

IEEE Recommended Practice for Routing Packets in IEEE 802.15.4™ Dynamically Changing Wireless Networks

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

IEEE
3 Park Avenue
New York, NY 10016-5997
USA

IEEE Std 802.15.10™-2017

IEEE Recommended Practice for Routing Packets in IEEE 802.15.4™ Dynamically Changing Wireless Networks

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 14 February 2017
IEEE-SA Standards Board

Abstract: The routing of packets in dynamically changing wireless networks is facilitated in this recommended practice. The result is an extension of the area of coverage as the number of nodes increase.

Keywords: IEEE 802.15.4™, layer 2 routing, mesh-under routing, mesh

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Copyright © 2017 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 21 April 2017. Printed in the United States of America.

IEEE is a registered trademark in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

Print: ISBN 978-1-5044-3848-3 STD22478
PDF: ISBN 978-1-5044-3849-0 STDPD22478

IEEE Prohibits discrimination, harassment, and bullying.

For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.

Important Notices and Disclaimers Concerning IEEE Standards Documents

IEEE documents are made available for use subject to important notices and legal disclaimers. These notices and disclaimers, or a reference to this page, appear in all standards and may be found under the heading “Important Notices and Disclaimers Concerning IEEE Standards Documents.” They can also be obtained on request from IEEE or viewed at <http://standards.ieee.org/IPR/disclaimers.html>.

Notice and Disclaimer of Liability Concerning the Use of IEEE Standards Documents

IEEE Standards documents (standards, recommended practices, and guides), both full-use and trial-use, are developed within IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (“IEEE-SA”) Standards Board. IEEE (“the Institute”) develops its standards through a consensus development process, approved by the American National Standards Institute (“ANSI”), which brings together volunteers representing varied viewpoints and interests to achieve the final product. IEEE Standards are documents developed through scientific, academic, and industry-based technical working groups. Volunteers in IEEE working groups are not necessarily members of the Institute and participate without compensation from IEEE. While IEEE administers the process and establishes rules to promote fairness in the consensus development process, IEEE does not independently evaluate, test, or verify the accuracy of any of the information or the soundness of any judgments contained in its standards.

IEEE Standards do not guarantee or ensure safety, security, health, or environmental protection, or ensure against interference with or from other devices or networks. Implementers and users of IEEE Standards documents are responsible for determining and complying with all appropriate safety, security, environmental, health, and interference protection practices and all applicable laws and regulations.

IEEE does not warrant or represent the accuracy or content of the material contained in its standards, and expressly disclaims all warranties (express, implied and statutory) not included in this or any other document relating to the standard, including, but not limited to, the warranties of: merchantability; fitness for a particular purpose; non-infringement; and quality, accuracy, effectiveness, currency, or completeness of material. In addition, IEEE disclaims any and all conditions relating to: results; and workmanlike effort. IEEE standards documents are supplied “AS IS” and “WITH ALL FAULTS.”

Use of an IEEE standard is wholly voluntary. The existence of an IEEE standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard.

In publishing and making its standards available, IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity nor is IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing any IEEE Standards document, should rely upon his or her own independent judgment in the exercise of reasonable care in any given circumstances or, as appropriate, seek the advice of a competent professional in determining the appropriateness of a given IEEE standard.

IN NO EVENT SHALL IEEE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO: PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE PUBLICATION, USE OF, OR RELIANCE UPON ANY STANDARD, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE AND REGARDLESS OF WHETHER SUCH DAMAGE WAS FORESEEABLE.

Translations

The IEEE consensus development process involves the review of documents in English only. In the event that an IEEE standard is translated, only the English version published by IEEE should be considered the approved IEEE standard.

Official statements

A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered or inferred to be the official position of IEEE or any of its committees and shall not be considered to be, or be relied upon as, a formal position of IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE.

Comments on standards

Comments for revision of IEEE Standards documents are welcome from any interested party, regardless of membership affiliation with IEEE. However, IEEE does not provide consulting information or advice pertaining to IEEE Standards documents. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Since IEEE standards represent a consensus of concerned interests, it is important that any responses to comments and questions also receive the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to comments or questions except in those cases where the matter has previously been addressed. For the same reason, IEEE does not respond to interpretation requests. Any person who would like to participate in revisions to an IEEE standard is welcome to join the relevant IEEE working group.

Comments on standards should be submitted to the following address:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854 USA

Laws and regulations

Users of IEEE Standards documents should consult all applicable laws and regulations. Compliance with the provisions of any IEEE Standards document does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

IEEE draft and approved standards are copyrighted by IEEE under U.S. and international copyright laws. They are made available by IEEE and are adopted for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making these documents available for use and adoption by public authorities and private users, IEEE does not waive any rights in copyright to the documents.

Photocopies

Subject to payment of the appropriate fee, IEEE will grant users a limited, non-exclusive license to photocopy portions of any individual standard for company or organizational internal use or individual, non-commercial use only. To arrange for payment of licensing fees, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Updating of IEEE Standards documents

Users of IEEE Standards documents should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect.

Every IEEE standard is subjected to review at least every ten years. When a document is more than ten years old and has not undergone a revision process, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE standard.

In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE-SA Website at <http://ieeexplore.ieee.org> or contact IEEE at the address listed previously. For more information about the IEEE SA or IEEE's standards development process, visit the IEEE-SA Website at <http://standards.ieee.org>.

Errata

Errata, if any, for all IEEE standards can be accessed on the IEEE-SA Website at the following URL: <http://standards.ieee.org/findstds/errata/index.html>. Users are encouraged to check this URL for errata periodically.

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 by the IEEE with respect to the existence or validity of any patent rights in connection therewith. If a patent holder or patent applicant has filed a statement of assurance via an Accepted Letter of Assurance, then the statement is listed on the IEEE-SA Website at <http://standards.ieee.org/about/sasb/patcom/patents.html>. Letters of Assurance may indicate whether the Submitter is willing or unwilling to grant licenses under patent rights without compensation or under reasonable rates, with reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to obtain such licenses.

Essential Patent Claims may exist for which a Letter of Assurance has not been received. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims, or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

At the time this standard was completed, the IEEE 802.15 Working Group had the following membership:

Robert F. Heile, *Working Group Chair*
Rick Alfvén, *Working Group Vice Chair*
Patrick Kinney, *Working Group Vice Chair, Working Group Secretary*
James P. K. Gilb, *Working Group Technical Editor*
Benjamin A. Rolfe, *Working Group Treasurer*

Clinton C Powell, *802.15.10 Chair*
Verotiana Rabarijaona, *802.15.10 Technical Editor*
Charles Perkins, *802.15.10 Secretary*

Mounir Archir
Keiji Akiyama
Arthur Astrin
Philip Beecher
Frederik Beer
Chandrashekhar P. S. Bhat
Kiran Bynam
Edgar Callaway
Chris Calvert
Radhakrishna Canchi
Kapsoek Chang
Soo-Young Chang
Clint Chaplin
Stephen Chasko
Paul Chilton
Sangsung Choi
Hendricus De Ruijter
Guido Dolmans
Igor Dotlic
Stefan Drude
Dietmar Eggert
Shahriar Emami
Andrew Estrada
David Evans
George Flammer
Kiyoshi Fukui
Matthew Gillmore
Tim Godfrey
Elad Gottlib
Jussi Haapola

Rainer Hach
Shinsuke Hara
Timothy Harrington
James Hartman
Marco Hernandez
Ken Hiraga
Iwao Hosako
Yeong Min Jang
Seong-Soon Joo
Akifumi Kasamatsu
Shuzo Kato
Toyoyuki Kato
Jeritt Kent
Jaehwan Kim
Youngsoo Kim
Shoichi Kitazawa
Tero Kivinen
Ryuji Kohno
Fumihide Kojima
Thomas Kuerner
Byung-Jae Kwak
Jae Lee
Myung Lee
Sangjae Lee
Huan-Bang Li
Liang Li
Qing Li
Michael Lynch
Itaru Maekawa
Hiroyuki Matsumura
Michael McLaughlin

Michael McInnis
Kenichi Mori
Robert Moskowitz
Jinesh Nair
Chiu Ngo
Paul Nikolich
John Notor
Hiroyo Ogawa
Taejoon Park
Glenn Parsons
Albert Petrick
Ivan Reede
Richard Roberts
Ruben Salazar Cardozo
Noriyuki Sato
Norihiro Sekine
Kunal Shah
Stephen Shellhammer
Shusaku Shimada
Gary Stuebing
Don Sturek
Mineo Takai
Kou Togashi
Billy Verso
Gabriel Villardi
Brian Weis
Makoto Yaita
Peter Yee
Yu Zeng
Chunhui Zhu

Major contributions were received from the following individuals:

Soo-Young Chang
Sangsung Choi
Kiyoshi Fukui
Kojima Fumihide
Hiroshi Harada

Young-Ae Jeon
Tero Kivinen
Jae-Beom Kim
Jaehwan Kim
Young-Bae Ko

Sangjae Lee
Charles Perkins
Verotiana Rabarijaona
Don Sturek
Noriyuki Sato

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

Thomas Alexander
Rick Alfvén
James D. Allen
Nobumitsu Amachi
Carol Ansley
Butch Anton
Michael Bahr
Tuncer Baykas
Mathilde Benveniste
Harry Bims
Gennaro Boggia
Riccardo Brama
Nancy Bravin
William Byrd
Edgar Callaway
Juan Carreon
Dave Cavalcanti
Minho Cheong
Charles Cook
Hendricus De Ruijter
Richard Doyle
Sourav Dutta
Richard Edgar
David Evans
Liu Fangfang
Avraham Freedman
Yukihiro Fujimoto
James P. K. Gilb
Tim Godfrey
Joel Goergen
Randall Groves

