

IEEE Standard for Local and metropolitan area networks— Part 15.6: Wireless Body Area Networks

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

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

IEEE Std 802.15.6™-2012

29 February 2012

IEEE Standard for Local and metropolitan area networks— Part 15.6: Wireless Body Area Networks

Sponsor

**LAN/MAN Standards Committee
of the
IEEE Computer Society**

Approved 6 February 2012

IEEE-SA Standards Board

Abstract: Short-range, wireless communications in the vicinity of, or inside, a human body (but not limited to humans) are specified in this standard. It uses existing industrial scientific medical (ISM) bands as well as frequency bands approved by national medical and/or regulatory authorities. Support for quality of service (QoS), extremely low power, and data rates up to 10 Mbps is required while simultaneously complying with strict non-interference guidelines where needed. This standard considers effects on portable antennas due to the presence of a person (varying with male, female, skinny, heavy, etc.), radiation pattern shaping to minimize the specific absorption rate (SAR) into the body, and changes in characteristics as a result of the user motions.

Keywords: BAN, body area network, IEEE 802.15.6

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

Copyright © 2012 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 29 February 2012. Printed in the United States of America.

IEEE and 802 are registered trademarks in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

PDF: ISBN 978-0-7381-7206-4 STD97213
Print: ISBN 978-0-7381-7227-9 STDPD97213

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.

Notice and Disclaimer of Liability Concerning the Use of IEEE Documents: IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. 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.

Use of an IEEE Standard is wholly voluntary. IEEE disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon any IEEE Standard document.

IEEE does not warrant or represent the accuracy or content of the material contained in its standards, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained in its standards is free from patent infringement. IEEE Standards documents are supplied “**AS IS**.”

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. 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 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.

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 the official position of IEEE or any of its committees and shall not be considered to be, nor 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 to ensure 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. Any person who would like to participate in evaluating comments or revisions to an IEEE standard is welcome to join the relevant IEEE working group at <http://standards.ieee.org/develop/wg/>.

Comments on standards should be submitted to the following address:

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

Photocopies: Authorization to photocopy portions of any individual standard for internal or personal use is granted by The Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, 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.

Notice to users

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

This document is copyrighted by the IEEE. It is made available 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 this document available for use and adoption by public authorities and private users, the IEEE does not waive any rights in copyright to this document.

Updating of IEEE 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. 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://standards.ieee.org/index.html> or contact the IEEE at the address listed previously. For more information about the IEEE Standards Association or the IEEE standards development process, visit IEEE-SA Website at <http://standards.ieee.org/index.html>.

Errata

Errata, if any, for this and all other standards can be accessed 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 submitted to the IEEE-SA Standards Board for approval, the 802.15 Working Group had the following membership:

Robert F. Heile, *Chair*
Rick Alfvén, *Vice Chair*
Patrick W. Kinney, *Vice Chair, Secretary*
James P. K. Gilb, *Technical Editor*
Clint Chaplin, *Treasurer*

Art Astrin, *Task Group 6 Chair*
Huan-Bang Li, *Task Group 6 Co-Vice Chair*
Ranjeet Patro, *Task Group 6 Co-Vice Chair*
Daniel Lewis, *Task Group 6 Editor*
Jin-Meng Ho, *Task Group 6 MAC and Security Editor*
Anuj Batra, *Task Group 6 Narrowband PHY Editor*
Marco Hernandez, *Task Group 6 UWB PHY Editor*
Seung-Hoon (Shannon) Park, *Task Group 6 HBC PHY Editor*
Igor Dotlic, *Task Group 6 Secretary*

Emad Afifi
Gahng-Seop Ahn
Roberto Aiello
Takamasa Asano
Taehan Bae
Michael Bahr
John Barr
Tuncer Baykas
Philip Beecher
Ghulam Bhatti
Gary Birk
Mathew Boytim
Peter Bradley
Nancy Bravin
David Britz
Monique Brown
Sverre Brubk
Brian Buchanan
John Buffington
Kiran Bynam
Brent Cain
Edgar Callaway
Chris Calvert
Radhakrishna Canchi
Ruben Cardozo
Russell Chandler
Kuo-Hsin Chang
Soo-Young Chang
Hind Chebbo
In-Kyeong Choi
Sangsung Choi
David Cypher
Matthew Dahl
David Davenport
Mark Dawkins
Hendricus De Ruijter
Gang Ding

Michael Dow
Dietmar Eggert
David Evans
Charles S. Farlow
John Farserotu
Kory Fifield
Stanislav Filin
Will Filippo
Michael Fischer
George Flammer
Ryosuke Fujiwara
Noriyasu Fukatsu
Kiyoshi Fukui
John Geiger
Gregory Gillooly
Paul Gorday
Elad Gottlib
Evan Green
Robert Hall
Shinsuke Hara
Hiroshi Harada
Shigekazu Harada
Timothy Harrington
Rodney Hemminger
Ryoichi Higashi
Garth Hillman
Wei Hong
Srinath Hosur
David Howard
Heqing Huang
Sterling Hughes
Jung-Hwan Hwang
Ichirou Ida
Tetsushi Ikegami
Akio Iso
Il Jang
Adrian Jennings

