is not relevant to the increment of HWMP sequence numbers. This is in contrast to the manner in which the result of comparing two HWMP sequence numbers is to be treated.

In order to ascertain that information about a destination is not stale, the MP compares its current numerical value for the sequence number with that obtained from the incoming HWMP message. This comparison shall be done using signed 32-bit arithmetic, this is necessary to accomplish sequence number rollover. If the result of subtracting the currently stored sequence number from the value of the incoming sequence number is less than zero, then the information related to that destination in the HWMP message shall be discarded, since that information is stale compared to the MP's currently stored information.

11A.6.3.4 Time-to-Live (TTL)

The TTL is set by the originator of an HWMP IEinformation element. It gives the maximum number of hops that the IE information element is allowed to be forwarded.

A receiving MP will decrement decrements the value of the TTL field before re-transmitting forwarding the IEinformation element. If the decremented value is zero, the IE will be information element is discarded.

11A.6.3.5 Re-transmission delay

11A.6.3.6 Forwarding delay

In general, retransmission forwarding of an HWMP IE information element is not subject to a delay. Exceptions exist for the RANN IE information element as described below.

11A.6.3.7 Forwarding Informationinformation

The forwarding information maintained by a an MP consists at least of a destination, the DSN of the destination, the next hop to the destination, the route path metric to the destination and the lifetime of this forwarding information.

Stored forwarding information can be active or inactive (invalidated). The latter means that the forwarding information is still known for future references but not used for forwarding.

11A.6.3.8 Creation and Update update of Forwarding Informationforwarding information

HWMP path selection information elements create or update the forwarding information in the MPs that process these information elements. The creation and update of forwarding information follows the same rules for RREQ, RREP, and RANN. These rules are given below. "HWMP_IE" stands for the IE information element under consideration (RREQ, RREP, or RANN). Exceptions are stated in the corresponding sections subclauses of the descriptions of the information elements. In the following text, the "→" indicates an entry into the local forwarding data base, i.e., the routing table.

1) If the MP does not have stored any active forwarding information to the originator of the HWMP_IE (hwmp_ie.originator_address), it creates this information from the field hwmp_ie.originator_address (→ destination), the transmitter address of the management frame containing the HWMP_IE (→ next hop), the accumulation of the value of field hwmp_ie.metric with the metric of the last link (→ path metric), and the value of field hwmp_ie.lifetime (→ lifetime).
2) If the MP does not have stored any active forwarding information to the transmitter of the HWMP_IE, it creates this information from the transmitter address of the management frame containing the HWMP_IE (→ destination and next hop), the metric of the last link (→ path metric), and the value of field hwmp_ie.lifetime (→ lifetime).
3) If the MP has stored any active forwarding information to the originator of the HWMP_IE (hwmp_ie.originator_address), then the MP updates this forwarding information with the transmitter address of the management frame containing the HWMP_IE (→ next hop), the accumulation of the value of field hwmp_ie.metric with the metric of the last link (→ path metric), and the larger one of the lifetime of the stored forwarding information and the value of field hwmp_ie.lifetime (→lifetime).
4) If the MP has stored any active forwarding information to the transmitter of the HWMP_IE and if the path metric of this information is larger than the metric of the last link, then the MP updates this forwarding information with the transmitter address of the management frame containing the HWMP_IE (→ next hop), the metric of the last link (→ path metric), and the larger one of the lifetime of the stored forwarding information and the value of field hwmp_ie.lifetime (→ lifetime).

11A.6.3.9 Metric of Last Linklast link

The term *metric of last link* specifies the current link metric between the transmitter of the IE information element under consideration and the mesh point MP that received this IEinformation element. The latter is the mesh point MP under consideration.

11A.6.4 Root Announcement (RANN)

This subclause describes the function, generation and processing of the *Root Announcement* IEinformation element.

11A.6.4.1 Function

This IE The RANN element, described in 7.3.2.70, is used for announcing the presence of a an MP configured as Root MP. RANN IEs are sent out periodically by the root MP.

The RANN IE information element propagates route path metric information across the network so that each MP can select a best metric path to the announced root MP. This mechanism allows bidirectional trees to be built, using a robust unicast procedure based on individually addressed frames initiated by the MPs. This ensures that the root MP is aware of all MPs in the mesh.

Receiving MPs shall propagate the RANN as described below.

11A.6.4.2 RANN information element

Figure s110—RANN Element

Octets: 1	1	1	1	1	6	4	4	4
Element ID	Length	Flags	**Hopcount**	**Time to Live**	**Originator Address**	**Destination Sequence Number**	**Lifetime**	**Metric**

Table s32—RANN Element Fields

Field	Value/description
ID	TBD
Length	Length of the IE
Flags	Bit 0: Portal Role (0 = non-portal, 1 = portal) Bit 1 – 7: Reserved
Hop Count	The number of hops from the Originator to the MP transmitting the request
Time to Live	Remaining number of times the RANN may be forwarded.
Originator Address	Root MAC address
Destination Sequence Number	A sequence number specific to the originator (root).
Lifetime	The time for which MPs receiving the RANN consider the forwarding information to be valid.
Metric	The cumulative metric from the Originator to the MP transmitting the announcement.

11A.6.4.3 Conditions for generating and sending a RANN

An MP will send sends out a RANN in the following cases:

Case A: Original transmission

All of the following applies:

- The MP is configured as a Root MP
- at every ROOT_ANNOUNCEMENT_INTERVAL

Content:

Content:

Field	Value
ID	RANN IE information element ID
Length	
Flags	Bit 0: Portal Role (0 = non-portal, 1 = portal) Bit 1 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this IE
Originator Address	Own MAC address
Destination Sequence Number	Last used DSN + 1
Lifetime	HWMP_PATH2ROOT_TIMEOUT (shall be greater than ROOT_ANNOUNCEMENT_INTERVAL)
Metric	initial metric value (0 for airtime metric)

Case B: Re-transmissionForwarding

All of the following applies:

- the MP has received and accepted a RANN – See 11A.6.4.4
- RANN_PROPAGATION_DELAY has expired – See 11A.6.4.4.2

Content:

Field	Value
ID	As received
Length	As received
Flags	As received
Hop Count	As received + 1
TTL	As received – 1
Originator Address	As received
Destination Sequence Number	As received
Lifetime	As received
Metric	As received + link metric to the transmitting MP

Field	Value
ID	As received
Length	As received
Flags	As received
Hop Count	As received + 1
TTL	As received – 1
Originator Address	As received
Destination Sequence Number	As received
Lifetime	As received
Metric	As received + link metric to the transmitting MP

11A.6.4.4 RANN Reception

Received RANN IEs are subject to certain acceptance criteria. Processing and actions taken depend on the content of the RANN and the forwarding information maintained by the receiving MP.
See also 11A.6.3: General rules for message processing

11A.6.4.4.1 Acceptance criteria

The RANN will be is discarded if any of the following is true:

— The DSN < previous DSN from this originator
— ((DSN = previous DSN DSN) AND (updated path metric >= previous path metric metric)

11A.6.4.4.2 Effect of receipt

The following applies only to RANN that has not been discarded.

1) The receiving MP shall initiate a timer for RANN_PROPAGATION_DELAY
2) The receiving MP may
    1) initiate a RREQ/RREP exchange with the Root root MP to set up or update a route to the Rootroot MP.
    See RREQ, when generated, case C (CHECK)C
    2) send a gratuitous RREP to the Rootroot MP. See: RREP, when generated, case D
3) The receiving MP shall record the Originator Address, together with the DSN, hopcount, metric.
4) The receiving MP shall transmit a RANN as defined in 11A.6.4.3, Case B
5) The receiving MP shall maintain one root MP as its active rootroot MP—this choice may be changed at any time, e.g. if the metric of the received RANN justifies this.

11A.6.5 Route Request (RREQ)

This subclause describes the function, generation and processing of the Route Request IEinformation element.

11A.6.5.1 Function

This IE is used for three purposes:
The Route Request (RREQ) element, described in 7.3.2.71, is used for three purposes:

— Discovering a route to one or more destinations
— Building a proactive (reverse) routing tree to the root MP
— Confirming a route to a destination (optional)

11A.6.5.2 RREQ information element

Figure s111—RREQ Element

Octets: 1	1	1	1	1	4	6	4	4
Element ID	Length	Flags	**Hopcount**	**Time to Live**	**RREQ ID**	**Originator Address**	**Originator Sequence Number**	**Lifetime**

4	1			6	4		1			6	4
Metric	Per Destination Flags			Destination Address #1	Destination Seq. Num.#1	....	Per Destination Flags			Destination Address #N	Destination Seq. Num.#N
	DO #1	RF #1	Reserved				DO #N	RF #N	Reserved		

Table s33—RREQ Element Fields

Field	Value/description
ID	TBD
Length	Length of the IE
Flags	Bit 0: Portal Role (0 = non-portal, 1 = portal) Bit 1: (0 = broadcast, 1 = unicast) (see 11A.6.2) Bit 2: RREP (0 = off, 1 = on) Bit 3 – 7: Reserved
Hop Count	The number of hops from the Originator to the MP transmitting the request
Time to Live	Maximum number of hops allowed for this IE
RREQ ID	Some unique ID for this RREQ
Originator Address	Originator MAC address
Originator's Destination Sequence Number	A sequence number specific for the originator
Lifetime	The time for which MPs receiving the RREQ consider the forwarding information to be valid.
Metric	The cumulative metric from the Originator to the MP transmitting the RREQ
Destination Count N	Gives the number of Destinations (N) contained in this RREQ
Per Destination Flags	Flags Bit 0: DO (Destination Only): If DO=0, an intermediate MP with active forwarding information to the corresponding destination shall respond to the RREQ with a unicast RREP; if DO=1, only the destination can respond with a unicast RREP. The default value is 1. Bit 1: RF (Reply-and-Forward):. The RF flag controls the forwarding of RREQ at intermediate MPs. When DO=0 and the intermediate MP has active forwarding information to the corresponding destination, the RREQ is not forwarded if RF=0 and forwarded if RF=1. The default value is 1. When DO=1, the RF flag has no effect.. Bit 2-7: Reserved
Destination Address	MAC address of the destination MP
Destination Sequence Number	The latest sequence number received in the past by the originator for any route towards the destination.

11A.6.5.3 Conditions for generating and sending a RREQ

An MP will shall send out a RREQ element in a RREQ frame in the following cases:

Case A: Original Transmission (Route Discovery)
All of the following applies:

- The MP needs to establish an on-demand route to one or more destinations for which there is no ongoing route discovery initiated by this MP.

- The MP has not sent more than (HWMP_RREQ_RATELIMIT – 1) route request messages during the last second. If this is the case, the transmission of the RREQ has to be postponed until this condition becomes true.
- The MP has not made more than (HWMP_MAX_RREQ_RETRIES – 1) repeated attempts at route discovery towards the destination of the RREQ.

Content:

Field	Value
ID	TBD
Length	27 + N*11
Flags	Bit 0: 0 (no portal role) Bit 1: 0 (broadcast) Bit 2: 0 (no proactive RREP applicable) Bit 3 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this information element, e.g., HWMP_NET_DIAMETER.
RREQ ID	Previous RREQ ID + 1
Originator Address	Own MAC address
Originator's Destination Sequence Number	Previous Originator DSN + 1. See Note 2
Lifetime	The time for which MPs receiving the RREQ consider the forwarding information to be valid, e.g. HWMP_ACTIVE_ROOT_TIMEOUT.
Metric	0
Destination Count	(N)
Per Destination Flags	DO flag, RF flag, as required
Destination Address	MAC address of requested destination
Destination Sequence Number	The latest sequence number received in the past by the originator for any route towards the destination.

Field	Value	
ID	Tbd	
Length	27 + N*11	
Flags	Bit 0: 0 (no portal role) Bit 1: 0 (broadcast) Bit 2: 0 (no proactive RREP applicable) Bit 3 – 7: Reserved	
Hop Count	0	

Time to Live	Maximum number of hops allowed for this IE, e.g., HWMP_NET_DIAMETER.	
RREQ ID	Previous RREQ ID + 1	
Originator Address	Own MAC address	
Originator's Destination Sequence Number	Previous Originator DSN + 1. See Note 2	
Lifetime	The time for which nodes receiving the RREQ consider the forwarding information to be valid, e.g. HWMP_ACTIVE_ROOT_TIMEOUT.	
Metric	0	
Destination Count	(N)	
Per Destination Flags	DO flag, RF flag, as required	
Destination Address	MAC address of requested destination	
Destination Sequence Number	The latest sequence number received in the past by the originator for any route towards the destination.	

Note 1: Repeated attempts by an MP at route discovery towards a given (set of) destination(s) shall be limited to HWMP_MAX_RREQ_RETRIES and utilize a binary exponential backoff between transmissions. The minimum waiting time for the RREP corresponding to an a RREQ is 2 * HWMP_RT_NETDIAMETER_TRAVERSAL_TIME.

Note 2: In order to improve route stability (and further reduce overhead), an MP may use the same Originator Destination Sequence Number (Originator DSN) for a certain time interval. The Originator DSN is incremented only after at least HWMP_RT_NETDIAMETER_TRAVERSAL_TIME has elapsed since the previous increment. This mechanism prevents MPs from changing the route frequently to the source every time the source sends a burst of RREQs within a very short time. This element of the protocol allows a source MP to immediately initiate on-demand route discovery to a new destination without affecting recently refreshed routes to the source in other MPs.

Case B: Original Transmission (Route Maintenance) (optional implementation enhancement)
All of the following applies:

- the MP has a route to a given destination that is not a Root MP
- the HWMP_ROUTE_MAINTENANCE_INTERVAL has expired

Content:

Field	Value
ID	TBD
Length	As required
Flags	Bit 0: 0 (no portal role) Bit 1: 0 (broadcast) Bit 2: 0 (no proactive RREP applicable) Bit 3 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this information element = HWMP_NET_DIAMETER.
RREQ ID	Previous RREQ ID +1
Originator Address	Own MAC address
Originator's Destination Sequence Number	Originator DSN + 1. See Note 2 under Case A.
Lifetime	The time for which MPs receiving the RREQ consider the forwarding information to be valid, e.g. HWMP_ACTIVE_ROOT_TIMEOUT.
Metric	0
Destination Count	N
Per Destination	Flags = DO flag = 1, RF flag = 0
	Address = Destination MAC Address
	Destination Sequence Number (for this destination)

Field	Value	
ID	Tbd	
Length	As required	
Flags	Bit 0: 0 (no portal role) Bit 1: 0 (broadcast) Bit 2: 0 (no proactive RREP applicable) Bit 3 – 7: Reserved	
Hop Count	0	
Time to Live	Maximum number of hops allowed for this IE = HWMP_NET_DIAMETER.	
RREQ ID	Previous RREQ ID +1	
Originator Address	Own MAC address	
Originator's Destination Sequence Number	Originator DSN + 1. See Note 2 under Case A.	
Lifetime	The time for which nodes receiving the RREQ consider the forwarding information to be valid, e.g. HWMP_ACTIVE_ROOT_TIMEOUT.	
Metric	0	

Destination Count	Gives the number of Destinations contained in this RREQ
Per Destination	Flags = DO flag = 1, RF flag = 0
	Address = Destination MAC Address
	Destination Sequence Number (for this destination)
Etc	

Case C: Root Path Confirmation.

All of the following applies:

- the MP has a route to a Root MP which broadcasts RANNs
- the HWMP_ROUTE_CONFIRMATION_INTERVAL has expired

Content:

Content:

Field	Value
ID	TbdTBD
Length	As required
Flags	Bit 0: 0 (no portal role) Bit 1: 1 (unicastindividually addressed) Bit 2: 0 (no proactive RREP applicable) Bit 3 – 7: Reserved
Hop Count	0
Time to Live	1
RREQ ID	Not used
Originator Address	Own MAC address
Originator's Destination Sequence Number	Originator DSN + 1. See Note 2 under Case A.
Lifetime	The time for which nodes MPs receiving the RREQ consider the forwarding information to be valid, e.g. HWMP_ACTIVE_ROOT_TIMEOUT.
Metric	0
Destination Count	1
Per Destination	Flags: DO flag = 1, RF flag = 0
	Address = Root MP MAC Address
	Destination Sequence Number (for this destination)
Etc	

Case D: RREQ Re-transmissionForwarding

Case D1 (destination count = 1, no RREP generation):

All of the following applies:

- the MP has received and accepted a RREQ – See 11A.6.5.4.1
- Destination_count = 1
- the MP is not the destination of the RREQ OR the destination of the RREQ is the MAC broadcast address (all 1's)
- the MP has no active forwarding information for the requested destination
- [Destination Only flag of the destination in the RREQ is ON (DO = 1)]
  OR
  [{Destination Only flag of the destination in the RREQ is OFF (DO = 0)} AND {MP has no active forwarding information for the requested destination rreq.destination_address}]

Content for D1:

Field	Value
ID	TbdTBD
Length	37
Flags	as received
Hop Count	As received + 1
Time to Live	As received – 1
RREQ ID	As received
Originator Address	As received
Originator's Sequence Number	As received
Lifetime	As received
Metric	As receiver + own metric towards transmitter of received RREQ
Destination Count	1
Per Destination flags #1	as received
Destination MAC address #1	as received
Destination Sequence Number #1	as received

Case D2 (destination count = 1, RREP generation as intermediate MP):

All of the following applies:

- the MP has received and accepted a RREQ – See 11A.6.5.4.1
- rreq.destination_count = 1

- the MP is not the destination of the RREQ
- the MP has active forwarding information for the requested destination rreq.destination_address
- Destination Only flag of the destination in the RREQ is OFF (DO = 0)
- Reply and Forward flag of the destination in the RREQ is ON (RF = 1)

Content for D2:

Field	Value
ID	TbdTBD
Length	37
Flags	as As received
Hop Count	As received + 1
Time to Live	As received – 1
RREQ ID	As received
Originator Address	As received
Originator's Sequence Number	As received
Lifetime	As received
Metric	As receiver + own metric towards transmitter of received RREQ
Destination Count	1
Per Destination flags #1	Bit 0 (DO): 1 (set to 1 before forwarding because MP sent a RREP) Bit 1 (RF): as received
Destination MAC address #1	as As received
Destination Sequence Number #1	as As received

Case D3 (destination count > 1): All of the following applies

- the MP has received and accepted a RREQ – See 11A.6.5.4.1
- there is at least one requested destination left after processing the RREQ according to 11A.6.5.4.

Content for D3:

Field	Value
ID	TBD
Length	26 + N * 11
Flags	As received
Hop Count	As received + 1
Time to Live	As received – 1
RREQ ID	As received
Originator Address	As received
Originator's Sequence Number	As received
Lifetime	As received
Metric	As receiver + own metric towards the transmitter
Destination Count	1 ≤ destination count ≤ received destination count received destination count less the number of requested destinations, for which the processing MP ² is the destination or ² has active forwarding information for the requested destination and the corresponding Destination Only flag is off (DO=0) and Reply and Forward flag is on (RF = 1)
For the per destination fields (per destination flags, destination MAC address, destination sequence number) assume the following: ² *destination #A: If destination A would have been the only requested destination, it would generate a RREQ for forwarding according to case **D1*** ² *destination #B: If destination B would have been the only requested destination, it would generate a RREQ for forwarding according to case **D2***	
Per Destination Flags #A	As received
Destination MAC address #A	As received
Destination Sequence Number #A	As received
Per Destination Flags #B	Bit 0 (DO): 1 (set to 1 because MP sent RREP) Bit 1 (RF): as received
Destination MAC address #B	As received
Destination Sequence Number #B	As received

Field	Value	
ID	Tbd	

Length		
Flags	as received	
Hop Count	As received + 1	
Time to Live	As received – 1	
RREQ ID	As received	
Originator Address	As received	
Originator's Sequence Number	As received	
Lifetime	As received	
Metric	As receiver + own metric towards the transmitter	
Destination Count	1 ≤ destination count ≤ received destination count received destination count less the number of requested destinations, for which the processing MP ² is the destination or ² has active forwarding information for the requested destination and the corresponding Destination Only flag is off (DO=0) and Reply and Forward flag is on (RF = 1)	
For the per destination fields (per destination flags, destination MAC address, destination sequence number) assume the following: ² *destination #A: If destination A would have been the only requested destination, it would generate a RREQ for retransmission according to case **D1*** ² *destination #B: If destination B would have been the only requested destination, it would generate a RREQ for retransmission according to case **D2***		
Per Destination Flags #A	as received	
Destination MAC address #A	as received	
Destination Sequence Number #A	as received	
Per Destination Flags #B	Bit 0 (DO): 1 (set to 1 because MP sent RREP) Bit 1 (RF): as received	
Destination MAC address #B	as received	
Destination Sequence Number #B	as received	

Case E: Proactive Root RREQ (original transmission)

All of the following applies:

- The Root MP is configured to send proactive root RREQs
- [The Root Announcement interval has expired]

Content:

Field	Value
ID	TBD
Length	37
Flags	Bit 0: As needed (portal role) Bit 1: 0 (broadcast) Bit 2: As needed (proactive RREP) Bit 3 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this information element, e.g. HWMP_NET_DIAMETER.
RREQ ID	Previous RREQ ID + 1
Originator Address	root MP MAC address
Originator's Destination Sequence Number	Previous DSN of root MP + 1
Lifetime	HWMP_PATH2ROOT_TIMEOUT (shall be greater than ROOT_ANNOUNCEMENT_INTERVAL)
Metric	0
Destination Count	1
Per Destination	DO = 1, RF = 1
	Address = All ones (broadcast address)
	Destination Sequence Number (none)

Field	Value	
ID	Tbd	
Length		
Flags	Bit 0: As needed (portal role) Bit 1: 0 (broadcast) Bit 2: As needed (proactive RREP) Bit 3 – 7: Reserved	
Hop Count	0	
Time to Live	Maximum number of hops allowed for this IE, e.g. HWMP_NET_DIAMETER.	
RREQ ID	Previous RREQ ID + 1	
Originator Address	Root MAC address	
Originator's Destination Sequence Number	Previous DSN of root + 1	
Lifetime	HWMP_PATH2ROOT_TIMEOUT (shall be greater than ROOT_ANNOUNCEMENT_INTERVAL)	
Metric	0	

Destination Count	1
Per Destination	DO = 1, RF = 1
	Address = All ones (broadcast address)
	Destination Sequence Number (none)

11A.6.5.4 RREQ processing

Received RREQ IEs are subject to certain acceptance criteria. Processing and actions taken depend on the contents of the RREQ and the information available to the receiving MP.
See also 11A.6.3: General rules for message processing

11A.6.5.4.1 Acceptance criteria

The RREQ will be is discarded if any of the following is true:

— The Originator DSN < previous Originator DSN
— ((DSN = previous DSN DSN) AND (updated path metric >= previous path metric metric)

Otherwise, the RREQ IE will be information element is accepted.

See also 11A.6.3: General rules for message processing

11A.6.5.4.2 Effect of receipt

The following applies only to a RREQ that has not been discarded:

1. The receiving MP shall record the RREQ ID, the Originator Address, and entries for each destination MAC Address and DSN
2. The receiving MP shall update the active forwarding information it maintains for the originator and previous hop MPs of the RREQ (see 11A.6.3.8)
3. If the MP is addressed by the RREQ it will initiate initiates the transmission of an a RREP to the originator (clause 11A.6.6.3 Case A)
4. If the MP has active forwarding information to any of the requested destinations and the DO flag for such a destination is OFF (DO=0), it will initiate initiates the transmission of an a RREP to each of these destinations (see clause 11A.6.6.3 Case A)
5. If there are destinations in the RREQ that have been not processed in steps 3 or 4 or that have been processed in step 4 but the corresponding Reply and Forward Flag is ON (RF = 1), the receiving MP shall retransmit forward the RREQ as defined in 11A.6.5.3, Case D

11A.6.6 Route Reply (RREP)

This subclause describes the function, generation and processing of the Route Reply IEinformation element.

11A.6.6.1 Function

The purpose of the Route Reply Information Element is
The purpose of the Route Reply (RREP) Information Element, described in 7.3.2.72, is

— to establish a forward route to a destination and
— to confirm that a destination is reachable.

11A.6.6.2 RREP information element

Octets: 1	1	1	1	1	6	4
ID	Length	Mode Flags	Hopcount	Time to Live	Destination Address	Destination Seq.Num.

4	4	6	4	1	6	4		6	4
Lifetime	Metric	Source Address #1	Source Seq. Num.	Dependent MP Count N	Dependent MP MAC Address #1	Dependent MP DSN #1	...	Dependent MP MAC Address #N	Dependent MP DSN #N

Figure s112—Route Reply Element

Table s34—RREP Element Fields

Field	Value/description
ID	TBD
Length	Length of the IE
Flags	Bit 0 – 7: Reserved
Hop Count	The number of hops from the route destination to the local MP
Time to Live	Maximum number of hops allowed for this IE
Destination Address	MAC address [of the destination for which a route is supplied]
Destination Sequence Number	DSN of the originator of the RREP
Lifetime	If applicable: reflects the Lifetime of the RREQ this RREP responds to
Metric	The cumulative metric from the route destination to the local MP.
Dependent MP Count N	Number of dependent MPs (N)
Dependent MP MAC Address #	MAC address of dependent MP
Dependent MP DSN #	Destination Sequence Number associated with MAC address of dependent MP

11A.6.6.3 Conditions for generating and sending a RREP

An MP will send sends out a RREP element in a RREP frame in the following cases:

Case A: Original transmission

A RREP is transmitted if the MP has received a RREQ fulfilling all of the following conditions:

a. One of the following applies:
   - The Destination Address of the RREQ is the same as MAC address of the receiving MP
   - The destination address of the RREQ = all 1's (broadcast) and the RREP flag is set to 1 ("Proactive RREP")

b. One of the following applies:
   - The Originator DSN of the RREQ (rreq.orig_dsn) is greater than the DSN of the last RREQ received from the same originator address (which includes the case that there is no route to the originating MP)
   - The Metric is better than the path selection metric currently associated with the Originator Address and the Originator DSN of the RREQ (rreq.orig_dsn) is equal to the DSN of the last RREQ received from the same originator address

The content of the generated RREP shall be:

Field	Value
ID	RREP information element ID
Length	As required
Flags	Bit 0 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this IE
Destination Address	Own MAC address
Destination Sequence Number	DSN of the originator of the RREP
Lifetime	As per the RREQ that triggered the transmission of this RREP
Metric	0
Dependents	As applicable: List of addresses of dependent MPs including MAC address and DSN

Field	Value	
ID	RREP IE ID	
Length	As required	
Flags	Bit 0 – 7: Reserved	
Hop Count	0	
Time to Live	Maximum number of hops allowed for this IE	
Destination Address	Own MAC address	
Destination Sequence Number	DSN of the originator of the RREP	
Lifetime	As per the RREQ that triggered the transmission of this RREP	
Metric	0	
Dependents	As applicable: List of addresses of dependent MPs including MAC address and DSN	

Note: the destination address of the action frame carrying the RREP IE information element is the Originator Address of the RREQ that triggered the RREP.

Case B: Re-transmission (forwarding)Forwarding

A RREP is transmitted if all of the following applies:

1. the MP has received and accepted the RREP – See 11A.6.6.4.1
2. the MP is not the destination of the RREP

Content:

Content:

Field	Value
ID	As received
Length	As received
Flags	As received
Hop Count	As received + 1
Time to Live	As received – 1
Destination Address	As received
Destination Sequence Number	As received
Lifetime	As received
Metric	As received + own metric towards the transmitting MP
Dependents	As received

Note: the destination address of the action frame carrying the RREP IE information element is the next hop to the Originator Address of the RREQ or RANN that triggered the RREP.

Case C: Intermediate reply

A RREP is transmitted if the MP has received a RREQ fulfilling all of the following conditions:

1. The RREQ Destination Only flag is set to 0
2. The receiving MP has active forwarding information with:
    a. A destination which is the same as the Destination Address of the RREQ
    b. A DSN which is less than or equal to the DSN of the RREQ (rreq.dest_dsn)
    c. A route path metric which is greater than or equal to the Metric or the RREQ
    d. A non-zero lifetime

The content of the generated RREP shall be:

Field	Value
ID	RREP information element ID
Length	As required
Flags	Bit 0 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this IE
Destination Address	Destination MAC address from the RREQ
Destination Sequence Number	DSN of the stored forwarding information of the Destination MAC address of the RREQ
Metric	0
Dependents	As applicable: List of addresses of dependent MPs including MAC address and DSN

Field	Value
ID	RREP IE ID
Length	As required
Flags	Bit 0 – 7: Reserved
Hop Count	0
Time to Live	Maximum number of hops allowed for this IE
Destination Address	Destination MAC address from the RREQ
Destination Sequence Number	DSN of the stored forwarding information of the Destination MAC address of the RREQ
Metric	0
Dependents	As applicable: List of addresses of dependent MPs including MAC address and DSN

11A.6.6.4 RREP processing

Received RREP IEs are subject to certain acceptance criteria. Processing and actions taken depend on the contents of the RREP and the information available to the receiving MP.

11A.6.6.4.1 Acceptance criteria

The RREP will be is discarded if any of the following is true:

— The DSN < previous DSN from this originator

— The Time to Live is 1 or less

11A.6.6.4.2 Effect of receipt

The following applies only to a RREP that has not been discarded

1) The receiving MP shall record the Originator Address, together with the DSN, hopcount and metric according to the rules defined in 11A.6.3.8.
2) The receiving MP may record the list of dependent MPs if present in the RREP.
3) If the receiving MP is not the final destination of the RREP, the RREP is re-transmitted forwarded as per Case B above.

11A.6.7 Route Error Information Element information element (RERR)

This subclause describes the function, generation and processing of the route error IEinformation element.

11A.6.7.1 Function

The RERR IE information element is used for announcing a broken link to all traffic sources that have an active path over this broken link. The active forwading forwarding information associated with the unreachable destinations should no longer be used for forwarding.

A RERR IE information element is propagated by MPs receiving a RERR if certain conditions are fulfilled.

An MP generating or receiving a RERR may decide to establish routes to unreachable destinations using any of the available HWMP mechanisms.

11A.6.7.2 Route Error Information Element

Octets: 1	1	1	1 (or 4)	6	4
ID	Length	Mode Flags	Num of Destinations	Destination Address	Destination MP Seq. Num

Figure s113—Route Error Element

Table s35—Route Error Element Fields

Field	Value/description
ID	TBD
Length	Length of the IE
Mode Flags	Bit 0 – 7: Reserved
Number of Destinations	Number of announced destinations in RERR (destination address and destination MP sequence number).
Destination Address	Detected unreachable destination MAC address
Destination Sequence Number	The sequence number of detected unreachable destination MP

11A.6.7.3 Conditions for generating and sending a RERR

A mesh point will send An MP sends out a RERR in the following cases:

Case A: Original transmission
All of the following applies:

- The MP detects a link break to the next hop of an active path in its stored forwarding information while transmitting packets frames to it. Note: the detection may be triggered by the fact that an MP is unable to forward a data frame to a next hop MP.
- The MP did not send more than HWMP_RERR_RATELIMIT – 1 RERR messages during the last second.