Rainer Hach
Robert F. Heile
Jerome Henry
Marco Hernandez
Werner Hoelzl
Noriyuki Ikeuchi
Sergiu Iordanescu
Atsushi Ito
Raj Jain
Anthony Jeffree
Michael Johas Teener
Adri Jovin
Piotr Karocki
Ruediger Kays
Stuart Kerry
Tero Kivinen
Yasushi Kudoh
Thomas Kuerner
Thomas Kurihara
James Lepp
Arthur H. Light
Michael Lynch
Elvis Maculuba
Roger Marks
Jeffery Masters
Michael McInnis
Ronald Murias
Michael Newman
Nick S. A Nikjoo
John Notor

Yoshihiro Ohba
Satoshi Oyama
Arumugam Paventhan
Clinton C Powell
Venkatesha Prasad
Verotiana Rabarijaona
Maximilian Riegel
Robert Robinson
Benjamin A. Rolfe
Noriyuki Sato
Frank Schewe
Reinhard Schrage
Kunal Shah
Daniel Smolinski
Ju-Hyung Son
Dorothy Stanley
Thomas Starai
Eugene Stoudenmire
Rene Struik
Walter Struppler
Mark Sturza
Patricia Thaler
Mark-Rene Uchida
Lorenzo Vangelista
Dmitri Varsanofiev
Prabodh Varshney
George Vlantis
Khurram Waheed
Hung-Yu Wei
Andreas Wolf
Dayin Xu
Oren Yuen

When the IEEE-SA Standards Board approved this standard on 14 February 2017, it had the following membership:

Jean-Philippe Faure, *Chair*
Vacant Position, *Vice Chair*
John D. Kulick, *Past Chair*
Konstantinos Karachalios, *Secretary*

Chuck Adams
Masayuki Ariyoshi
Ted Burse
Stephen Dukes
Doug Edwards
J. Travis Griffith
Gary Hoffman

Michael Janezic
Thomas Koshy
Joseph L. Koepfinger*
Kevin Lu
Daleep Mohla
Damir Novosel
Ronald C. Petersen
Annette D. Reilly

Robby Robson
Dorothy Stanley
Adrian Stephens
Mehmet Ulema
Phil Wennblom
Howard Wolfman
Yu Yuan

*Member Emeritus

Introduction

This introduction is not part of IEEE Std 802.15.10-2017, IEEE Recommended Practice for Routing Packets in IEEE 802.15.4™ Dynamically Changing Wireless Networks.

Devices that implement the IEEE 802.15.4 technology are widely deployed and support a variety of applications. The latest PHYs and features of the IEEE 802.15.4 technology enable support for networks operating over a diverse range of geographical areas.

The use of mesh network topologies features heavily in many of these applications. This recommended practice aims at enabling mesh routing at the link layer on these types of IEEE 802.15.4 networks.

Contents

1.	Overview.....	14
1.1	Scope.....	14
1.2	Purpose.....	14
2.	Normative references.....	14
3.	Definitions, acronyms, and abbreviations.....	15
3.1	Definitions	15
3.2	Acronyms and abbreviations	16
4.	General description	18
4.1	Introduction.....	18
4.2	Format conventions.....	18
4.3	Network topology	19
4.4	Architecture	21
4.5	Functional overview	22
4.5.1	L2R mesh and routing.....	22
4.5.2	Addressing modes.....	25
5.	L2R functional description	26
5.1	Starting and joining an L2R mesh	26
5.1.1	Starting and stopping an L2R mesh.....	26
5.1.1.1	Procedure to start an L2R mesh.....	26
5.1.1.2	Procedure to update the L2R mesh topology parameters	27
5.1.1.3	Procedure to stop an L2R mesh	27
5.1.2	Joining an L2R mesh	28
5.1.2.1	Mesh discovery procedure	28
5.1.2.1.1	Discovery of L2R meshes with a specific mesh ID.....	29
5.1.2.1.2	Discovery of all existing meshes	30
5.1.2.1.3	Discovery within a PAN.....	31
5.1.2.2	Procedure to join an L2R mesh.....	31
5.1.2.2.1	Mesh selection by the L2R sublayer.....	31
5.1.2.2.2	Mesh selection by the next higher layer	33
5.1.2.3	Procedure to rejoin an L2R mesh	36
5.1.2.4	Procedure to leave an L2R mesh	36
5.1.2.5	Short address management	37
5.1.2.5.1	Initial short address assignment.....	37
5.1.2.5.2	Short address maintenance.....	38
5.1.2.5.3	Short address release.....	39
5.2	L2R mesh formation and route establishment	40
5.2.1	L2R mesh construction	40
5.2.2	L2R route metrics	47
5.2.2.1	Received Signal Weakness (RSW).....	48
5.2.2.2	ETX.....	50
5.2.2.3	Expected airtime metric	51
5.2.3	US route establishment	51
5.2.4	DS route establishment	53

5.2.4.1	Hop-by-hop DS route establishment.....	53
5.2.4.2	DS source route establishment.....	54
5.2.5	Route establishment for TMCTP.....	54
5.2.5.1	MPO US route establishment	54
5.2.5.2	MPO DS route establishment	55
5.2.6	Multicast route establishment	55
5.2.7	P2P route establishment.....	57
5.2.7.1	Storing mode.....	58
5.2.7.2	Non-storing mode	63
5.3	L2R route maintenance	65
5.3.1	L2R US route maintenance.....	65
5.3.2	L2R DS route maintenance.....	65
5.3.3	L2R Route maintenance with MPO.....	66
5.4	L2R routing capabilities.....	66
5.4.1	Unicast routing.....	66
5.4.1.1	Source routing.....	67
5.4.1.2	Hop-by-hop routing	67
5.4.1.3	L2R Routing with MPO.....	69
5.4.1.4	Hop-by-hop retransmission.....	70
5.4.1.5	End to end acknowledgment.....	71
5.4.1.6	DCat.....	72
5.4.2	Multicast routing.....	73
5.4.2.1	When L2R Multicast is set to TRUE in an L2R mesh in a regular PAN	73
5.4.2.2	When L2R Multicast is set to FALSE	77
5.4.3	Broadcast routing.....	77
5.4.3.1	General case.....	77
5.4.3.2	Broadcast routing with indirect transmission in non-beacon enabled PAN ..	77
5.4.3.3	Broadcast routing with MPO	78
5.5	L2R security aspects	78
5.5.1	Bootstrapping.....	78
5.5.1.1	Non secured mode bootstrapping	78
5.5.1.2	Out of band mode bootstrapping	78
5.5.1.3	Bootstrapping with KMP	79
5.5.1.4	Securing the L2R-D IE	79
5.5.2	Securing frames	80
6.	L2R extensions to MAC frame formats.....	81
6.1	Short nested IEs	81
6.1.1	L2R-D IE	82
6.1.1.1	Descriptor field	82
6.1.1.2	Mesh ID Information field.....	83
6.1.1.3	Mesh Root Address field	84
6.1.2	TC IE.....	84
6.1.2.1	Descriptor field	84
6.1.2.2	Mesh Root Address field	85
6.1.2.3	Tx Extended Address field	85
6.1.2.4	Tx Short Address field.....	85
6.1.2.5	L2R Max Depth field.....	85
6.1.2.6	Depth field	85
6.1.2.7	MSN field	85
6.1.2.8	TC IE Interval field.....	85
6.1.2.9	MPO Descriptor field.....	85
6.1.2.10	DCat Buffering Time field.....	86

6.1.2.11	PQM field	86
6.1.3	NLM IE	88
6.1.3.1	Number of Neighbors field	88
6.1.3.2	Number of Containers field	88
6.1.3.3	NLM IE Interval field	88
6.1.3.4	Neighbor Metric Container field.....	88
6.1.4	SRA IE	89
6.1.4.1	Vendor Specific Usage field	89
6.1.4.2	SRA SN or Vendor Specific field.....	89
6.1.4.3	Source Address field	89
6.1.4.4	Number of Intermediate Addresses field	90
6.1.4.5	Intermediate Address List field.....	90
6.1.5	P2P-RQ IE	90
6.1.5.1	Descriptor field	90
6.1.5.2	P2P DA field	91
6.1.5.3	P2P SA field.....	91
6.1.5.4	PSN field.....	91
6.1.5.5	TTL field.....	91
6.1.5.6	PQM field	91
6.1.5.7	Number of Intermediate Addresses field	91
6.1.5.8	Intermediate Address List field.....	91
6.1.6	P2P-RP IE	91
6.1.6.1	Address Mode field.....	92
6.1.6.2	Mesh Root Address field	92
6.1.6.3	P2P DA field	92
6.1.6.4	P2P SA field.....	92
6.1.6.5	PSN field.....	92
6.1.6.6	PQM field	92
6.1.6.7	TTL	92
6.1.6.8	Number of Intermediate Addresses field	92
6.1.6.9	Intermediate Address List field.....	92
6.1.7	SLR IE	93
6.1.7.1	Source Address field.....	93
6.1.7.2	Destination Address field.....	93
6.1.7.3	SLSN field	93
6.1.7.4	Number of Intermediate Addresses field	93
6.1.7.5	Intermediate Address List field.....	93
6.1.8	E2E ACK IE	93
6.1.9	A-RLS IE	94
6.1.9.1	Releasing Device Extended Address field.....	94
6.1.9.2	Short Address field	94
6.1.10	AA-RQ IE	94
6.1.10.1	Joining Device Extended Address field.....	94
6.1.10.2	Allocated Address field	94
6.1.10.3	Expiration Time field.....	95
6.1.11	AA-RP IE.....	95
6.1.11.1	Status field	95
6.1.11.2	Joining Device Extended Address field.....	95
6.1.11.3	Allocated Address field	96
6.1.11.4	Expiration Time field.....	96
6.2	L2R Long nested IE (LLIE).....	96
6.2.1	RA IE	97
6.2.1.1	Descriptor field	97
6.2.1.2	Mesh Root Address field	97