Wuncheol Jeong
Jorjeta Jetcheva
Steven Jillings
Seong-Soon Joo
Tae-Gyu Kang
Shuzo Kato
Jeritt Kent
Prithpal Khakuria
Dae Kim
Dong-Sun Kim
Jinkyong Kim
Youngsoo Kim
Yunjoo Kim
Jeffrey King
Ryuji Kohno
Fumihide Kojima
Bruce Kraemer
Raymond Krasinski
Thomas Kuerner
Masahiro Kuroda
John Lampe
Zhou Lan
Khanh Tuan Le
Cheolhyo Lee
Hoosung Lee
Hyungsoo Lee
Myung Lee
Yuro Lee
Liang Li
Sang-Kyu Lim
Jeremy Link
Lu Liru
Michael Lynch
Robert Mason
Jeff McCullough
Michael McInnis
Michael McLaughlin

Charles Millet
Dino Miniutti
Siamak Mirnezami
Rishi Mohindra
Emmanuel Monnerie
Robert Moskowitz
Hamilton Moy
Peter Murray
Theodore Myers
Ken Naganuma
Chiu Ngo
Estelle Nguyen
Paul Nikolich
John Notor
Jong-Ee Oh
Okundu Omeni
Laurent Ouvry
Hyung-II Park
Taejoon Park
Albert Petrick
Dalibor Pokrajac
Daniel Popa
Steve Pope
Clinton Powell
Richard Powell
Sridhar Rajagopal
Jayaram Ramasastry
Marc Reed
Ivan Reede

Emmanuel Riou
Richard Roberts
Craig Rodine
June Roh
Benjamin Rolfe
Didier Sagan
Kentaro Sakamoto
H. Sanderford
Naotaka Sato
Kamran Sayrafian-Pour
Timothy Schmidl
Michael Schmidt
Jean Schwoerer
Cristina Seibert
Kunal Shah
Steve Shearer
Stephen Shellhammer
Jie Shen
Shusaku Shimada
Chang Sub Shin
CheolHo Shin
Michael Sim
Jonathan Simon
Jaeseung Son
Ho-Jin Song
Paul Stadnik
Rene Struik
Chin-Sean Sum
Hui-Hsia Sung
Gu Sungi

Ronald Tabroff
Kenichi Takizawa
Hirokazu Tanaka
Larry Taylor
James Tomcik
Ichihiko Toyoda
Jana van Greunen
Hartman Van Wyk
Billy Verso
Bhupender Virk
Khurram Waheed
Joachim Walewski
Junyi Wang
Xiang Wang
Andy Ward
Scott Weikel
Nicholas West
Mark Wilbur
Ludwig Winkel
Eun Tae Won
Alan Wong
Tao Xing
Wen-Bin Yang
Yang Yang
Kazuyuki Yasukawa
Kamy Yazdandoost
Kaoru Yokoo
Mu Zhao
Bin Zhen

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

Jon Adams	Noriyuki Ikeuchi	Jayaram Ramasastry
Rajni Agarwal	Paul Isaacs	Ivan Reede
Rick Alfvén	Akio Iso	Maximilian Riegel
Nobumitsu Amachi	Atsushi Ito	Robert Robinson
Butch Anton	Raj Jain	Benjamin Rolfe
Art Astrin	Oyvind Janbu	Randall Safier
Taehan Bae	Junghoon Jee	Didier Sagan
Jay Bain	Tal Kaitz	Kazuyuki Sakoda
Anuj Batra	Shinkyō Kaku	Shigenobu Sasaki
Nancy Bravin	Piotr Karocki	Naotaka Sato
Vern Brethour	Ruediger Kays	Bartien Sayogo
Sverre Brubk	Stuart J. Kerry	Timothy Schmidl
John Buffington	Brian Kiernan	Cristina Seibert
Kiran Bynam	Yongbum Kim	Peretz Shekalim
William Byrd	Patrick W. Kinney	Jie Shen
Edgar Callaway	Bruce Kraemer	Shusaku Shimada
Dave Cavalcanti	Yasushi Kudoh	Gil Shultz
Clint Chaplin	Thomas Kurihara	Jaeseung Son
Keith Chow	Jeremy Landt	Kapil Sood
Charles Cook	Khanh Tuan Le	Robert Soranno
David Davenport	David G. K. Li	Thomas Starai
Mark Dawkins	Jan-Ray Liao	Rene Struik
Joseph Decuir	Arthur Light	Walter Struppler
Wael Diab	Lu Liru	Mark Sturza
Patrick Diamond	HaiTao Liu	Chin-Sean Sum
Thomas Dineen	William Lumpkins	Steven Thoen
Sourav Dutta	Greg Luri	Ichihiko Toyoda
Richard Edgar	Michael Lynch	Anna Urta
John Egan	Elvis Maculuba	Sunil Vadgama
Charles S. Farlow	Michael McInnis	Dmitri Varsanofiev
Avraham Freedman	Michael McLaughlin	Prabodh Varshney
Monisha Ghosh	Apurva Mody	John Vergis
James P. K. Gilb	Kenichi Mori	Billy Verso
Randall C. Groves	Peter Murray	Bhupender Virk
Michael Gundlach	Michael Newman	George Vlantis
Jose Gutierrez	Paul Nikolich	Khurram Waheed
C. Guy	John Notor	Stanley Wang
Rainer Hach	Satoshi Obara	Xiang Wang
Hiroshi Hamano	Okundu Omeni	Andy Ward
Timothy Harrington	Satoshi Oyama	Hung-Yu Wei
Rodney Hemminger	Ranjeet Patro	Andreas Wolf
Marco Hernandez	Clinton Powell	Eun Tae Won
Jin-Meng Ho	Richard Powell	Ariton Xhafa
Srinath Hosur	Venkatesha Prasad	Tao Xing
Heqing Huang	Mohammad Azizur Rahman	Yang Yang
Ichirou Ida	Sridhar Rajagopal	Oren Yuen
Tetsushi Ikegami		Daidi Zhong