Actions before sending the RERR:

- The destination sequence numbers in all valid stored forwarding information of unreachable destinations announced in this RERR is incremented.
- The stored forwarding information for each unreachable destination announced in this RERR is invalidated.

Content:

Field	Value
ID	RERR information element ID
Length	2 + N * 10
Mode Flags	Bit 0 – 7: Reserved
Number of Destinations	Number of announced unreachable destinations in RERR. A destination is unreachable if its next hop in the stored forwarding information is an unreachable neighbor.
Destination Address	MAC address of detected unreachable destination #1
Destination Sequence Number	Last used DSN for Destination Address #1 + 1

Field	Value
ID	RERR IE ID
Length	
Mode Flags	Bit 0 – 7: Reserved
Number of Destinations	Number of announced unreachable destinations in RERR. A destination is unreachable if its next hop in the stored forwarding information is an unreachable neighbor.
Destination Address	MAC address of detected unreachable destination #1
Destination Sequence Number	Last used DSN for Destination Address #1 + 1

Case B: Re-transmissionForwarding

All of the following applies:

- The MP received a RERR from a neighbor for one or more of its active paths in its stored forwarding information.
- The MP has not sent or forwarded more than HWMP_RERR_RATELIMIT – 1 RERR messages during the last second.

Content:

Field	Value
ID	RERR information element ID
Length	2 + N * 10
Mode Flags	Bit 0 – 7: Reserved
Number of Destinations	Number of announced unreachable destinations in RERR (≤ received value) A destination is unreachable if its next hop in the corresponding stored forwarding information is the transmitter of the received RERR.
Destination Address	MAC address of detected unreachable destination #1 (as received, but maybe at different position in destination list)
Destination Sequence Number	as received (but maybe at different position in destination list)

Field	Value
ID	RERR IE ID
Length	
Mode Flags	Bit 0 – 7: Reserved
Number of Destinations	Number of announced unreachable destinations in RERR (≤ received value) A destination is unreachable if its next hop in the corresponding stored forwarding information is the transmitter of the received RERR.
Destination Address	MAC address of detected unreachable destination #1 (as received, but maybe at different position in destination list)
Destination Sequence Number	as received (but maybe at different position in destination list)

11A.6.7.4 RERR Reception

Received RERR IEs are subject to certain acceptance criteria. Processing and actions taken depend on the contents of the RERR and the information available to the receiving MP.

See also 11A.6.3: General rules for message processing

11A.6.7.4.1 Acceptance criteria

The RERR is not discarded if the following applies:

— The MP that receives the RERR has forwarding information stored where

- the destination is contained in the list of unreachable destinations of the RERR and
- the next hop is the transmitter of the received RERR

11A.6.7.4.2 Effect of receipt

The following applies only to RERR that has not been discarded.

a) The destination sequence numbers in all valid stored forwarding information of unreachable destinations announced in this RERR is set to the values taken from the corresponding fields of the RERR (rerr.dsnX) if the latter is greater than the destination sequence number currently stored in the forwarding information (→ destination sequence number).

b) The stored forwarding information of unreachable destinations announced in this RERR is invalidated.

c) The receiving MP shall transmit a RERR as defined in 11A.6.4.3, Case B

3A.5.1 HWMP parameters

To be added.

See T.3 for recommended parameter values.

11A.7 Radio Aware OLSR Path Selection Protocol path selection protocol (Optional)

11A.7.1 Introduction

Radio Aware Optimized Link State Routing (RA-OLSR) protocol is a proactive, link-state wireless mesh path selection protocol based on Optimized Link State Routing (OLSR) protocol [IETF RFC 3626] with extensions from Fisheye State Routing (FSR) protocol and uses radio-aware metrics in forwarding path and multipoint relay (MPR) set calculation. RA-OLSR enables the discovery and maintenance of optimal routes based on a predefined metric, given that each node MP has a mechanism to determine the metric cost of a link to each of its neighbors. In order to propagate the metric link cost information between nodesMPs, a metric field is used in RA-OLSR control messagesinformation elements. In disseminating topology information over the network, RA-OLSR adopts the following approaches in order to reduce the related control overhead:

— It uses only a subset of nodes MPs in the network, called MPRs, in flooding process;

— It optionally controls (and thereby reduces) the message exchange frequencies based on the Fisheye scopes.

The current RA-OLSR protocol specifications protocols also include an association discovery and maintenance protocol to support non-mesh STAs both internal (associated with MAPs) and external (with MPPs): When a non-mesh STA is a source or a destination of an IEEE 802 data link, RA-OLSR protocol

sets up a routing path to the MAP or the MPP associated with that STA by complementing routing information among MPs with STA association information.

11A.7.2 Overview

RA-OLSR is an optimization over a classical link-state routing protocol, tailored for WLAN mesh networks; as such, it inherits all the benefits of link-state routing protocols, including the immediate availability of routes when needed, which greatly reduces the initial route setup delay.

The optimization in RA-OLSR mainly focuses on the minimization of flooding overhead: First, in RA-OLSR only a selected subset of 1-hop neighbor MPs (i.e., MPRs) is used by an MP in retransmitting control messagesforwarding information elements. The MPRs are selected such that a broadcast message, retransmitted forwarded by these MPRs, can reach all 2-hop neighbor MPs of the selecting MP (i.e., MPR selector). The information required to perform MPR selection is acquired through the periodic exchange of HELLO messages. In addition, RA-OLSR can also optionally control the message exchange frequencies based on the fisheye scopes to further reduce the amount of control messages information elements exchanges. These techniques significantly reduce the number of retransmissions forwarding transmissions required to flood a message to all MPs in the network. Second, RA-OLSR requires only partial link state to be flooded in order to provide shortest best path routes. The minimal set of link state information required is that all the nodes MPs selected as MPRs must shall declare the links to their MPR selectors. Additional topological information, if present, may be utilized, e.g., for redundancy purposes.

In wireless mesh networks, unlike traditional mobile ad hoc networks where all mobile nodes MPs are directly participating in routing procedures, legacy STAs that do not support WLAN mesh services are indirectly participating in routing procedures through their associated MAPs or MPPs. To support these STAs, RA-OLSR provides information repositories for STA association at MAPs/MPPs — Local Association Base (LAB) and Global Association Base (GAB) — with efficient advertisement and management schemes, which complements a routing table at each MP during path selection process. This is detailed in clause 11A.7.13.

RA-OLSR may optimize the reactivity to topological changes by reducing the maximum time interval for periodic control message information element transmissions. Furthermore, as RA-OLSR continuously maintains routes to all destinations in the network, the protocol is beneficial for traffic patterns where a large subset of MPs are communicating with another large subset of MPs, and where the [source, destination] pairs are changing over time. The protocol is particularly suited for large and dense networks, as the optimization done using MPRs works well in this context. The larger and denser is a network, the more optimization can RA-OLSR achieves than the classic link-state algorithm.

RA-OLSR is designed to work in a completely distributed manner and does not depend on any central entity. The protocol does not require reliable transmission of control messagesinformation elements: Each MP sends control messages information elements periodically, and can therefore sustain a reasonable loss of some of such messages. Such losses occur frequently in radio networks due to collisions or other transmission problems.

Also, RA-OLSR does not require sequenced delivery of messages. Each control message information element contains a sequence number which is incremented for each message. Thus the The recipient of a control message an information element can, if required, easily identify which information is more recent — even if messages have been re-ordered while in transmission.

Given a network with only single-interface MPs, an MP may deduct the neighbor set directly from the information exchanged as part of link sensing: the “Main Address” of a single interface MP is, by definition, the address of the only interface on that MP. In a network with multiple-interface MPs, additional information is required in order to map interface addresses to main addresses (and, thereby, to MPs). This

additional information is acquired through multiple interface declaration (MID) messages, described in clause 11A.7.5.

11A.7.2.1 Terminology

Terminology	Description
Main Address	Primary MAC address of the MP, if it has more than one radio interfacePHY
Originator Address	The main address of a NODEan MP, which sent a given message
Sender Interface Address	The address of the interface, over which the message was last transmitted
Receiving Interface	The interface, over which the message was received
Associated Station	A station which is associated with a an MAP
Local Association Base (LAB)	The table of the associated stations to a given MAP
Global Association Base (GAB)	The information base which maintains the list of the associated stations to all the MAPs in the WLAN Mesh (in other terms, all the Local Association Base of all the MAPs of the WLAN Mesh)
Association Tuple	The information about associated stations is maintained in entries called "association tuples", either "Local Association Tuple", in the LAB, or "Global Association Tuple", in the GAB
Checksum	The value obtained by applying a hash function to the information in the LAB / GAB (or subsets of these Association Bases)

11A.7.3 Message Processing processing and Forwardingforwarding

One or more RA-OLSR messages (i.e., IEs; see clause 7.3.2.74 for details) are carried in an RA-OLSR frame (see clause 7.4.9.13 for its format). Here we describe a general rule for processing and forwarding messages included in an RA-OLSR frame.

11A.7.3.1 Message Processing processing and Floodingflooding

Upon receiving an RA-OLSR management frame, an MP examines each of the included messages (i.e., IEs). Based on the "ID" value, the MP can determine the further processing of the message.

An MP may receive the same message several times in a wireless mesh network. Therefore, each MP maintains a "Duplicate Set" where the MP records information about the most recently received messages, by which any duplicate processing of those messages can be avoided. For such a message, an MP records a "Duplicate Tuple" as follows:

(D_addr, D_seq_num, D_retransmittedD_forwarded, D_iface_list, D_time)

Field	Description
D_addr	The originator address of the message
D_seq_num	The message sequence number of the message
D_retransmitted D_forwarded	A boolean indicating if the message has already been retransmittedforwarded
D_iface_list	A list of the addresses of the interfaces on which the message has been received
D_time	Expiration time of the tuple when it must shall be removed

In an MP, the set of Duplicate Tuples are denoted the “Duplicate set”.

Thus, upon Upon receiving an RA-OLSR management frame, an MP must shall perform the following tasks for each encapsulated message:

a) If the RA-OLSR frame contains no messages, the frame must shall be silently discarded.

b) If the TTL of the message is less than or equal to zero, or if the message was sent by the receiving MP (i.e., the Originator Address of the message is the main address of the receiving MP), then the message must shall be silently dropped.

c) Processing condition:

1) If there exists a tuple in the duplicate set, where:

D_addr == Originator Address, AND

D_seq_num == Message Sequence Number.

Then the message has already been processed and must shall not be processed again.

2) Otherwise, if the MP implements the message type (i.e., “ID”) of the message, the message must shall be processed according to the specifications for the message type.

d) Forwarding condition:

3) If there exists a tuple in the duplicate set, where:

D_addr == Originator Address, AND

D_seq_num == Message Sequence Number, AND

the receiving interface (address) is in D_iface_list.

Then the message has already been considered for forwarding and should not be retransmitted forwarded again.

4) Otherwise:

i) If the MP implements the message type of the message, the message must shall be considered for forwarding according to the specifications for the message type.

i) Otherwise, if the MP does not implement the message type of the message, the message should be processed according to the default forwarding algorithm described in clause 11A.7.3.3.

11A.7.3.2 Default Forwarding Algorithm

11A.7.3.3 RA-OLSR default forwarding algorithm

The default forwarding algorithm is the following:

a) If the sender interface address of the message is not detected to be in the symmetric 1-hop neighborhood of the MP, the forwarding algorithm must shall silently stop here (and the message must shall not be forwarded).

b) If there exists a tuple in the duplicate set where:

D_addr == Originator Address, AND

D_seq_num == Message Sequence Number.

Then the message will be is further considered for forwarding if and only if:

D_retransmitted D_forwarded is false, AND

the (address of the) interface which received the message is not included among the addresses in D_iface_list.

c) Otherwise, if such an entry doesn't exist, the message is further considered for forwarding.

If the message is not considered for forwarding after steps 1 through 3, the processing of this clause stops here (i.e., steps 4 to 8 are ignored). Otherwise, the following algorithm is used:

d) If the sender interface address is an interface address of an MPR selector of this MP AND the TTL of the message is greater than '1', the message must shall be retransmitted forwarded (as described later in steps 6 to 8).

e) If there exists an entry in the duplicate set with the same Originator Address and Message Sequence Number, the entry is updated as follows:

D_time = current time + DUP_HOLD_TIME

The receiving interface MAC address is added to D_iface_list.

D_retransmitted D_forwarded is set to true if and only if the message will be retransmitted is forwarded according to step 4.

Otherwise a new entry is added to the duplicate set with:

D_addr == Originator MAC Address.

D_seq_num == Message Sequence Number.

D_time = current time + DUP_HOLD_TIME.

D_iface_list contains the receiving interface address.

D_retransmitted D_forwarded is set to true if and only if the message will be retransmitted is forwarded according to step 4.

If, and only if, according to step 4, the message must shall be retransmitted forwarded then:

f) The TTL of the message is decremented by one.

g) The hop-count of the message is increased by one.

h) The message is broadcast on all interfaces.

11A.7.3.4 Considerations on Processing processing and Forwardingforwarding

It should be noted that the processing and the forwarding of messages are two different actions, conditioned by different rules. Processing is related to using the content of the messages, while forwarding is related to retransmitting forwarding the same message for other MPs in the network.

Notice that this specification standard includes a description for both the forwarding and the processing of each known message type (i.e., "ID"). Messages with known message types must shall not be forwarded "blindly" by the algorithm described in clause 11A.7.3.3. Forwarding (and setting the correct message field in the forwarded, known message) is the responsibility of the algorithm specifying how the message is to be handled and, if necessary, retransmittedforwarded. This enables a message type to be specified such that the message can be modified while in transit (e.g., to reflect the route the message has taken). In this regard, it would be possible to bypass the MPR flooding mechanism if for some reasons, classical flooding is required for a message type; in such a case, the forwarding algorithm for that message type will specify specifies simple rebroadcasting of received messages, regardless of MPRs.

NOTE--By defining a set of message types, which must shall be recognized by all implementations of RA-OLSR, it will be is possible to extend the protocol through introduction of additional message types, while still being able to maintain compatibility with older implementations.

11A.7.3.5 Message Emission emission and Jitterjitter

As a basic implementation requirement, the synchronization of control messages information elements should be avoided. As a consequence, RA-OLSR control messages information elements should be emitted such that they avoid synchronization.

Emission of control traffic from neighbor MPs may — for various reasons (mainly timer interactions with packet frame processing) — become synchronized such that several neighbor MPs attempt to transmit control traffic simultaneously. This may lead to collisions and hence message loss, possibly the loss of several subsequent messages of the same type.

To avoid such synchronizations, the following simple strategy for emitting control messages information elements is recommended. A node An MP should add an amount of jitter to the interval at which messages are generated. The jitter must shall be a pseudo-random value for each message generated. Thus, for For an MP utilizing jitter:

Actual message interval = MESSAGE_INTERVAL – jitter

where jitter is a pseudo-random value in [0, MAXJITTER] and MESSAGE_INTERVAL is the value of the message interval specified for the message being emitted (e.g., HELLO_INTERVAL for HELLO messages, TC_INTERVAL for TC-messages, etc.).

Jitter SHOULD also may be introduced when forwarding messages. The following simple strategy may be adopted: When a message is to be forwarded by an MP, it should be kept in the MP during a short period of time:

Keep message period = jitter

where jitter is a pseudo-random value in [0, MAXJITTER]. Notice that when the MP sends a control messagean information element, the opportunity to piggyback other messages (before their keeping period is expired) maybe taken to reduce the number of frame transmissions.

11A.7.4 Information Repositoriesrepositories

Through the exchange of RA-OLSR control messagesinformation elements, each node MP accumulates information about the network. This information is stored according to the descriptions in this sectionsubclause.

11A.7.4.1 Link Setset

A node An MP records a set of "link tuples":

(L_local_iface_addr, L_neighb_iface_addr, L_time, L_link_metric)

describing links between local and remote interfaces. The tuples in this set are maintained by some other component of 802.11 this standard and populated using Neighbor Table Entry whose state is subordinate link up or Superordinate, link up as described in clause 11A.5.1.

Field	Description
L_local_iface_addr	The interface on the local MP
L_neighb_iface_addr	The interface on the remote MP, with which there exists a symmetric link
L_time	The expiration time of this tuple
L_link_metric	The value representing the metric cost of the link. An example is the Airtime cost given in 11A.5

Field	Description
L_local_iface_addr	The interface on the local MP
L_neighb_iface_addr	The interface on the remote MP, with which there exists a symmetric link
L_time	The expiration time of this tuple
L_link_metric	The value representing the metric cost of the link. An example is the Airtime cost given in clause 11A.5

11A.7.4.2 Neighbor Setset

An MP records a set of "neighbor tuples":

(N_neighb_main_addr, N_willingness)

describing neighbors.

Field	Description
N_neighb_main_addr	The interface on the local MP
N_willingness	An integer, between 0 and 7, specifying the MPs willingness to carry traffic on behalf of other MPs

Field	Description
N_neighb_main_addr	The interface on the local MP
N_willingness	An integer, between 0 and 7, specifying the nodes willingness to carry traffic on behalf of other MPs

11A.7.4.3 Interface Association Setassociation set

For each destination in the network, “Interface Association Tuples” are recorded:

(I_iface_addr, I_main_addr, I_time)

Field	Description
I_iface_addr	An interface address of a nodean MP
I_main_addr	The main address of a nodean MP
I_time	The expiration time of the tuple

11A.7.4.4 2-hop Neighbor Setneighbor set

A node An MP records a set of “2-hop tuples”:

(N_neighb_main_addr, N_2hop_addr, N_time)

describing symmetric links between its neighbors and the symmetric 2-hop neighborhood.

Field	Description
N_neighb_main_addr	The main address of a neighbor
N_2hop_addr	The main address of a neighbor, which has a symmetric link to N_neighbor_main_addr
N_time	The expiration time of the tuple

Field	Description
N_neighb_main_addr	The main address of a neighbor
N_2hop_addr	The main address of a neighbor, which has a symmetric link to N_neighbor_main_addr
N_time	The expiration time of the tuple

11A.7.4.5 MPR Setset

An MP maintains a set of neighbors which are selected as MPR. Their main addresses are listed in the MPR Set.

11A.7.4.6 MPR Selector Setselector set

An MP records a set of "MPR-selector tuples":

(MS_main_addr, MS_time)

describing the neighbors which have selected this MP as an MPR.

Field	Description
MS_main_addr	The main address of an MPR-selector
MS_time	The expiration time of the tuple

Field	Description
MS_main_addr	The main address of an MPR-selector
MS_time	The expiration time of the tuple

11A.7.4.7 Topology Setset

Each node MP in the network maintains topology information about the network. This information is acquired from TC-messages and is used for routing table calculations. Thus, for For each destination in the network, at least one "Topology Tuple":

(T_dest_addr, T_last_addr, T_seq, T_time, T_link_metric)

is recorded.

Field	Description
T_dest_addr	The main address of an MP, which may be reached in one hop from the MP with main address T_last_addr
T_last_addr	An MP which can reach T_dest_addr in one hop
T_seq	A sequence number
T_time	The expiration time of the tuple
T_link_metric	The value representing the metric cost of the link. If more than one link exists, the minimum cost (i.e., the cost of the link with the best quality) should be used. An example of link metric is the Airtime cost given in 11A.5.

Field	Description
T_dest_addr	The main address of a node, which may be reached in one hop from the node with main address T_last_addr
T_last_addr	An MP which can reach T_dest_addr in one hop
T_seq	A sequence number
T_time	The expiration time of the tuple
T_link_metric	The value representing the metric cost of the link. If more than one link exists, the minimum cost (i.e., the cost of the link with the best quality) should be used. An example of link metric is the Airtime cost given in clause 11A.5.

11A.7.4.7.1 Local Association Base (LAB)

Each MAP, as a result of the 802.11 association protocol, MAP keeps a set of associated stations, denoted Local Association Base" (LAB) in this documentstandard, which holds "Local Association Tuple", one for each associated station. Additionally, to provide support for a large number of stations, each MAP divides its LAB into blocks of local association tuples: each block is a subset of the LAB, and the LAB is the union

of those subsets. The blocks of the LAB are numbered; hence each block is identified by an integer, the "block index".

Each element of the LAB, a "Local Association Tuple" of an associated station, has the following fields:

(block_index, station_address, station_sequence_number)

Field	Description
block_index	Index of the block the "Local Association Tuple" belongs to
station_address	The Mac address of the associated station. A station address does not include the "Group MAC address bit" in the 48 bit MAC address
station_sequence_number	The sequence number (i.e., the whole 2-octet "Sequence Control" field) of the management frame that a STA sent to the MAP when the STA is associated with or disassociated with the MAP

Field	Description
block_index	Index of the block the "Local Association Tuple" belongs to
station_address	The Mac address of the associated station. Note that a station address does not include the "Group MAC address bit" in the 48 bit MAC address
station_sequence_number	The sequence number (i.e., the whole 2-octet "Sequence Control" field) of the management frame that a STA sent to the MAP when the STA is associated with or disassociated with the MAP

11A.7.4.7.2 Global Association Base (GAB)

Each MAP maintains information concerning which station is associated to which MAP in the entire WLAN Mesh, in a "Global Association Base" (GAB) which is union of all LAB of each MAP in the WLAN Mesh. It is a set of "global association tuples" containing

(block_index, AP_address, station_address, station_sequence_number, expiration_time)

Field	Description
block_index	Index of block associated to the station address by the MAP in its LAB
AP_address	The MAC address of the MAP
station_address	The MAC address of the associated STA. Note that a A STA address does not include the "Group MAC address bit" in the 48 bit MAC address
station_sequence_number	The sequence number (i.e., the whole 2-octet "Sequence Control" field) of the management frame that a STA sent to the MAP when the station associated with or disassociated with the MAP
expiration_time	The time at which the entry is no longer valid

11A.7.5 Multiple Interfacesinterfaces

The relationship between RA-OLSR interface addresses and main addresses is defined through the exchange of Multiple Interface Declaration (MID) messages. This section subclause describes how MID messages are exchanged and processed.

Each node MP with multiple interfaces MUST shall announce, periodically, information describing its interface configuration to other MPs in the network. This is accomplished through flooding a MID message to all MPs in the network.

11A.7.5.1 MID Message Generationmessage generation

An MID message is sent by an MP in the network to declare its multiple interfaces (if any) i.e., the MID message contains the list of interface addresses which are associated to its main address. 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 interface set from an MP MUST shall be complete within a certain refreshing period (MID_INTERVAL). The information diffused in the network by these MID messages will help each MP to calculate its routing table. An MP which has only a single interface address participating in the WLAN mesh network running RA-OLSR, MUST NOT shall not generate any MID message.

A node An MP with several interfaces, where more than one is participating in the WLAN Mesh and running RA-OLSR MUST shall generate MID messages as specified.

The "Vtime" and "TTL" fields of the MID message shall be set to constant MID_HOLD_TIME and 255, respectively.

11A.7.5.2 MID Message Forwardingmessage forwarding

MID messages are broadcast and retransmitted forwarded by the MPRs in order to diffuse the messages in the entire network. The "default forwarding algorithm" MUST shall be used for forwarding of MID messages.

11A.7.5.3 MID Message Processingmessage processing

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 shall be computed from the Vtime field of the message header (as described in clause 11A.7.3). Then the Multiple Interface Association Information Base should be updated as follows:

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

b) For each interface address listed in the MID message:

   1) If there exists a tuple in the interface association set where:

      I_iface_addr == interface 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 == interface address,

      I_main_addr == originator address,

      I_time = current time + validity time.

11A.7.5.4 Mapping Interface Addresses interface addresses and MP Addressesaddresses

In networks with multiple interface MPs operating within a common RA-OLSR area, it is required to be able to map any interface address to the corresponding main address.

Given an interface address:

a) if there exists some tuple in the interface association set where:

   I_iface_addr == interface address.

   Then the result of the main address search is the originator address I_main_addr of the tuple.

b) Otherwise, the result of the main address search is the interface address itself.

11A.7.6 HELLO Message Generationmessage generation, Forwarding forwarding & Processingprocessing

HELLO messages are exchanged between neighbor MPs, and serve the purpose of populating the 2-hop neighbor set as well as carry MPR signaling.

11A.7.6.1 HELLO Message Generationmessage generation

This involves transmitting the Neighbor Set and the MPR Set. These are transmitted periodically, and are scoped for the nodeMP's neighborhood only (TTL=1). The lists list of addresses declared in a HELLO message is a list of neighbor interface addresses computed as follows: For each non-expired tuple in the Link Set, where L_local_iface_addr is the interface where the HELLO is to be transmitted, L_neighb_iface_addr is advertised. If the MP, to which L_neighb_iface_addr is associated is selected as MPR (i.e., a corresponding tuple exists in the MPR-set), a neighbor-tlv with type = MPR is associated with this address. The originator address for a HELLO message is set to the main address of the MP, generating the HELLO message.

11A.7.6.2 HELLO Message Forwardingmessage forwarding

Each HELLO message generated is broadcast by the node MP on one interface to its neighbors (i.e., the interface for which the HELLO was generated). HELLO messages SHALL not be forwarded.

11A.7.6.3 HELLO Message Processingmessage processing

An MP processes incoming HELLO messages for the purpose of conducting link sensing (detailed in clause 11A.7.11), neighbor detection and MPR selector set population (detailed in clause 11A.7.10). Only messages from subordinate and superordinate links (as described in clause 11A.1.5) are accepted

11A.7.7 Populating the Neighbor Setneighbor set

An MP maintains a set of neighbor tuples, based on the link tuples. This information is updated according to changes in the Link Set.

The Link Set keeps the information about the links, while the Neighbor Set keeps the information about the neighbors. There is a clear association between those two sets, since an MP is a neighbor of another MP if and only if there is at least one link between the two MPs. Only messages from subordinate and superordinate link (as described in 11A.1.5) are accepted

In any case, the formal correspondence between links and neighbors is defined as follows:

The “associated neighbor tuple” of a link tuple, is, if it exists, the neighbor tuple where:

N_neighb_main_addr == main address of L_neighb_iface_addr

The “associated link tuples” of a neighbor tuple, are all the link tuples, where:

N_neighb_main_addr == main address of L_neighb_iface_addr

The Neighbor Set MUST shall be populated by maintaining the proper correspondence between link tuples and associated neighbor tuples, as follows:

Creation

Each time a link appears, that is, each time a link tuple is created, the associated neighbor tuple MUST shall be created, if it doesn't already exist, with the following values:

N_neighb_main_addr == main address of L_neighb_iface_addr (from the link tuple)

Removal

Each time a link is deleted, that is, each time a link tuple is removed, the associated neighbor tuple MUST shall be removed if it has no longer any associated link tuples.

These rules ensure that there is exactly one associated neighbor tuple for a link tuple, and that every neighbor tuple has at least one associated link tuple.

11A.7.7.1 HELLO Message Processingmessage processing

The “Originator Address” of a HELLO message is the main address of the MP, which has emitted the message. Likewise, the “willingness” MUST shall be computed from the Willingness field of the HELLO message.

Upon receiving a HELLO message, an MP SHOULD should update its Neighbor Set as follows:

If the Originator Address is the N_neighbor_main_addr from a neighbor tuple included in the Neighbor Set:

Then, the neighbor tuple SHOULD should be updated as follows:

N_willingness == willingness from the HELLO message

11A.7.8 Populating the 2-hop Neighbor Setneighbor set

The 2-hop neighbor set describes the set of MPs which have a symmetric link to a symmetric neighbor. This information set is maintained through periodic exchange of HELLO messages as described in this section-subclause.

11A.7.8.1 HELLO Message Processingmessage processing

The "Originator Address" of a HELLO message is the main address of the MP, which has emitted the message.

Upon receiving a HELLO message, an MP SHOULD should update its 2-hop Neighbor Set. Notice that a HELLO message MUST shall neither be forwarded nor be recorded in the duplicate set.

Upon receiving a HELLO message, the "validity time" MUST shall be computed from the Vtime field of the message header (see clause 11A.7.3).

The 2-hop Neighbor Set SHOULD should be updated as follows:

a) For each address (henceforth: 2-hop neighbor address), listed in the HELLO:

1) If the main address of the 2-hop neighbor address == main address of the receiving MP:
silently discard the 2-hop neighbor address.
(in other words: an MP is not its own 2-hop neighbor).

2) Otherwise, a 2-hop tuple is created with:
N_neighb_main_addr == Originator address;
N_2hop_addr == main address of the 2-hop neighbor;
N_time = current time + validity time.

This tuple may replace an older similar tuple with same N_neighb_main_addr and N_2hop_addr values.

11A.7.9 Populating the MPR set

MPRs are used to flood control messages information elements from an MP into the network while reducing the number of retransmissions forwarding transmissions that will occur in a region. Thus, the The concept of MPR is an optimization of a classical flooding mechanism.

Each MP in the network selects, independently, its own set of MPRs among its 1-hop neighborhood. The neighbor interfaces, which are selected as MPR, are advertised with an associated neighbor-tlv with type=MPR in HELLO messages.

The MPR set MUST shall be calculated by an MP in such a way that it, through the neighbors in the MPR-set, can reach all strict 2-hop neighbors with minimum radio-aware metric. (Notice that an MP, a, which is a direct neighbor of another MP, b, is not also a strict 2-hop neighbor of MP b). This means that the union of the symmetric 1-hop neighborhoods of the MPR MPs contains the strict 2-hop neighborhood. MPR set recalculation should occur when changes are detected in the symmetric neighborhood or in the symmetric strict 2-hop neighborhood. MPR set recalculation should also occur when the change of a radio-aware metric is larger greater than a defined threshold.

MPRs are computed per interface, the union of the MPR sets of each interface makes up the MPR set for the MP.

While it is not essential that the MPR set is minimal, it is essential that all strict 2-hop neighbors can be reached through the selected MPR MPs. An MP SHOULD should select an MPR set such that any strict 2-hop neighbor is covered by at least one MPR MP. Keeping the MPR set small ensures that the overhead of the protocol is kept at a minimum.

The MPR set can coincide with the entire neighbor set. This could be the case at network initialization (and will correspond to classic link-state routing).

The heuristic for the selection of MPRs in the original OLSR does not take into account the radio-aware metric. It computes an MPR set with minimal cardinality and therefore links with lower radio-aware metric can be omitted. Consequently, the path calculated between two nodes MPs using the known partial topology is not optimal (in terms of radio-aware metric) in the whole network.

The decision of how each node MP selects its MPRs is essential to determinate the optimal radio-aware metric path in the network. In the MPR selection, links with low radio-aware metric SHOULD should not be omitted.