6.2.1.3	Depth field	97
6.2.1.4	LSN field.....	97
6.2.1.5	RA IE Interval field	97
6.2.1.6	Source Address field	97
6.2.1.7	Multicast Subscription field.....	98
6.2.1.8	Number of Intermediate Addresses field	98
6.2.1.9	Intermediate Address List field.....	98
6.2.2	L2R Routing IE.....	98
6.2.2.1	Descriptor field	98
6.2.2.2	Mesh Root Address field	100
6.2.2.3	Addressing fields	100
6.2.2.4	LSN field.....	101
6.2.2.5	TTL field.....	101
6.2.2.6	Retry Limit field	101
6.2.2.7	Number of Intermediate Addresses field	101
6.2.2.8	Intermediate Address List field.....	101
6.2.3	DCat IE	101
7.	L2R services	102
7.1	L2R management service.....	102
7.1.1	Primitives to manage an L2R mesh	102
7.1.1.1	L2RLME-PAN-SCAN.request.....	103
7.1.1.2	L2RLME-PAN-SCAN.confirm.....	104
7.1.1.3	L2RLME-PAN-SCAN.indication.....	106
7.1.1.4	L2RLME-MESH-START.request.....	106
7.1.1.5	L2RLME-MESH-START.confirm.....	108
7.1.1.6	L2RLME-MESH-STOP.request.....	109
7.1.1.7	L2RLME-MESH-STOP.confirm.....	110
7.1.1.8	L2RLME-JOIN-MESH.request.....	110
7.1.1.9	L2RLME-JOIN-MESH.confirm.....	111
7.1.1.10	L2RLME-MESH-DISCOVERY.request.....	112
7.1.1.11	L2RLME-MESH-DISCOVERY.confirm.....	113
7.1.1.12	L2RLME-MESH-SELECT.request.....	113
7.1.1.13	L2RLME-MESH-SELECT.confirm.....	113
7.1.1.14	L2RLME-LEAVE-MESH.request	114
7.1.1.15	L2RLME-LEAVE-MESH.confirm	115
7.1.1.16	L2RLME-DISCONNECT-MESH.indication.....	115
7.1.1.17	L2RLME-INTERMEDIATE-ADDR-LIST.indication	116
7.1.1.18	L2RLME-NOTIFY.indication	117
7.1.2	Primitives to manage a short address assignment.....	118
7.1.2.1	L2RLME-AA-RQ.request	118
7.1.2.2	L2RLME-AA-RQ.confirm	118
7.1.2.3	L2RLME-AA-RQ.indication.....	120
7.1.2.4	L2RLME-AA-RP.request.....	120
7.1.2.5	L2RLME-AA-RP.confirm.....	121
7.1.2.6	L2RLME-ARLS.request.....	122
7.1.2.7	L2RLME-ARLS.confirm.....	122
7.1.2.8	L2RLME-ARLS.indication	123
7.1.3	Primitives to handle multicast route establishment	123
7.1.3.1	L2RLME-MULTICAST-SUBSCRIPTION.request	123
7.1.3.2	L2RLME-MULTICAST-SUBSCRIPTION.confirm	124
7.2	L2R data service	125
7.2.1	Unicast and broadcast data service	125

7.2.1.1	L2R-DATA.request	126
7.2.1.2	L2R-DATA.confirm	128
7.2.1.3	L2R-DATA.indication	129
7.2.2	Multicast data service	130
7.2.2.1	L2R-MC-DATA.request.....	131
7.2.2.2	L2R-MC-DATA.confirm.....	132
7.2.2.3	L2R-MC-DATA.indication	132
7.3	L2IB attributes	134
Annex A (normative) NLM IE usage and mutual link metric calculation		140
Annex B (informative) Bibliography.....		141

IEEE Recommended Practice for Routing Packets in IEEE 802.15.4™ Dynamically Changing Wireless Networks

1. Overview

1.1 Scope

This recommended practice defines a protocol that routes packets in a dynamically changing IEEE 802.15.4 network with minimal impact from route management. The result is an extension of the area of coverage as the number of nodes increases.

1.2 Purpose

This recommended practice will facilitate the routing of packets in dynamically changing wireless networks. In particular, it will provide for automatic handling of route related capabilities such as the following:

- Discovery and addition of new nodes
- Route establishment
- Dynamic route reconfiguration
- Re-establishment of broken/lost routes
- Purging of inactive routes
- Real time link status information exchange
- Single hop appearance at the networking layer (not breaking standard Layer 3 mechanisms)¹
- Support for broadcast
- Support for multicast
- Frame forwarding

2. Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

IEEE Std 802[®], IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.^{2, 3}

¹Refer to the network layer of the OSI model as described in “OSI Reference Model - The ISO Model of Architecture for Open Systems Interconnections”

²The IEEE standards or products referred to in this clause are trademarks of The Institute of Electrical and Electronics Engineers, Inc.

³IEEE publications are available from The Institute of Electrical and Electronics Engineers (<http://standards.ieee.org/>)

IEEE Std 802.15.4™, IEEE Standard for Low-Rate Wireless Networks.

IEEE Std 802.15.9™, IEEE Recommended Practice for Transport of Key Management Protocol (KMP) Datagrams.

3. Definitions, acronyms, and abbreviations

3.1 Definitions

For the purposes of this document, the following terms and definitions apply. The *IEEE Standards Dictionary Online* should be consulted for terms not defined in this clause.⁴

ancestor: Neighbor at a depth at least one less than the current device.

best upstream next hop: Neighbor selected as a next hop towards the mesh root based on the current next hop selection scheme.

broadcast address: Broadcast short address defined as 0xffff or the 64-bit broadcast MAC address defined in IEEE Std 802[®] as 0xffff ffff ffff ffff.

coordinator: Full function device allowing other devices to associate to the personal area network (PAN) to which it is itself associated.

current transmitter: Device currently forwarding a frame whose address is found in the Source Address field of the MAC Header of the frame.

descendant: Neighbor at least of one depth more than the current device.

depth: Position of a device relative to the mesh root with respect to the path quality metric (PQM), increasing as the PQM worsens.

destination address: Address of the final destination device of a frame.

direct connection: Connection between a mesh root and a personal area network (PAN) coordinator through an external network or by being located in the same device.

enhanced beacon: Beacon frame that may contain one or more information elements (IE) defined in IEEE Std 802.15.4™.

enhanced beacon request: MAC command frame used to request an enhanced beacon (EB) defined in IEEE Std 802.15.4™.

extended address: A 64-bit extended unique identifier (EUI-64) or a 64-bit local MAC address.

NOTE—EUI-64's are defined by IEEE Std 802-2014 and assigned by the IEEE Registration Authority. Interested applicants should contact the IEEE Registration Authority, <http://standards.ieee.org/develop/regauth/>.

intermediate hop: Device on a route between the source device and the destination device of a frame.

link quality metric: Value of a metric used to evaluate the quality of a link between two neighbors.

⁴The *IEEE Standards Dictionary Online* is available <http://dictionary.ieee.org>.

mesh root: Device in the layer 2 routing mesh with depth zero that manages the mesh.

multicast routing: Routing of a frame from a device to a group of devices belonging to a multicast group.

mutual link metric: Value of the link quality metric between two devices, used when links are assumed to be asymmetric.

neighbor: Device within the coverage area of the current device.

next hop: Neighbor to which the transmitter forwards the current frame in order to reach the final destination.

parent PAN: Personal area network (PAN) with which a PAN coordinator is associated before starting its own PAN in a television white space multichannel cluster tree personal area network (TMCTP).

path quality metric: Value of a metric used to evaluate the quality of a path between a device and the mesh root.

short address: An assigned IEEE 802.15.4™ 16-bit address.

sibling: Neighbor with the same depth as the current device.

small scale layer 2 routing mesh: L2R mesh where the mesh root is located at the PAN coordinator operating with short nested information elements.

small scale personal area network: Personal area network upon which a small scale layer 2 routing mesh is deployed.

source address: Address of the original source device of a frame.

super PAN coordinator: First PAN coordinator starting a television white space multichannel cluster tree personal area network as defined in IEEE Std 802.15.4™.

transmitting address: Address of the current transmitter.

unicast routing: Routing of a frame from a unique source to a unique destination.