When the IEEE-SA Standards Board approved this standard on 6 February 2012, it had the following membership:

Richard H. Hulett, *Chair*
John Kulick, *Vice Chair*
Robert M. Grow, *Past Chair*
Judith Gorman, *Secretary*

Masayuki Ariyoshi
William Bartley
Ted Burse
Clint Chaplin
Wael Diab
Jean-Philippe Faure
Alexander Gelman
Paul Houzé

Jim Hughes
Joseph L. Koepfinger*
David J. Law
Thomas Lee
Hung Ling
Oleg Logvinov
Ted Olsen

Gary Robinson
Jon Walter Rosdahl
Sam Sciacca
Mike Seavey
Curtis Siller
Phil Winston
Howard L. Wolfman
Don Wright

*Member Emeritus

Also included are the following nonvoting IEEE-SA Standards Board liaisons:

Richard DeBlasio, *DOE Representative*
Michael Janezic, *NIST Representative*
Satish K. Aggarwal, *NRC Representative*

Michelle D. Turner
IEEE Standards Program Manager, Document Development

Patricia A. Gerdon
IEEE Standards Program Manager, Technical Program Development

Introduction

This introduction is not part of IEEE Std 802.15.6-2012, IEEE Standard for Local and metropolitan area networks—Part 15.6: Wireless Body Area Networks.

With the decreasing size and increasing capability of electronic devices, thanks to the Moore's Law, it was inevitable that small and portable devices would be developed for communications around human bodies. Some devices are wearable and some are implementable for medical purposes. These devices need to communicate with their remote controllers. IEEE Std 802.15.6-2012 is a standard for short-range, wireless communications in the vicinity of, or inside, a human body (but not limited to humans). It uses ISM and other bands as well as frequency bands in compliance with applicable medical and communication regulatory authorities. It allows devices to operate on very low transmit power for safety to minimize the specific absorption rate (SAR) into the body and increase the battery life. It supports quality of service (QoS), for example, to provide for emergency messaging. Since some communications can carry sensitive information, it also provides for strong security.

Contents

1. Overview	1
1.1 Scope	1
1.2 Purpose	1
2. Normative references.....	2
3. Definitions, acronyms, and abbreviations	2
3.1 Definitions	2
3.2 Special terms.....	7
3.3 Acronyms and abbreviations	7
4. General framework elements.....	10
4.1 General	10
4.2 Network topology	10
4.3 Reference model	11
4.4 Time base.....	12
4.5 MAC and security state diagrams	12
4.6 Security paradigm.....	15
5. MAC frame formats	16
5.1 Conventions	16
5.2 General format.....	17
5.3 Management type frames.....	27
5.4 Control type frames	55
5.5 Data type frames	60
5.6 MAC/PHY Capability fields.....	61
5.7 Information elements.....	65
6. MAC functions	77
6.1 General	77
6.2 Frame processing.....	78
6.3 Access classification and division	88
6.4 BAN creation/operation and node connection/disconnection.....	90
6.5 Random access	92
6.6 Improvised access and unscheduled access	98
6.7 Scheduled access and scheduled-polling access	107
6.8 Access continuation, termination, and timeout.....	110
6.9 MICS band communication.....	116
6.10 Two-hop star topology extension	121
6.11 Clock synchronization and guard time provisioning	132
6.12 Power management.....	140
6.13 Coexistence and interference mitigation.....	142
6.14 MAC/PHY capability handling/interaction and Application Specific IE usage	147
6.15 MAC sublayer parameters	148
7. Security services.....	151
7.1 Security association and disassociation	151
7.2 PTK creation and GTK distribution.....	163
7.3 Message security.....	165
7.4 Optional cipher functions	172