11A.7.9.1 MPR Computation

The following specifies a recommended heuristic for selection of MPRs. It constructs an MPR-set that enables an MP to reach any MP in the strict 2-hop neighborhood through relaying by one MPR with willingness different from WILL_NEVER. The heuristic MUST shall be applied per interface: The MPR set for an MP is the union of the MPR sets found for each interface. The following terminology will be is used in describing the heuristics:

Terminology	Description
Neighbor of an interface	An MP is a "neighbor of an interface" if the interface (on the local MP) has a link to any one interface of the neighbor MP.
2-hop neighbors reachable from an interface	The list of 2-hop neighbors of the MP that can be reached from neighbors of this interface.
MPR set of an interface	A (sub)set of the neighbors of an interface with a willingness different from WILL_NEVER, selected such that through these selected MPs, all strict 2-hop neighbors reachable from that interface are reachable.
N	N is the subset of neighbors of the MP, which are neighbor of the interface I.
N2	The set of 2-hop neighbors reachable from the interface I, excluding: (i) the MPs only reachable by members of N with willingness WILL_NEVER (ii) the MP performing the computation (iii) all the neighbors: the MPs for which there exists a link to this MP on some interface
D(y)	The degree of a 1-hop neighbor MP y (where y is a member of N), is defined as the number of symmetric neighbors of MP y, EXCLUDING all the members of N and EXCLUDING the MP performing the computation.

The recommended heuristic is as follows:

a) Start with an MPR set made of all members of N with N_willingness equal to WILL_ALWAYS.

b) Calculate D(y), where y is a member of N, for all MPs in N.

c) Add to the MPR set those MPs in N, which are the *only* only MPs to provide reachability to an MP in N2. For example, if MP 'b' in N2 can be reached only through a symmetric link to MP 'a' in N, then add MP 'a' to the MPR set. Remove the MPs from N2 which are now covered by an MP in the MPR set.

d) While there exist MPs in N2 which are not covered by at least one MP in the MPR set:

1) For each MP in N, calculate the reachability, i.e., the number of MPs in N2 which are not yet covered by at least one MP in the MPR set, and which are reachable through this 1-hop neighbor;

2) Select as an MPR the MP with the highest N_willingness among the MPs in N with non-zero reachability. In case of multiple choices, we use tie-breakers in the following order:

Maximum reachability (i.e., reachability to the maximum number of MPs in N2);

Maximum degree (D(.))

Minimum radio-aware metric (i.e., the best link quality according to the radio-aware metric);

Remove the MPs from N2 which are now covered by an MP in the MPR set.

e) An MP's MPR set is generated from the union of the MPR sets for each interface. As an optimization, process each MP 'y', in the MPR set in increasing order of N_willingness. If all MPs in N2 are still covered by at least one MP in the MPR set excluding MP 'y', and if N_willingness of MP 'y' is smaller than WILL_ALWAYS, then MP 'y' MAY may be removed from the MPR set.

Other algorithms, as well as improvements over this algorithm, are possible. For example, assume that in a multiple-interface scenario there exists more than one link between MPs 'a' and 'b'. If MP 'a' has selected MP 'b' as MPR for one of its interfaces, then MP 'b' can be selected as MPR without additional performance loss by any other interfaces on MP 'a'.

11A.7.10 Populating the MPR Selector Setselector set

The MPR selector set of an MP, n, is populated by the main addresses of the MPs which have selected n as MPR. MPR selection is signaled through HELLO messages.

11A.7.10.1 HELLO Message Processingmessage processing

Upon receiving a HELLO message, if an MP finds one of its own interface addresses in the list with a Neighbor Type equal to MPR_NEIGH, information from the HELLO message must shall be recorded in the MPR Selector Set.

The "validity time" MUST shall be computed from the Vtime field of the message header (see clause 11A.7.3). The MPR Selector Set SHOULD should then be updated as follows:

a) If there exists no MPR selector tuple with:

MS_main_addr == Originator address

then a new tuple is created with:

MS_main_addr == Originator address

b) The tuple (new or otherwise) with

MS_main_addr == Originator address

is then modified as follows:

MS_time = current time + validity time.

Deletion of MPR selector tuples occurs in case of expiration of the timer or in case of link breakage as described in the "Neighborhood and 2-hop Neighborhood Changes".

11A.7.10.2 Neighborhood and 2-hop Neighborhood Changesneighborhood changes

A change in the neighborhood is detected when:

— The L_SYM_time field of a link tuple expires. This is considered as a neighbor loss if the link described by the expired tuple was the last link with a neighbor MP (on the contrary, a link with an interface may break while a link with another interface of the neighbor MP remains without being observed as a neighborhood change).

— A new link tuple is inserted in the Link Set with a non expired L_SYM_time or a tuple with expired L_SYM_time is modified so that L_SYM_time becomes non-expired. This is considered as a neighbor appearance if there was previously no link tuple describing a link with the corresponding neighbor MP.

A change in the 2-hop neighborhood is detected when a 2-hop neighbor tuple expires or is deleted according to clause 11A.7.8.

The following processing occurs when changes in the neighborhood or the 2-hop neighborhood are detected:

— In case of neighbor loss, all 2-hop tuples with N_neighb_main_addr == Main address of the neighbor MUST shall be deleted.

— In case of neighbor loss, all MPR selector tuples with MS_main_addr == Main address of the neighbor MUST shall be deleted.

— The MPR set MUST shall be re-calculated when a neighbor appearance or loss is detected, or when a change in the 2-hop neighborhood is detected.

— An additional HELLO message MAY may be sent when the MPR set changes.

11A.7.11 Topology Discoverydiscovery

The link sensing and neighbor detection part of the protocol basically offers, to each MP, a list of neighbors with which it can communicate directly and, in combination with the Packet Frame Format and Forwarding part, an optimized flooding mechanism through MPRs. Based on this, topology information is disseminated through the network. The present section subclause describes which part of the information given by the link sensing and neighbor detection is disseminated to the entire network and how it is used to construct routes.

Routes are constructed through advertised links and links with neighbors. An MP must shall at least disseminate links between itself and the MPs in its MPR-selector set, in order to provide sufficient information to enable routing.

11A.7.11.1 Advertised Neighbor Setneighbor set

A TC message is sent by an MP in the network to declare a set of links, called advertised link set which MUST shall include at least the links to all MPs of its MPR Selector set, i.e., the neighbors which have selected the sender MP as an MPR.

The sequence number (ANSN) associated with the advertised neighbor set is also sent with the list. The ANSN number MUST shall be incremented when links are removed from the advertised neighbor set; the ANSN number SHOULD should be incremented when links are added to the advertised neighbor set.

11A.7.11.2 TC Message Generationmessage generation

In order to build the topology information base, each MP, which has been selected as MPR, broadcasts Topology Control (TC) messages. TC messages are flooded to all MPs in the network and take advantage of MPRs. MPRs enable better scalability in the distribution of topology information.

The list of addresses can be partial in each TC message (e.g., due to message size limitations, imposed by the network), but parsing of al TC messages describing the advertised link set of an MP MUST shall be complete within a certain refreshing period (TC_INTERVAL). The information diffused in the network by these TC messages will help each MP calculate its routing table.

When the advertised link set of an MP becomes empty, this MP SHOULD should still send (empty) TC-messages during the duration equal to the "validity time" (typically, this will be is equal to TOP_HOLD_TIME) of its previously emitted TC-messages, in order to invalidate the previous TC-messages. It SHOULD should then stop sending TC-messages until some MP is inserted in its advertised link set.

An MP MAY may transmit additional TC-messages to increase its reactiveness to link failures. When a change to the MPR selector set is detected and this change can be attributed to a link failure, a TC-message SHOULD should be transmitted after an interval shorter than TC_INTERVAL.

In order to realize the frequency control, when a node an MP initiates a TC message, RA-OLSR sets different maximum hop count value for the packet frame according to the different frequencies. With the different maximum hop count values, the packets frames can only reach different designed scopes. The "TTL" field of the RA-OLSR message for the TC messages is used to control the scope. In the default behavior, the TTL is alternatively set to 2, 4, and the maximum TTL value in every TC_INTERVAL.

11A.7.11.3 TC Message Forwardingmessage forwarding

TC messages are broadcast and retransmitted forwarded by the MPRs in order to diffuse the messages in the entire network. TC messages MUST shall be forwarded according to the "default forwarding algorithm" (described in clause 11A.7.3).

11A.7.11.4 TC Message Processingmessage processing

Upon receiving a TC message, the "validity time" MUST shall be computed from the Vtime field of the message header (see clause 11A.7.3). The topology set SHOULD should then be updated as follows:

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

b) If there exist some tuple in the topology set where:

T_last_addr == originator address, AND

T_seq > ANSN,

then further processing of this TC message MUST NOT shall not be performed and the message MUST shall be silently discarded (case: message received out of order).

c) All tuples in the topology set where:

T_last_addr == originator address, AND

T_seq < ANSN

MUST shall be removed from the topology set.

d) For each of the advertised neighbor main address addresses received in the TC message:

1) If there exist some tuple in the topology set where:

T_dest_addr == advertised neighbor main address, AND

T_last_addr == Originator address,

then the holding time of that tuple MUST shall be set to:

T_time = current time + validity time.

2) Otherwise, a new tuple MUST shall be recorded in the topology set where:

T_dest_addr == Advertised neighbor main address,

T_last_addr == Originator address,

T_seq == ANSN,

T_time = current time + validity time

11A.7.12 Routing Table Calculationtable calculation

11A.7.12.1 General

Each MP maintains a routing table which allows it to route data destined for the other MPs in the network. The routing table is based on the information contained in the local link information base and the topology set. Therefore, if any of these sets are changed, the routing table is recalculated to update the route information about each destination in the network. The route entries are recorded in the routing table in the following format:

1. R_dest_addr R_next_addr R_dist R_metric R_iface_addr

2. R_dest_addr R_next_addr R_dist R_metric R_iface_addr

3. ,, ,, ,, ,, ,,

Each entry in the table consists of R_dest_addr, R_next_addr, R_dist, R_metric and R_iface_addr. Such entry specifies that the MP identified by R_dest_addr is estimated to be R_dist hops away from the local MP with the cumulative radio-aware metric equal to R_metric, that the symmetric neighbor MP with interface address R_next_addr is the next hop MP in the route to R_dest_addr, and that this symmetric neighbor MP is reachable through the local interface with the address R_iface_addr. Entries are recorded in the routing table for each destination in the network for which a route is known. All the destinations, for which a route is broken or only partially known, are not recorded in the table.

More precisely, the routing table is updated when a change is detected in either:

— the link set,

— the neighbor set,

— the 2-hop neighbor set,

— the topology set,

— the Multiple Interface Association Information Base,

More precisely, the routing table is recalculated in case of neighbor appearance or loss, when a 2-hop tuple is created or removed, when a topology tuple is created or removed or when multiple interface association information changes. The update of this routing information does not generate or trigger any messages to be transmitted, neither in the network, nor in the 1-hop neighborhood.

11A.7.12.2 Path Selection Algorithmselection algorithm

To construct the routing table of MP X, a shortest path algorithm is run on the directed graph containing the arcs X -> Y where Y is any symmetric neighbor of X (with Neighbor Type equal to SYM), the arcs Y -> Z where Y is a neighbor MP with willingness different of WILL_NEVER and there exists an entry in the 2-hop Neighbor set with Y as N_neighb_main_addr and Z as N_2hop_addr, and the arcs U -> V, where there exists an entry in the topology set with V as T_dest_addr and U as T_last_addr.

The optimal path will be is selected through the following procedure using the neighbor set, the link set, the 2-hop neighbor set and the topology set:

a) All the entries from the routing table are removed. Clear the list of candidate mesh pointsMPs. Initialize the shortest-path tree to only the rootroot MP.

b) Call the mesh point MP just added to the tree mesh point MP V. For each mesh point MP W which is the one-hop neighbor and the link between V and W is the SYM link, calculate the link cost (the sum of radio-aware metric) D of the resulting path from the root MP to W. D is equal to the sum of the link cost of the (already calculated) shortest path to vertex V and the advertised cost of the link between vertices V and W. If D is:

- Greater than or equal to the value that already appears for vertex W on the candidate list, then examine the next mesh pointMP.
- Less than the value that appears for W on the candidate list, or if W does not yet appear on the candidate list, then set the entry for W on the candidate list to indicate D from the rootroot MP. The next hop that result results from the candidate path for W accordingly is set to the same as the next hop of V.

c) If at this step the candidate list is empty, the shortest-path tree (of transit vertices) has been completely built and the algorithm terminates. Otherwise, choose the mesh point MP belonging to the candidate list that is closest to the rootroot MP, and add it to the shortest-path tree (removing it from the candidate list in the process).

d) The new route entries for the destination MP W is recorded in the routing table.

e) Iterate the algorithm by returning to Step 2.

11A.7.13 Associated Station Discoverystation discovery

There are two different operation modes for associated station discovery — (1) “Full Base Diffusion” mode and (2) “Checksum” mode; in the first mode a an MAP advertises the full contents of its LAB, while in the second mode it advertises only the checksums of blocks in the LAB for possible detection of inconsistent blocks in GAB at receiving MPs.

11A.7.13.1 Associated Station Discovery station discovery in “Full Base Diffusion” mode

As explained previously, in “Full Base Diffusion” mode, the MAP broadcasts periodically messages (LABA messages), representing the full content of its Local Association Base; the other MPs receive them and populate their Global Association Table by this means.

11A.7.13.2 Local Association Base Advertisement (LABA) Message Generationmessage generation

A Mesh AP with associated station generates periodically LABA messages containing the pairs of (Associated Station Addresses, Station Sequence Number), with proper block index, corresponding to the currently associated stations in its Local Associated Station Base. LABA messages are broadcast to the entire WLAN Mesh.

The Mesh AP must shall generate LABA messages before the entries of the previous LABA message(s) reach their expiration time (as advertised by expire time).

As an optimization to increase responsiveness of the station discovery protocol to changes in the association tables, a Mesh AP may generate an earlier LABA message in case of change (which may include information for all blocks as previous LABA messages, or only the blocks that have changed).

11A.7.13.3 LABA Message Forwardingmessage forwarding

LABA messages are broadcast and retransmitted forwarded by the MPRs in order to diffuse the messages to the entire network. LABA messages MUST shall be forwarded according to the "default forwarding algorithm" (described in clause 11A.7.3.3)

11A.7.13.4 LABA Message Processingmessage processing

Upon receiving an LABA Message, each Mesh Point MP updates its Global Association Base and its Local Association Base, as described in the following sectionssubclauses.

11A.7.13.4.1 Populating the Global Association Base Populationassociation base population

Having received an LABA Message, the receiver Mesh Point will use MP uses the following information:

— the originator address in the RA-OLSR message header of the LABA message: the address of the Mesh AP which generated the message

— the list of associated STAs addresses (with their sequence numbers)

It will perform performs the actions of:

— ensuring that any new association is added to the Global Association Base

— ensuring that any terminated association is removed from the Global Association Base

— updating properly the expiration times of the global association tuple with the one from the message received

More precisely, it updates the Global Association Base as follows:

a) Global Association Base Block cleaning for originator Mesh AP: for each advertised block, the received information will include includes all the content of the block as present in the last Local Association Base (of the originator Mesh AP): all preexisting tuples are now obsolete, and hence the following step is performed:

b) The receiving Mesh Point MP deletes any global association tuple where GA_AP_address == LABA message originator, and where GA_block_index is one of the block indices in the LABA message.

c) Update of information for advertised associated stations:

For each station entry in the list of associated STAs, potential corresponding information in the Global Association Table is added or updated:

Table is added or updated:

1) If there exists a global association tuple in global association base where:

GA_AP_address == LABA Message originator, AND

GA_station_address == station address, AND

GA_station_sequence_number <= Station sequence number of the entry

Then, it is updated as follows:

GA_station_sequence_number == station sequence number in frame

GA_block_index == block index

GA_expiration_time = current time + validity time

2) Otherwise, if there exists a station tuple in the global association bases where:

GA_AP_address == LABA Message originator, AND

GA_station_address == station address of the current entry, AND

GA_station_sequence_number > Station sequence number of the entry

Then, the information in the message relative to the station is ignored because it is considered obsolete

d) Otherwise, a new station tuple is recorded in the Global Association Base where:

GA_AP_address == LABA Message originator,

GA_station_address == station address of the current entry

GA_block_index == block index,

GA_station_sequence_number == station sequence number of the entry

GA_expiration_time = current time + validity time

11A.7.13.4.2 Populating the Local Association Baselocal association base: Updateupdate

Because a Mesh AP may not always be informed directly of a disassociation of a STA, it can be the case that it detects indirectly a disassociation by discovering that another Mesh Point MP advertises a STA believed to be associated, and with a newer sequence number.

To handle the occurrence of such event, the following processing is done:

— For each added entry in the list of associated STAs, if there exists a tuple in the Local Association Base where:

LA_station_address == associated station address of the entry

LA_station_sequence_number <= station sequence number of the entry

Then the Local Association Tuple is deleted and the corresponding STAs are assumed to be disassociated (and other IEEE 802.11 protocols might be informed of such an event).

11A.7.13.5 Associated Station Address Search station address search and Population population of the Routing Table routing table

The intent of the maintenance of the different Association Tables is that when an MP needs a Mesh Point would need a mesh path for an associated station address, it would proceeds as follows: it would search searches in its Local Association Base, and in its Global Association Base, for association tuples with the required address. If more than one were found to be present, the Mesh Point would choose MP chooses the association tuple with the highest GA_station_sequence_number.

Because the routes are proactively computed, instead of performing this search, the Mesh Point will complement MP complements the routing table calculation described in another section subclause in an equivalent way with the following step:

For each entry in the Global Association Base where:

— There is no local association tuple in the LAB, nor global association tuple in the GAB with the same station_address, and a higher station_sequence_number, AND

— There exists a recorded routing entry such that:

R_dest_addr == GA_AP_address

Then a route entry is created with:

R_dest_addr == GA_station_address

R_next_addr == R_next address (of the recorded routing entry)

R_dist == R_dist (of the recorded routing entry)

R_iface_addr == R_iface_addr (of the recorded routing entry)

11A.7.13.6 Associated Station Discovery station discovery in "Checksum Diffusion" mode

11A.7.13.6.1 Overview

Because sending periodically the full Local Association Base can be expensive, especially in the presence of numerous Mesh APs, another mode of operation is presented here, with the following changes in the message generation and processing.

— Message Generation by a Mesh AP: rather than sending the whole Local Association Base, a checksum is computed for each block of the Base, and the list of checksums is sent in a Local Association Base Checksum Advertisement (LABCA) message.

— Message Processing by receiver Mesh PointsMPs: as a result, the receiver Mesh PointsMPs, can no longer populate their Global Association Table for the originator Mesh AP, naturally. But they will verify that the checksum of their Global Association Table matches instead.

— Block Request by a receiver Mesh PointMP: when a checksum mismatch is detected, a receiver Mesh Point will issue MP issues an Association Base Block Request (ABBR) message, with the list of indices of mismatching blocks. It will be unicast is transmitted using an individually addressed frame to the Mesh AP.

— Block Request processing: the indices of mismatching blocks are recorded, and will be is advertised in the next LABA/LABCA generation.

The policy for choosing to generate LABCA messages instead of LABA messages is left as a choice to the Mesh AP, but the following is suggested:

— If there have been no changes in the Local Association Table for a duration greater than STABLE_LOCAL_ASSOCIATION_TIME, the Mesh AP generates LABCA messages.

— If there is a change (i.e., station association or disassociation), it switches back to sending LABA.

As a side note, remark that a prerequisite for application of this mode of operation is that all Mesh Points MPs in the WLAN Mesh support it.

11A.7.13.6.2 Detailed Message Generation message generation and Message Processing message processing

The processing of the LABA messages is unchanged compared to the operating mode "Full Base Diffusion".

11A.7.13.6.3 LABCA Message Generationmessage generation, Forwarding forwarding and Processingprocessing

The generation of LABCA messages can replace the generation of LABA messages, and is straightforward from the LABCA message format description, and from the checksum calculation described in a following sectionsubclause.

The forwarding of LABCA messages is done according to the "default forwarding algorithm".

The processing of LABCA messages is also simple: as indicated, a verification of all the Checksum Status Values is performed for all blocks.

In case of one mismatch or more mismatches an ABBR message is generated. However, for each block which matches, the GA_expiration_time of tuples in the matching blocks is updated, and also each entry in a block which has disappeared is deleted:

a) For each tuple in the Global Association Base where:

- GA_AP_address == Originator Address, AND
- GA_block_index == in the range from 0 to the number of blocks-1 in the message

The field GA_expiration_time is updated as follows:

- GA_expiration_time = current time + validity time (from message header)

b) For each tuple in the Global Association Base where:

- GA_AP_address == Originator address, AND
- GA_block_index is greater than last block index in the LABCA message,

The tuple is deleted.

11A.7.13.6.4 ABBR Message Generationmessage generation, Forwarding forwarding and Processingprocessing

As indicated, at the receiver Mesh PointMP, an ABBR message is generated each time are detected by the receiver Mesh PointMP, for the blocks of the Global Association Base corresponding the originator Mesh AP address of the LABCA message. The ABBR message includes the list of the block indices for which there is a mismatch. It is sent to the Mesh AP originator of the mismatching LABCA message, with an RA-OLSR Message header in a RA-OLSR packet frame format, as a unicast packet an individually addressed frame with destination MAC address set to the address of the Mesh AP.

The ABBR message forwarding is simple: because the ABBR message is unicastindividually addressed, the message is forwarded in a similar way the data message are forwarded, i.e., hop by hop and using the routing table to reach the destination Mesh AP. Note however that the hop count and the TTL fields of the RA-OLSR message header may be updated at each hop.

The processing of ABBR messages at the destination Mesh AP consists in recording all the blocks indices for which there is a mismatch. At the next LABA/LABCA generation, the Mesh AP will either:

— Broadcast a Full LABA message (with all Local Association Base).

— Broadcast a LABA message with all the content of blocks for which there was (at least) one request, and broadcast one LABCA message in addition.

11A.7.13.6.5 Checksum Calculationcalculation

The other sections subclauses use the fact that a checksum is calculated and verified. In this clause, a procedure for performing checksum calculation is provided: For a block with a given index,

— All the corresponding association tuples, appropriately either in LAB or in the GAB, are retrieved.

— They are converted in sequence of octets (8): (MAC address, sequence number)

— They are sorted by increasing order.

— They are concatenated as a sequence of octets.

— A hashing function[3] is applied to the sequence.

— The result of the hashing function is used to represent the checksum

11A.7.14 Recommended Values values for Constantsconstants

This clause lists the values for the constants used in the description of the protocol.

11A.7.14.1 Setting emission intervals and holding times

The recommended value for constant C is the following:

C = 1/16 seconds (equal to 0.0625 seconds)

[3] In the specification of LABCA IE information element in clause 7.3.2.74.5, we assume that MD5 is used for checksum calculation.

C is a scaling factor for the "validity time" calculation ("Vtime" and "Htime" fields in message headers, see clause 11A.7.3 and 11A.7.6). The "validity time" advertisement is designed such that MPs in a network may have different and individually tunable emission intervals, while still interoperates fully. For protocol functioning and interoperability to work:

— The advertised holding time MUST shall always be greater than the refresh interval of the advertised information. Moreover, it is recommended that the relation between the interval (from clause 11A.7.5), and the hold time is kept as specified in clause 11A.7.5, to allow for reasonable packet frame loss.

— The constant C SHOULD should be set to the suggested value. In order to achieve interoperability, C MUST shall be the same on all MPs.

— The emission intervals, along with the advertised holding times (subject to the above constraints MAY may be selected on a per MP basis.

Note that the The timer resolution of a given implementation might not be sufficient to wake up the system on precise refresh times or on precise expire times: the implementation SHOULD should round up the 'validity time' ("Vtime" and "Htime" of packetsframes) to compensate for coarser timer resolution, at least in the case where "validity time" could be shorter than the sum of emission interval and maximum expected timer error.

11A.7.14.2 Emission Intervalsintervals

HELLO_INTERVAL = 2 seconds

REFRESH_INTERVAL = 2 seconds

TC_INTERVAL = 5 seconds

MID_INTERVAL = TC_INTERVAL

HNA_INTERVAL = TC_INTERVAL

11A.7.14.3 Holding Timetime

NEIGHB_HOLD_TIME = 3 x * REFRESH_INTERVAL

TOP_HOLD_TIME = 3 x * TC_INTERVAL

DUP_HOLD_TIME = 30 seconds

MID_HOLD_TIME = 3 x * MID_INTERVAL

HNA_HOLD_TIME = 3 x * HNA_INTERVAL

The Vtime in the message header (see clause 11A.7.3), and the Htime in the HELLO message (see clause 11A.7.6) are the fields which hold information about the above values in mantissa and exponent format (rounded up). In other words:

Value = C*(1+a/16)*2b [in seconds]

where a is the integer represented by the four highest bits of the field and b the integer represented by the four lowest bits of the field.

Notice, that for the previous recommended value of C, (1/16 seconds), the values, in seconds, expressed by the formula above can be stored, without loss of precision, in binary fixed point or floating point numbers with at least 8 bits of fractional part. This corresponds with NTP time-stamps and single precision IEEE Standard 754 floating point numbers.

Given one of the above holding times, a way of computing the mantissa/exponent representation of a number T (of seconds) is the following:

— Find the largest integer ‘b’ such that: T/C >= 2b

— Compute the expression 16*(T/(C*(2b))-1), which may not be a integer, and round it up.

This results in the value for ‘a’

— if ‘a’ is equal to 16: increment ‘b’ by one, and set ‘a’ to 0 now,

— ‘a’ and ‘b’ should be integers between 0 and 15, and the field will be is a byte holding the value a*16+b

For instance, for values of 2 seconds, 6 seconds, 15 seconds, and 30 seconds respectively, ‘a’ and ‘b’ would be: (a=0, b=5), (a=8,b=6), (a=14,b=7) and (a=14,b=8) respectively.

11A.7.14.4 Message Typestypes

HELLO_MESSAGE = 1
TC_MESSAGE = 2
MID_MESSAGE = 3
HNA_MESSAGE = 4

11A.7.14.5 Neighbor Typestypes

NOT_NEIGH = 0
SYM_NEIGH = 1
MPR_NEIGH = 2

11A.7.14.6 Willingness

WILL_NEVER = 0
WILL_LOW = 1
WILL_DEFAULT = 3
WILL_HIGH = 6
WILL_ALWAYS = 7

The willingness of an MP may be set to any integer value from 0 to 7, and specifies how willing an MP is to be forwarding traffic on behalf of other MPs. MPs willMPs, by default, have a willingness WILL_DEFAULT. WILL_NEVER indicates an MP which does not wish to carry traffic for other MPs, for example due to resource constraints (like being low on battery). WILL_ALWAYS indicates that an MP always should be selected to carry traffic on behalf of other MPs, for example due to resource abundance (like permanent power supply, high capacity interfaces to other MPs).

An MP may dynamically change its willingness as its conditions change.

One possible application wouldis, for example, be for an MP, connected to a permanent power supply and with fully charged batteries, to advertise a willingness of WILL_ALWAYS. Upon being disconnected from the permanent power supply (e.g., a PDA being taken out of its charging cradle), a willingness of WILL_DEFAULT is advertised. As battery capacity is drained, the willingness would be further reduced. First to the intermediate value between WILL_DEFAULT and WILL_LOW, then to WILL_LOW and

finally to WILL_NEVER, when the battery capacity of the MP does no longer support carrying foreign traffic.

11A.7.14.7 Misc. Constantsconstants

MAXJITTER = HELLO_INTERVAL / 4

11A.7.15 Sequence Numbersnumbers

Sequence numbers are used in RA-OLSR with the purpose of discarding "old" information, i.e., messages received out of order. However with a limited number of bits for representing sequence numbers, wrap-around (that the sequence number is incremented from the maximum possible value to zero) will occuroccurs. To prevent this from interfering with the operation of the protocol, the following MUST shall be observed:

The sequence number S1 is said to be "greater than" the sequence number S2 if:

S1 > S2 AND S1 - S2 <= MAXVALUE/2 OR

S2 > S1 AND S2 - S1 > MAXVALUE/2

where the term MAXVALUE designates the largest possible value for a sequence number.

Thus when When comparing two messages, it is possible — even in the presence of wrap-around — to determine which message contains the most recent information.

11A.8 Intra-Mesh Congestion Controlmesh congestion control

11A.8.1 Motivation (Informative)

The original 802.11 MAC and all the recent MAC enhancements (e.g., 11e, 11i, 11k) are designed primarily for one-hop wireless networks. One of the key distinctions of mesh networks is their multi-hop data forwarding nature. IEEE 802.11 DCF and EDCA provide no end to end consideration or coordination beyond a single hop at all. The Mesh Points MPs in a mesh network get fair share of the channel access on a nodean MP-by-node MP basis without any coordination among the nodesMPs. More specifically, each mesh point MP contends for the channel independently, without any regard for what is happening in the upstream or downstream nodesMPs. One of the consequences is that a sender with backlogged traffic may rapidly inject many packets frames into the network which would result in local congestion for nodes MPs downstream. Local congestion is defined as the condition when an intermediate mesh point MP receives more packets frames than it can transmit out in a pre-defined time window. The result of local congestion is that the local buffer gets filled up quickly, and eventually the buffer may become full and packets frames will have to be dropped from the buffer.

Local congestion exists in both wired and wireless networks, but the performance degradation it causes in wireless networks is even worse than in wired network, because of the very nature of wireless medium being a shared resource. In a wired network, the neighboring links across a flow can be regarded as independent resources and so bandwidth consumed by one link does not adversely affect the bandwidth available on another link. This is not true in a multi-hop wireless mesh networks any morenetworks. Typically multiple mesh nodes MPs use the same channel to stay connected to the network, and so when a node an MP is transmitting, other nodes MPs in the vicinity have to refrain from transmitting. Therefore, if an aggressive upstream sender just blindly uses its share of the channel access time slots to inject packets frames into the network, while the downstream intermediate mesh points MPs cannot effectively forward the packets frames to the final destinations, the channel access time slots used by the sender not only are wasted but also further reduce the channel access time slots available to the downstream nodes MPs and hence adversely reduce the end to end throughput even more. Furthermore, hot spots may occur in a wireless mesh network where

bandwidth consumed on one link in a neighborhood may adversely affect the bandwidth available to other links. This may cause unfairness due to the fact that the contention level seen by the nodes MPs in a neighborhood may be different depending on each nodeMP's location. Therefore in wireless networks congestion can occur even after the admission control is applied at the connection level.

Wired networks have been dealing with congestion control for a long time, and one of the effective tools to combat congestion has been end-to-end flow control implemented at the higher layers of the network stack. The TCP sliding window has been the primary example for end-to-end flow control at L4. Then why can't we rely on higher layer flow control like TCP to solve this problem?

In wireless mesh networks, neighboring nodes share the same medium and so scheduling across neighboring links on a multi-hop path is important to maximize the network throughput. Research has shown that TCP congestion control does not work well across a multi-hop wireless network, so simply relying on TCP alone is not a viable solution.