3.2 Acronyms and abbreviations

AA	address assignment
AA-RP	address assignment reply
AA-RQ	address assignment request
AR	acknowledgment required
A-RLS	address release
BI	beacon interval
BO	beacon order
BOP	beacon only period
CAP	contention access period
CFP	contention free period
CSL	coordinated sampled listening
DA	destination address

DBS	dedicated beacon slot
DCat	data concatenation
DS	downstream
E2E	end-to-end
EB	enhanced beacon
EBR	enhanced beacon request
ETT	expected transmission time
ETX	expected transmission count
EUI-64	64-bit extended unique identifier
FFD	full-function device
ID	identifier
IE	information element
KMP	key management protocol
L2IB	layer 2 routing information base
L2R	layer 2 routing
L2R-D	layer 2 routing discovery
L2RLME	layer 2 routing sublayer management entity
LQM	link quality metric
LQT	link quality threshold
LSN	layer 2 routing sequence number
MCPS	medium access control common part sublayer
MHR	medium access control header
MLME	medium access control sublayer management entity
MP	multipurpose
MPO	multi-personal area network operation
MSN	mesh sequence number
MT	mesh table
MTU	maximum transmission unit
NHA	next hop address
NLM	neighbor link metric
NT	neighbor table
P2P	peer-to-peer
P2P-RP	peer-to-peer reply
P2P-RQ	peer-to-peer request
PAN	personal area network
PanC DC	personal area network coordinator direct connection
PIB	personal area network information base
PQM	path quality metric
PSN	peer-to-peer sequence number
RA	route announcement
RFD	reduced-function device
RIT	receiver initiated transmission
RL	retry limit
RSW	received signal weakness
RvS	routing via a sibling
SA	source address

SHR	synchronization header
SL2R	small scale personal area network using layer 2 routing
SLR	short layer 2 routing
SLSN	small scale personal area network layer 2 routing sequence number
SO	superframe order
SPC	super personal area network coordinator
SRA	short route announcement
SSPAN	small scale personal area network
TMCTP	television white space multichannel cluster tree personal area network
TxA	transmitting address
TC	topology construction
TTL	time to live
US	upstream

4. General description

4.1 Introduction

Layer 2 routing (L2R) defines a routing protocol for IEEE 802.15.4 networks. An L2R mesh allows increasing the range of a wireless personal area network (WPAN) while maintaining low power transmission. The L2R protocol provides routing for small, medium and large scale networks with tens to thousands of devices.

L2R routing may be used in diverse applications such as smart metering, smart city, and other control and monitoring applications.

The following functionalities are provided in this recommended practice:

- Topology construction
- L2R mesh discovery, join, update, and recovery
- Hop-by-hop retransmission
- Unicast, multicast, and broadcast routing
- Data concatenation (DCat)
- Short address assignment
- Security

L2R enables routing over the MAC and PHY layers available in IEEE Std 802.15.4. The L2R sublayer is deployed over a working IEEE 802.15.4 network. Any device wishing to be part of an L2R mesh shall be part of a personal area network (PAN). The procedures to manage a PAN (start, join, leave, scan...) are described in IEEE Std 802.15.4.

4.2 Format conventions

This document follows the formatting conventions as defined in Clause 4 of IEEE Std 802.15.4-2015.

Unless otherwise stated, the word “integer” refers to “unsigned integer” throughout the document.

4.3 Network topology

An L2R mesh may operate in a variety of topologies. The L2R mesh may operate within a single PAN. It may also operate across multiple PANs in a television white space multichannel cluster tree personal area network (TMCTP).

The L2R mesh is built around a mesh root device which enables access to a service such as a data collection service, control and monitoring service, etc. An L2R mesh is identified by a mesh identifier (ID), and the service(s) available are attached to the mesh ID and known to the higher layers of a device. The mesh root acts as the controller of the L2R mesh and defines the features and the metrics in use. The mesh root is implemented in an FFD, and may optionally also be implemented in the same device as the PAN coordinator.

The mesh root has a direct connection with the PAN coordinator (PanC DC) if it is able to communicate with the PAN coordinator through a network external to the L2R mesh or by being located in the same device. The PanC DC allows communication between a device part of the L2R mesh and the PAN coordinator via the mesh root. In this document, the PanC DC is used to exchange short address assignment (AA) related IEs between L2R devices and the PAN coordinator, or is used by the mesh root to deliver PAN broadcast frames to the PAN coordinator.

An example of the basic topology of an L2R mesh is illustrated in Figure 4-1, where the L2R mesh with the mesh ID L2rMeshId1 enables access to an internet connection service, and the mesh root has a PanC DC.

An L2R mesh may also be deployed over a small scale PAN (SSPAN) defined as a PAN where the mesh root and the PAN coordinator are implemented in the same device and the L2R mesh is connected to a unique service. An SSPAN aims at low overhead L2R operations performed using small sized nested IEs for downstream (DS) route establishment and for routing.




Figure 4-1—L2R mesh topology

Several L2R meshes, with their respective mesh roots, may be deployed over a PAN. A mesh root may provide a connection to several services and the connection to a service may also be provided by several mesh roots. Routing uses extended addresses or short addresses. The mechanism to ensure the uniqueness of extended addresses when they are 64-bit local MAC addresses is out of the scope of this document. As L2R is to be used above the IEEE Std 802.15.4 that operates with 64-bit extended unique identifiers (EUI-64), it is recommended that the implementer also uses EUI-64 for L2R. Short addresses, when used, are allocated and managed by the PAN coordinator. A device may simultaneously belong to different L2R meshes. The device chooses which L2R mesh to use depending on the service it requires. Figure 4-2 illustrates an example of multiple L2R meshes deployed over a PAN providing two services: internet and data collection. The L2R mesh with the ID L2rMeshId1 managed by mesh root 1 provides both services while the L2R mesh with the ID L2rMeshId2 managed by mesh root 2 only provides the data collection service.




Figure 4-2—Multiple L2R meshes deployed over a PAN

An L2R mesh may also span multiple PANs organized in different neighboring clusters such as in a TMCTP defined in IEEE Std 802.15.4. An L2R mesh deployed over multiple PANs is described in 5.4.1.3. EUI-64 addresses are used in a TMCTP to avoid duplicate addresses across the different PANs.

4.4 Architecture

This document specifies the L2R sublayer above the MAC sublayer defined in IEEE Std 802.15.4. The L2R sublayer interacts with the MAC sublayer through the MAC sublayer management entity (MLME) service access point (SAP) and through the MAC common part sublayer (MCPS)-SAP. This document also provides two interfaces to allow the higher layers to interact with the L2R sublayer: the L2R sublayer management entity (L2RLME)-SAP to initiate operations such as discovering, starting or joining an L2R mesh, and the L2RDATA-SAP for data routing services. The L2R sublayer does not interact with the PHY layer.

Figure 4-3 illustrates the system architecture.




Figure 4-3—System architecture

4.5 Functional overview

4.5.1 L2R mesh and routing

An L2R mesh may be a part of a PAN or may span several PANs. The formation of a PAN is described in IEEE Std 802.15.4.

An L2R mesh enables access to one or more services corresponding to their mesh ID. A device that can provide access to a particular service may optionally become a mesh root. A device that needs to communicate with a service joins the corresponding L2R mesh. Then in turn, the device informs its neighbors of the existence of the L2R mesh it has just joined.

An L2R mesh is organized hierarchically. The hierarchy is defined by the depth of each device in the L2R mesh, in other words the relative distance of a device to a mesh root (which has a depth 0) with respect to a metric. Ancestor-descendant relationships and sibling relationships are established based on the depth of the devices. Topology construction is described in 5.2.1.

This recommended practice allows the following types of routing:

- Unicast routing: routing from one source device to one destination device. Unicast routing can be performed in the following ways:
 - a) Upstream (US): routing from a device to the mesh root
 - b) Downstream (DS): routing from the mesh root to a device
 - c) Peer-to-peer (P2P): routing between two devices not necessarily traversing the mesh root
- Multicast routing: routing from one source device to multiple destination devices belonging to a multicast group.
- Broadcast routing: routing from one source device to all the devices within a mesh, or within a PAN (if PAN broadcast is enabled)

US routing is performed on a hop-by-hop basis by forwarding a frame from a device to an ancestor. A device only uses information about its neighbors' depth and metric to perform US routing. US route establishment is described in 5.2.3.

DS routing is performed by forwarding a frame from a device to a descendant based on routing information advertised beforehand. A device selects a next hop towards the destination by evaluating its neighbors based on a metric determined by the mesh root. Source routing may also be used for DS route establishment. DS route establishment is described in 5.2.4.

A device able to perform routing via a sibling (RvS) may optionally use a sibling for US and DS routing if allowed by the device originating the frame and if the device supports loop avoidance.

P2P routing may be achieved through a combination of US and DS routing. When a device needs to send a frame to another device other than the mesh root, it first checks if it has a route to the destination within its DS routing information. If a DS route is available, the frame is forwarded using the found path. If no DS route is found, the frame is forwarded US until the first L2R router, including the mesh root, that has a DS route to the destination. The frame is then forwarded DS to the destination.

On-demand P2P routing may also be performed. In this case the source device starts a path discovery towards the destination. If a path is established between the source and destination, the path is used to route frames between the two devices. Otherwise, the frame is routed through US and DS routing.

Multi-PAN operation (MPO) is used if an L2R mesh spans multiple PANs, each operating in a single channel.

Hop-by-hop retransmission of a data frame at the L2R sublayer may optionally be performed in case of failure at the MAC sublayer.

A source device may request an end to end acknowledgment from a destination device to ensure the successful transmission of a data frame. Otherwise, the source device assumes that all its transmissions are successful.

Data concatenation may be used for the transmission of multiple data frames intended for the same destination. This allows traffic reduction in the mesh and is recommended for applications such as data collection or monitoring where a limited number of devices collect information from the rest of the devices in an L2R mesh.

The different types of PAN, the route establishment features and the routing features are summarized in Table 4-1, Table 4-2 and Table 4-3 respectively. The relevant clause for each feature is enclosed in parenthesis.