8. Narrowband PHY specification.....	172
8.1 Data-rate-dependent parameters	173
8.2 PLCP preamble.....	176
8.3 PLCP header.....	178
8.4 PSDU	181
8.5 Constellation mapping	185
8.6 General requirements.....	187
8.7 PHY layer timing.....	188
8.8 Transmitter specifications.....	190
8.9 Receiver specifications	194
9. Ultra wideband PHY specification	196
9.1 Definition of hubs and devices	196
9.2 Modes of operation.....	197
9.3 Rules for use of modes and options	197
9.4 Pulse shape option	198
9.5 UWB PHY frame format.....	198
9.6 PSDU construction	198
9.7 PHR construction.....	201
9.8 Synchronization header	204
9.9 IR-UWB symbol structure.....	206
9.10 UWB modulations	209
9.11 IR-UWB PSDU timing parameters.....	215
9.12 Operating frequency bands.....	217
9.13 Transmit spectral mask.....	218
9.14 IR-UWB pulse shapes.....	219
9.15 Type II hybrid ARQ mechanism	224
9.16 FM-UWB.....	228
9.17 General UWB PHY requirements.....	231
9.18 General radio specifications	232
10. Human body communications PHY specification.....	235
10.1 General	235
10.2 HBC packet structure.....	235
10.3 HBC transmitter.....	236
10.4 PLCP Preamble.....	237
10.5 Start frame delimiter and rate indicator	239
10.6 PHY Header.....	242
10.7 PSDU.....	244
10.8 Transmitter specifications.....	247
10.9 Receiver specifications	248
10.10 General requirements.....	248
10.11 PHY layer timing.....	249
Annex A (informative) Bibliography	250
Annex B (informative) Coexistence applicability guide	251
Annex C (informative) Ultra wideband.....	252
Annex D (informative) Features of human body communication	256

IEEE Standard for Local and metropolitan area networks— Part 15.6: Wireless Body Area Networks

IMPORTANT NOTICE: IEEE Standards documents are not intended to ensure safety, health, or environmental protection, or ensure against interference with or from other devices or networks. Implementers 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.

This IEEE document is made available for use subject to important notices and legal disclaimers. These notices and disclaimers appear in all publications containing this document and may be found under the heading “Important Notice” or “Important Notices and Disclaimers Concerning IEEE Documents.” They can also be obtained on request from IEEE or viewed at <http://standards.ieee.org/IPR/disclaimers.html>.

1. Overview

1.1 Scope

This is a standard for short-range, wireless communication in the vicinity of, or inside, a human body (but not limited to humans). It uses existing industrial scientific medical (ISM) bands as well as frequency bands approved by national medical and/or regulatory authorities. Support for quality of service (QoS), extremely low power, and data rates up to 10 Mbps is required while simultaneously complying with strict non-interference guidelines where needed. This standard considers effects on portable antennas due to the presence of a person (varying with male, female, skinny, heavy, etc.), radiation pattern shaping to minimize specific absorption rate (SAR) into the body, and changes in characteristics as a result of the user motions.

1.2 Purpose

The purpose is to provide an international standard for a short-range (i.e., about human body range), low power, and highly reliable wireless communication for use in close proximity to, or inside, a human body. Data rates, typically up to 10Mbps, can be offered to satisfy an evolutionary set of entertainment and healthcare services. Current *personal area networks* (PANs) do not meet the medical (proximity to human tissue) and relevant communication regulations for some application environments. They also do not support the combination of reliability, QoS, low power, data rate, and noninterference required to broadly address the breadth of body area network (BAN) applications.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used, so each referenced document is cited in text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

ETSI EN 301 839-1, V1.3.1 (2009), Electromagnetic Compatibility and Radio Spectrum Matters (ERM); Short Range Devices (SRD); Ultra Low Power Active Medical Implants (ULP-AMI) and Peripherals (ULP-AMI-P) operating in the frequency range 402 MHz to 405 MHz; Part 1: Technical characteristics and test methods.¹

FIPS Pub 186-3 (2009), Digital Signature Standard (DSS).²

FIPS Pub 197 (2001), Advanced Encryption Standard (AES).

IEEE Std 802[®]-2001, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.^{3, 4}

IEEE Std 1363[™]-2000, IEEE Standard Specifications for Public-Key Cryptography.

ISO/IEC 10646:2003, Information Technology—Universal Multiple-Octet Coded Character Set (UCS). Amendment 1, November 2005; Amendment 2, July 2006; Amendment 3, February 2008.⁵

NIST Special Publication 800-38B (2005), Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication.⁶

NIST Special Publication 800-38C (2004), Recommendation for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality.