To provide a good tradeoff between complexity and performance, we recommend a simple hop-by-hop congestion control mechanism is defined to address the problem. This mechanism should be implemented at each mesh pointMP, and it includes three basic elements: local congestion monitoring, congestion control signaling, and local rate control. The basic idea is that each mesh point shall MP actively monitor monitors its local channel utilization condition so that it can detect local congestion when it happens; three new mesh action frames are defined ("Congestion Control Request", "Congestion Control Response" and "Neighborhood Congestion Announcement") to support the necessary hop-by-hop signaling between neighboring Mesh PointsMPs; upon receiving "Congestion Control Request" from a downstream Mesh PointMP, the upstream neighbors shall employ local rate control to help relieve the congestion being experienced downstream and upon receiving "Neighborhood Congestion Announcement" from a neighbor mesh pointMP, the neighbors shall employ local rate control to help relieve the congestion being experienced in the neighborhood. The exact mechanism for local congestion monitoring and the exact trigger for congestion control signaling is outside the scope of this specification standard and entirely up to the implementation; but the frame formats for congestion signaling are specified here.

11A.8.2 Local Congestion Monitoring congestion monitoring (Informative)

In order to effectively control or even avoid congestion in the network, each Mesh Point has to monitor MP monitors its local/neighborhood congestion condition so that when necessary, it can notify the neighborhood/upstream nodes MPs of congestion, by transmitting a broadcast "Neighborhood Congestion Announcement" and/or a unicast an individually addressed "Congestion Control Request". How the monitoring and congestion detection are implemented is a pure local implementation matter and is considered outside the scope of this specificationstandard. For the sake of completeness here two different monitoring and congestion detection mechanisms are provided as examples:

One way for detecting congestion is for each Mesh Point MP to keep track of its own effective MAC transmission and receiving rate for the packets frames to be forwarded (excluding the received packets frames destined for this Mesh PointMP), and to monitor the backpressure which is the difference between the aggregate receive and transmit rates. The goal of rate control is to maintain near zero backpressure at the local nodeMP. When the backpressure builds up significantly at the local nodeMP, the node MP informs its previous hop nodes MPs by sending "Congestion Control Request" messages so that the recipient nodes MPs can decrease their transmission rate accordingly by local rate control mechanism. Furthermore, the Mesh Point will inform MP informs its neighbors of the congestion level by sending broadcast signaling messages "Neighborhood Congestion Announcement", so that its neighbors shall regulate their transmission rate depending on service differentiation criteria. The rate regulation should be limited by an expiration timer.

Another possible method for congestion detection is based on queue size. Following is an illustrative example: In this method we select upper and lower queue size thresholds. If the queue size is above the upper threshold, the node MP informs its previous hop neighbors by sending unicast individually addressed

signaling messages "Congestion Control Request Messages", so that its previous hop neighbors can decrease their transmission rate to it accordingly by local rate control mechanisms. If the queue size is between the lower and the upper thresholds, the node MP again declares congestion, but this time, sends the unicast individually addressed signaling messages "Congestion Control Request" messages to the upstream nodes MPs with a probability given by

$$\frac{queuesize - lowerthreshold}{upperthreshold - lowerthreshold}$$

Furthermore, the possible methods described above for congestion detection may be based on queue size or difference between the MAC transmit and receive rates per upstream node MP as the downstream node MP may use different forwarding links for different upstream nodesMPs' traffic. This method may help to reduce the congestion problem caused by a particular forwarding link by rate limiting only the upstream nodes MPs whose traffic is forwarded through this link.

The cause of the congestion may be estimated during the local congestion monitoring process in order to trigger the appropriate congestion control signaling and local rate control mechanism. For example, the channel access rate of a downstream node MP may be high within a neighborhood with a small channel load and still congestion may occur at this node MP that has high retry count due to forward link characteristics or its downstream node MP capacity. In this case, only Only the upstream Mesh Point MP should be informed of the congestion for the application of local rate control while the neighbors' transmissions should not be affected. On the other hand, a high channel load value in the neighborhood where the congested Mesh Point MP determined the channel to be busy (CCA/NAV indications) in a manner that adversely affects its transmission attempts may be an indication to inform the neighbor Mesh Points MPs about the congestion for the application of local rate control.

11A.8.3 Congestion Control Signaling control signaling

This specification standard does not define the exact conditions that would trigger the congestion control signaling. Such signaling can be done periodically or non-periodically. Note that the The inclusion of "Expiration Timer" in the "Congestion Control Request" and "Neighborhood Congestion Announcement" frames provides enough flexibility for both cases.

"Congestion Control Request" message is transmitted as a unicast an individually addressed message. The frame contains the "Target Transmission Rate Element" which specifies the target data rate that should not be exceeded by the upstream Mesh Point MP transmitting to this nodeMP. When a node an MP experiences local congestion, this target rate is generally already exceeded by the previous hop. The purpose of "Congestion Control Request" message is to notify the previous hop node MP of the congestion condition and to request from it to rate limit its transmission to help remove the congestion. The target rate is specified per AC per upstream neighbor so that the flow control can be done separately for each AC at each upstream neighbor. For example, this allows the nodes MPs to request rate control of higher throughput traffic flows (such as bulk data or video traffic) while maintaining the QoS requirements of delay-sensitive flows (such as voice traffic) by leaving them unaffected by the congestion control scheme.

Nodes MPs that receive a "Congestion Control Request" message must shall enact the congestion control request and immediately adjust their MAC transmission rate to meet the target rate specified in the "Congestion Control Request" message. They should also reply to the congestion control request with an "Offered Traffic Load" element. This message includes the offered traffic load of each access category at the nodeMP. This information can be used by the congested node MP to compute the target rate.

The "Neighborhood Congestion Announcement" message is transmitted as a broadcast message. The purpose of this message is to notify the neighbor Mesh Points MPs of the congestion condition due to the high channel load in the neighborhood and to request that all neighbors rate limit their transmission in order to eliminate the congestion.

11A.8.4 Target Rate Computation rate computation (Informative)

In this section subclause informative text and guidelines are provided on how nodes MPs can compute the target rate to be communicated to the upstream nodes MPs in two cases of channel congestion and channel underutilization.

In the case that a node an MP detects congestion on the channel, it measures its outgoing traffic rate as an upper bound for the allowed aggregate incoming traffic to this nodeMP. Then the node MP computes the share of each upstream node MP of this rate. In order to compute the individual upstream nodesMPs' shares, the node MP can take into account any additional information it might have, for example the offered load of upstream nodes MPs or system fairness policies, when computing the target rate. In the simplest case, where there is no additional information available, the share of each upstream node MP with traffic for this congested node MP is less than or equal to the measured outgoing traffic divided by number of such nodesMPs. Note that it It is recommended that the congested nodes MPs utilize a more conservative approach in computing the target rate, as it would allow additional resources for initialization of new traffic flows.

When a node an MP detects that the channel is underutilized, it needs to notify the upstream nodes MPs of its available resources. Such notification is especially important in cases where the nodes MPs reduce their traffic due to existing congestion and later on the source of congestion disappears. In order to ensure efficient utilization of the network resources, it is important that the upstream nodes MPs be notified of the excess resources. However, the increase in the rate should happen in a way that does not result in congestion in the system. Computation of the target rate in this case requires additional information including the channel utilization (can be measured in the form of channel idle time), probability of collision (an indication of number of nodes MPs contending for the channel), and upstream nodesMPs' offered load (to ensure efficient allocation of resources). Moreover, an accurate computation of target rate for each node MP can only be done if additional information is available regarding the channel quality (achievable transmission rate), average data packet frame size, and the MAC policy (for example whether RTS/CTS is used, in order to compute the MAC overhead per pkt frame transmission).

Assuming a network consisting of only a one-hop flow and an underutilized channel, which is idle for t time units per measurement window on average, one can compute the target increase rate n (packets frames per measurement window) as $n = \frac{tC}{P + CT_{oh}}$, where C denotes average channel capacity, P denotes average packet frame size, and T_{oh} denotes average overhead per packet frame in time units.

In the case that there exist more than one node MP contending for the channel, this upper bound target rate needs to be adjusted accordingly.

11A.8.5 Local Rate Control Mechanism rate control mechanism (Informative)

Upon receiving either a "Congestion Control Request" or a "Neighborhood Congestion Announcement" message, the receiving node MP needs to reduce its effective MAC transmission rate accordingly by locally rate limiting its traffic. If the received message is a "Congestion Control Request" message, the node MP is only required to rate limit the traffic that it sends out to the neighbor requesting the congestion control. Note that reducing Reducing the MAC transmission rate does not mean changing the PHY layer modulation and encoding scheme to reduce the transmission rate of the radio. Typically a node an MP can adjust its effective MAC transmission rate by internal scheduling algorithms that purposely delay the transmission of packets frames in a particular AC to a particular neighbor. How the rate control is implemented is an implementation choice and is hence outside the scope of this specificationstandard. Although there are many different ways of implementing a local rate controller, the objective of any implemented mechanism should be to keep the service differentiation and fairness in the neighborhood.

Here we mention two examples of an example rate control mechanismsmechanism. First example of the local rate control mechanism is based on dynamically adjusting EDCA parameters depending on the congestion condition in the node MP and/or the neighborhood. The EDCA parameters to be adjusted can be AIFSN, CWmin, or both. Use of such a mechanism would be especially effective when the source of a congestion-causing flow is a STA associated with a an MAP. Although the BSS communication between a an MAP and associated STAs is beyond the scope of this WLAN Mesh specificationstandard, note that a an MAP implementation may adjust the EDCA parameters for the BSS to alleviate congestion due to traffic generated by associated STAs. Thus, STAs do not require explicit knowledge of the congestion control scheme.

Note that while While this specification standard focuses only on intra-mesh congestion control, it is important to point out that when some of the Mesh Points MPs are also Mesh APs, it is critical for MAP to implement effective rate control mechanism for the BSS traffic in conjunction with rate control mechanism for mesh traffic. When a an MAP receives Congestion Control Request or Neighborhood Congestion Announcement from its neighbor, it should effectively control both the mesh traffic transmission rate by itself and the BSS traffic transmission from its STAs so that fairness is maintained. The exact mechanism to control BSS traffic is out of scope for this specificationstandard. While one can dynamically adjust EDCA parameters to control mesh traffic, one can also employ the same method for MAP to regulate QSTA traffic by setting different EDCA parameters for the BSS.

Similarly, if a an MP itself generates local application traffic to be forwarded into the mesh, it is also important to regulate the local traffic injection such that when it needs to ask its upstream neighbor to regulate the mesh traffic transmission rate, it also needs to regulate its own local traffic injection from its applications. In another word, it should treat its own local application traffic as if it is coming from another neighbor over the air, and it should not allow an aggressive local application saturate the queues when it asks other neighbor nodes MPs to slow down.

11A.9 Mesh Beaconing beaconing and Synchronizationsynchronization

11A.9.1 Synchronization

Synchronization and beacon generation services in a WLAN mesh are based upon the procedures defined in clause 11.1 for Infrastructure and IBSS modes of operation.

It is optional for an MP to support synchronization. An MP supporting synchronization may choose to be either synchronizing or unsynchronizing based on either its own requirements or the requirements of its peerspeer MPs. MP's synchronization behavior is communicated through the "synchronization capability field" within the WLAN Mesh Capability element. The synchronizing behaviour for the two classes is defined as follows.

11A.9.2 Unsynchronizing MPs

An unsynchronizing MP is a an MP that maintains an independent TSF timer and does not update the value of its TSF timer based on time stamps and offsets received in beacons Beacon frames or probe responses Probe Response frames from other MPs. An unsynchronizing MP may start its TSF timer independently of other MPs. The "Synchronizing with peer MP" bit in the "Synchronization Capability" field of the WLAN Mesh Capability element, when set to 0, indicates that an MP is currently an unsynchronizing MP. A An MP that supports synchronization may elect to be an unsynchronizing MP if it is communicating with peers that are not requesting synchronization.

11A.9.2.1 Synchronizing MPs (Optional)

A synchronizing MP is an MP that updates its timer based on the time stamps and offsets (if any) received in beacons Beacon frames and probe responses Probe Response frames from other synchronizing MPs. The

"Synchronized Synchronizing with peer MP" bit in the "Synchronization Capability" field of the WLAN Mesh Capability element, when set to 1, indicates that the MP is currently a synchronizing MP.

Synchronizing MPs should attempt to maintain a common TSF time called the Mesh TSF time. An MAP maintains the mesh TSF in terms of its TSF timer and its self TBTT offset such that the sum of the self TSF timer and the self TBTT offset equals the mesh TSF time. All beacons Beacon frames and probe responses Probe Response frames by such synchronizing MAPs carry the Beacon Timing IE information element to advertise its self offset value relative to the Mesh TSF time.

Synchronizing MPs translate the received time stamps and offsets (if any) from beacons Beacon frames and probe responses Probe Response frames from other synchronizing MPs to their own timer base, and update their timer as described as follows:

Translated time stamp = Received time stamp + Received offset (if any) – Receiver's offset (if any);

Any synchronizing MP will adopt adopts the translated time stamp as its own if it is later than the timer value of self as described for IBSS mode of synchronizationsynchronization in 11.1.1.2.

Synchronizing MPs may optionally choose to update their offsets instead of their timers. The offset update process in this case is as below.

If (received time stamp + received offset) > (self time + self offset)

New self offset value = received time stamp + received offset – self time.

Note that the The "Received offset" above is the "self offset" in the received Beacon Timing IE information element from the neighbor MP, and the "Receiver's offset" is the receiving MP's own self offset.

11A.9.2.2 Interaction between synchronizing and unsynchronizing MPs

A synchronizing MP may or may not request synchronization from its peerspeer MPs by setting the "Requests Synchronization from Peer" bit in the Mesh Capability element in Beacon or Probe Response frames. However, if an MP requests synchronization from its peers, it has to shall be a synchronizing MP at that time. Initially, an MP may be in unsynchronized state, but it may switch to synchronized state and vice-versa based on either its own requirements or the requirements of peers.

An unsynchronizing MP may change into a synchronizing MP if it is capable of synchronizing, by setting its "Synchronizing with peer MP" bit to 1.

A MP that associates with an unsynchronizing MP and intends to enter power save mode may need to maintain additional information to wake up at the neighboring MP's DTIM interval during its PS mode operations as described below.

An unsynchronizing MP or an MP that has an established peer link with an unsynchronizing MP shall maintain information to wake up at the neighboring MP's Mesh DTIM beacon timing when it is in power save mode, as described below.

11A.9.3 Beaconing

Any An MP may choose to beacon either as defined in the IBSS mode (clause 11.1.2.2) or as defined in the infrastructure mode of operation (clause 11.1.2.1).

11A.9.3.1 Beaconing by unsynchronizing MPs

Unsynchronizing MPs generate beacons Beacon frames according to the beacon generation procedures defined in clause 11.1.2.1 1. Unsynchronizing MPs may choose their own beacon interval and TSF independent of other MPs.

Unsynchronizing MPs may implement utilize the beacon collision avoidance mechanism defined in clause 11A.9.4 to reduce the chances chance that it will transmit beacons transmits Beacon frames at the same time as one of its neighbors or neighbor's neighbors.

Unsynchronizing MPs shall treat any associated MPs operating in PS mode identical to a STA and the receive operation procedures of the associated MP in PS mode are as defined in clause 11.2.1.6 for STAs in the PS mode.

11A.9.3.2 Beaconing by synchronizing MPs

Synchronizing MAPs generate beacons Beacon frames according to the beacon generation procedures described in clause 11.1.2.1. The value of the aBeaconPeriod attribute used by synchronizing MAPs shall equal a sub-multiple of the Mesh DTIM interval. Synchronizing MAPs shall use and advertise a non-zero self TBTT offset value using the Beacon Timing element.

Synchronizing non-AP MPs generate beacons Beacon frames according to the beacon generation procedures described for IBSS operation in clause 11.1.2.2, unless acting as a designated beacon broadcaster (see clause 11A.9.3.3). A non-AP MP that receives a beacon from an MP with the Mesh ID the same as its own after TBTT and before being able to send its own beacon may cancel that beacon transmission. Specifically, the following rules apply for beacon transmission.

a) If the BSS mesh has a Beacon Broadcasterbeacon broadcaster, all MPs, except the BB suspend the transmission of the beacons Beacon frames upon reception of a beacon or a Connectivity Report indicating received beacon during the `dot11BBBeaconRecoveryTimeOut` + `dot11BBBeaconRecoveryAddition` previous MESH Mesh DTIM transmissions. If the MP is able to receive Connectivity Reports from all MPs in the BSSmesh, it suspends the transmissions of beacons Beacon frames if it has received a beacon or a Connectivity Report indicating received beacon during the `dot11BBBeaconRecoveryTimeOut` previous MESH mesh DTIM transmissions.

b) Suspend the decrementing of backoff timers for any non Beacon traffic.

c) Calculate a pseudo-random delay uniformly distributed over the range of zero and twice aCWmin X aSlot time. The CWmin is as used for AC_VO.

d) Wait for the period of pseudo-random delay, decrementing the pseudo-random delay timer using the same algorithm as for back off.

e) A An MP that receives a beacon from another MP which belongs to the same mesh and before being able to send its own beacon may cancel its beacon transmission

f) Send a beacon if the pseudo-random delay has expired and no beacon has arrived during the delay period.

Each synchronizing MP can select its own Beacon interval, but all synchronizing MPs need to share a common Mesh DTIM interval. The Beacon intervals selected by MP must shall always be a submultiple of the Mesh DTIM interval. A synchronizing MP that establishes a Mesh selects its Beacon interval and the MP DTIM period and establishes the common Mesh DTIM interval of the mesh. Mesh DTIM interval equals the product of the beacon interval and the MP DTIM period. A synchronizing MP that joins an existing mesh needs to should adopt the Mesh DTIM interval of the mesh.

Note that the The Mesh DTIM interval and the BSS DTIM interval in MAPs do not have to be identical. MPs use the Mesh DTIM IE information element to advertise the Mesh DTIM interval, whereas the TIM IE

information element is used for advertising the DTIM interval in a BSS. The DTIM Period of these IEs do not have to be identical since one will be is used for the AP service while the other for the Mesh service.

MPs supporting PS mode operation may use the optional designated beacon broadcaster approach described in clause 11A.9.3.3. In case an MP has more than one designated beacon broadcaster in range, its it should follow the ATIM windows for each of them.

An MP operating in Power Save mode will set sets its MP DTIM period to 1 prior to leaving the active state (meaning the beacon interval is the same as the mesh DTIM interval) and would therefore attempt to beacon only once on every Mesh DTIM interval.

11A.9.3.3 Designated Beacon Broadcaster

It The designated beacon broadcaster mechanism enables a mesh to have a designated MP perform beaconing for a defined period of time while all other MPs defer from sending beaconsBeacon frames.

An MP that serves as the designated beacon broadcaster will transmit transmits its beacon using the procedure as described for infrastructure AP operation clause 11.1.2.1 (i.e., will does not use pseudo-random backoff). The general operation of the power management in a mesh network is discussed in clause 11A.10 and the beacon broadcaster is discussed in clause 11A.9.3.3.

An MP supporting this option that received at least one beacon from an MP that is marked as Designated Beacon Broadcaster designated beacon broadcaster in the last 2 Mesh DTIM intervals will does not schedule a Beacon for transmission.

By default the beacon broadcaster in a mesh is rotating between the MPs. MAPs may need to send a beacon even if another MP had already sent a beacon on that same TBTT.

MPs that support power save operation must also may support the designated beacon broadcasting selection as described in this sectionsubclause.

An MP not supporting this service will use uses a pseudo-random backoff procedure for sending of Beacons Beacon frames as described for IBSS operation in clause 11.1.2.2. An MP that receives a Beacon after TBTT and before being able to send its own beacon may cancel that beacon transmission. MPs that do implement this service and are set to be the Beacon Broadcaster beacon broadcaster (BB) will transmit Beacons transmits Beacon frames immediately after TBTT or after a PIFS interval after the clearing of CCA in case CCA was active during at the TBTT.

An MP that supports this option and starts a mesh will set sets itself to be the Beacon broadcaster.

11A.9.3.4 Reporting to the Beacon Broadcaster using Connectivity Reports

The connectivity reports are broadcasted to share information of the received beacons Beacon frames and transmissions from other terminalsneighbors. The reports are exchanged during mesh DTIM ATIM periods, so that each terminal neighbor that is in active state may receive the report. All nodes MPs send connectivity reports according to a periodicity as specified by the BB. The connectivity reports should be transmitted quite seldomly in order to avoid congestions during the ATIM periods and to keep ATIM periods short.

The MPs send connectivity reports between the integer number of MESH mesh DTIM Intervals as specified in Connectivity Reporting Interval field in MP Control field of the neighbor Info element in beacon of the BB. If the Connectivity Reporting Interval is set to zero, the connectivity reporting is not used. The MP may select the pseudo-random transmission time within one Connectivity Reporting Interval for the transmission of the first report. The BB shall maintain the same Connectivity Reporting Interval as specified for the BSS-mesh.

The MP lists the beacon transmitters during the previous Connectivity Reporting Interval. The list of the received beacons Beacon frames is used to give a notification to BB that BSS the mesh has multiple STAs MPs transmitting a beacon. The BB may switch the role to new MP if it receives connectivity reports that indicate multiple beaconsBeacon frames.

The MAC Addresses of Connectivity Reporting MPs field in the Connectivity Report contains the addresses of the MPs, where the MP has received a Connectivity Report within `dot11BBConnectivityReportTimeout`. If no Connectivity Report from the MP is received within the timeout, the address is removed from the neighbor list Mesh Neighbor List element. The BB shall list to Neighbor List all MAC Addresses in the received Connectivity Reports within the `dot11BBConnectivityReportTimeout`. Thus, the The Connectivity Report frames are used to distribute network topology information in the BSSmesh.

A Connectivity Report indicating a received beacon suspends beaconing of the MPs, except the BB continues to beacon. The suspencion suspension of the beaconing of the neighboring MPs avoids the collisions of the beacon transmission. The MPs are allowed to beacon, if the Connectivity Reports from all neighboring MPs indicate no received beacon.

Recovery mechanisms:

a) If connectivity reports indicate that no transmitter of the connectivity report has received a beacon within `dot11BBBeaconRecoveryTimeOut` DTIM beacon intervals and a connectivity reporting interval from the last received beacon from any terminal neighbor has elapsed, the nodeMP(s) that has a connectivity report from all MPs listed in the last received beacon starts to compete on the beacon transmission opportunity

b) After `dot11BBBeaconRecoveryTimeOut` + `dot11BBBeaconRecoveryAddition` DTIM beacon period from the last received beacon from any terminals neighbors has elapsed, the other nodes MPs may start to compete on the beacon transmission opportunities.

11A.9.3.5 Change of Beacon broadcaster

The beacon broadcaster role may be changed periodically. An MP needs to should relinquish its role as the designated beacon broadcaster after no more than MAX_CONT_BB Mesh DTIM intervals. A suggested value of MAX_CONT_BB is 32.

In every Mesh DTIM interval, the current beacon broadcaster sets the BB switch bit to 1 if it wants to change the beacon broadcaster to another MP (see clause 11A.9.3.3). In this case the The neighbor that is first in the neighbor list Mesh Neighbor List shall start acting as the beacon broadcaster and shall send the next Mesh DTIM beacon.

The BB has to make sure should ensure that the neighbor appearing it places at the head of the list is supporting supports frame transmission to MPs in power save mode and Designated Beacon broadcasting.

If a neighbor assigned to be the beacon broadcaster fails to transmit its beacon (shuts down or goes out of range), other MPs will attempt attempts to take over its role. A An MP supporting designated beacon broadcasting will start starts attempting to send beacons Beacon frames using the standard backoff procedure with a Mesh Neighbor list IE List information element designating itself as the BB as soon as it fails to receive 3 consecutive beacons Beacon frames from the last designated BB.

If another MP already transmitted a Beacon with a neighbor list IE Mesh Neighbor List information element designating itself as the BB, then the MP will cancel cancels its pending Beacon.

The selection of the next Beacon Broadcaster beacon broadcaster may be based on Connectivity Reports allowing the current BB to select the next BB in such a manner that network connectivity is ensured. The beacon broadcaster may obtain the connectivity status of the network and apply this information to select the

next BB in the network. The use of information in connectivity reports may aid to detect multiple beacon transmissions in the same coverage and guide the recovery process of merging the beacon transmissions. There may be multiple attributes, like power management mode, connectivity or device capabilities that affect to selection of the BB.

11A.9.4 Mesh Beacon Collision Avoidance (MBCA) mechanism

MPs may optionally adjust their TSF timers to reduce the chances that they will transmit beacons Beacon frames at the same time as one of their neighbors.

Individual MPs may take steps either prior to, or during association, with a WLAN mesh to select a TBTT that does not conflict with its mesh neighbors. An MP may adjust its TSF timer if it discovers that its TBTT may repeatedly collide with the TBTT of a neighbor. Options a an MP has for adjusting its TSF include advancing or suspending the TSF for a period of time.

An MP may collect and report information about the TBTT of neighboring synchronizing MAPs and unsynchronizing MPs using a variety of techniques. The following describes options to receive such information from neighboring MPs.

a) Information from synchronizing neighbors

   MPs that are synchronizing collect the beacon timing information of their neighbors and report it through the beacon timing IEinformation element. This IE information element may be transmitted in selected beaconsBeacon frames, and in action frames responding to requests for such information. Synchronizing MPs may choose any frequency of including the beacon timing information in the Beacon Timing IE information element in their beaconsBeacon frames. The beacon timing information may also be requested via action frames (described in clause 11A.9.4.1), with the response through the beacon timing IE information element in action frames. The action frame approach is especially useful for use by synchronizing MAPs to proactively detect and avoid beacon collisions. Synchronizing MPs are required to be able to respond to requests for such information using the beacon timing IEinformation element. Synchronizing non-AP MPs, when using the Beacon Timing IEinformation element, set the Self TBTT Offset field in the Self Beacon Timing field to 0.

b) Information from unsynchronizing neighbors

   Unsynchronizing MPs may optionally collect and report beacon timing information of their neighbors. Since unsynchronizing MPs do not track the mesh TSF, they report beacon time offsets relative to their self TSF. This information either may be periodically transmitted in beacons Beacon frames at whatever periodicity the MP chooses, or it may be transmitted based on a request response approach through action frames. The Beacon Timing IE information element is used to report this information in beacons Beacon frames as well as in action frames as a response to request action frames. The Self Beacon Timing field in the Beacon Timing IE information element is set to all zeros in this case.

In addition, 802.11k beacon reports may be used by MPs to exchange beacon timing information of their neighbors, with the usage as defined by 802in 11.10.8.11k1.

As an option, Synchronizing MAPs may occasionally delay their beacons Beacon frames after their TBTTs for a pseudo-random time. The pseudo-random delay is chosen so that the transmission time is interpreted by the MAP as not colliding with other beaconsBeacon frames. This behavior further helps in discovery of neighbors through beacons Beacon frames in case they choose colliding offsets. The MBCA mechanism may then be used for choosing non-colliding offsets, in case any colliding offsets are observed

11A.9.4.1 Action frames for beacon timing request and response

Neighbors' Beacon Timing information may be requested and provided through action frames as described above. Clause 7.4.9.10 and Clause 7.4.9.11 show the frame formats — including category and action field values — for the "Beacon Timing Request" and "Beacon Timing Response" frames, respectively.

11A.10 Power Management management in a Mesh mesh (Optional)

11A.10.1 Overview

The need for power save in a mesh environment depends on specific scenarios of operation. In certain scenarios where the MPs are all MAPs or only carry backbone traffic, the devices may not be expected to be power constrained. Specifically MAPs are expected to be awake all the time. However, in scenarios with lightweight and non-forwarding MPs, battery powered MPs power save can be useful. Specifically, MPs that since this class of devices are lightweight or non-forwarding may be expected to be power constrained. Thus, power Power saving in MPs is specified as an optional feature in this documentstandard. The expectation is that devices manufactured to operate in specific scenarios will choose to implement power save mechanism, while other devices may be spared the additional overhead of supporting it.

Some aspects of optional power save support are as follows. The capability to support power save is advertised by MPs. In case a neighbor of an MP does not support power save, the MP may take one of the following two approaches. It may choose not to communicate with that particular neighbor and still go into power save, or it may choose to not use power save mechanism and continue communication with that neighbor. An MP supporting power save may reject a peer link establishment attempt from another MP if this MP does not supporting support power save. MPs supporting power save may operate in power save mode only if all the MPs they have established peer links with support power save. Lightweight MPs may communicate with neighbors without association. If they choose to operate in power save mode, they are aware that communication with non supporting neighbors is not possible. The decision of whether to go in enter power save mode or not has to shall be made considering the power versus communication constraints. Such a decision can be changed dynamically.

In certain scenarios, devices may also choose to operate in STA mode and use the power save service through an AP. While such functionality is beyond the scope of this specificationstandard, this example is included here for completeness. Such a scenario is particularly attractive in the case power save support from mesh point MP neighbors is not available, but a an MAP is available in vicinity. It should be noted that the choice of mesh versus non-mesh device class or role can be made dynamic; that is, a consumer electronic device such as a camera could configure itself as a mesh device when AC powered, but may configure itself as a simple client STA when operating from a battery.

The power save mode operation in a mesh environment has to shall be coordinated with the neighbor discovery mechanisms and route discovery mechanisms, both of which are salient but essential mechanisms for mesh operation. MPs The behavior of an MP in a power save mode shall clearly behave varies depending on whether or not it intends to be a part of utilizes these mechanisms or not. MP's power save mode behavior will be diverse depending upon these intentionsmechanisms. It should be noted that the MP's inappropriate behavior in terms of power save mode may degrade not only the performance of local MP but also the performance of entire mesh.

11A.10.2 Power management and neighbor discovery mechanism

There are possibly 2 two possible cases for an MP in a power save mode in terms of the power save mode coordination with neighbor discovery, as listed in Table s36.

Table s36—Power Save Neighbor Discovery Coordination States

Neighbor Discovery Mode	Description
neighbor discovery active PS MP	Cooresponding to an MP in power save mode which intends to establish a new peer link.
neighbor discovery passive PS MP	Cooresponding to an MP in power save mode which does not intend to establish a new peer link.

neighbor discovery active PS MP	MP in power save mode which intends to establish a new peer link.
neighbor discovery passive PS MP	MP in power save mode which does not intends to establish a new peer link.

Neighbor A neighbor discovery active PS MP shall keep on beaconing transmit beacons periodically, so that the neighboring MPs can determine the existence of the MP. Neighbor A neighbor discovery active PS MP shall transmit a beacon frame at least once in a mesh DTIM period, though it may enter the doze state if a series of conditions are fulfilled as described later in this subclause. Neighbor A neighbor discovery active PS MP shall process the probe request frame if it is received, and respond with a probe response frame.