The notations used in Table 4-1, Table 4-2 and Table 4-3 are interpreted as follows:

- Y: the feature shall be used
- O: the use of the feature is optional
- —: the feature is not relevant to the routing mode
- O.1, O.2: One or the other of these features shall be used

Table 4-1—PAN type

PT1	Regular PAN	O.1
PT2	SSPAN	O.1
PT3	TMCTP	O.1

Table 4-2—Features for route establishment

Function		Mesh root centralized path establishment	P2P	Multicast
RE1	TC IE based US route establishment	M (5.2.3)	—	—
RE2	RA IE based DS route establishment (RE1 required to use this function)	SM1: O (5.2.4) SM2: M (5.2.4)	—	O (5.2.6)
RE3	L2R Routing IE based DS route establishment	SM1: M (5.2.4) SM2: —	—	—
RE4	P2P route establishment	—	M (5.2.7)	—
SM1	Storing mode	O.2 (5.2.4.1)	O.2 (5.2.7.1)	—
SM2	Non-storing mode	O.2 (5.2.4.2)	O.2 (5.2.7.2)	—

NOTE — Broadcast is not listed in Table 4-2 as broadcast routing does not require route establishment.

Table 4-3—Features for routing

Function		Unicast US	Unicast DS	P2P	Multicast	Broadcast
R1	Source routing	—	SM1: — SM2: M (5.4.1.1)	SM1: — SM2: M (5.4.1.1)	—	—
R2	Hop-by-hop routing	M (5.4.1.2)	SM1: M (5.4.1.2) SM2: —	SM1: M (5.4.1.2) SM2: —	—	—
R3	L2R multicast routing	—	—	—	PT1: O (5.4.2.1) PT2, PT3: —	—
R4	Flooding	—	—	—	PT1: O (5.4.2.2) PT2, PT3: M (5.4.2.2)	M (5.4.3)
RvS	Routing via a sibling	O (5.4.1.2)	O, only when used in US routing (5.4.1.2)	—	—	—
HR	Hop-by-hop retransmission	O (5.4.1.4)	—	—	—	—

Table 4-3—Features for routing (continued)

Function		Unicast US	Unicast DS	P2P	Multicast	Broadcast
E2EA	End to end acknowledgment	PT1, PT3: O (5.4.1.5) PT2: —			—	—
DCat	Data concatenation	O (5.4.1.6)			—	—
MPO	Multi-PAN operation	O (5.4.1.3)			—	O (5.4.3.3)

The L2R mesh may be used in an energy constrained network where the devices work on a duty cycle. A device may use a low energy scheme such as coordinated sampled listening (CSL) or receiver initiated transmission (RIT). A device should synchronize with the awake time of a duty cycling next hop in order to forward a frame.

4.5.2 Addressing modes

This document specifies the use of short and extended addresses as found in Table 4-4.

Table 4-4—Addressing mode usage

Routing mode	Network type			
	PAN (1 or more meshes)		SSPAN (1 mesh)	TMCTP (1 mesh)
	Mesh operating with short addresses	Mesh operating with extended addresses		
Unicast (US/DS/P2P)	Short address	Extended address	Short address OR extended address	Extended address
Multicast	Short address (0xff00 – 0xfffd)		Short address (0xff00 – 0xfffd) ^a	Short address (0xff00 – 0xfffd) ^b
Broadcast	Short broadcast address	Short or extended broadcast address	Broadcast address in the same mode as that of the mesh root	Extended broadcast address

^aMulticast routing is performed through flooding with filtering by the next higher layer as described in 5.4.2.2 and is only supported in SL2R meshes operating with short addresses.

^bMulticast routing is performed through flooding with filtering by the next higher layer as described in 5.4.2.2.

5. L2R functional description

5.1 Starting and joining an L2R mesh

Aside from end devices and routers, an L2R mesh comprises two special components, namely a PAN coordinator and an L2R mesh root.

A PAN coordinator is defined as described in IEEE Std 802.15.4.

An L2R mesh root is the primary controller of an L2R mesh. A PAN coordinator usually acts as the L2R mesh root. Alternatively, the function of L2R mesh root may be implemented in one or more devices other than the PAN coordinator.

There are two types of joining devices: the L2R router device and the L2R end device. The L2R router device has the ability to forward frames and is implemented in an FFD as specified in IEEE Std 802.15.4. The L2R end device does not have the ability to route frames and it may be implemented in an FFD or a reduced-function device (RFD). An L2R router device may optionally become the L2R mesh root.

The mesh root may start an L2R mesh in one of the PAN types listed in Table 4-1 in 4.5.1. A joining device discovers the type of PAN on which an L2R mesh is running during the discovery phase described in 5.1.2.1.

5.1.1 Starting and stopping an L2R mesh

5.1.1.1 Procedure to start an L2R mesh

Before starting an L2R mesh, the next higher layer of a device initiates the discovery procedure described in 5.1.2.1. The device may decide to become a mesh root if it did not discover an existing L2R mesh with the same mesh ID as the L2R mesh it intends to start. In this case, the device first associates with a PAN as described in IEEE Std 802.15.4 before starting the L2R mesh. Otherwise, the device joins the existing L2R mesh. A device should become the mesh root of only one L2R mesh at a time.

If the L2R mesh to start is an SSPAN L2R (SL2R) mesh, the PAN coordinator is the only device allowed to initiate its construction. If a device is already a part of an SL2R mesh, it shall not start another L2R mesh.

An L2R mesh is started when the L2R sublayer in the L2R router that is to become the mesh root, receives the L2RLME-MESH-START.request primitive from the next higher layer. The L2R sublayer then initializes a mesh table (MT) described in Table 5-1. The mesh root address is set to the device's own address. The mesh ID and the mesh ID encoding are set to the values of MeshId and MeshIdEncoding from the primitive respectively. My depth is set to zero. The mesh sequence number (MSN) is set to an initial value between 0xf0 and 0xff as described in 5.2.1. The local neighbor table (NT) is initialized as an empty table. The value of My PQM is initialized to zero. If On-demand P2P route discovery is TRUE, the P2P path list is initialized as an empty table. Other parameters in the MT are set as indicated by the parameters in the primitive. The L2R sublayer adds a new L2R mesh descriptor to *l2rMeshDescriptorList* for the L2R mesh the device is starting. The elements of the new L2R mesh descriptor that are not set from the values of the L2RLME-START-MESH.request primitive are set to default values. At the end of the procedure to start the L2R mesh, the next higher layer may set these elements to different values. The L2R sublayer then starts the periodic transmission of enhanced beacons (EB) containing a topology construction information element (TC IE) with the frequency indicated by *l2rTcIeInterval*. If the L2R router has a direct connection to the PAN coordinator (PanC DC), it sets the PanC DC field in the Descriptor field of the TC IE to 1. The Descriptor field of the TC IE is illustrated in Figure 6-5. Upon successful completion of this start-up procedure, the L2R router becomes the mesh root and the L2R sublayer sends an L2RLME-MESH-START.confirm primitive, with a status SUCCESS, to the next higher layer. The TC IE format is described in 6.1.2; the

L2RLME-MESH-START.request and L2RLME-MESH-START.confirm primitives are described in 7.1.1.4 and 7.1.1.5 respectively. The procedure is illustrated in Figure 5-1.

A single address mode is used in the L2R mesh, set by the L2R mesh root as indicated by the AddressingMode in the L2RLME-MESH-START.request primitive. If short addressing is to be used in an L2R mesh and short address assignment (AA) is not managed by a higher layer, the short addresses should be managed according to the process described in 5.1.2.5 and the corresponding L2R mesh root should have a short address. If its short address has expired, the mesh root should stop its L2R mesh.




Figure 5-1—Message sequence chart to start an L2R mesh

5.1.1.2 Procedure to update the L2R mesh topology parameters

If a mesh root needs to modify the parameters currently in use in its L2R mesh (path quality metric (PQM), storing mode, use of DCat, service(s) available, etc.), the next higher layer informs the L2R sublayer of the desired modifications with an L2RLME-MESH-START.request primitive. The MT is updated according to the parameters of the primitive and the mesh root transmits an updated TC IE with an MSN between 0xf0 and 0xff. The L2R sublayer reports the result of the restart procedure with the L2RLME-MESH-START.confirm primitive.

5.1.1.3 Procedure to stop an L2R mesh

A mesh root should stop its mesh if it is no longer able to provide its services. To stop an L2R mesh, the next higher layer issues the L2RLME-MESH-STOP.request primitive. If SoftStop in the primitive is TRUE, the L2R sublayer transmits the next scheduled TC IE with the Depth field set to 0xff and suspends the operation. The mesh root may optionally transmit several new TC IEs successively with the same MSN to ensure the notification of all its descendants before suspending the operation. The number of TC IEs to be sent is

determined by the next higher layer and is out of scope of this specification. If SoftStop is FALSE, the updated TC IE is transmitted and the operation is suspended immediately. Upon successful completion of the stop procedure, the L2R sublayer sends an L2RLME-MESH-STOP.confirm primitive, with a SUCCESS status, to the next higher layer. This procedure is illustrated in Figure 5-2. The TC IE is described in 6.1.2. The L2RLME-MESH-STOP.request and the L2RLME-MESH-STOP.confirm primitives are described in 7.1.1.6 and 7.1.1.7, respectively.




Figure 5-2—Message sequence chart to stop an L2R mesh

5.1.2 Joining an L2R mesh

5.1.2.1 Mesh discovery procedure

A device intending to join or to start an L2R mesh performs the mesh discovery procedure summarized in Figure 5-3.