3. Definitions, acronyms, and abbreviations

For the purposes of this document, the following terms and definitions apply. The *IEEE Standards Dictionary: Glossary of Terms & Definitions* [B1] should be consulted for terms not defined in this clause.^{7, 8}

3.1 Definitions

abbreviated address: A one-octet identifier (ID) selected as an address for a node, hub, or body area network (BAN) in frame exchanges.

active state: An internal power management state that is ready for frame reception and transmission.

¹ ETSI publications are available from the European Telecommunications Standards Institute (<http://www.etsi.org>).

² FIPS publications are available from the National Technical Information Service (<http://www.ntis.gov/>).

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

⁴ The IEEE standards or products referred to in Clause 2 are trademarks owned by The Institute of Electrical and Electronics Engineers, Incorporated.

⁵ ISO/IEC publications are available from the ISO Central Secretariat (<http://www.iso.org/>). ISO publications are also available in the United States from the American National Standards Institute (<http://www.ansi.org/>).

⁶ NIST publications are available from the National Institute of Standards and Technology (<http://www.nist.gov/>).

⁷ The *IEEE Standards Dictionary: Glossary of Terms & Definitions* is available at <http://shop.ieee.org/>.

⁸ The numbers in brackets correspond to those of the bibliography in Annex A.

active superframe: A superframe in which frame transmission typically occurs within the body area network (BAN) of the hub announcing such a superframe.

allocation: One or more time intervals that a node or a hub obtains using an access method for initiating one or more frame transactions. An allocation comprises one or more allocation intervals. Reference to allocation of a node means that the node is the sender or recipient in the allocation.

allocation interval: A continuous time interval in an allocation, comprising one or more consecutive allocation slots. Reference to allocation interval of a node means that the node is the sender or recipient in the allocation interval.

allocation slot: A time unit used to designate the lengths of medium access related time intervals, such as beacon period (superframe) and allocation interval.

beacon: A frame transmitted by a hub to facilitate network management, such as the coordination of medium access and power management of the nodes in the body area network (BAN) of the hub, and to facilitate clock synchronization therein.

beacon period: A repetitive time interval to which medium access is referenced and in which a beacon is transmitted when appropriate, comprising the same number of time units (allocation slots) of equal duration.

bilink: A communications link for transfer of management and data traffic from a hub to a node or/and vice versa.

bilink allocation: An allocation with allocation interval(s) in which a hub or a node initiates one or more frame transactions to transmit management and data traffic to a node or a hub, respectively, and the recipient returns acknowledgment if required, with the provision that the node initiates frame transaction(s) only after receiving a poll from the hub.

connected node: A node that has a connection with a hub.

connection: A relationship between a node and a hub in a body area network (BAN), substantiated by a connected node identification assigned to the node by the hub and by a wakeup arrangement between them, and optionally by one or more scheduled allocations or unscheduled bilink allocations between them.

contended allocation: A non-reoccurring time interval, within a random access phase (RAP) or a contention access phase (CAP), that a node obtains using random access for initiating a frame transaction. A contended allocation is an uplink allocation, suitable for servicing “unpredictable” uplink traffic (for example, due to data rate variations and/or channel impairments).

contention access: An access method, based on carrier sense multiple access with collision avoidance (CSMA/CA) or slotted Aloha access but not both, whereby a node obtains a time interval in a contention access phase (CAP) for initiating one or more frame transactions. As an access method, contention access is synonymous with random access.

contention access phase (CAP): A time span set aside by a hub and announced via a preceding non-beacon frame for contention access to the medium by the nodes in the body area network (BAN) of the hub.

downlink: A communications link for transfer of management and data traffic from a hub to a node.

downlink allocation: An allocation with allocation interval(s) in which a hub initiates one or more frame transactions to transmit management and data traffic to a node and the node returns acknowledgment if required.

entity authentication: Corroboration of the identity of the node or the hub in a security association procedure.

exclusive access phase (EAP): A time span set aside by a hub in a beacon period (superframe) for transfer of the traffic of the highest user priority (UP) (for emergency or medical implant event report).

frame: An uninterrupted sequence of octets delivered by the medium access control (MAC) sublayer to the physical (PHY) layer, or vice versa, within a node or a hub.

hub: An entity that possesses a node's functionality and coordinates the medium access and power management of the nodes in its body area network (BAN).

hub identifier (HID): An abbreviated address of a hub.

improvised access: An access method, based on impromptu polling or posting by a hub, whereby a hub grants to a node or itself a polled or posted allocation, typically outside scheduled downlink and uplink allocations, for initiating one or more frame transactions by the node or hub, respectively.

inactive state: An internal power management state that is not ready for frame reception and transmission.

inactive superframe: A superframe in which no frame transmission occurs within the body area network (BAN) of the hub announcing such a superframe.

information element (IE): An optional part, with variable but self-identifiable length, of some frames.