An MP which tries to establish a new peer link with MPs in power save mode shall perform passive scanscanning. Since MPs in power save mode may transmit beacon beacons at a low rate, an MP which tries to establish a new peer link with MPs in power save mode should perform passive scan scanning for a relatively longer time compared to passive scan scanning in BSS infrastructrure mode operationnetworks. MPs An MP in power save mode which that sets longer mesh DTIM interval intervals may not be found discovered by the neighbor MPs due to the shorter scan duration. MPs in power save mode should set shorter mesh DTIM intervalintervals, if it intends they intend to establish a new peer link links with MPs with higher probability.

Neighbor A neighbor discovery passive PS MP may not transmit beacon periodically, beacons periodically since it does not have to let the newly arrived neighboring MPs to determine the existence of itselfdiscover it. However, it is recommended that an MP transmit some frames to neighbors periodically in order to let the neighboring MPs to perform maintain the maintenance of the peer link, link if the MP is a mobile device and it has a risk of physical topology changes.

In summary, ,an MP shall select selects how to beacon depending on the MP's intention for the new establishment of the peer linklinks. While MP's The trigger to change the beaconing behavior may transit of an MP from one mode to another, the trigger to change the beaconing profile another is beyond the scope of this standard.

11A.10.3 Power management and route discovery mechanism

There are 2 cases for MP in a power save mode in terms of the power save mode coordination with route discovery, as shown in Table s37.

Table s37—Power save route discovery coordination states

Route Discovery Mode	Description
route discovery active PS MP	MP in power save mode which intends to act as a forwarding MP.
route discovery passive PS MP	MP in power save mode which does not intend to act as a forwarding MP.

route discovery active PS MP	MP in power save mode which intends to act as a forwarding MP.
route discovery passive PS MP	MP in power save mode which does not intends to act as a forwarding MP.

Route discovery passive PS MP may not receive nor process the route discovery related mesh management frames. Neighboring MPs of route discovery passive PS MP may not deliver the route discovery related mesh discovery frames and shall try to establish a new route without a participation of route discovery passive PS MP.

A An MP whose neighbors are all in the power save mode or MP which find the destination MP of the path is in power save mode shall deliver the route discovery related mesh management frames to the MP in power save mode.

Since this amendment standard assumes multiple route discovery mechanisms and some of the route discovery mechanim assumes immediate packet frame delivery, all the forwarding MPs should assume that forwarding neighbor MPs are in active state all the time.

11A.10.4 Basic approach

A mesh point An MP supporting power save mode may either operate in an awake state or a doze state. The mesh point will MP shall advertise its power management mode to all neighboring mesh points MPs by using it's beacons setting the "Current Power Management Mode" bit in the power save capability field of the Mesh Capability element in its Beacons and by sending a frame with the Power Management bit in the frame control field set.

Mesh points MPs in power save mode shall periodically listen for mesh DTIM beaconsBeacon frames. A mesh point An MP waking to transmit or receive a beacon will stay stays in the awake state for a minimum period of ATIM window as indicated in their beaconsBeacon frames, before returning to the doze state.

Mesh points MPs in power save mode shall also wakeup according to any negotiated schedule as part of TSPEC setup with other mesh pointsMPs. The mesh point will remain MP remains in the awake state until the end of service period period.

Mesh point An MP wishing to communicate with mesh points MPs that are in power save mode would buffer the traffic targeted for these mesh pointsMPs. They could deliver the traffic to the mesh point MP in one of the following 6 ways:

a) Send traffic to these mesh points MPs only on agreed schedules as negotiated as part of APSD (Automatic Power Save Delivery) TSPEC setup

b) Send directed or broadcast ATIM packets frames to mesh point an MP in power save mode during ATIM window in order to signal them to remain awake and wait for further traffic

c) Send a single null data packet frame to mesh point an MP in power save mode during their the ATIM window in order to reactivate a flow that has been suspended or to signal power management mode change.

d) Send a beacon frame with TIM IE information element indicating the existence of the traffic buffered to the mesh point MP in the power save mode, in order to signal them to remain in awake state and send PS-poll frames to retrieve this buffered traffic.

e) Send a single, short broadcast or multicast frame to mesh points MPs during the ATIM window as described in 11A.10.7.

f) Broadcast Connectivity reports during ATIM window.

MPs operating in power save mode shall wake up and try to receive all the mesh DTIM beacons Beacon frames of the neighboring MPs with which it establishes a peer relationship. Alternatively, a lightweight MP may associate with a an MAP as a simple STA if it intends to enter PS mode.

The PS mode operation of an unsynchronizing MAP is based upon IEEE 802.11 infrastructure mode BSS power save mode operation. In particular, STAs (including MPs) changing Power Management mode shall inform the MAP of this fact using the Power Management bit within the Frame Control field of transmitted frames. Unsynchronizing MAP shall not arbitrarily transmit MSDUs to MP operating in a PS mode, but shall buffer MSDUs and only transmit them at designated times (during BSS DTIM intervals). A An MP that intends to operate in power save mode using this mechanism shall attempt to send a PS-poll frame when it receives a beacon frame whose TIM IE information element notifies it of the existence of buffered traffic. This mechanism is also used for an unsynchronizing MP wishing to communicate with MPs in power save mode or a an MP wishing to communicate with unsynchronizing MPs in power save mode.

11A.10.5 Initialization of power management within a mesh

The following procedure shall be used to initialize power management within a new mesh or on joining an existing mesh.

a) An unsynchronizing mesh point MP that creates/joins a mesh shall set the values of ATIM window, beacon interval, mesh DTIM interval, and power management mode. These values are advertised in its beacon Beacon frames.

b) A synchronizing mesh point MP that creates a mesh or joins a mesh where all the neighbors are unsynchronizing MPs shall set the values of ATIM window, beacon interval, mesh DTIM interval, and power management mode. These values are advertised in its beacon Beacon frames.

c) A synchronizing mesh point MP that joins a mesh where one or more neighbors are synchronizing MPs shall update its ATIM window and mesh DTIM interval according to the values of received beacons Beacon frames from the synchronizing MP and set beacon interval and power management mode. These values are advertised in its beacon Beacon frames.

d) The start of the ATIM window will be is measured from TBTT TBTT.

e) A power management capable unsynchronizing MP assigns AID to every peer MP through the peer link establishment procedure. AID 0 is reserved to indicate the presence of buffered group-addressed MSDUs. This AID will be used so that the MP identifies those peer MPs for which it is prepared to deliver buffered MSDUs by setting bits in the TIM's partial virtual bitmap that correspond to the appropriate AIDs.

11A.10.6 Mesh point MP power state transitions

A mesh point An MP may change its power management mode to the power save mode only if the following conditions are fulfilled:

a) The mesh point MP supports power save mode operation

b) All of the mesh points MPs that the mesh point MP is connected to (has peer relationships) are capable of transmitting traffic to mesh points MPs operating in power save mode

Note that a MAP never changes its power management mode to power save mode since it also provides the functionality of an AP. A mesh point An MP changing power management mode to power save mode will inform informs all it's its mesh neighbors of the change by sending a Null-Data frame to a broadcast address with the power management bit in its frame control header set. The packet will be frame is sent during the ATIM window of a mesh DTIM beacon and will be is repeated at least twice on two different mesh DTIM intervals. If a beacon is received with the "“neighbor power management mode" bit of the specific MP in the Mesh Neighbor list List not updated, the MP will continue continues to send the Null-Null Data packet frame on the next mesh DTIM. The mesh point will include MP includes a Mesh Power Save Capability field of the Mesh PS IE Capability element with a value of power save mode in its following beaconsBeacon frames.

A mesh point An MP changing power management mode to active mode will inform informs all its mesh neighbors of the change by sending a Null-Null Data frame to a Broadcast address with the power management bit in its frame control header cleared. The packet will be frame is sent during the ATIM window of a mesh DTIM beacon and will be is repeated twice on two consecutive mesh DTIM intervals. The mesh point will include MP includes a Mesh PS IE information element with a value of Active in its following beaconsBeacon frames. The mesh point will transfer MP transfers to the active state immediately with no relation to when the Beacons will be Beacon frames are sent.

A mesh point An MP operating in power save mode will set sets the power management bit in the frame control header of every outgoing frame. A mesh point An MP operating in active mode will clear sets the power management bit to 0 in the frame control header of every out going frame.

A mesh point An MP in power save mode will transition transitions between awake and doze states according to the following rules:

a) A mesh point An MP shall enter awake state prior to every TBTT that matches the mesh DTIM interval beacon timing of its own or of its neighbor MPs with which it maintains a peer relationship.

b) A mesh point An MP that entered the awake state due to the mesh DTIM TBTT event and had not sent an ATIM, a broadcast frame, or a multicast frame, and did not receive a directed unicast or multicast ATIM, a broadcast frame, or a multicast frame, ATIM may return to the Doze state following the end of the ATIM window

c) If a mesh point an MP received an ATIM frame, broadcast frame, or multicast frame, it may return to doze state after receiving a packet frame with the more bit in the control field cleared from all the sources that sent an ATIM, broadcast, or multicast frame during the ATIM window.

d) A mesh point An MP receiving a broadcast or multicast frame during the ATIM window with the more data bit of its control field cleared may return to the doze state either following the end of the ATIM window or after receiving a packet frame with the more bit in the control field cleared from all other active sources, whichever comes later.

e) A mesh point An MP that transitions to the awake state may transmit a beacon, but this would not prevent it from returning to doze state following the ATIM window

f) In addition a mesh point will enter an MP enters awake state prior to every agreed schedule as negotiated as part of a periodic APSD TSPEC exchange with other mesh pointsMPs

g) A mesh point An MP entering awake state may return to doze state after receiving and/or sending a directed unicast frame to/from the specific flow for which this schedule was set with EOSP bit set or with the expiration of the maximal service interval for that flow.

h) A mesh point An MP may transition to awake state if it has traffic to transmit at any given point of time

11A.10.7 Frame transmission

The following description applies to mesh points MPs that support power save mode operation. Mesh points MPs that do not support this capability do not have to track the power save mode of other mesh points neighbor MPs and will may only use the standard procedure for frame transmission.

A mesh point will An MP supporting PS shall store information on the power save mode of all its neighbors by monitoring their beacon or connectivity report Mesh PS IE information element and by extracting information from the neighbor list IE Mesh Neighbor List information element in beacons Beacon frames or connectivity reports of other MPs.

A mesh point An MP considers the mesh to operate in a power save mode if any one of its neighbors is operating in the power save mode. In a mesh that operates in the Active state, frames may be sent at any time to other mesh pointsMPs. For a mesh that operates in a power save scheme the following rules apply for transmission:

a) All broadcast and multicast traffic will shall be buffered by mesh points MPs that perceive the mesh to operate in power save mode. These packets will be frames are transmitted only on immediately after the mesh DTIM intervalsbeacon transmission.

b) All unicast individually addressed traffic targeted to mesh points MPs in power save mode will be is buffered. These packets will be frames are transmitted upon the reception of PS-poll frames or upon the positive acknowledgement reception to the frames transmitted during the ATIM window.

c) The only types of frames mesh points MPs may transmit during the ATIM window of synchronizing MPs are ACK, CTS, RTS, ATIM, Beacon, broadcast or multicast MPDU and Null Data frames. A An MP may transmit any type of frames during the ATIM window of unsynchronizing MPs.

d) A mesh point An MP may transmit one short broadcast or multicast MPDU during the ATIM window if the MAC frame length of the MPDU is less than dot11shortMulticastFrameLengthLimit. If the mesh point MP has more than one broadcast or multicast frame to transmit, it should set the more data bit of the broadcast or multicast frame transmitted during the ATIM window and contend for the channel following the end of the ATIM window to transmit the additional frames.

e) Mesh points MPs that transmit to mesh points MPs in power save mode (including broadcast and multicast) will shall set the More data bit in frame control headers to indicate if more frames are to be transmitted for the specific destination.

f) All other aspects of ATIM based transmission are as defined in 802.11 specification clause 11.2.2.4

g) For traffic that belongs to a flow for which an APSD TSPEC and schedule was setup with another mesh pointMP, the transmission will be is performed according to the agreed schedule.

11A.10.8 Power management operation with APSD

Power management operation with APSD in a mesh is very similar to the operation in a BSS.

The following are the modifications compared to the description provided in clause 11.2.1.4

The Mesh supports periodic and aperiodic APSD operation modes. The periodic APSD mode is similar to Scheduled APSD. The aperiodic APSD is similar to Unscheduled APSD. The periodic and aperiodic APSDs APSD modes use the same signaling as scheduled and unscheduled APSDsAPSD modes, but they can be used only with neighbors in a mesh that support it..

Use of aperiodic APSD in a mesh may be limited. It may only be used if one of the mesh points MPs in the link is Active and the other in Power Save and EDCA is used for transport of the data. In this case the mesh point The MP that operates in power save mode will be is the one to initiate the data exchange by sending packets frames to the other mesh point MP triggering it to send the traffic back.

Periodic APSD can be used for all cases (i.e., both mesh points MPs in power save or only one, EDCA or HCCA access).

The ADDTS request is modified to include a Schedule element that describes the requested schedule from the mesh pointMP. The ADDTS response will include includes the Schedule that can be supported by the other mesh point MP and that should be used for this flow. If this schedule is not acceptable to the originating mesh point MP it may reattempt the ADDTS request with modified schedule or tear down the flow. An unsynchronizing MP shall set the service start time field of the TSPEC element to the four lower octets of the TSF timer value of the MP which initiates the ADDTS request.

For periodic APSD both sides can initiate transactions as long as they are sent within the service interval defined. The service interval will last lasts up to the maximal service duration as defined in the schedule IE information element or if EOSP is declared in frames sent/received for that flow. For a unidirectional flow the originating mesh point MP sets the EOSP when it wants to end the service interval. The interval will be is considered terminated once the ACK is received for that packet frame (if ACK is required). For a bi-directional bidirectional flow the service period will end ends only after both ends of the flow send a packet frame with EOSP bit set and the matching ACK packets frames are received.

11A.10.9 Power Save parameters selection (Informative)

The power save operation of a mesh point is controlled by a set of global parameters. The following are the global mesh parameters with their default recommended values:

Beacon Period: 100TU

DTIM period: 10

ATIM Window: 10TU

Mesh points may wish to use other parameters but doing so may effect the power save efficiency and also delay the service initiation in the mesh.

EDITORIAL NOTE—Power save parameters selection moved to Annex T per CID 3444

11A.10.10 TS Reinstatement

A mesh point An MP wishing to reinstate a TS with another mesh point MP that is operating in Power save mode will send sends a QoS-Null frame to the mesh point MP in power save mode during its ATIM window.

11A.10.11 Beacon broadcaster power save mode

The beacon broadcaster can enter the power save mode where it sends only mesh DTIM beacons Beacon frames and stays awake for the next beacon period. In this case, MPs shall not send any frames to the beacon broadcaster during normal beacon periods.

If the beacon broadcaster is in power save mode it sets the BB power management mode bit of the MP control field in the Neighbor List element to 1.

11A.10.12 Naive Mesh operation (Informative)

This section describes the operation of a naive mesh that does not include any Mesh APs and is also not supporting any routing capabilities.

This type of mesh would be mainly useful in cases were all Mesh point are maintaining a neighbor relationship with each other, as such it does not need to follow the association procedures and may use the three address format to directly exchange frames between the mesh members.

For this specific case the Mesh point does not have to support any of the route messages and link state announcements defined in this specification.

The mesh point will include in its WLAN Mesh Capability IE a Peer Capability field with only bit 15 set. This would signal that the MP is not supporting association with any other MPs and does not support any 802.1X capabilities.

Since the mesh point is not going to receive any association request, and it has no need to initiate one itself, it does not have to support the association messages as well as any of congestion control messages that can be exchanged only between associated peers.

The power save operation is an optional feature and may be implemented by the MPs of the Naive mesh.

A naive mesh supporting power save may use the basic power save scheme for simple data exchange and may optionally extend to include the APSD support for real time power save stream delivery.

11A.11 Examples (Informative)

11A.11.1 Mesh Point Boot Sequence

At power up, a configured MP shall perform the following sequence of operations:

a) Passive or Active scanning to discover other MP

b) Channel selection

c) Begin mesh beaconing.

d) Neighbor MP link establishment[4]

   1) 802.11 open authentication

   2) Association

   3) 802.11i authentication and key exchange

e) Local link state measurement

f) Path selection initialization

g) AP initialization (optional – if MAP)

This sequence is illustrated in Figure s114.

Link establishment may be performed with a number of nodes but not all of the links will become active – that depends on the outcome of the link state measurements and on the routing initialization. The final (optional) step of the boot sequence is AP service initialization. This may be observed as a Mesh AP (MAP) starting to transmit a valid SSID information element in beacons. Before this point, that is, before the path selection specific state has been initialized, the MAP will not accept association requests or probe requests from simple STA devices. After this time, however, the MAP will continue to perform maintenance operations on the path selection state.

[4] Neighbor MP link establishment may be repeated multiple times if there are multiple neighbor MPs

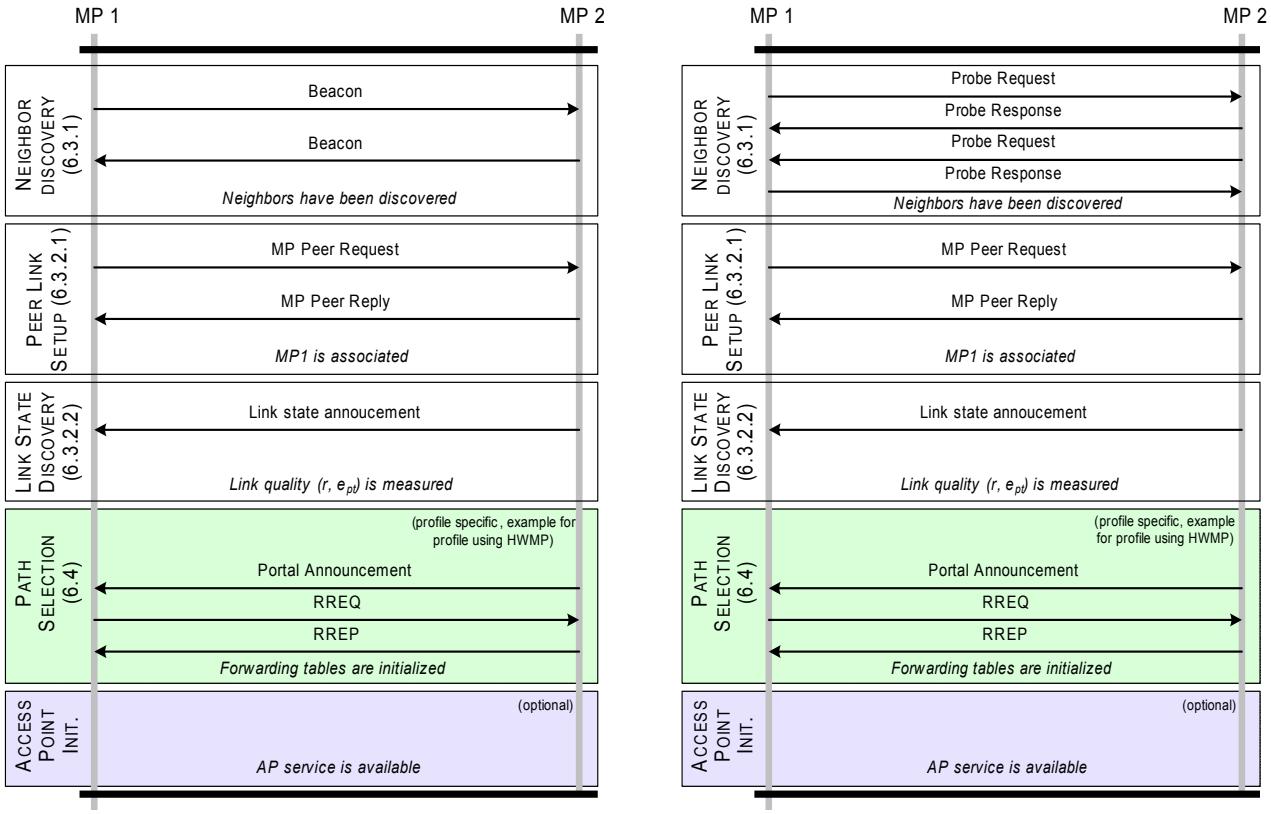

(a) Using Passive Scanning **(b) Using Active Scanning**

Figure s114—Mesh Point Boot Sequence

11A.11.2 Mesh Point Tables

11A.11.2.1 MP Neighbor Table

An MP shall maintain a table containing an entry for each discovered neighbor MP. By definition, all neighbor MPs have the same Mesh ID. Each entry shall contain the information shown in Table s38. The r and e_{pt} fields are specific to the default airtime metric defined in clause 11A.5. An implementation using a different path selection metric may require fields other than r and e_{pt}.

Table s38—MP Neighbor Table Entry

Value	Description
Neighbor MAC address	MAC address of the neighbor MP radio interface
Primary MAC address	Primary MAC address of the MP, if it has more than one radio interface
State	State of the association with the neighbor
Directionality	Directionality value in previous association request
c_o	Operating channel number
p_l	Channel precedence value
r	Reference bit rate (modulation mode)
e_{pt}	PER for the reference frame size at the reference bit rate
Q	Received signal strength or quality (internal units)

The state of the association with the neighbor shall take one of the values shown in Table s39, and shall be initialized on discovery to *neighbor* or *candidate peer* based on beacon or probe response contents as described in clause 7.2.3.1 and 7.2.3.9.

Table s39—State Values

State	Description
Neighbor	Discovered, no peer capability
Candidate peer	Has peer capability, no association established
Association pending	Association sent, reply not received
Subordinate, link down	Association established with this node as the subordinate, link not yet measured
Subordinate, link up	Association established with this node as the subordinate, link measured and active
Superordinate, link down	Association established with peer as the subordinate, link not yet measured
Superordinate, link up	Association established with peer as the subordinate, link measured and active

The neighbor MAC address is the address of the neighbor MP's radio interface that was discovered. The state information for the link to the MP stored in the MP neighbor table entry is with respect to this advertised address.

The primary MAC address of a neighbor MP is the primary unique address of the MP. In the case where the neighbor MP has only one radio interface, the primary MAC address will be equal to the radio interface MAC address. In the case where the neighbor MP has more than one radio interface, the primary MAC address is typically the radio interface MAC address with the smallest address value. (Note: more than one table entry may be created for a given neighbor primary MAC address).

The operating channel number is the channel on which the beacon was received from the node.

The channel precedence value is a number chosen by all MP radio interfaces in a given Mesh. It is contained in the beacon transmitted by the neighbor MP. It is used when merging disjoint networks and for the purpose of supporting DFS.

The bit rate and PER values are created by the local link state discovery procedures, described in Clause 11A.5.

The received signal strength or quality may represent any convenient quality measure; this value is never presented at an exposed interface, but rather is used for comparisons.

11A.11.2.2 MP Proxy Table

Each MP maintains a proxy table for the devices outside of the WLAN Mesh. The format of a logical proxy table is shown in Table s40.

Table s40—A logical proxy table maintained at each MP (the information can be derived from other sources).

Value	Description
MAC Address	MAC address of a given STA
inMesh	If the destination is within the mesh
isProxied	If the destination is proxied by MAP/MPP
Owner	MAC address of its proxy

An example for proxy registration procedure (Figure s115):

During the initialization phase, a STA first associates with MAP3 by using standard IEEE 802.11 procedures. Once associated, MAP3 may initiate a proxy registration procedure on behalf of STA towards the MPP. To do so, it sends a proxy registration request message on behalf of STA, to the mesh portal. MAP1 after receiving the registration request message updates its own proxy registration table for STA and forwards it towards mesh portal. The MPP thereby learns that STA is being proxied by MAP3. The MPP may then update its proxy registration table with new/updated entry for STA with owner set to MAP3 and inMesh, isProxied flags set to true. The MPP may then create a proxy registration reply message which is sent towards MAP3. Once MAP3 receives a registration confirmation for STA, proxy routing for STA is established. (Further optimization can be done using selective proxying).

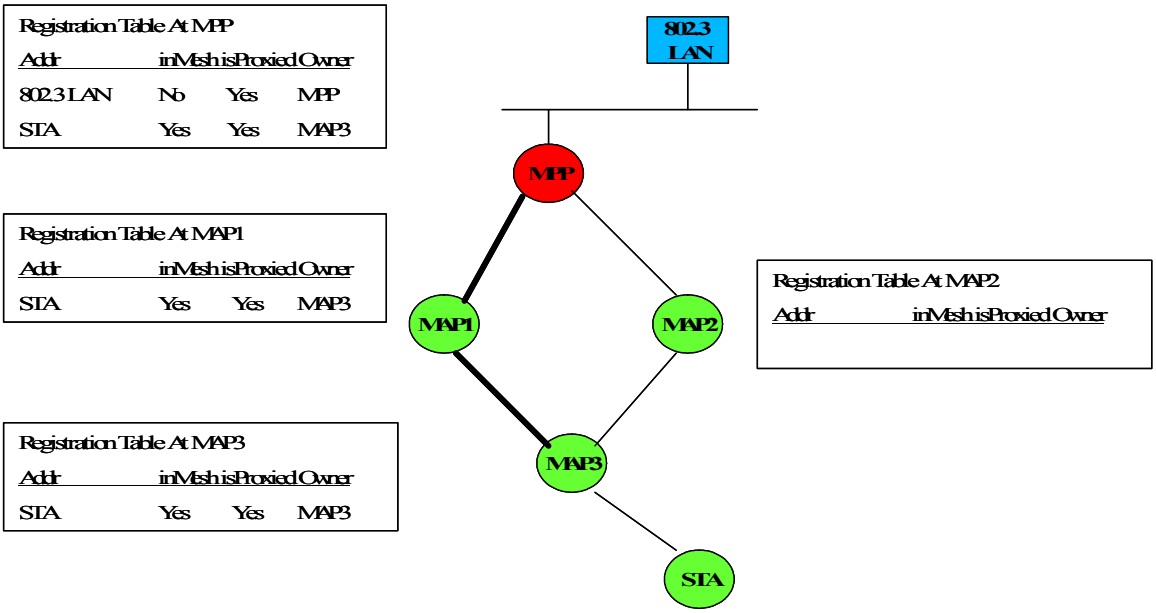

Figure s115—Example of optional proxy registration procedure to MPP.

EDITORIAL NOTE—Naive Mesh operation moved to Annex T per CID 3445

EDITORIAL NOTE—MP boot sequence examples moved to Annex T per CID 3447

EDITORIAL NOTE—MP table examples moved to Annex T per CID 3447

Annex D (normative) ASN.1 encoding of the MAC and PHY MIB

Insert the following at the end of Annex D:

```
************************************************************************

* dot11MeshPointConfig TABLE

************************************************************************

dot11BBConnectivityReportTimeout OBJECT-TYPE
              SYNTAX INTEGER (0..1000)
              MAX-ACCESS read-write
              STATUS current
            DESCRIPTION
              "This attribute shall specify the amount of MESH Mesh DTIM
              intervals, when no beacon or connectivity report indicating
              received beacon is received before the MP is removed from the
              neighbor list Mesh Neighbor List element in beacon or in
              connectivity report"
            ::= { dot11MeshPointConfigEntry  1 }

dot11BBBeaconRecoveryTimeOut OBJECT-TYPE
              SYNTAX INTEGER (0..1000)
              MAX-ACCESS read-write
              STATUS current
            DESCRIPTION
              "This attribute shall specify the amount of MESH Mesh DTIM
              intervals, when no beacon or connectivity report indicating
              received beacon is received before the terminal MP starts to
              transmit a beacon"
            ::= { dot11MeshPointConfigEntry  2 }

dot11BBBeaconRecoveryAddition OBJECT-TYPE
              SYNTAX INTEGER (0..100)
              MAX-ACCESS read-write
              STATUS current
            DESCRIPTION
              "This attribute shall specify extra MESH Mesh DTIM intervals that
              is used to for MPs that have not received connectivity reports from
              all other MPs. The MPs shall wait for  with beacon or connectivity
              report indicating received beacon is received before the terminal
              MP starts to transmit a beacon"
            ::= { dot11MeshPointConfigEntry  3 }

************************************************************************

* End of dot11MeshPointConfig TABLE

************************************************************************
```

Annex P WLAN Mesh Annex

EDITORIAL NOTE—Amendments prior to TGn define Annexes up to Q. TGn defines R and S. Next available is T.

Insert the following new Annex after Annex S:

Annex T Mesh Annex (Informative)

T.1 Example WLAN Mesh related entities with different levels of functionality

T.1.1 Overview

As described in clause 5.2.9.1 different WLAN Mesh related entities (MPs and entities that exhibit MP functionalitiesfunctions, such as MAPs) may operate with different levels of functionality. Table s20 describes applicability of different WLAN Mesh related functionalities functions to different entities. Note that this This table is not intended to classify mandatory vs. optional WLAN Mesh features; such classifications are made elsewhere in this specificationstandard.

Table s20—Comparison of different example mesh related entities according to an example set of functions

	MAP	MP	non-forwarding MP	Lightweight -MP
Uses 4 Address frame format	A	A	A	A
Association (providing the DS-service)	A	A	-	-
Association (using the DS-service)	A	A	A	-
To/From DS: 00	-	-	-	-
To/From DS: 01	U	-	-	-
To/From DS: 10	-	-	-	-
To/From DS: 11	U	U	U	U
Mesh Link security	A	A	A	A
MDA	A	A	A	A
DFS (5GHz Europe)	A	A	A	A
Discovery & peer link establishment	A	A	A	A

Table s20—Comparison of different example mesh related entities according to an example set of functions

Path selection	A	A	A	-
Forwarding	A	A	-	-
Interworking	A	A	A	-
Congestion Control	A	A	A	-
Beaconing	A	A	A	A
Synchronization	A	A	A	A
Power Management (provide, support neighbors in doing so)	A	A	A	A
Power Saving (actively going to doze/ sleep mode)	-	A	A	A

Legend:	”A”	= Applicable
	“-“	= Not applicable
	“U”	= Uses/Makes use of

Table s38—Comparison of different example WLAN Mesh related entities according to an example set of functionalities

	MAP	MP	non-forwarding MP	Lightweight-MP
Uses 4 Address frame format	A	A	A	A
Association (providing the DS-service)	A	A	-	-
Association (using the DS-service)	A	A	A	-
To/From DS: 00	-	-	-	-
To/From DS: 01	U	-	-	-
To/From DS: 10	-	-	-	-
To/From DS: 11	U	U	U	U
Mesh Link security	A	A	A	A
MDA	A	A	A	A
DFS (5GHz Europe)	A	A	A	A
Discovery & peer link establishment	A	A	A	A
Path selection	A	A	A	-
Forwarding	A	A	-	-
Interworking	A	A	A	-
Congestion Control	A	A	A	-
Beaconing	A	A	A	A
Synchronization	A	A	A	A
Power Managment (provide, support neighbors in doing so)	A	A	A	A
Power Saving (actively going to doze/ sleep mode)	-	A	A	A

Legend:"A" = Applicable
" - " = Not applicable
"U" = Uses/Makes use of

T.1.2 Lightweight mesh point operation

Lightweight mesh points (LWMPs) are minimal functionality mesh pointsMPs. They support a subset of mesh point MP functionality and are able to communicate only with their neighbors. Such MPs can have extremely lightweight implementation. This functionality can be achieved as a subset of full functionality, e.g. by having such LWMPs adopt "Null" routing profiles. This indicates to neighbors that these MPs are unable to provide functions such as routing. The choice of not using all MP functions does not require any modification to mesh services specification. Table s20 summarizes applicability of various WLAN Mesh related functions to lightweight mesh pointsMPs.