Figure 5-3—Message sequence chart to discover an L2R mesh

The next higher layer of a joining device issues the L2RLME-PAN-SCAN.request primitive to request the broadcast of an enhanced beacon request (EBR) with an L2R Discovery (L2R-D) IE where the Content field is set as described in 5.1.2.1.1 or 5.1.2.1.2. The L2R-D IE is defined in 6.1.1. The scan procedure is performed on the channels indicated in L2RLME-PAN-SCAN.request primitive. The next higher layer indicates which L2R mesh it wishes to discover with the parameter MeshId of the primitive. The L2R-D IE is sent in an EBR with the Destination PAN Identifier and the Destination Address fields set to 0xffff. The L2RLME-PAN-SCAN.request primitive is described in 7.1.1.1.

A device that is to use L2R should associate with a PAN that supports L2R as indicated by L2R-D IEs received from at least one of the existing devices. An L2R mesh may be deployed over a regular PAN, an SSPAN or a TMCTP. The PAN coordinator, or the super PAN coordinator (SPC) and the PAN coordinators in a TMCTP, should be configured to know the type of PAN it is running if L2R is to be used. In an SSPAN, the PAN coordinator is the only device allowed to start an L2R mesh. In a TMCTP, the SPC is the only device allowed to start an L2R mesh. Therefore, when an FFD that can act as a coordinator (including the PAN coordinator) that does not belong to any L2R mesh receives the EBR with the L2R-D IE, it replies with an EB with an L2R-D IE where the Mesh Root Address field is omitted. In a regular PAN, all the flags of the Descriptor field of the L2R-D IE are set to 0. In an SSPAN, the SSPAN field is set to 1. In a TMCTP, the MPO field is set to 1. This is to allow the joining devices to know the type of PAN that is running before association.

5.1.2.1.1 Discovery of L2R meshes with a specific mesh ID

If MeshId is a string with a nonzero length, the L2R sublayer should attempt to discover the L2R mesh identified by MeshId. The device sends an EBR with an L2R-D IE where the Mesh ID Present field is set to 1 and the Mesh ID Information field is set according to MeshId and MeshIdEncoding from the

L2RME-PAN-SCAN.request primitive. The remaining flags of the Descriptor field are set to 0 and ignored by the receiver.

When a device receives the L2R-D IE in an EBR, the value of the Mesh ID field is compared to the MeshId of the L2R mesh(es) that the receiving device has joined. If the device belongs to an L2R mesh with the requested MeshId, it responds with an L2R-D IE in an EB. The device sets the Mesh ID Present field and the Mesh Root Present field to 1 and sets the Mesh ID field to the requested MeshId. The remaining fields are set based on the information from the corresponding MT. If the device does not belong to an L2R mesh with the requested MeshId, it discards the L2R-D IE.

If the *macAutoRequest* MAC PIB attribute is set to FALSE, the L2R sublayer is notified with an MLME-BEACON-NOTIFY.indication primitive upon receiving each EB frame. In this case, the L2R sublayer creates an MT corresponding to each distinct L2R mesh with the requested MeshId if one has not been created yet. The L2R sublayer issues an L2RLME-PAN-SCAN.indication primitive to the next higher layer after receiving each MLME-BEACON-NOTIFY.indication primitive from the MAC sublayer. After the scan is completed the L2R sublayer issues the L2RLME-PAN-SCAN.confirm primitive with an empty ScanResultList.

If *macAutoRequest* is set to TRUE, the L2R sublayer is notified of all the scan results with the MLME-SCAN.confirm primitive from the MAC sublayer at the end of the scan. If the L2R sublayer receives one or more L2R-D IEs, it creates an MT for each distinct L2R mesh with the requested MeshId. The L2RLME-PAN-SCAN.confirm primitive is returned with a Status SUCCESS and with a ScanResultList corresponding to the mesh ID of interest.

The joining device records a list of mesh root addresses corresponding to each mesh ID of the L2R mesh(es) discovered based on the received L2R-D IEs in the *l2rMeshIdMeshRootList*.

If a device receives an L2R-D IE where the Mesh ID field does not match the requested MeshId, it then discards the L2R-D IE.

If an L2R mesh identified by MeshId is not found at the end of the scan, the L2RLME-PAN-SCAN.confirm primitive is returned with a Status MESH_NOT_FOUND and with an empty ScanResultList.

5.1.2.1.2 Discovery of all existing meshes

If MeshId is a string with a zero length, the L2R sublayer discovers all the existing meshes in the device's vicinity. The device sends an EBR with an L2R-D IE where all the flags of the Descriptor field are set to 0. All neighbors receiving the IE respond with an EB containing an L2R-D IE set according to their respective L2R meshes. If a device belongs to more than one mesh, it responds with the same number of L2R-D IEs successively.

If the *macAutoRequest* MAC PIB attribute is set to FALSE, the L2R sublayer is notified with an MLME-BEACON-NOTIFY.indication primitive upon receiving each EB frame. In this case, the L2R sublayer creates an MT corresponding to each distinct mesh root address. The L2R sublayer then issues an L2RLME-PAN-SCAN.indication primitive to the next higher layer after receiving each MLME-BEACON-NOTIFY.indication primitive from the MAC sublayer. After the scan is completed the L2R sublayer issues the L2RLME-PAN-SCAN.confirm primitive with an empty ScanResultList.

If *macAutoRequest* is set to TRUE, the L2R sublayer is notified of all the scan results with the MLME-SCAN.confirm primitive from the MAC layer at the end of the scan. If the L2R sublayer receives one or more L2R-D IEs, it creates an MT for each distinct mesh root address. The L2R sublayer then submits the L2RLME-PAN-SCAN.confirm primitive with the available coordinator candidates in the ScanResultList parameter. After the discovery, the device associates with a PAN containing at least one L2R mesh providing

the desired service. This procedure is illustrated in Figure 5-3. The L2RLME-PAN-SCAN.confirm and the L2RLME-PAN-SCAN.indication are described in 7.1.1.2 and 7.1.1.3, respectively.

5.1.2.1.3 Discovery within a PAN

If a device is already associated with a PAN, it may also discover the L2R mesh(es) deployed within its PAN. In this case, the device sends the L2R-D IE on the channel of the current PAN within an EBR where the Destination PAN Identifier field is set to the current PAN ID and the Destination Address field is set to 0xffff to allow a response from all potential neighbors. This procedure is illustrated in Figure 5-4.




Figure 5-4—Message sequence chart to discover an L2R mesh within the associated PAN

5.1.2.2 Procedure to join an L2R mesh

A mesh root manages one L2R mesh at a time; therefore once the discovery procedure described in 5.1.2.1 is completed, the L2R mesh is identified by the mesh root address used in the joining procedure and for all other L2R operations. A device joins an L2R mesh after the association following the discovery procedure.

5.1.2.2.1 Mesh selection by the L2R sublayer

If *l2rMeshSelection* is TRUE, mesh selection is handled by the L2R sublayer. When a device wishes to join a mesh, the next higher layer issues the L2RLME-JOIN-MESH.request primitive to request the L2R sublayer to join a mesh managed by one of the mesh roots found in the MeshRootList indicated in the primitive. Upon reception of this primitive, the L2R sublayer initiates an enhanced active scan and broadcasts an EBR with a TC IE with an empty Content field. The TC IE is defined in 6.1.2. When an L2R router receives the EBR with an empty TC IE, it immediately replies with an EB containing a non-empty TC

IE, and then resumes its regular periodic TC IE transmissions. When the joining device receives the response TC IE and if the mesh root address therein corresponds to the mesh root address of a recorded MT, the device computes its own depth and PQM as described in 5.2.1 and updates an MT entry related to the L2R mesh advertised in the TC IE. The device also creates or updates the global NT entry for the neighbor transmitting the TC IE. At the end of the scan, the L2R sublayer selects the L2R mesh providing the best PQM value among the L2R meshes whose mesh root are found in the MeshRootList. If the device discovers an L2R mesh operating with short addresses where the mesh root does not have a PanC DC, as indicated by the PanC DC field of the TC IE set to 0, the device should join this L2R mesh only if it has already been assigned with a short address. If multiple meshes with the same mesh ID are discovered during the procedure described in 5.1.2.1, and have different PQM IDs, the algorithm to select the L2R mesh is out of the scope of this document. The L2R sublayer adds a new L2R mesh descriptor to *l2rMeshDescriptorList* for the L2R mesh the device is joining. The elements of the L2R mesh descriptor that are not set from the content of the TC IE are set to default values. After the joining procedure the L2R mesh, the next higher layer may set these elements to different values. The device is allowed to join an L2R mesh if its depth does not exceed the value in the L2R Max Depth field of the TC IE. The device deletes unnecessary MTs, MT entries or MT entry elements and global NT entries or elements according to the condition to record each element as described in 5.2.1. The device then transmits its own TC IE if it is an L2R router as indicated by a non-zero *l2rTcleInterval*. The L2R sublayer sends an L2RLME-JOIN-MESH.confirm primitive with a SUCCESS Status to the next higher layer. This procedure is illustrated in Figure 5-5.

If no TC IE is received during the scan or if no mesh satisfies the requirements set by the L2RLME-JOIN-MESH.request primitive, the L2R sublayer may reattempt to trigger an enhanced active scan to find the desired L2R mesh up to *l2rMaxScanRetry* times. The L2RLME-JOIN-MESH.request and L2RLME-JOIN-MESH.confirm primitives are described in 7.1.1.8 and 7.1.1.9 respectively.