managed access phase (MAP): A time span set aside by a hub for improvised access, scheduled access, and unscheduled access to the medium by the hub and the nodes in the body area network (BAN) of the hub.

master key (MK): A secret bit string activated or established between a node and a hub in a security association procedure and used to create a pairwise temporal key (PTK) shared between them.

message authentication: Corroboration of the origin of a message in a message transfer.

multi-periodic (m-periodic) allocation: A scheduled allocation or an unscheduled bilink allocation that has allocation intervals reoccurring in every m th beacon period (superframe) with m being an integer larger than one. An m -periodic scheduled allocation is an uplink allocation, a downlink allocation, or a bilink allocation, suitable for servicing low duty cycle periodic or quasi-periodic traffic on a committed schedule. An m -periodic unscheduled allocation is a bilink allocation, suitable for servicing low duty cycle periodic or quasi-periodic traffic on a best-effort basis.

node: An entity that contains a medium access control (MAC) sublayer, a physical (PHY) layer, and that optionally provides security services.

node identifier (NID): An abbreviated address of a single node or of a logical group of nodes.

nonce: A number that is unique per instantiation of a cryptography protocol as part of a measure to thwart cryptanalytic and other cryptographic attacks.

non-secure frame: A term that is interchangeable with unsecured frame.

one-periodic (1-periodic) allocation: A scheduled allocation that has allocation intervals reoccurring in every beacon period (superframe), or an unscheduled bilink allocation that has allocation intervals reoccurring in every beacon period (superframe) or in round-robin together with the allocation intervals of other 1-periodic unscheduled bilink allocations. A 1-periodic scheduled allocation is an uplink allocation, a

downlink allocation, or a bilink allocation, suitable for servicing high duty cycle periodic or quasi-periodic traffic on a committed schedule. A 1-periodic unscheduled bilink allocation is a bilink allocation, suitable for servicing high duty cycle periodic or quasi-periodic traffic on a best-effort basis.

pairwise temporal key (PTK): A secret bit string shared between a node and a hub and used to secure frames transferred between them.

pairwise temporal key (PTK) creation: A procedure used to create a PTK between a node and a hub based on a master key (MK) shared between them, and to confirm possession of a shared MK between the node and the hub.

poll: A control type frame or its variant sent by a hub to grant an immediate polled allocation to the addressed node or to inform the node of a future poll or post.

polled allocation: A non-reoccurring time interval that a hub grants to a node using polling access for initiating one or more frame transactions by the node. A polled allocation is an uplink allocation interval, suitable for servicing “ordinary,” “unexpected,” or “extra” uplink traffic (for example, due to data rate variations and/or channel impairments). A polled allocation is also called a *polled allocation interval*.

polling access: An access method, based on impromptu or scheduled polling by a hub, whereby a hub grants to a node a polled allocation for initiating one or more frame transactions by the node.

post: A management or data type frame sent by a hub to a node within its body area network (BAN). A post starts a posted allocation.

posted allocation: A non-reoccurring time interval that a hub grants to itself using posting access for initiating a frame transaction. A posted allocation is a downlink allocation interval, suitable for servicing “unexpected” or “extra” downlink traffic (for example, due to network management needs, data rate variations, and/or channel impairments).

posting access: An access method, based on impromptu or scheduled posting by a hub, whereby a hub grants to itself a posted allocation, typically outside scheduled uplink allocations, for initiating one or more frame transactions by the hub.

random access: An access method, based on carrier sense multiple access with collision avoidance (CSMA/CA) or slotted Aloha access but not both, whereby a node obtains a time interval in a random access phase (RAP) for initiating one or more frame transactions.

random access phase (RAP): A time span set aside by a hub and announced via a beacon frame for random access to the medium by the nodes in the body area network (BAN) of the hub.

relayed node: A node that communicates with a hub through another node.

relaying node: A node through which another node communicates with a hub.

scheduled access: An access method, based on advance reservation and committed scheduling, whereby a node and a hub obtain scheduled reoccurring time intervals for initiating frame transactions.

scheduled allocation: One or more scheduled reoccurring time intervals that a node and a hub obtains using scheduled access for initiating frame transactions. A scheduled allocation is an uplink allocation, a downlink allocation, or a bilink allocation, suitable for servicing high or low duty cycle periodic or quasi-periodic traffic on a committed schedule.

scheduled-polling access: A combination of scheduled access and polling access, whereby a node and a hub obtain scheduled reoccurring time intervals, wherein the hub grants to the node and/or itself non-reoccurring time intervals for initiating frame transactions on uplink and/or downlink.

secure frame: A term that is interchangeable with secured frame.

secured communication: Exchange of secured frames.

secured frame: A frame that is secured with authenticity, integrity, confidentiality if required, and replay protection.

security association: A procedure used to identify a node and a hub to each other and establish a new master key (MK) shared between them or activate an existing MK pre-shared between them.

star network: A logical network partition comprising a hub and zero or more nodes whose medium access and power management are coordinated by the hub.