T.2 Radio Metric AODV Example and FlowCharts

T.2.1 An Example

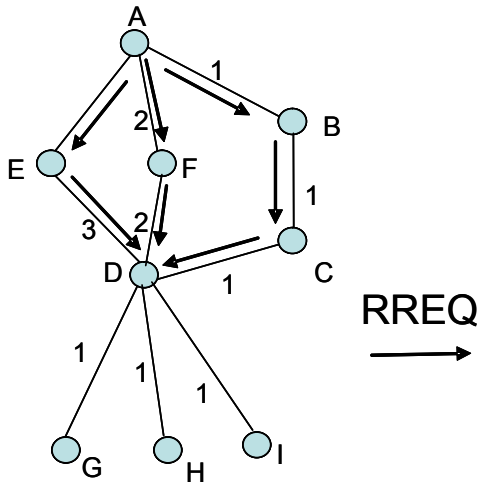

Figure s112—Example network.

Consider an example network shown in Figure s112. In the figure, lines between nodes indicate that they are neighbors (i.e) they are in radio-range with each other. The numbers beside the lines indicates the metric for the corresponding link. In this example, the link qualities are assumed to be symmetric. Radio Metric AODV should determine routes between nodes that will minimize the end-to-end metric.

Assume that node A attempts to find a route to destination node D and sends a RREQ (OSN=2, RREQID=2, Dest=D, DSN= 0, DO for D= 1). Assuming there is no packet loss, node D would get 3 copies of this RREQ each traversing a different path: A-E-D, A-F-D and A-B-C-D. Let us assume that the three RREQs reached D in the following order: A-E-D, A-F-D, A-B-C-D. On receiving the RREQ through A-E-D, node D first creates a route to node A through E. At this point, the upstream route to the source (node A) would have been established in node E and D. Then, node D sends a RREP (Dest=D, DSN=2) along the route D-E-A. The RREP establishes the downstream route to node D on nodes E and A. Table s39, Table s40, and Table s41 show the routing tables in nodes A, E and D respectively at this point.

When node D receives the RREQ that came along A-F-D, it determines that this RREQ came along a path with a better metric to A than the current route (4 vs. 5). Therefore, node D modifies the nexthop from E to F and the metric from 5 to 4. D sends a unicast RREP (D, 3) back to A via F. Similarly, when D receives the RREQ that came along A-B-C-D, it modifies the nexthop to A from F to C, as this RREQ came along a better route than the current route. A unicast RREP is sent via node C. The RREP establishes the route to D in the intermediate nodes C and B as well as the source node A. Table s42, Table s43, Table s44 and Table s45 show the routing tables in nodes A, B, C and D respectively at this point.

Assume that intermediate node E already has a valid route E-D to the destination node D with metric 3. Furthermore, the DO flag for D = 0 in the RREQ. When intermediate node E receives the RREQ, it creates a reverse route to the source node from which it receives the RREQ as the next hop (source node A) of the reverse route. Intermediate node E responds to the RREQ with a unicast RREP because it has a valid route to the destination D and the DO flag for D = 0. The RREP establishes a forward route to destination node A in source node A. As soon as source node A creates the route to destination node D with the RREP from intermediate node E, source node A can start sending data frames to destination node D via route A-E-D. If the RF flag for D = 1, the intermediate node E sets the DO flag for D = 1 in the RREQ message and forwards it further. Destination D will process this RREQ the same way as described above. The generated RREP will

refresh the path metric of the downstream route to D on nodes E and A. Assume that intermediate nodes F, B and C do not have valid routes to the destination node D. When intermediate nodes F, B and C receive the flooded RREQ messages, they create the reverse route to the source node A and forward the RREQ messages further the same way as described above.

In order to maintain an optimal route between A and D, A sends maintenance RREQs. The DO flag is set to 1 in maintenance RREQs. Intermediate nodes do not reply to a maintenance RREQ with a RREP. They forward the RREQ message. Assume that A sends a maintenance RREQ (3, 3, D, 2) and that the RREQ reaches D in the order A-E-D, A-F-D, and A-B-C-D. Since D has a better route to A through C and since A's sequence number in the RREQ (3) is not newer than what is in D's route table by at least HWMP_RREQ_LOSS_THRESHOLD, node D will not modify its current route to A. Instead, it stores this route (through E) as an alternate route to A, and starts a *rreq_wait_alarm* for this RREQ that would go off in current_time+HWMP_NETDIAMETER_TRAVERSAL_TIME. It is assumed that HWMP_NETDIAMETER_TRAVERSAL_TIME is configured large enough for a RREQ to traverse one-way along the diameter of the network. This would guarantee that node D receives all copies of a RREQ originating from A within that time. In particular, node D would receive the RREQs that traversed the paths A-F-D and A-B-C-D before the timer goes off. When node D receives the RREQ that came along A-F-D, it does not modify the current route as before. But since this RREQ offers a better route to A than the current alternate route, it changes the alternate route to go through F. When the RREQ arrives via A-B-C-D, node D keeps the current route, since it is still the optimal one, and changes the sequence number on the route to the one in the RREQ. Node D also cancels the rreq_wait_alarm for this RREQ and sends a unicast RREP back to A using the current route via C. The rest of this section considers and illustrates a few hypothetical situations that may happen in the network from this point onwards.

Example Scenario1

Consider a scenario when the metric along the path A-B-C-D became worse. In particular, say the metric along the link A-B changed from 1 to 4. When A sends the next maintenance RREQ (4, 4, D, 3), the processing of the RREQs that come along A-E-D and A-F-D in node D will be similar to the processing of the previous maintenance RREQ (3, 3, D, 2). But when D receives the RREQ that traversed A-B-C-D, it finds that the current route's metric has deteriorated (from 3 to 6) and modifies the route to use the alternate route (via nexthop F) since its offered metric is better than the current route (4 vs. 6). As done before, when the route is modified, node D cancels the rreq_wait_alarm for this RREQ, forwards the RREQ and sends a unicast RREP (D, 4) through the current route to A (via nexthop F).

Example Scenario2

Assume that A sends the maintenance RREQ (4, 4, D, 3) and that the RREQ reaches node D along the paths A-E-D and A-F-D, but never along the path A-B-C-D (gets lost). In this case, processing the received RREQs will be similar to the processing of the previous maintenance RREQ along the same paths. But since the RREQ never comes along the current route, the rreq_wait_alarm does not get cancelled. When the timer expires, node D sends a unicast RREP (D, 4) along the current route.

Example Scenario3

Assume that HWMP_RREQ_LOSS_THRESHOLD=2. Let us further assume that the maintenance RREQs from A with OSN 3 and 4 got lost somewhere along the path A-B-C-D. Now when node D receives the RREQ (5, 5, D, 4) along the path A-E-D, it will modify its route to A to go through nexthop E even though the metric on this route is worse than the current route through C. This is because, the sequence number difference between what is in the received RREQ and the current route entry is greater than HWMP_RREQ_LOSS_THRESHOLD, which indicates that too many RREQs were lost along the current route.

Example Scenario4

Assume that node A generated 3 RREQs, one each for destinations G, H, and I. Also, let us say that A incremented its sequence number every time it sent a RREQ. So, the RREQs for G, H, and I would be RREQ4 (4, 4, G, 0), RREQ5 (5, 5, G, 0), and RREQ6 (6, 6, I, 0) respectively. Assume that all the three RREQs reach D along the path A-E-D even before RREQ4 reaches D along the path A-B-C-D. In this case, D would change its route to A to go through nexthop E after processing RREQ6, since the A's sequence in the RREQ indicates that it has not received more than HWMP_RREQ_LOSS_THRESHOLD RREQs along its current route to A. But shortly after processing this, all the three RREQs (4, 5, and 6) arrive at D along A-B-C-D. Now D would change its route to A back to its original route through nexthop C. This route flapping can happen every time a burst of RREQs originate from a node with unique and increasing originating sequence numbers. Radio Metric AODV avoids this situation by imposing that nodes not increment their current sequence numbers if the time at which they sent the first RREQ with the current sequence number is less than HWMP_NETDIAMTER_TRAVERSAL_TIME than the current_time. In the current example, node A would have sent the burst of RREQs with the same sequence number, but with different RREQ IDs. This would prevent node D from switching its route to the inferior route through E even if the burst of RREQs arrive at D along A-E-D.

EDITORIAL NOTE—Recommendationf for use of EDCA in MPs moved from 9.9.1.7 to Annex T per CID 3100

T.3 Recommendations for use of EDCA in MPs

T.3.1 General

EDCA is used as the basis for the Mesh Media Access Mechanism. A set of recommendations on how to optimize EDCA for use by MPs without changing the basic Medium Access Mechanism is presented here.

T.3.2 Forwarding and BSS traffic interaction

Since an MP is a logical entity, it is possible that an MP may be physically collocated with an Access Point. It is also possible to have an MP implemented on a device that also acts as an Application End Point, that is, in addition to participating in the mesh and forwarding frames on behalf of other MPs, it also generate its own application traffic. In both of these cases, one single device has to forward a mixture of mesh traffic (with 4-address frame formats) and BSS traffic (with 3-address frame formats). How these two different kinds of traffic are handled within a single device can have a profound impact on overall network performance. For example, the forwarding traffic tends to traverse the network through multi-hop paths and hence may have already consumed significant amount of network resources before reaching a certain MP. Dropping such traffic basically renders all resource previously used to forward the traffic as being wasted. The frames originated from a local BSS and destined to the mesh have only traversed one hop, and so dropping such traffic only has local impact. It is also possible that an aggressive STA with heavy traffic backlog in the BSS can potentially starve the neighboring MPs in the network. Such traffic prioritization may have different implication from the point of view of fairness and the prioritization policy may very well depend on the mesh network deployment scenario and the business model used. There are many implementation choices as how to best support such traffic prioritization within a single device like MAP. For example, one may choose to employ multiple PHYs to separate the BSS traffic and mesh forwarding traffic into different PHYs operating at different channels. Such choice is entirely an implementation matter, outside the scope of this standard, but it is highly recommended that consideration is taken into account to separate BSS traffic and mesh forwarding traffic as much as possible and regulate the interaction between the traffic when complete separation is not possible. For example, it is recommended in 11A.8 that a BSS traffic rate control mechanism be used in conjunction with intra-mesh congestion control to ensure the overall network performance.

Table s39—Routing Table in node A

Dest.	DSN	Nexthop	Metric	Alternate Route
E	0	E	2	Invalid
D	2	E	5	Invalid

Table s40—Routing Table in node E

Dest.	DSN	Nexthop	Metric	Alternate Route
A	2	A	2	Invalid
D	2	D	3	Invalid

Table s41—Routing: Table in node D

Dest.	DSN	Nexthop	Metric	Alternate Route
E	0	E	2	Invalid
A	2	E	5	Invalid

Table s42—Routing Table in node A

Dest.	DSN	Nexthop	Metric	Alternate Route
B	0	B	1	Invalid
D	2	B	3	Invalid
E	0	E	2	Invalid
F	0	F	2	Invalid

Table s43—Routing: Table in node B

Dest.	DSN	Nexthop	Metric	Alternate Route
A	2	A	1	Invalid
C	0	C	1	Invalid
D	2	C	2	Invalid

Table s44—Routing: Table in node C

Dest.	DSN	Nexthop	Metric	Alternate Route
A	2	B	2	Invalid
B	0	B	1	Invalid
D	2	D	1	Invalid

Table s45—Routing: Table in node D

Dest.	DSN	Nexthop	Metric	Alternate Route
C	0	C	1	Invalid
A	2	C	3	Invalid
E	0	E	2	Invalid
F	0	F	2	Invalid

Table s46—Routing: Table in node D.(case 2)

Dest.	DSN	Nexthop	Metric	Alternate Route
C	0	C	1	Invalid
A	3	C	3	DSN=3 Nexthop=F Metric=4

Example Scenario 5

Station management and multiple interface operation.

Consider an example network shown in Figure s113 where MAP with multiple radio interfaces has stations. Assume the station X transmits packets to station Z. MAP A transmits RREQ frame on behalf of station X, because station X does not have routing functionality. In the same way, MAP G transmits RREP frame on behalf of station Z because station Z cannot transmits RREP frame. Routing table in MAP A is shown in Table 27. Destination Sequence number is managed in MAP G. If station X moves from MAP A to MAP C, MAP A receives the disassociation frame or detects station X association timeout, then MAP A sends RERR message to MAP B in precursor list.

Example Scenario 6

Assume that station Y transmits data frame to station Z when MAP A has a path from X to Z. There are two options in this protocol. One is that MAP A forwards the packets from Y by using the path from X to Z. The routing table 27 is used in this case. Another is that MAP A searches another path by transmitting RREQ frame (Elements #5). In this case, routing table has source field entry as shown in Table 28. The

second case option is useful when the destination is a device that has heavy traffic with many stations such as GW. Each MAP can use the different interface to transmit data frame to the same destination if the source of packet is different, which takes an advantage to improve the load balance in each radio.

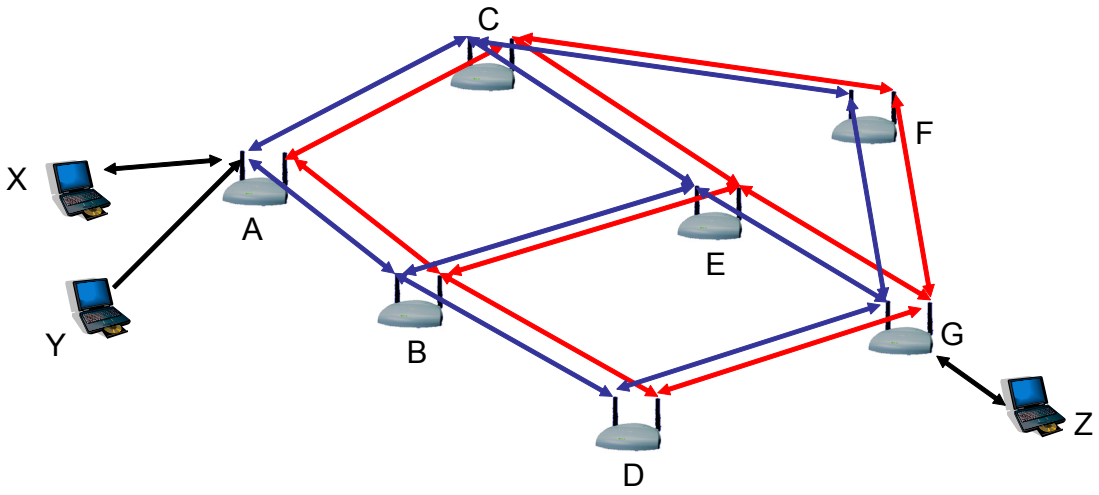

Figure s113—An example of network with multiple interface MAP and stations.

Table s47—Routing Table in MAP A

Dest.	DSN	Nexthop	Metric	Alternate Route
Z	0	B#1	1	Invalid
B	0	B#1	0.33	Invalid
C	0	C#1	0.33	Invalid

Table s48—Routing Table in MAP A (source-destination pair)

Dest.	DSN	Source	Nexthop	Metric	Alternate Route
Z	0	Y	B#1	1	Invalid
Z	0	X	B#1	1	Invalid

T.3.3 Radio Metric AODV Algorithm Flowchart

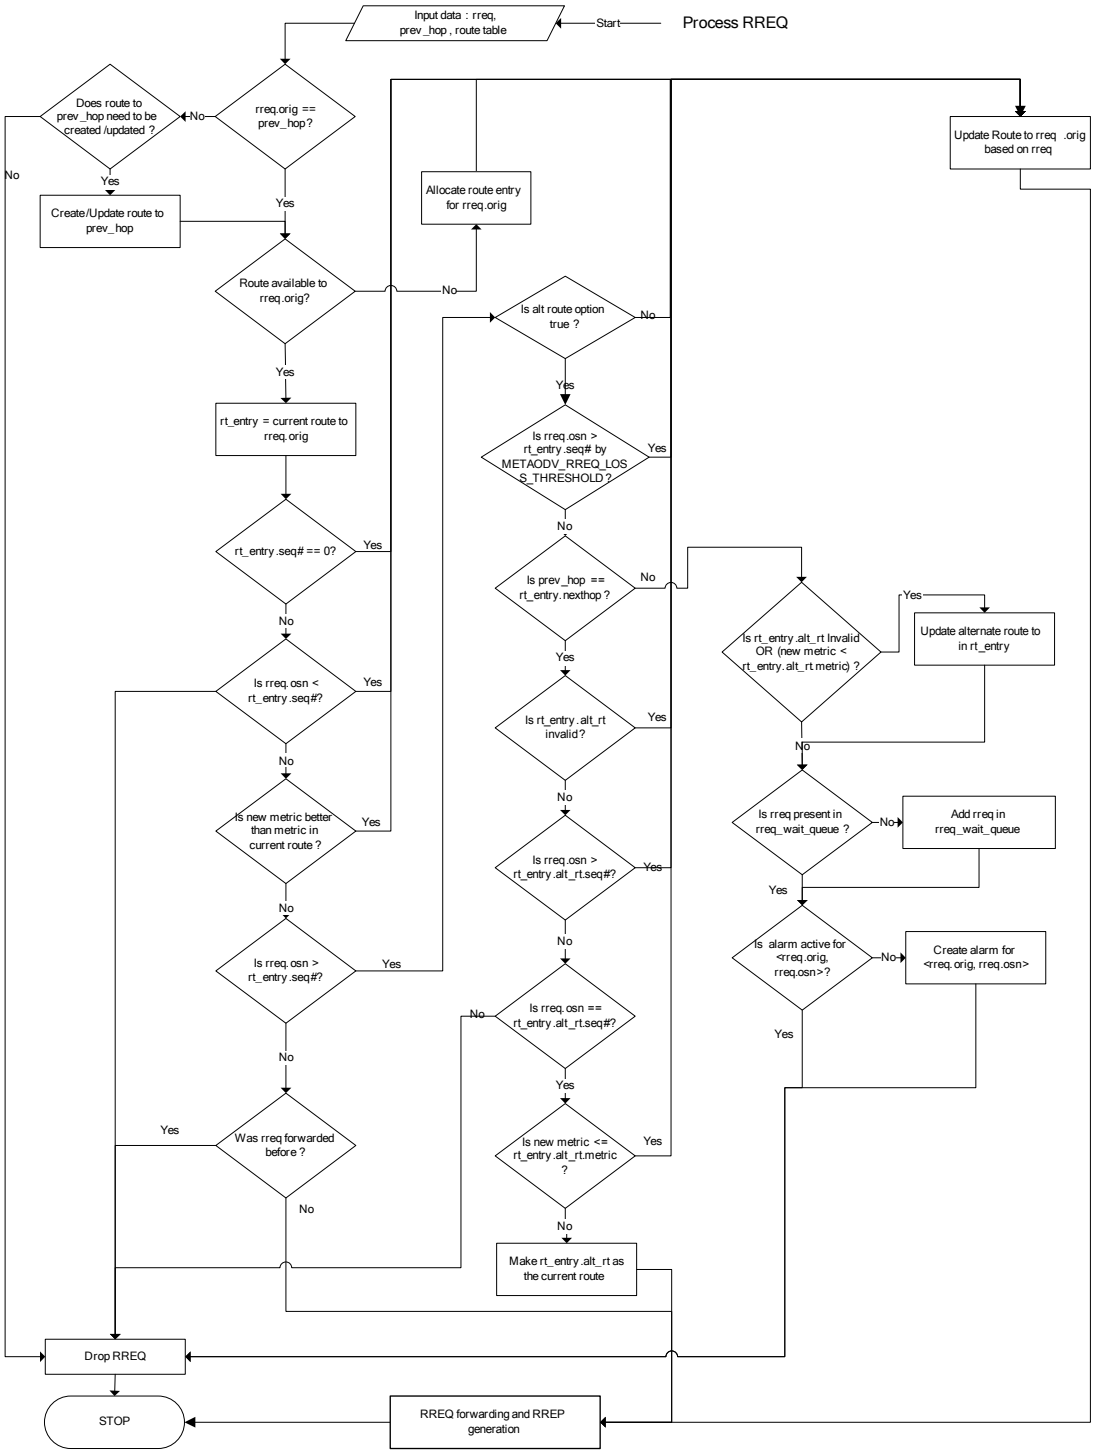

Figure s114—Flowchart for processing a RREQ

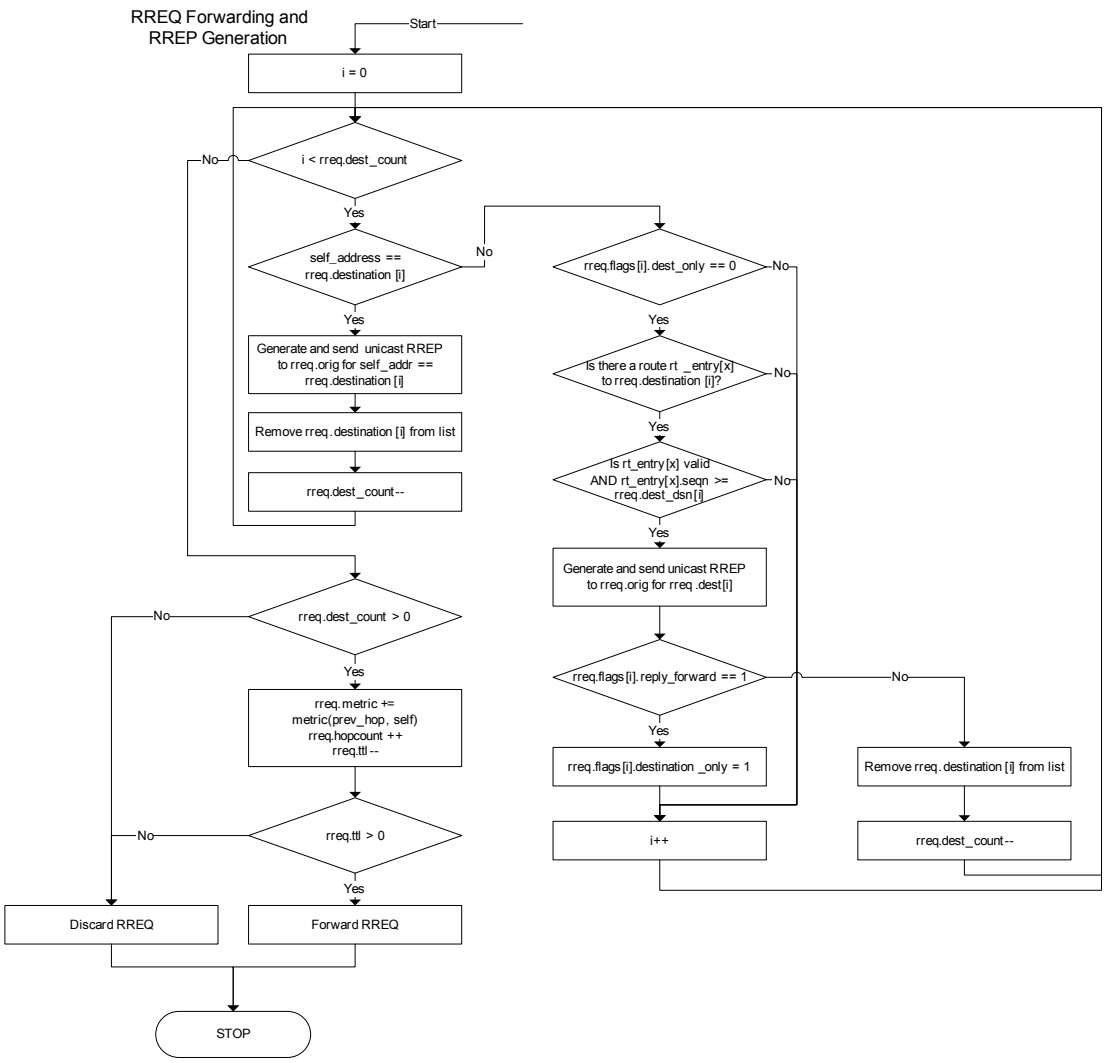

Figure s115—Flowchart for RREQ forwarding and RREP generation

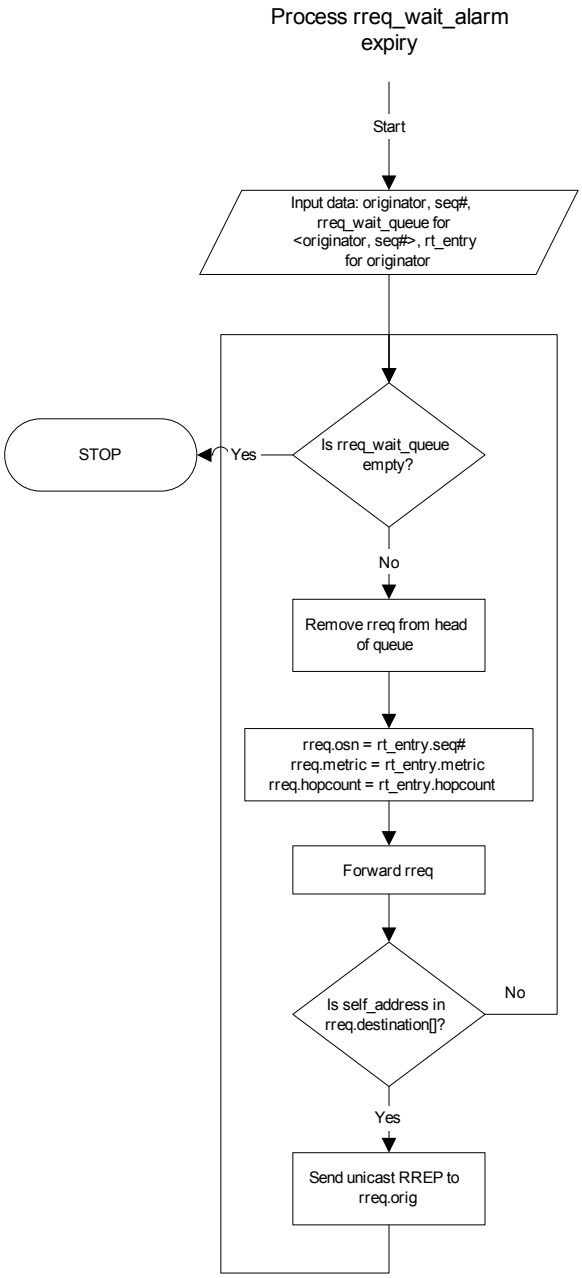

Figure s116—Flowchart for processing expiry of a RREQ wait alarm

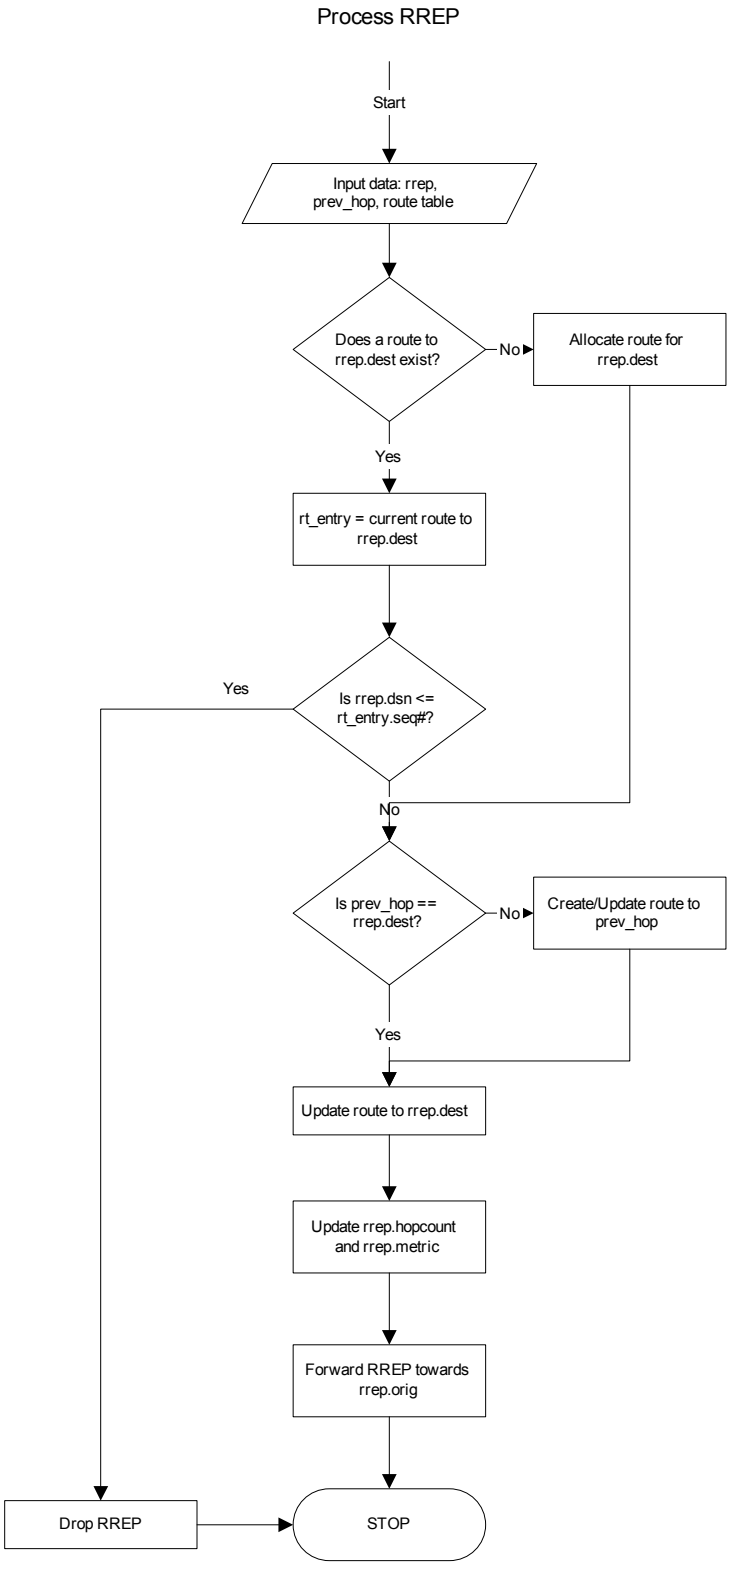

Figure s117—Flowchart for processing a RREP

EDITORIAL NOTE—Radio Metric AODV example and flowcharts removed per CID 4454

T.4 Recommended HWMP Default Values

```
HWMP_RREQ_REFRESH_PERIOD              = 15 seconds
HWMP_ROUTE_LOSS_THRESHOLD             = 2
HWMP_ACTIVE_ROUTE_TIMEOUT             = 5000 milliseconds
HWMP_RREQ_RATELIMIT                   = 2
HWMP_NET_DIAMETER                     = 20
HWMP_NODE_TRAVERSAL_TIME HWMP_MP_TRAVERSAL_TIME             = 40
milliseconds
HWMP_NETDIAMTER_TRAVERSAL_TIME        =

(HWMP_NET_DIAMETER*HWMP_NODE_TRAVERSAL_TIMEHWMP_MP_TRAVERSAL_TIME)
HWMP_RT_NETDIAMETER_TRAVERSAL_TIME    =
(2*HWMP_NETDIAMTER_TRAVERSAL_TIME)
HWMP_MAX_RREQ_RETRIES                 = 3
```