If *l2rSwitchMesh* is TRUE, after having joined an L2R mesh, if a device receives a TC IE that is not encrypted or that is encrypted with a known key from one of the L2R meshes recorded in an *l2rMeshIdMeshRootList* entry with the same mesh ID as one of the currently joined L2R meshes with a better PQM, the L2R sublayer disconnects from the current L2R mesh and joins the new one.

If *l2rSwitchMesh* is FALSE, a device should not switch between L2R meshes.

If the joining device is an L2R router, the L2R sublayer starts the periodic transmission of EBs with a TC IE at an interval of *l2rTcleInterval*. The value of *l2rTcleInterval* may differ from device to device depending on their sleeping or duty cycling patterns. These patterns are determined by the low energy mechanism used at the MAC sublayer defined in IEEE Std 802.15.4.

If the RA IE Required field of the TC IE Descriptor field is set to 1, L2R routers and end devices should transmit MP frames with a route announcement IE (RA IE) or a short route announcement IE (SRA IE) periodically at an interval of *l2rRaleInterval* to establish DS routes. *l2rTcleInterval* and *l2rRaleInterval* have a common unit set by *l2rIntervalUnit*.




Figure 5-5—Message sequence chart to join an L2R mesh

5.1.2.2.2 Mesh selection by the next higher layer

If *l2rMeshSelection* is FALSE, when a device wishes to join a mesh, the next higher layer issues the L2RLME-MESH-DISCOVERY.request primitive to request the L2R sublayer to discover the L2R meshes around the joining device. Upon reception of this primitive, the joining device initiates an enhanced active scan and broadcasts an EBR with a TC IE with an empty Content field. When an L2R router receives the EBR with a TC IE with an empty Content field, it replies with an EB containing a non-empty TC IE. When the joining device receives the response TC IE and if the mesh root address therein corresponds to the mesh root address of a recorded MT, it computes its own depth and PQM as described in 5.2.1. The device updates an MT entry related to the L2R mesh advertised in the TC IE. The device also creates or updates the global NT entry for the neighbor transmitting the TC IE. At the end of the scan, the L2R sublayer sends an L2RLME-MESH-DISCOVERY.confirm primitive with a SUCCESS Status to the next higher layer. The next higher layer selects the L2R mesh to join based on the information in the MT stored in the L2R sublayer and informs the L2R sublayer by issuing the L2RLME-MESH-SELECT.request primitive. If the device discovers an L2R mesh operating with short addresses where the mesh root does not have a PanC DC, the next higher layer should join this L2R mesh only if the device has already been assigned with a short address. The device deletes unnecessary MTs, MT entries or MT entry elements, and global NT entries or

elements according to the condition to record each element as described in 5.2.1. The L2R sublayer adds a new L2R mesh descriptor to *l2rMeshDescriptorList* for the L2R mesh to join. The elements of the L2R mesh descriptor that are not set from the content of the TC IE are set to default values. At the end of the procedure to join the L2R mesh, the next higher layer may set these elements to different values. The device then transmits its own TC IE. The L2R sublayer sends an L2RLME-MESH-SELECT.confirm primitive with a SUCCESS Status to the next higher layer. This procedure is illustrated in Figure 5-6.

If *l2rSwitchMesh* is TRUE, whenever a device receives and is able to decrypt a TC IE from one of the L2R meshes recorded in an *l2rMeshIdMeshRootList* entry with the same mesh ID as one of the currently joined L2R meshes with a better PQM, the L2R sublayer informs the next higher layer with the L2RLME-NOTIFY.indication primitive where the Notification is set to BETTER_MESH_DETECT. Then the next higher layer decides whether or not to request the L2R to disconnect from the current L2R mesh and join to the new L2R mesh. The L2RLME-MESH-DISCOVERY.request, L2RLME-MESH-DISCOVERY.confirm, L2RLME-MESH-SELECT.request and L2RLME-MESH-SELECT.confirm are described in 7.1.1.10, 7.1.1.11, 7.1.1.12 and 7.1.1.13, respectively.

If the joining device is an L2R router, the L2R sublayer starts the periodic transmission of EBs with a TC IE at an interval of *l2rTcleInterval*. The value of *l2rTcleInterval* may differ from device to device depending on their sleeping or duty cycling patterns.

If the RA IE Required field of the TC IE Descriptor field is set to 1, L2R routers and end devices transmit MP frames with an RA IE or an SRA IE periodically at an interval of *l2rRaleInterval*.




Figure 5-6—Message sequence chart for mesh selection by the next higher layer

5.1.2.3 Procedure to rejoin an L2R mesh

If a device has no neighbor left in the local NT of an MT, it determines that it has been disconnected from the corresponding L2R mesh. Upon detection of the disconnection, the L2R sublayer issues an L2RLME-DISCONNECT-MESH.indication primitive to the next higher layer. The L2RLME-DISCONNECT-MESH.indication primitive is described in 7.1.1.16.

The device should rediscover an L2R mesh according to the procedure described in 5.1.2.1.1 or 5.1.2.1.2 and associate with the appropriate PAN.

If the device wishes to remain within the same PAN, it should try to rediscover the L2R meshes within its PAN according to the procedure described in 5.1.2.1.3.

Then the device tries to join the L2R mesh according to the procedure described in 5.1.2.2.1 if *l2rMeshSelection* is TRUE, or according to the procedure described in 5.1.2.2.2 if *l2rMeshSelection* is FALSE.

If the device has lost connection with its coordinator, it should reestablish the connection to the PAN before trying to rejoin the L2R mesh. This procedure is illustrated in Figure 5-7.




Figure 5-7—Message sequence chart to rejoin an L2R mesh

5.1.2.4 Procedure to leave an L2R mesh

The next higher layer issues the L2RLME-LEAVE-MESH.request primitive to disconnect from an L2R mesh. The L2R sublayer transmits its next scheduled EB with a TC IE with a depth 0xff. Upon successful

completion of the leave procedure, the L2R sublayer sends an L2RLME-LEAVE-MESH.confirm primitive, with a status SUCCESS, to the next higher layer. The procedure is illustrated in Figure 5-8. The L2RLME-LEAVE-MESH.request and the L2RLME-LEAVE-MESH.confirm primitives are described in 7.1.1.14 and 7.1.1.15, respectively.




Figure 5-8—Message sequence chart to leave an L2R mesh

5.1.2.5 Short address management

If an L2R mesh operates using short addresses but short addresses are not managed by a higher layer, they should be managed using the procedures described in this clause. As the maximum *macResponseWaitTime* allowed between an Association Request frame and an Association Response frame at the MAC sublayer does not provide enough time to request a short address to the PAN coordinator over multiple hops, short address assignment is conducted after association to the PAN. Short address assignment (AA) management is carried out using AA request (AA-RQ), AA reply (AA-RP) and address release (A-RLS) IEs. These IEs are transmitted within an MP frame along with an L2R Routing IE in order to route them to the mesh root. The AA-RQ IE, the AA-RP IE and the A-RLS IE are described in 6.1.9, 6.1.10, and 6.1.11, respectively. Short addresses are managed by the PAN coordinator in order to avoid duplicate addressing. Therefore, all three aforementioned AA IEs shall be transmitted to a mesh root with a PanC DC as indicated by the PanC DC with a value TRUE in the MT.

If multiple PANs are present in the L2R mesh, L2R routing uses extended addresses.

5.1.2.5.1 Initial short address assignment

If a device wants to be assigned a short address, its next higher layer sends an L2RLME-AA-RQ.request primitive to the L2R sublayer. Upon receiving the primitive, the L2R sublayer issues the MCPS-DATA.request primitive to trigger the transmission of an MP frame with an address assignment request (AA-RQ) IE along with an L2R Routing IE to the coordinator with which it is associated, which is also an L2R router within the L2R mesh. The device sets the Allocated Address field of the AA-RQ IE to a short address that it wishes to be assigned or that it has already been assigned. The Allocated Address field is set to 0xffff if no specific address is requested. The Expiration Time field is set to the duration for which the device expects to use the short address, or to 0x00 for an indefinite duration.

If the coordinator receives the AA-RQ IE, the L2R sublayer issues the MCPS-DATA.request primitive to initiate the transmission of an MP frame with the received AA-RQ IE along with an L2R Routing IE to a mesh root with a PanC DC. When the mesh root receives the AA-RQ IE, the L2R sublayer issues an L2RLME-AA-RQ.indication to the next higher layer which delivers the AA-RQ IE to the PAN coordinator which manages and assigns all the short addresses. The mesh root may communicate with the PAN coordinator using the mesh or through another communication channel. If the short address retrieved from the Allocated Address field is available, the PAN coordinator records the AA. If the requested short address is not available, the PAN coordinator allocates a new short address and records the AA. If there is no available short address, an error notification is sent to the mesh root. If the expiration time requested by the source of the AA-RQ IE is allowed, the PAN coordinator records the expiration time of the AA, otherwise, the expiration time is set to the maximum expiration time allowed. The result of the AA is sent back to the mesh root whose next higher layer issues the L2RLME-AA-RP.request primitive to trigger the transmission of an AA reply (AA-RP) IE to the coordinator with which the source of the AA-RQ IE is associated. If the AA-RP IE is successfully transmitted, the L2R sublayer issues an L2RLME-AA-RP.confirm to the next higher layer with status SUCCESS. Upon receiving the AA-RP IE, the coordinator forwards the IE to the device identified by the address found in the Joining Device Extended Address field. RFD devices are only able to communicate with their respective coordinator as defined in IEEE Std 802.15.4, therefore if the joining device is an RFD, the coordinator records the assigned short address for subsequent communication with the device. When the joining device receives the AA-RP IE from its coordinator, the L2R sublayer issues an L2RLME-AA-RQ.confirm primitive with a status SUCCESS to the next higher layer. This procedure is illustrated in Figure 5-9. The L2RLME-AA-RQ.request and the L2RLME-AA-RQ.confirm primitives are described in 7.1.2.1 and 7.1.2.2, respectively. The L2RLME-AA-RQ.indication is described in 7.1.2.3. The L2RLME-AA-RP.request and the L2RLME-AA-RP.confirm primitives are described in 7.1.2.4 and 7.1.2.5, respectively.