superframe: A term that is interchangeable with beacon period used especially when no beacons are transmitted.

type-I polled allocation: A polled allocation the length of which is specified in terms of the duration of time granted for transmission.

type-I polling access: Polling access that provides type-I polled allocations.

type-II polled allocation: A polled allocation the length of which is specified in terms of the number of frames granted for transmission.

type-II polling access: Polling access that provides type-II polled allocations.

unconnected node: A node that does not have a connection with a hub in a body area network (BAN).

unscheduled access: A combination of best-effort scheduled access and polling access, whereby a node and a hub obtain unscheduled reoccurring time intervals, wherein the hub grants to the node and/or itself non-reoccurring time intervals for initiating frame transactions in an uplink and/or downlink.

unscheduled bilink allocation: One or more unscheduled reoccurring time intervals that a node and a hub obtains using unscheduled access for initiating frame transactions. An unscheduled bilink allocation is a bilink allocation, suitable for servicing high or low duty cycle periodic or quasi-periodic traffic in an uplink and/or downlink on a best-effort basis.

unsecured communication: Exchange of unsecured frames.

unsecured frame: A frame that is not secured with authenticity, integrity, confidentiality, or replay protection.

uplink allocation: An allocation with allocation interval(s) in which a node initiates one or more frame transactions to transmit management and data traffic to a hub and the hub returns acknowledgment if required.

uplink: A communications link for transfer of management and data traffic from a node to a hub.

3.2 Special terms

association: A term synonymous to security association in absence of the qualifier “security.”

appropriate: Subject to the rules specified in Clause 5, Clause 6, and Clause 7.

frame transaction: All management or data type frames of the same frame subtype, and an acknowledgment frame if required, that are separated in time by an appropriate interframe space and transmitted between a node and a hub. Examples of a frame transaction: (a) a management type frame and an immediate acknowledgment frame; (b) all data type frames of the same frame subtype after a block acknowledgment frame, and the block acknowledgment frame that follows; (c) a management or data type frame, which is followed by another management or data type frame of a different frame subtype; (d) a management or data type frame, which is separated by more than an appropriate interframe space by another management or data type frame of the same frame subtype.

implant: An entity that is placed inside a human body for medical purposes.

low power low duty cycle (LP/LDC): Power, duty cycle, and transmission (transmissions/per hour) limits defined by standards and regulations for implants transmitting within the band 403.5 MHz to 403.8 MHz without coordination with a hub.

medical implant event: A term referenced in regulatory documents governing the use of the Medical Device Radiocommunication Service (MedRadio), which includes the Medical Implant Communications Service (MICS) from 402 MHz to 405 MHz.

3.3 Acronyms and abbreviations

AES	Advanced Encryption Standard
ARQ	automatic repeat request
AWGN	additive white Gaussian noise
B-Ack	block acknowledgment
BAN	body area network
BCH code	Bose, Ray-Chaudhuri, Hocquenghem code
BPSK	binary phase shift keying
CAP	contention access phase
CBC	cipher block chaining
CCA	clear channel assessment
CCM	counter mode for message encryption and cipher block chaining (CBC) mode for message authentication
CMAC	(block) cipher-based message authentication code algorithm
CP	contention probability
CP-BFSK	continuous-phase binary frequency shift keying
CRC	cyclic redundancy check

CSMA/CA	carrier sense multiple access with collision avoidance
CW	contention window (for carrier sense multiple access with collision avoidance)
D8PSK	differential 8-phase-shift keying
DBPSK	differential (or differentially encoded) binary phase-shift keying
DQPSK	differential (or differentially encoded) quadrature phase-shift keying
DRF	data rate field
EAP	exclusive access phase
ED	energy detection
EFC	electric field communication
EIRP	equivalent isotropically radiated power
EUI	extended unique identifier
EVM	error vector magnitude
FCS	frame check sequence
FEC	forward error correction code
FM-UWB	frequency modulation ultra-wideband
FSDT	frequency selective digital transmission
FS-Spreader	frequency selective spreader
G-Ack	group acknowledgment
GF	Galois field
GFSK	Gaussian frequency-shift keying
GPPM	group pulse position modulation
GT	guard time
GTK	group temporal key
HARQ	hybrid automatic repeat request
HBC	human body communications
HCS	header check sequence
HID	hub identifier
HME	hub management entity
I-Ack	immediate acknowledgment
IE	information element
IR-UWB	impulse radio ultra-wideband
ISM	industrial scientific medical
ISO	International Organization for Standardization