T.5 Radio Aware OLSR Flowchartsflowcharts

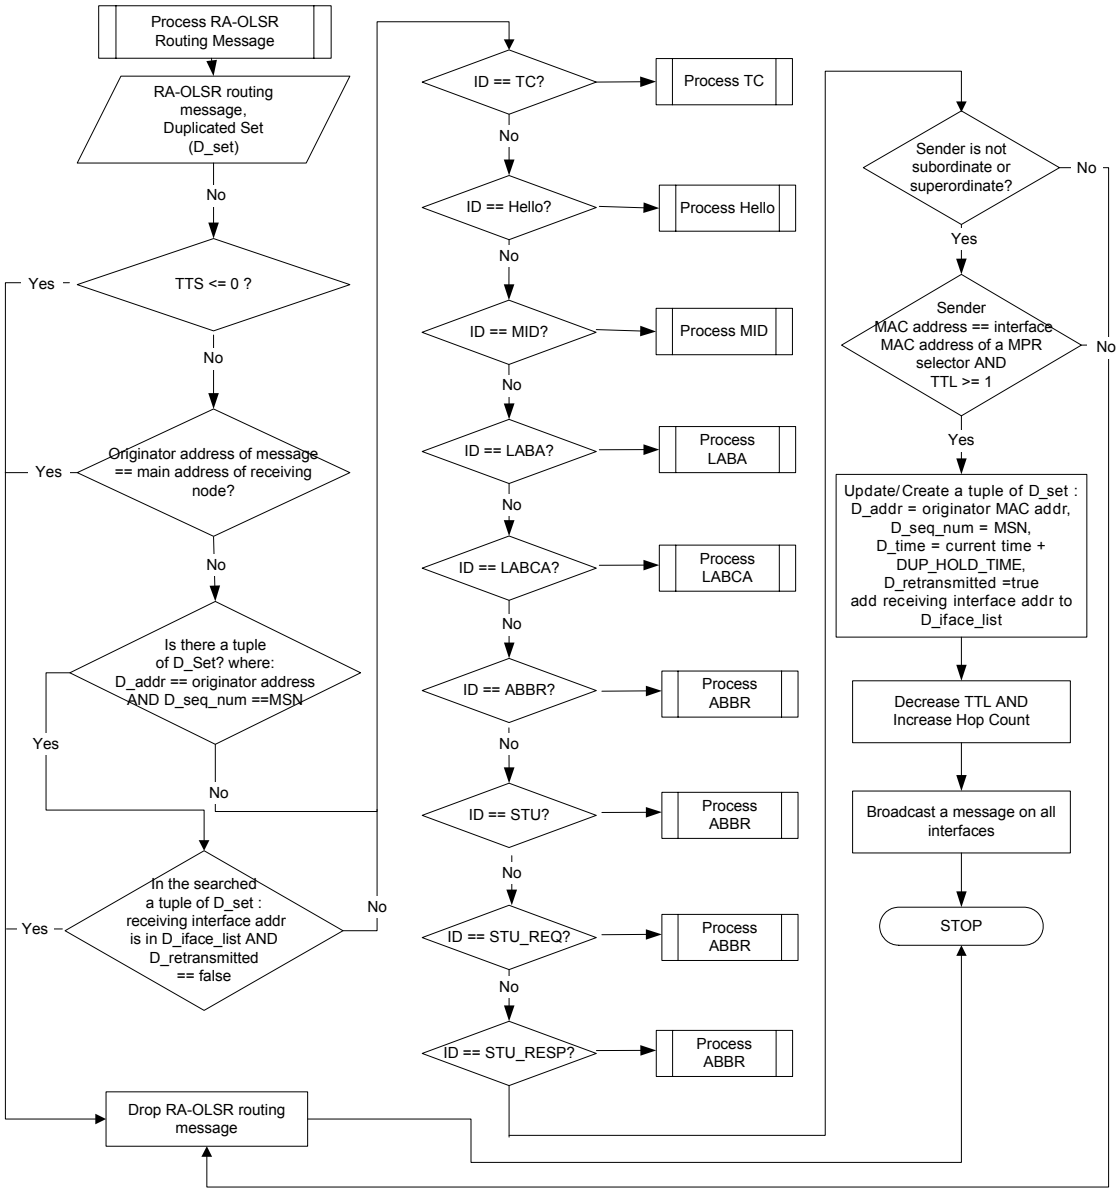

Figure s108—Flowchart for processing an RA-OLSR routing message.

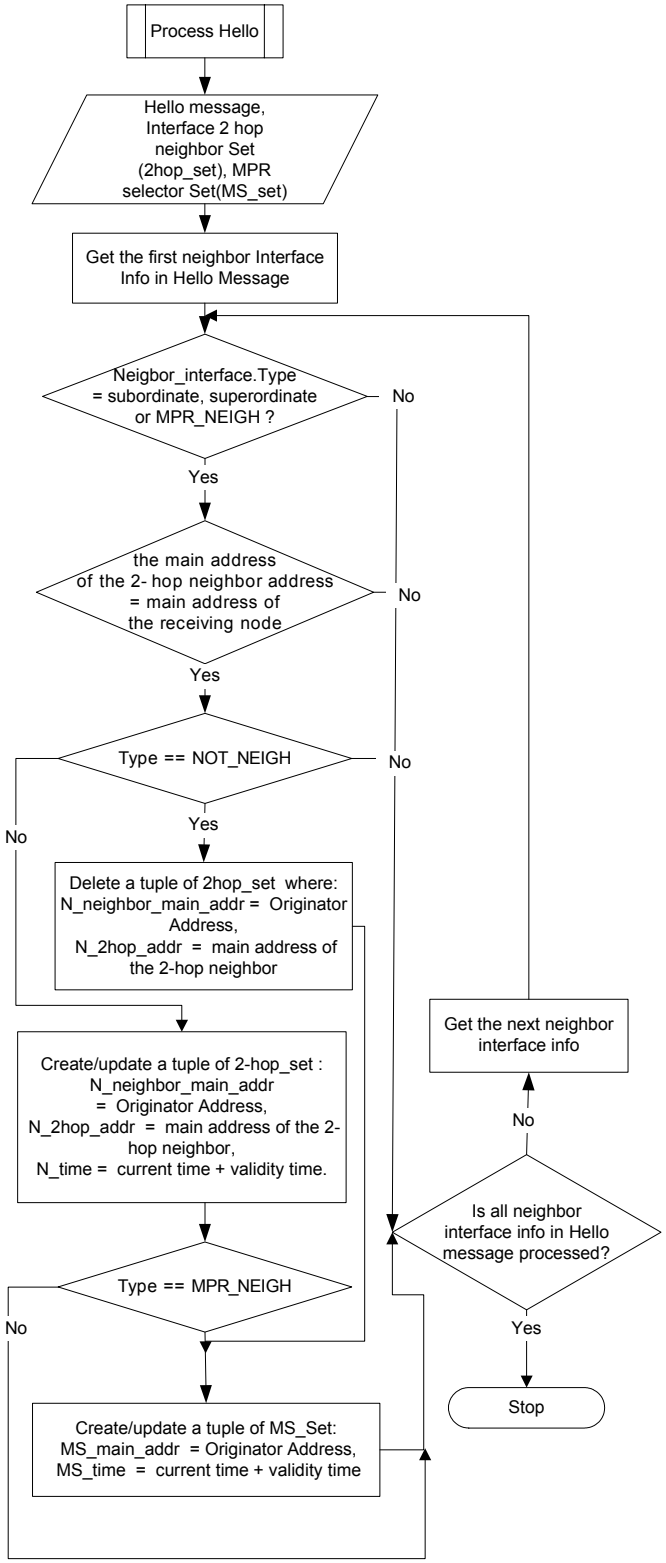

Figure s109—Flowchart for processing a HELLO message.

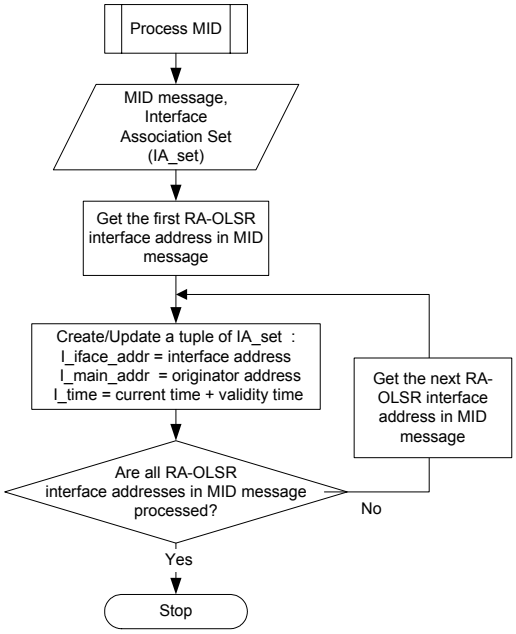

Figure s110—Flowcharts for processing an MID message.

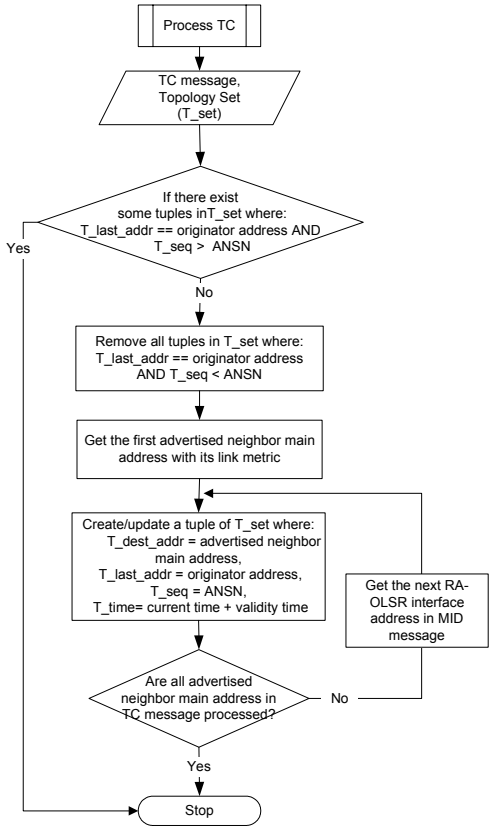

Figure s111—Flowcharts for processing a TC message.

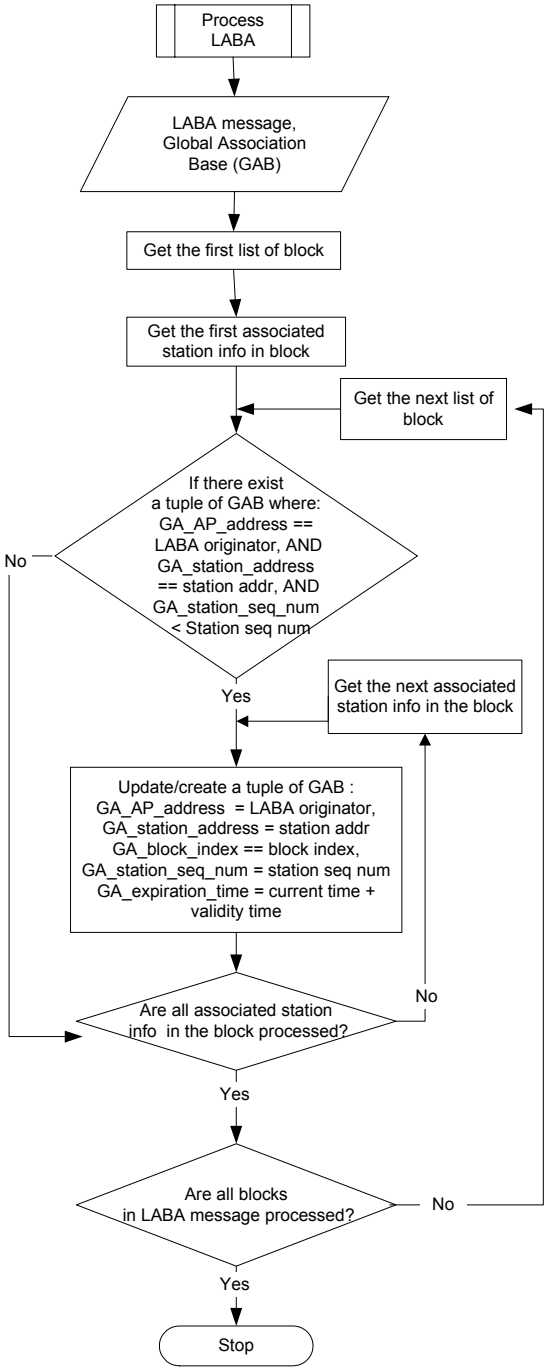

Figure s112—Flowchart for processing LABA message.

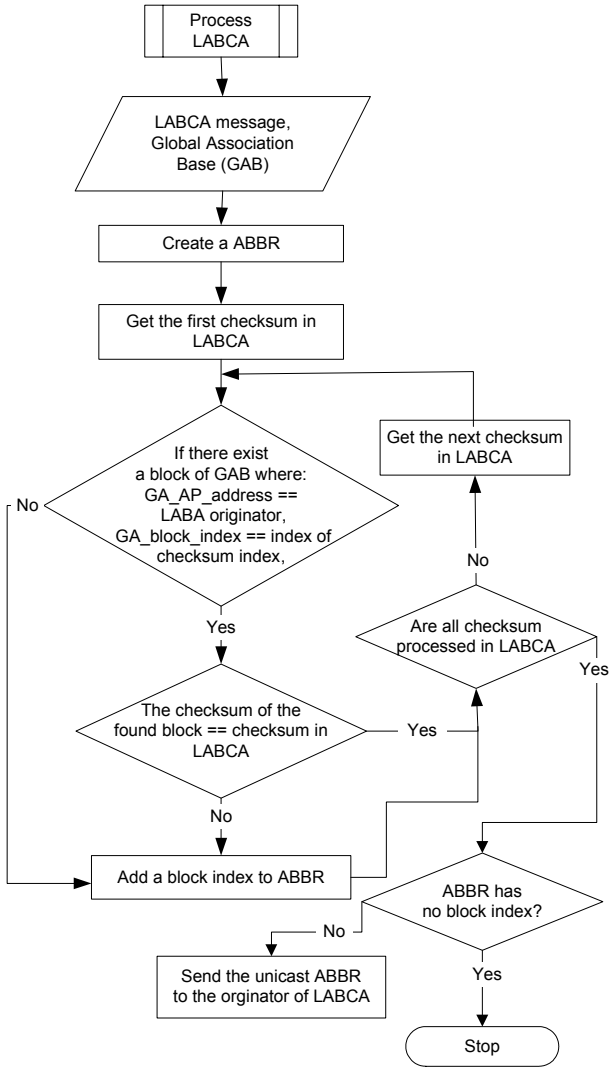

Figure s113—Flowcharts for processing an LABCA message.

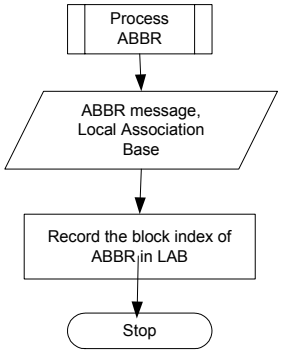

Figure s114—Flowcharts for processing an ABBR message.

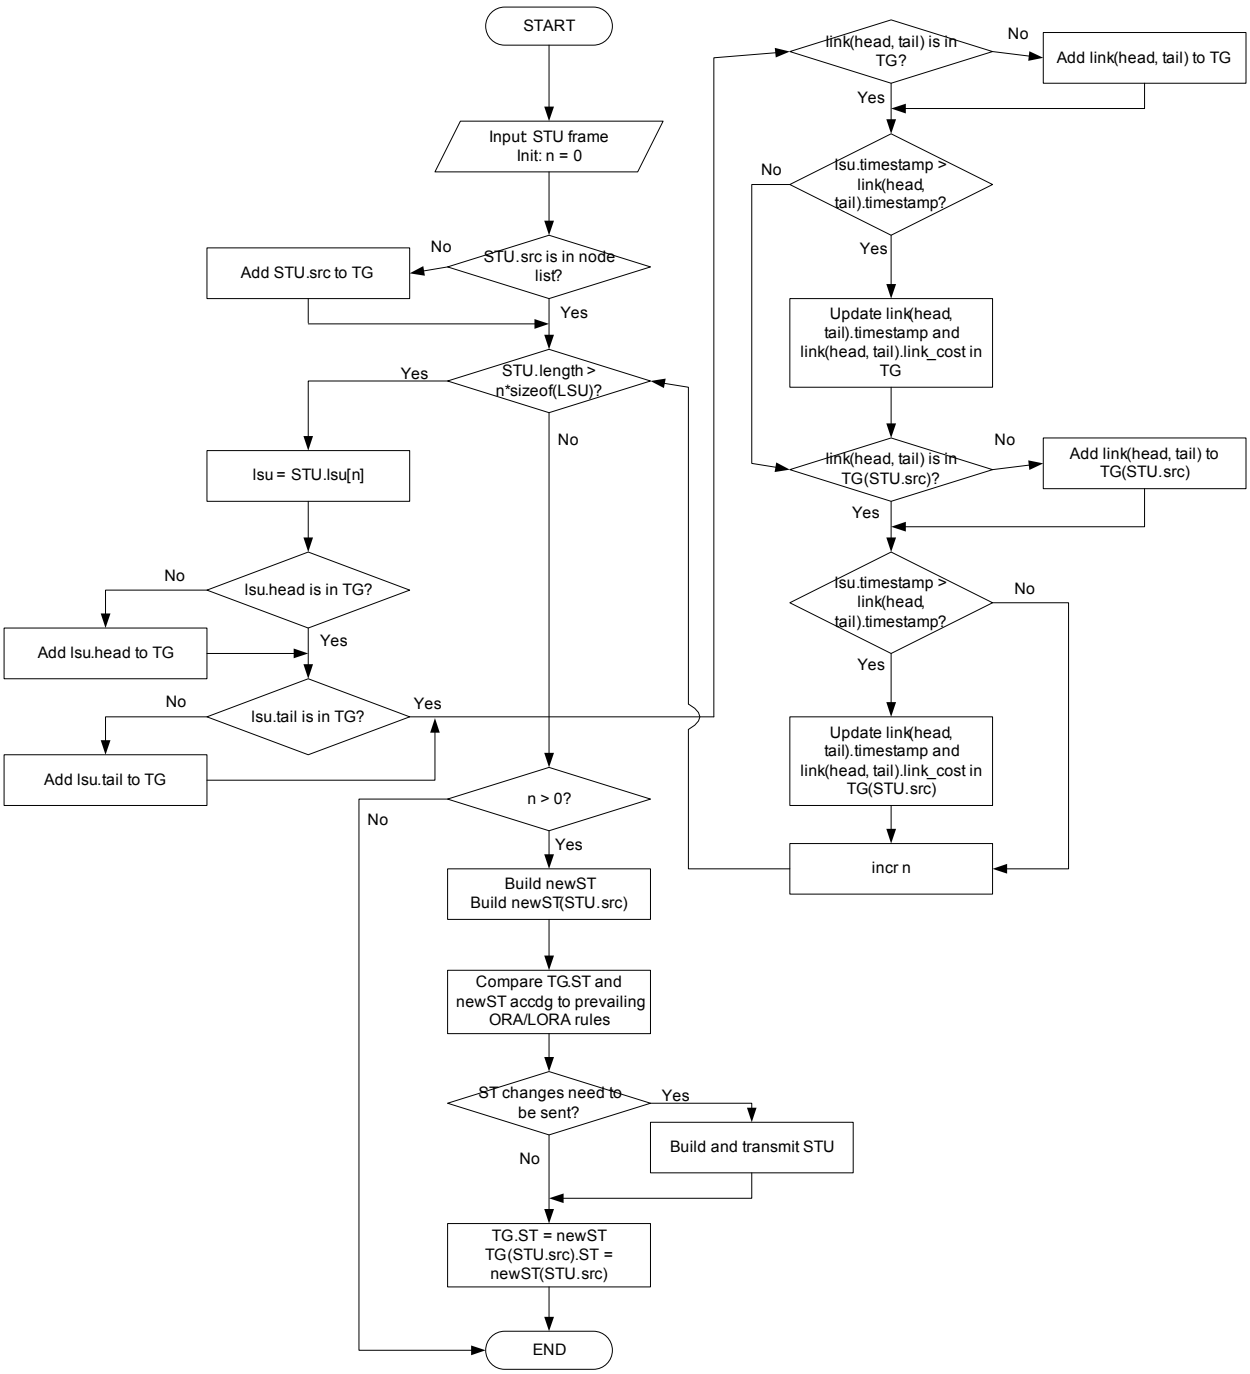

Figure s115—Flowchart for processing an optional STU message.

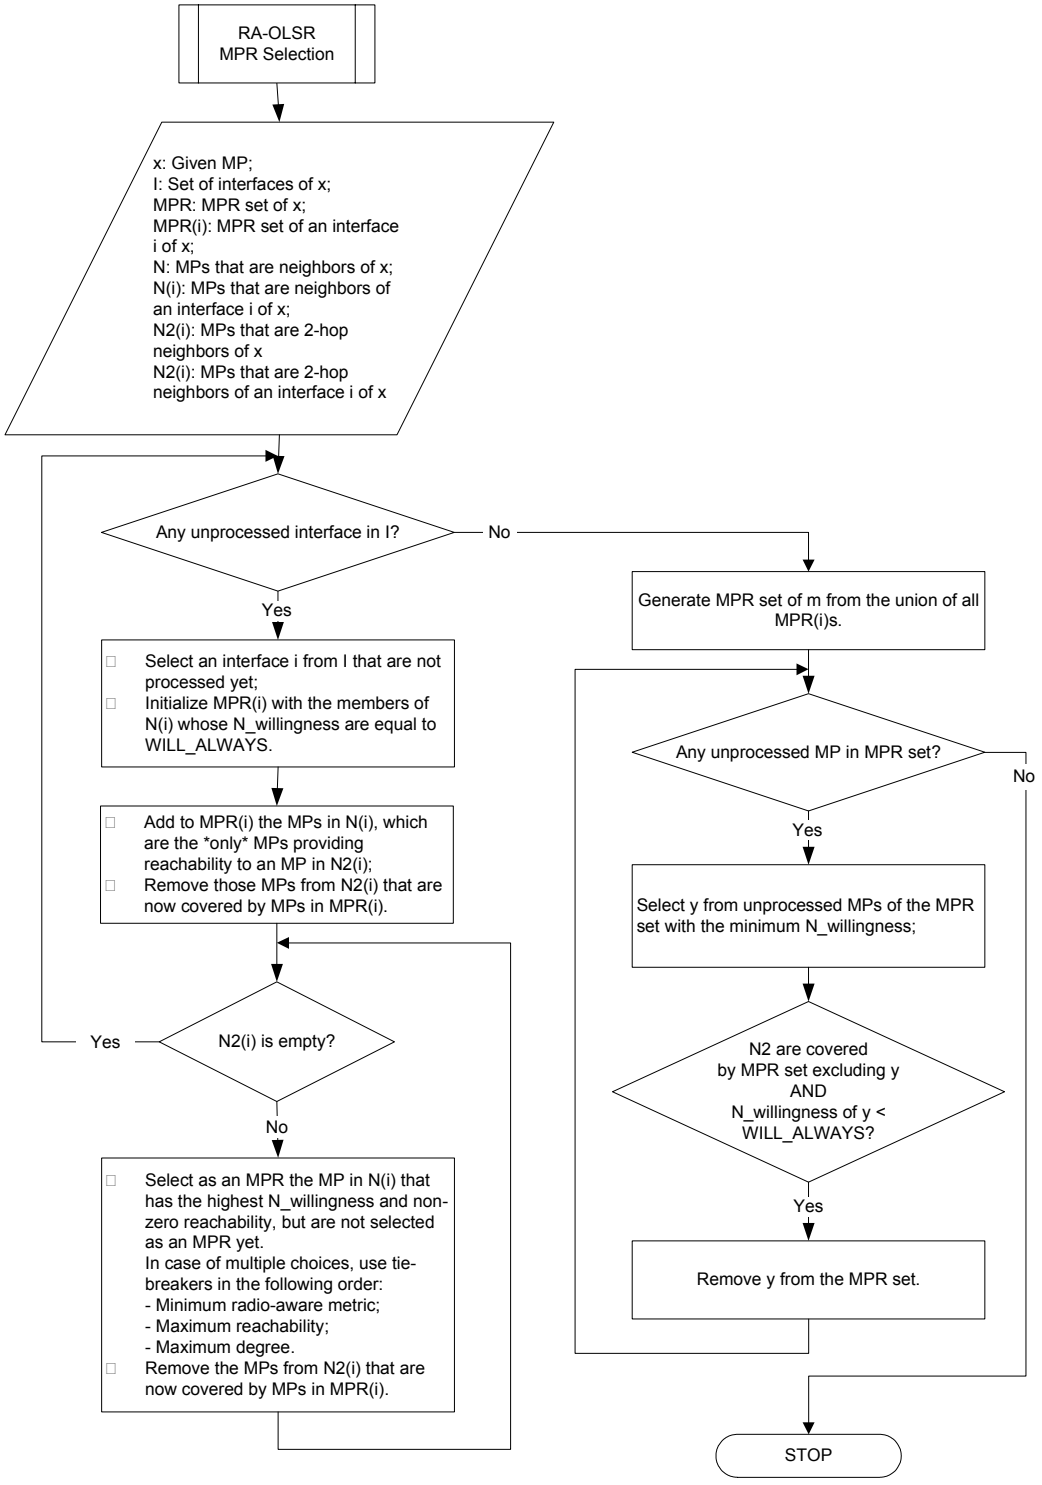

Figure s116—Flowchart for selection of MPRs.

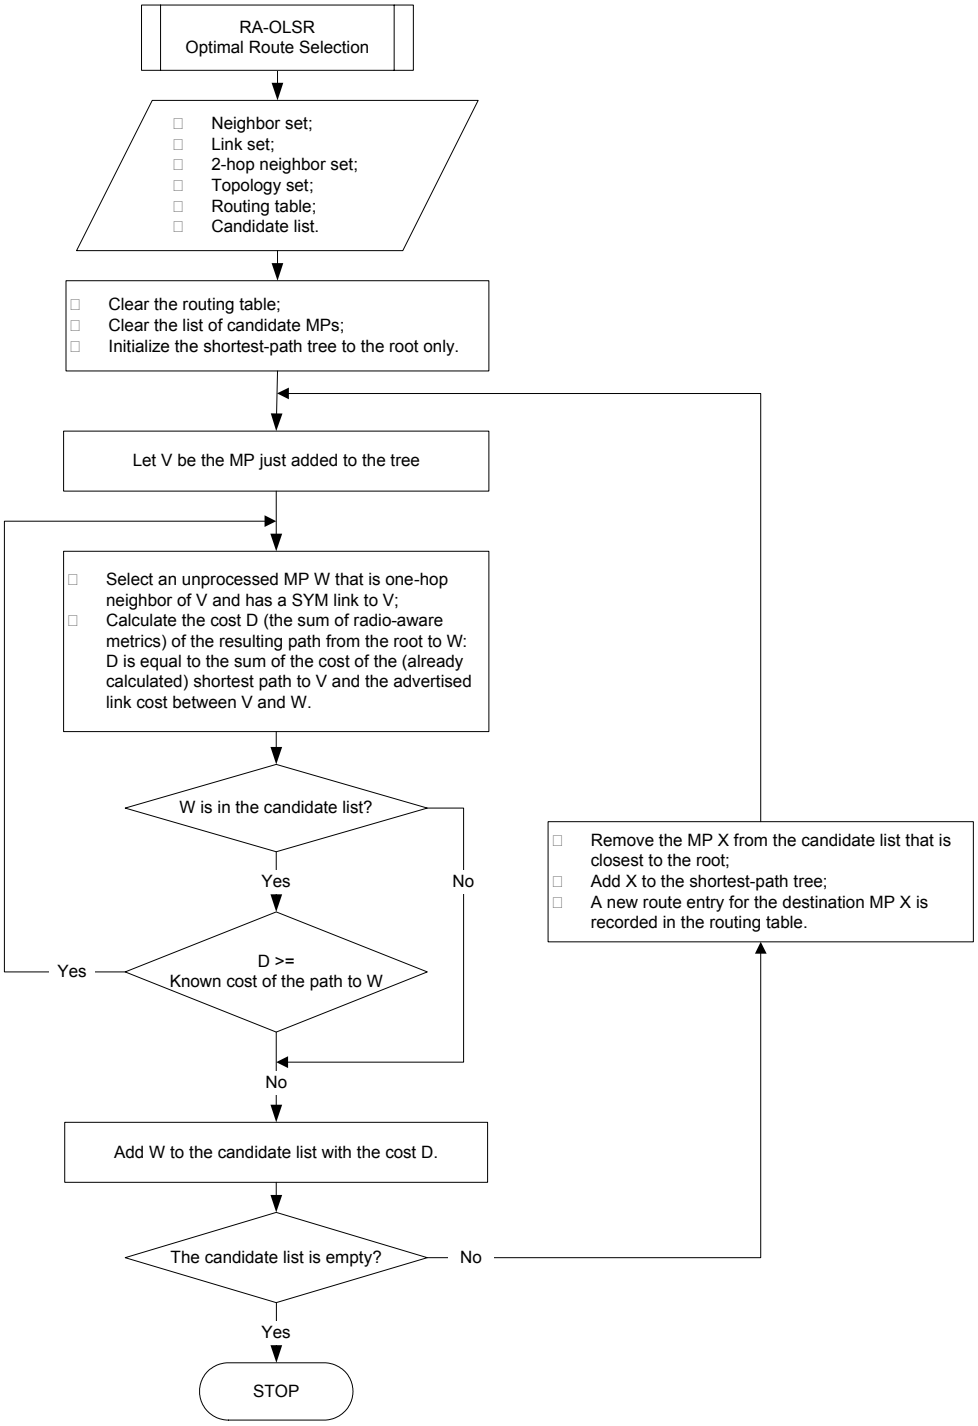

Figure s117—Flowchart for selection of optimal routes.