Figure 5-9—Message sequence chart for initial short AA

5.1.2.5.2 Short address maintenance

If a device wishes to continue to use a short address, it repeats the short AA procedure described in 5.1.2.5.1 before the expiration of the short address. However, the AA-RQ IE and AA-RP IE are exchanged between the device and the mesh root directly without requiring the coordinator with which the device associated when joining the PAN to transfer the IEs. The mesh root forwards the AA request to the PAN coordinator and awaits the response before transmitting the AA-RP IE to the device requesting the address maintenance. The choice of delivery method for this additional step is out of scope. This procedure is illustrated in

Figure 5-10. If the short address has not been maintained by a device before its expiration, the device should not operate on the mesh any longer, and the short address is marked as available by the PAN coordinator.




Figure 5-10—Message sequence chart for short address maintenance

5.1.2.5.3 Short address release

If a device no longer needs a previously assigned short address, it informs the PAN coordinator by transmitting an address release (A-RLS) IE through a mesh root with a PanC DC.

The next higher layer of the device issues an L2RLME-ARLS.request primitive to trigger the transmission of an A-RLS IE by the L2R sublayer. If the A-RLS IE is successfully transmitted, the L2R sublayer issues an L2RLME-ARLS.confirm primitive with a status SUCCESS to the next higher layer. Upon receiving the A-RLS IE, the L2R sublayer of the mesh root informs its next higher layer with an L2RLME-ARLS.indication primitive. The mesh root informs the PAN coordinator of the address release. This procedure is illustrated in Figure 5-11. The L2RLME-ARLS.request and the L2RLME-ARLS.confirm primitives are described in 7.1.2.6 and 7.1.2.7. The L2RLME-ARLS.indication primitive is described in 7.1.2.8.




Figure 5-11—Message sequence chart for short address release

5.2 L2R mesh formation and route establishment

Each device generates its own L2R sequence number (LSN) regardless of the number of unique devices with which it wishes to communicate. The LSN is incremented before the transmission of a new IE used by the L2R sublayer requiring a sequence number with the exception of the TC IE.

When the L2R mesh is deployed over an SSPAN, the LSN is one octet long and is indicated by SLSN in the IEs, and SRA IEs use a 7 bit SRA SN if the first bit of the IE is set to 0. Otherwise, the LSN is two octets long.

If on-demand P2P routing is enabled, each device generates a P2P sequence number (PSN), incremented before the transmission of a P2P-RQ IE or a P2P-RP IE as the original source device.

The LSN, the SLSN, the SRA SN (in an SL2R), and the PSN (if on-demand P2P routing is enabled) are permitted to roll over.

5.2.1 L2R mesh construction

An L2R mesh is managed based on the information retrieved from the layer 2 routing information base (L2IB) attributes found in Table 7-38 in 7.3, in the L2R Discovery (L2R-D IE) and in the TC IE. The relevant information is stored in an MT as shown in Table 5-1. A device manages as many MTs as the number of L2R meshes it has joined. A device may optionally manage the MTs other than that of the L2R mesh it has joined if the corresponding TC IEs are not encrypted, or if they are encrypted with a key known by the device. The parameters in each MT are recorded or omitted depending on the condition to record as follows:

- “J” indicates that the parameter is recorded only if the device is a member of the mesh.
- “M” indicates that recording the parameter depends on the *l2rMeshRecordMode* L2IB attribute.
 - If *l2rMeshRecordMode* is set to JOINED_MESH, the parameter is recorded only when the device is a part of the mesh.
 - If *l2rMeshRecordMode* is set to SAME_MESH_ID, the parameter is also recorded in MTs of L2R meshes with the same mesh ID as one of the L2R meshes the device has joined.
 - If *l2rMeshRecordMode* is set to ALL_MESH, the parameter is recorded for all existing L2R meshes. If *l2rMeshRecordMode* is set to ALL_MESH, recording the local neighbor table (NT) depends on *l2rNeighborRecordMode*. If *l2rNeighborRecordMode* is set to JOINED_MESH, the local NT is recorded only in the MT of the L2R mesh the device has joined. If *l2rNeighborRecordMode* is set to SAME_MESH_ID, the local NT is also recorded in the MT of the L2R mesh(es) with the same mesh ID as that of the L2R mesh(es) the device as joined. If *l2rNeighborRecordMode* is set to ALL_NEIGHBORS, the local NT is recorded in all MTs.
- “P” indicates that the parameter is recorded when on-demand P2P routing is in use in the L2R mesh.

Table 5-1—Elements of an MT Entry

Name	Type	Valid range	Description	Retrieved from IE	Condition to record
Mesh ID encoding	ENUMERATION	UTF-8, UNSPECIFIED	Specifies the encoding of Mesh ID.	L2R-D IE	M
Mesh ID	String	Any string	L2R mesh identifier.	L2R-D IE	M

Table 5-1—Elements of an MT Entry (continued)

Name	Type	Valid range	Description	Retrieved from IE	Condition to record
Mesh address mode	Enumeration	SHORT, EXTENDED	Indicates the address mode used in the mesh. If SHORT, short addresses are used. If EXTENDED, extended addresses are used.	L2R-D IE, TC IE	M
Mesh root address	Address	Short or extended address as specified by Mesh address mode	Address of the mesh root of the L2R mesh	L2R-D IE, TC IE	M
MSN	Integer	0x00 – 0xff	Identifies the latest TC IE, used as described in 5.2.1.	TC IE	M
SSPAN	Boolean	TRUE, FALSE	Indicates whether the L2R mesh is an SL2R mesh.	L2R-D IE	M
Local NT	—	As specified in Table 5-2	Table of neighbors belonging to the mesh indicated by Mesh root address.	—	M
PanC DC	Boolean	TRUE, FALSE	Indicates whether the mesh root has a direct connection to the PAN coordinator. (See example in Figure 4-1 or Figure 4-2).	L2R-D IE, TC IE	M
Metric ID	Integer	0x0 – 0xf	Identifies the metric in use in the mesh. The metric identifier values are listed in Table 6-5 in 6.1.2.11.	TC IE	J
LQT	Depends on the metric ID	Depends on the metric ID	Indicates the threshold of the metric that a link shared with a next hop should satisfy.	TC IE	J
My PQM	Depends on the metric ID	Depends on the metric ID	Value of the metric of the current device.	—	J
L2R max depth	Integer	0x00 – 0xff	Indicates the maximum depth allowed in an L2R mesh.	TC IE	M
My depth	Integer	0x00 – 0xfe	Depth of the current device.	—	J
Key exchange mode	Integer	0x00 – 0x02	As described in Table 6-2 in 6.1.1.1.	L2R-D IE	M
DCat	Boolean	TRUE, FALSE	Indicates whether DCat is allowed in the L2R mesh.	L2R-D IE	J

Table 5-1—Elements of an MT Entry (continued)

Name	Type	Valid range	Description	Retrieved from IE	Condition to record
DCat buffering unit	Enumeration	20 MILLISECONDS, SECOND, MINUTE, HOUR	Indicates the unit of the DCat Buffering time.	TC IE	J
DCat buffering time	Integer	0x00 – 0x3f	Duration a frame may be buffered for DCat in the unit indicated by DCat buffering time.	TC IE	J
Multi-PAN operation	Boolean	TRUE, FALSE	Indicates whether MPO is used.	L2R-D IE	M
RA IE Required	Boolean	TRUE, FALSE	Indicates whether all devices are required to send a Route Announcement (RA) IE to build downstream routes.	L2R-D IE	M
Storing mode	Boolean	TRUE, FALSE	Indicates whether the L2R mesh is in storing mode.	L2R-D IE	M
L2R multicast	Boolean	TRUE, FALSE	Indicates whether multicast routing is handled by the L2R sublayer and RA IEs may optionally contain a Multicast Subscription field.	L2R-D IE	M
On-demand P2P route discovery	Boolean	TRUE, FALSE	Indicates whether on-demand P2P discovery is allowed in the L2R mesh.	L2R-D IE	M
P2P path list	List of P2P paths	List of the elements of a P2P path defined in Table 5-3	List of P2P paths available through the current neighbor. Omitted if on-demand P2P route discovery or storing mode are disabled.	P2P-RQ IE, P2P-RP IE	P

When a mesh root initiates an L2R mesh, it transmits a TC IE with the MSN set to any value between 0xf0 to 0xff. The mesh root increments the MSN each time it transmits a new TC IE. The MSN is set and permitted to roll over in the following ways:

- 0xf0 – 0xff: These values are used when the mesh root is initialized. The MSN rolls over to 0x00 after 0xff is used.
- 0x00 – 0xef: These values are used outside the initialization phase. The MSN rolls over to 0x00 after 0xef is used.

If a device receives a TC IE with a MSN between 0xf0 and 0xff, the device determines that the mesh root is in initialization phase. If a device receives a TC IE with a MSN between 0xf0 and 0xff after a TC IE with a MSN between 0x00 and 0xef, the device determines that the mesh root has been reinitialized. The device updates the MT with the matching mesh root address based on the TC IE. If the metric used in the L2R mesh