T.6 Co-located Mesh Point mesh point and Station station functionality

The architecture and scenarios presented in this document standard introduce additional services in the IEEE 802.11 environment that enable generalized multi-hop wireless networks. Mesh services introduce a logical MAC interface that is independent of the 802.11 MAC interface. Any devices 802.11 entities that support mesh services are Mesh Points (MPs). In a particular device the The mesh interface and the BSS interface may be independently and individually invoked. This allows for devices that are both APs and mesh pointsMPs, or both STAs and mesh pointsMPs. The operation of mesh APs or MAPs, the APs with a mesh interface, has a significant bearing on the definition of mesh services, and is described in the documentstandard. On the other hand, the operation of devices that are STAs and MPs at the same time does not have any bearing on the standard specification. A brief description of such implementation specific operation is described in this appendix.

A special type of mesh point MP may be referred to as a mesh point an MP station (MPS). Such a device has a separate logical MAC interface that functions as a STA, along with a logical MAC interface that functions as a an MP. The internal communication between STA and MP interface are implementation dependent. However, given that STAs are end user client devices, bridging functionality within the IEEE 802 domain is not expected. A usage scenario for a mesh point an MP station is shown in Figure s118. In Figure s118, the MPS is connected to two logical networks. The station STA interface connects it to a wired distribution system, and the mesh MP interface connects it to a mesh. Such a scenario may be useful, for example, if connectivity is expected with different security profiles. One interface may be secure (for example the STA interface in Figure s118:), while the other (for example the mesh interface in the figure) may allow insecure access to limited services.

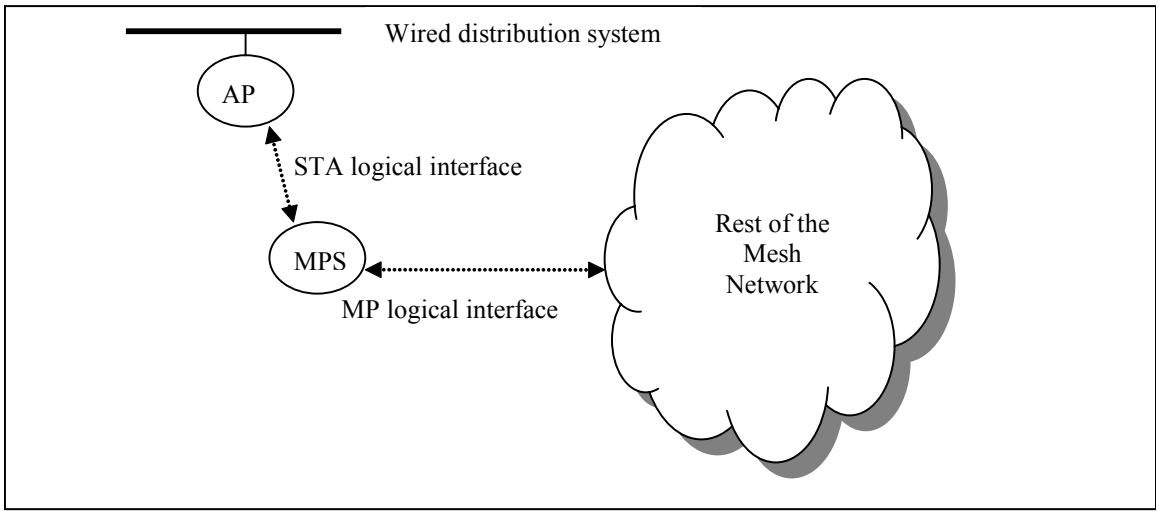

Figure s118—An example usage scenario for a mesh point an MP station (MPS) with both the MP and the STA logical interfaces

T.7 Interworking Support Example support example and Flowchartsflowcharts

T.7.1 General Interworking Example Topologies

T.7.2 General interworking example topologies

In order for a WLAN Mesh to behave as a traditional 802-style LAN, it must should be possible to interconnect the mesh with other networks using both layer 2 bridging and layer 3 internetworking. Figure s119 (a) illustrates an example network where two WLAN Mesh LANs are bridged with 802.3 LAN segments. In this example, each Mesh Point MP collocated with a mesh Portal (MPP) acts as a bridge, connecting the mesh to another LAN using standard bridge protocols (e.g., 802.1D). This configuration effectively creates a single logical layer 2 subnet LAN spanning both meshes and two 802.3 LAN segments. Figure s119 (b) illustrates an example network where the two WLAN Mesh LANs are internetworked with 802.3 LAN segments using layer 3 routing (e.g., IP). In this example, the the devices where MPP is implemented also includes IP gateway functionality, resulting in a network with multiple interconnected subnet LANs.

One or more meshes may be connected to each other through Mesh Portals (see Figure s119 (b)). This may be useful, for example, when different meshes are running different routing protocols, or are configured differently.

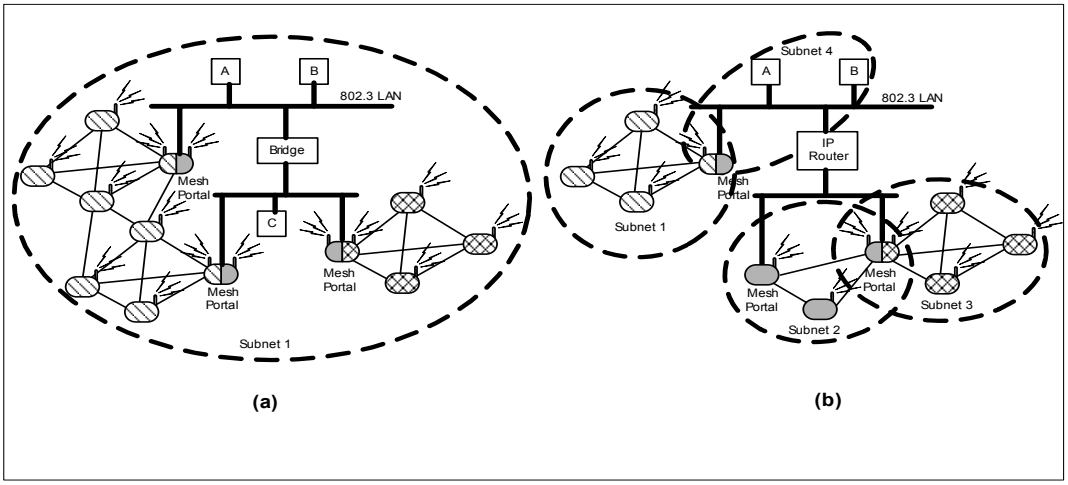

Figure s119—Connecting a WLAN Mesh with other LANs via Mesh Portals. (a) Layer 2 bridging. (b) Layer 3 internetworking.

T.7.3 An Exampleexample

Consider the network in Figure s120 consisting of two wired LAN segments connected by a wireless Mesh. Nodes MPs 1 through 11 make up the Mesh. MPPs A and B act as transparent layer-2 bridges. We assume that AODV is used for discovery of paths for unicast route discoveryindividually addressed frames.

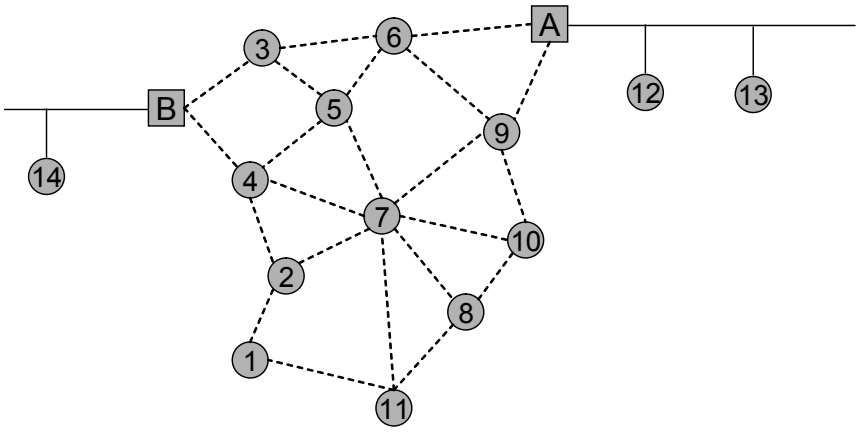

Figure s120—An example bridged network containing two wired segments and a wireless Mesh.

If node MP 11 wants to send a message to node MP 6, it looks in its route table, and finding no route, it initiates a route request. Eventually a route reply propagates back through the network, filling in route table entries along a path between 11 and 6. It is likely that the route request will be received by MPPs A and B, which will add node MP 11 to their bridging tables. It is also possible that the route reply will be overheard by MPP A, in which case bridge table entries will be added for all nodes MPs in the route. Once the route reply reaches node MP 11, data packets can be unicast individually addressed from node MP 11 to node MP 6.

If node MP 14 sends a packet to node MP 2, it will be promiscuously received by MPP B, which is acting as a transparent bridge. MPP B will look in its routing table, and finding no entry, it initiates a route request for node MP 2. As before, node MP A will receive the route request and learn that node MP B is in the Mesh. Eventually, node MP 2 will receive the route request and generate a route reply. The route reply will propagate back to MPP B, creating route entries at nodes MPs along the path. MPP B will also create bridge entries for node MP 2 and the nodes MPs along the path. Subsequent packets received by MPP B to node MP 2 will use this route.

If node MP 3 sends a packet to node MP 12, it will create a route request. Nodes MPs A and B will receive the route request and add node MP 12 to their bridging tables. Ultimately, the route request will timeout, and node MP 3 will add an entry to its routing table for node MP 12 containing the broadcast address as the next hop. Data packets will then be sent via flood, with node MP 12 as the ultimate destination. When MPPs A and B receive the flood, they will repeat the packet on their wired LANs, allowing node MP 12 to receive the packet. Eventually, node MP 12 will send a packet to node MP 3 (most application protocols are bidirectional), allowing MPP A to learn that node MP 12 is on its wired LAN. MPP A will flood a *portal update add* message over the Mesh, allowing all nodes MPs to learn that MPP A is the correct MPP for node MP 12 (by adding an entry to their routing tables). [Note that the The route request sent by node MP A for node MP 3 to deliver the packet for node MP 12 could also allow nodes MPs to learn that MPP A is the right way to get to node MP 12, eliminating the need for the *portal update* flood.] Subsequent packets from node MP 3 to node MP 12 will be unicastindividually addressed. Node MP 3 would look in its routing table for node MP 12 and find MPP A. It would then generate a route request for node MP A, eventually establishing a route. At this point, all nodes MPs along the path from node MP 3 to MPP A know that MPP A is the right way to reach node MP 12, and they know the correct next hop to reach A. Thus, each node MP can forward a unicast an individually addressed packet with 12 as the ultimate destination and the appropriate next hop to node MP A.

T.7.4 Interworking Support Flowchartssupport flowcharts

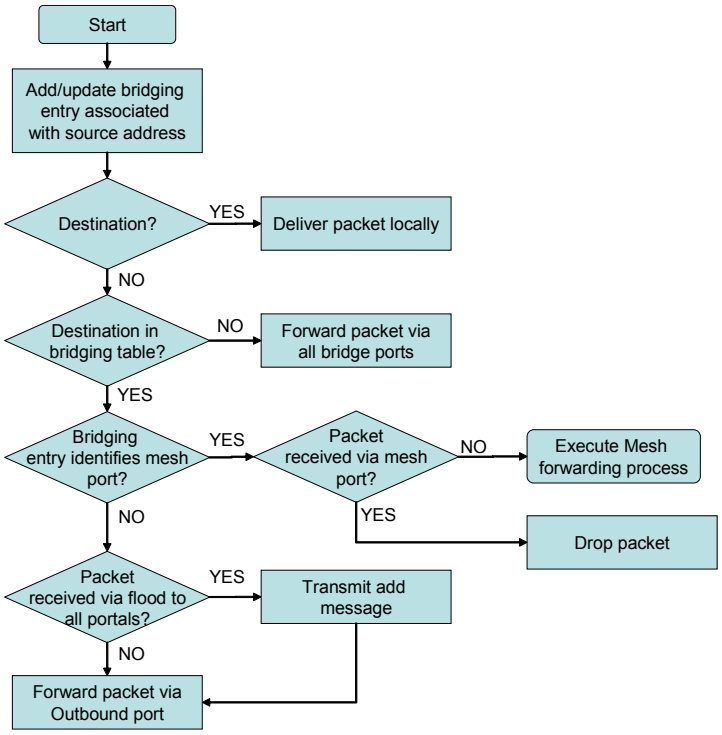

Figure s121—The unicast individually addressed packet forwarding procedure for MPPs.

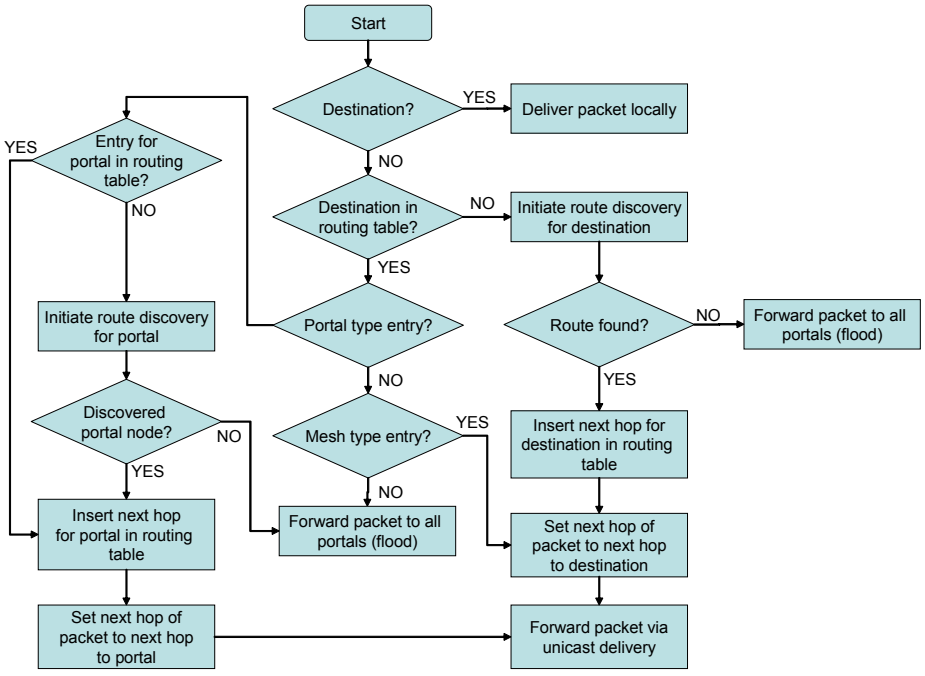

Figure s122—The unicast individually addressed packet forwarding procedure for Mesh nodes MPs with reactive routing.

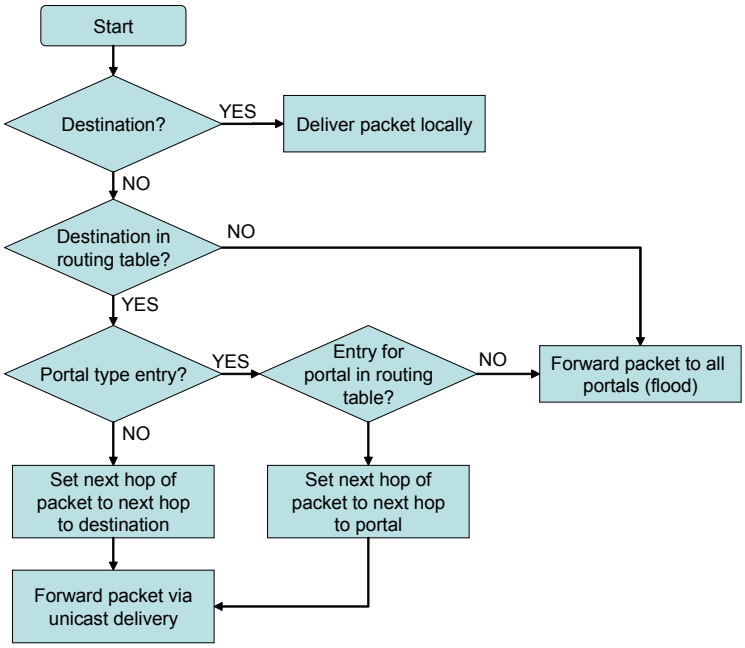

Figure s123—The unicast individually addressed packet forwarding procedure for Mesh nodes MPs with proactive routing.

EDITORIAL NOTE—Operational considerations for interworking moved from 11A.3 to Annex T per CID 3218

T.8 Operational considerations for interworking

T.8.1 Formation and maintenance of the IEEE 802.1D spanning tree

No special action is required to support formation of the IEEE 802.1D spanning tree. Spanning tree control messages are typically delivered to bridges in multicast frames. These messages are data frames from the point of view of the Mesh.

T.8.2 MP mobility

— MP mobility in a bridged network can be within or between physical LANs. Four cases can occur:

— *Mobility of an MP within the mesh.* This kind of mobility is handled through the mesh path selection mechanisms.

— *An MP may move from one LAN outside the Mesh to another LAN outside the Mesh.* The MPPs through which the MP can be reached by MPs in the mesh may change. This case occurs in typical bridged networks and can be handled through bridge learning and timing out of old bridge table entries.

— *An MP may move from inside the Mesh to outside the Mesh.* When an on-demand routing protocol is used, the movement is detected through the route maintenance mechanisms of the protocol, which triggers route repair procedures. When a proactive routing protocol is used, MP failure and information on the new whereabouts of an MP are disseminated during triggered and periodic route update rounds.

— *An MP may move from outside the Mesh to inside the Mesh.* See 11A.4.3 above.

EDITORIAL NOTE—MP boot sequence example moved from 11A.10 per CID 3447

T.9 MP boot sequence example

At power up, a configured MP should perform the following sequence of operations:

a) Passive or Active scanning to discover other MP
b) Channel selection
c) Begin mesh beaconing.
d) Neighbor MP link establishment[1]
    1) IEEE 802.11 open authentication
    2) Association
    3) IEEE 802.11i authentication and key exchange
e) Local link state measurement
f) Path selection initialization
g) AP initialization (optional – if MAP)

This sequence is illustrated in Figure s124.

Link establishment may be performed with a number of MPs but not all of the links become active – that depends on the outcome of the link state measurements and on the routing initialization. The final (optional) step of the boot sequence is AP service initialization. This may be observed as a Mesh AP (MAP) starting to transmit a valid SSID information element in Beacon frames. Before this point, that is, before the path selection specific state has been initialized, the MAP does not accept authentication requests, association requests or probe requests from STAs. After this time, however, the MAP will continue to perform maintenance operations on the path selection state.

[1] Neighbor MP link establishment may be repeated multiple times if there are multiple neighbor MPs

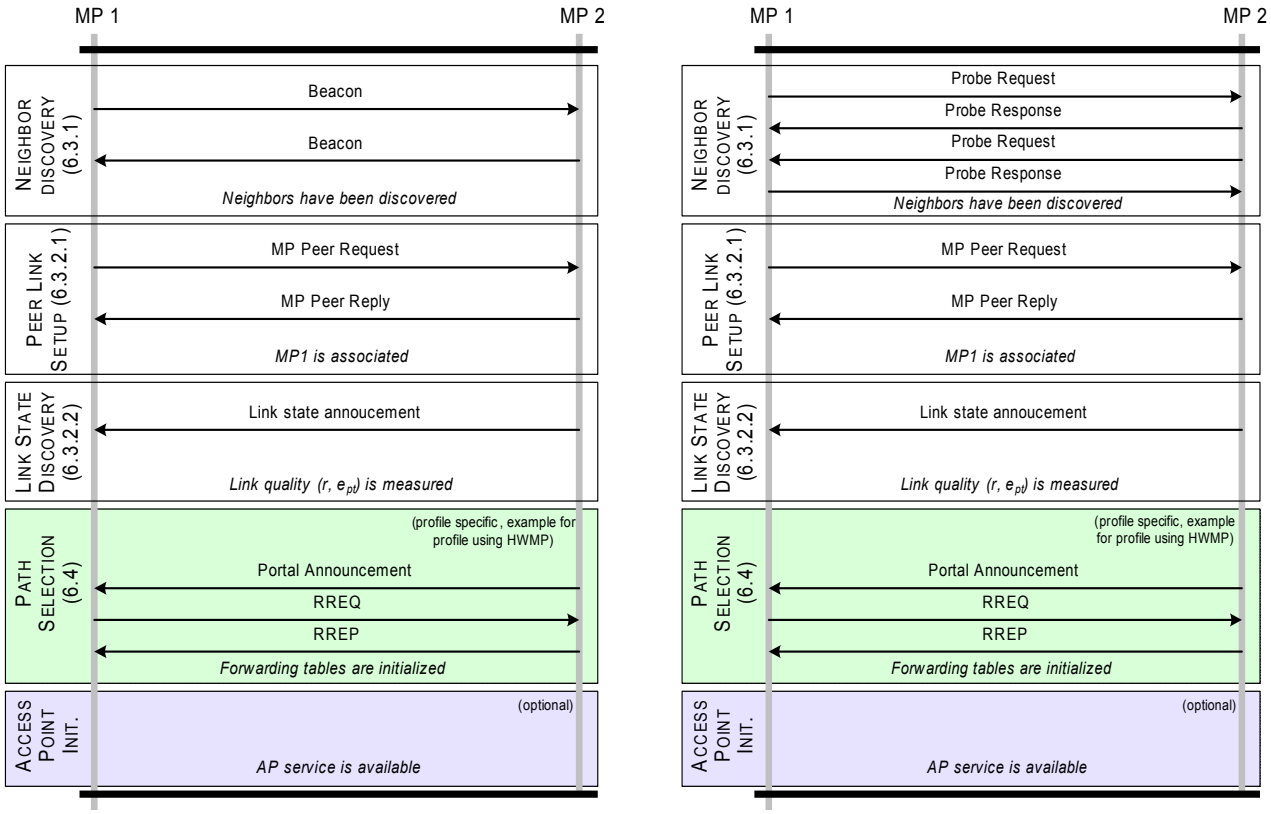

(a) Using Passive Scanning **(b) Using Active Scanning**

Figure s124—MP Boot Sequence

EDITORIAL NOTE—MP table examples moved from 11A.10 per CID 3447

T.10 MP Table Examples

T.10.1 MP neighbor table

An MP should maintain a table containing an entry for each discovered neighbor MP. By definition, all neighbor MPs have the same Mesh ID. Each entry should contain the information shown in Table s49. The r and e_f fields are specific to the default airtime metric defined in 11A.5. An implementation using a different path selection metric may require fields other than r and e_f.

The state of peer link establishment with the neighbor should take one of the values shown in Table s50, and should be initialized on discovery to *neighbor* or *candidate peer* based on beacon or probe response contents as described in 7.2.3.1 and 7.2.3.9.

The neighbor MAC address is the address of the neighbor MP's PHY that was discovered. The state information for the link to the MP stored in the MP neighbor table entry is with respect to this advertised address.

The primary MAC address of a neighbor MP is the primary unique address of the MP. In the case where the neighbor MP has only one PHY, the primary MAC address is equal to the PHY MAC address. In the case where the neighbor MP has more than one PHY, the primary MAC address is typically the PHY MAC

Table s49—MP Neighbor Table Entry

Value	Description
Neighbor MAC address	MAC address of the neighbor MP PHY
Primary MAC address	Primary MAC address of the MP, if it has more than one PHY
State	State of the association with the neighbor
Directionality	Directionality value in previous association request
c_o	Operating channel number
p_l	Channel precedence value
r	Reference bit rate (modulation mode)
e_f	Frame error rate for the reference frame size at the reference bit rate
Q	Received signal strength or quality (internal units)

Table s50—State Values

State	Description
Neighbor	Discovered, no peer capability
Candidate peer	Has peer capability, no association established
Association pending	Association sent, reply not received
Subordinate, link down	Association established with this MP as the subordinate, link not yet measured
Subordinate, link up	Association established with this MP as the subordinate, link measured and active
Superordinate, link down	Association established with peer as the subordinate, link not yet measured
Superordinate, link up	Association established with peer as the subordinate, link measured and active

address with the smallest address value. (Note: more than one table entry may be created for a given neighbor primary MAC address).

The operating channel number is the channel on which the beacon was received from the MP.

The channel precedence value is a number chosen by all MP PHYs in a given Mesh. It is contained in the beacon transmitted by the neighbor MP. It is used when merging disjoint networks and for the purpose of supporting DFS.

The bit rate and frame error rate values are created by the local link state discovery procedures, described in 11A.5.

The received signal strength or quality may represent any convenient quality measure; this value is never presented at an exposed interface, but rather is used for comparisons.

T.10.2 MP proxy table

Each MP maintains a proxy table for the devices outside of the Mesh. The format of a logical proxy table is shown in Table s51.

Table s51—A logical proxy table maintained at each MP (the information can be derived from other sources).

Value	Description
MAC Address	MAC address of a given STA
inMesh	If the destination is within the mesh
isProxied	If the destination is proxied by MAP/MPP
Owner	MAC address of its proxy

An example for proxy registration procedure (Figure s125):

During the initialization phase, a STA first associates with MAP3 by using standard IEEE 802.11 procedures. Once associated, MAP3 may initiate a proxy registration procedure on behalf of STA towards the MPP. To do so, it sends a proxy registration request message on behalf of STA, to the mesh portal. MAP1 after receiving the registration request message updates its own proxy registration table for STA and forwards it towards mesh portal. The MPP thereby learns that STA is being proxied by MAP3. The MPP may then update its proxy registration table with new/updated entry for STA with owner set to MAP3 and inMesh, isProxied flags set to true. The MPP may then create a proxy registration reply message which is sent towards MAP3. Once MAP3 receives a registration confirmation for STA, proxy routing for STA is established. (Further optimization can be done using selective proxying).

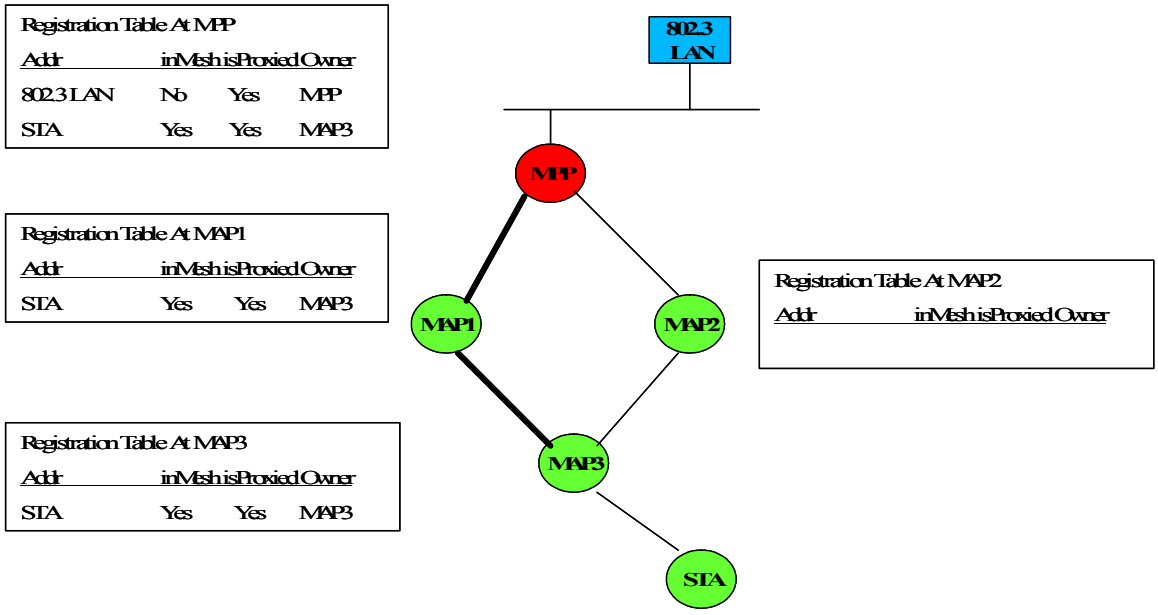

Figure s125—Example of optional proxy registration procedure to MPP.

EDITORIAL NOTE—Power Save parameters selection moved from 11A.9.9 per CID 3444

T.11 Power Save parameters selection

The power save operation of an MP is controlled by a set of global parameters. The following are the global mesh parameters with their default recommended values:

Beacon Period: 100TU

mesh DTIM period: 10

ATIM Window: 10TU

MPs may wish to use other parameters but doing so may effect the power save efficiency and also delay the service initiation in the mesh.

EDITORIAL NOTE—Naive mesh operation moved from 11A.9.12 per CID 3445

T.12 Naive mesh operation

This subclause describes the operation of a naive mesh that does not include any Mesh APs and is also not supporting any routing capabilities. A naive mesh consists of lightweight MPs only.

This type of mesh would be mainly useful in cases were all MP are maintaining a neighbor relationship with each other, as such it does not need to follow the association procedures and may use the three address format to directly exchange frames between the mesh members.

For this specific case the MP does not have to support any of the route messages and link state announcements defined in this standard.

The MP includes in its Mesh Capability information element a Peer Capability field with only bit 15 set. This would signal that the MP is not supporting association with any other MPs and does not support any IEEE 802.1X capabilities.

Since the MP is not going to receive any association request, and it has no need to initiate one itself, it does not have to support the association messages as well as any of congestion information elements that can be exchanged only between associated peers.

The power save operation is an optional feature and may be implemented by the MPs of the Naive mesh.

A naive mesh supporting power save may use the basic power save scheme for simple data exchange and may optionally extend to include the APSD support for real time power save stream delivery.

T.13 Non-forwarding mesh point operation (Informative)

These are normal mesh points (MPs) configured/functioning to be selfish in the mesh. The operation has no bearing on the standard, and is completely an implementation issue. Consider an MP that has a single neighbor only. Such a 'leaf' MP never forwards any other MPs MP's data in the mesh, but can still communicate with the rest of the mesh through its single neighbor. Such behaviour behavior can be extended to a scenario when there are multiple neighbors of an MP. Non-forwarding MPs are such MPs that 'lie' to their neighbors that they cannot reach/do not have any other neighbors. Thus, they never receive data to be forwarded. Such MPs do not advertise any routes or reachability to other MPs to any of their neighbors. They can possibly communicate with all of the rest of the mesh through any of their neighbors.

While a functionality similar to a non-forwarding mesh point MP can be achieved through the STA functionality, the non non-forwarding MP operation allows the added flexibility of communicating over the mesh interface even if no access points are available in the vicinity.

T.14 Informative references[2]

EDITORIAL NOTE—These informative references were relocated from Clause 2 since this draft does not use RFCs 3561 and 3626.

IETF RFC 3561, "Ad hoc On-Demand Distance Vector (AODV) Routing", C. Perkins, E. Belding-Royer, S. Das, July 2003. (status: experimental)

IETF RFC 3626, "Optimized Link State Routing Protocol (OLSR)", T. Clausen and P. Jacquet, October 2003. (status: experimental)

[2]Internet RFCs are available from the Internet Engineering Task Fource at http://www.ietf.org/.