KCK	key confirmation key
KMAC	key message authentication code
L-Ack	late acknowledgment (acknowledgment later)
LFSR	linear feedback shift register
LP/LDC	low power low duty cycle
LSB	least significant bit
MAC	media access control
MAP	managed access phase
MCU	micro controller unit
MIC	message integrity code
MICS	medical implant communications service
MIFS	minimum interframe space
MK	master key
MPDU	medium access control protocol data unit
MSB	most significant bit
MSDU	media access control service data unit
MUX	multiplexer
N-Ack	no acknowledgment
NID	node identifier
NME	node management entity
OSI	open systems interconnection
PER	packet error rate
PHR	physical layer header
PHY	physical or physical layer
PLCP	physical layer convergence protocol
PN	pseudo-random noise
PPDU	physical layer protocol data unit
P.PRF	peak pulse repetition frequency
PRF	pulse-repetition-frequency
PSD	power spectral density
PSDU	physical layer service data unit
PSK	phase shift keying
PTK	pairwise temporal key

QPSK	quadrature phase shift keying
RAP	random access phase
RI	rate indicator
RX	receive or reception
S2P	serial-to-parallel
SAP	service access point
SAR	specific absorption rate
SF	spreading factor
SFD	start-of-frame delimiter
SHR	synchronization header
SIFS	short interframe spacing
SRRC	square-root raised cosine
TK	temporal key
TX	transmit or transmission
UP	user priority
UWB	ultra wideband

4. General framework elements

4.1 General

This clause provides the basic framework of nodes and hubs. The framework serves as a prerequisite to supporting the functions of nodes and hubs and their interactions specified later in detail. It covers the following aspects: the network topology used for medium access, the reference model used for functional partitioning, the time base used for access scheduling, the state diagrams used for frame exchange, and the security paradigm used for message protection.

4.2 Network topology

All nodes and hubs are to be organized into logical sets, referred to as *body area networks* (BANs) in this specification, and coordinated by their respective hubs for medium access and power management as illustrated in Figure 1. There is to be one and only one hub in a BAN, whereas the number of nodes in a BAN is to range from zero to $mMaxBANSize$. In a one-hop star BAN, frame exchanges are to occur directly between nodes and the hub of the BAN. In a two-hop extended star BAN, the hub and a node are to exchange frames optionally via a relay-capable node.

Access coordination at the MAC sublayer between BANs is not specified in this standard. Optional mechanisms for coexistence and interference mitigating between adjacent or overlapping BANs are provided (in 6.13). Nodes referenced in this standard are in the context of a given BAN, unless noted otherwise.




Figure 1—Network topology

4.3 Reference model

All nodes and hubs are internally partitioned into a physical (PHY) layer and a medium access control (MAC) sublayer, in accordance with the IEEE 802[®] reference model, as shown in Figure 2. Direct communications between a node and a hub are to transpire at the PHY layer and MAC sublayer as specified in this standard; the PHY layer and MAC sublayer of a node or a hub are to use only one operating channel at any given time. Message security services are to occur at the MAC sublayer, and security key generations are to take place inside and/or outside the MAC sublayer.




Figure 2—Reference model

Within a node or a hub, the MAC provides its service to the MAC client (higher layer) through the MAC service access point (SAP) located immediately above the MAC sublayer, while the PHY provides its service to the MAC through the PHY SAP located between them. On transmission, the MAC client passes MAC service data units (MSDUs) to the MAC sublayer via the MAC SAP, and the MAC sublayer passes MAC frames (also known as MAC protocol data units or MPDUs) to the PHY layer via the PHY SAP. On reception, the PHY layer passes MAC frames to the MAC sublayer via the PHY SAP, and the MAC sublayer passes MSDUs to the MAC client via the MAC SAP. Both MAC SAP and PHY SAP are not exposed and their specifications are beyond the scope of this standard.

There may be a logical node management entity (NME) or hub management entity (HME) that exchanges network management information with the PHY and MAC as well as with other layers. The HME is a superset of the NME in terms of the management functionality they each support. However, the presence of the NME or HME and the partitioning between the NME or HME and the MAC or the PHY is not mandated, nor is the behavior of the NME or HME specified, in this standard.

4.4 Time base

All nodes and hubs are to establish a time reference base, as shown in Figure 3, if their medium access is to be scheduled in time, where the time axis is divided into beacon periods (superframes) of equal length and each beacon period (superframe) is composed of allocation slots of equal length and numbered from 0, 1, ..., s , where $s \leq 255$. An allocation interval may be referenced in terms of the numbered allocation slot comprising it, and a point of time may be referenced in terms of the numbered allocation slots preceding or following it as appropriate.




Figure 3—Time reference base

If time reference is needed for access scheduling in its BAN, the hub is required to choose the boundaries of beacon periods (superframes) and hence of the allocation slots therein. In beacon mode operation for which beacons are transmitted, the hub needs to communicate such boundaries by transmitting beacons at the start or other specified locations of beacon periods (superframes), and optionally timed frames (T-Poll frames) containing their transmit time relative to the start time of current beacon period (superframe). In non-beacon mode operation for which beacons are not transmitted but time reference is needed, the hub is required to communicate such boundaries by transmitting timed frames (T-Poll frames) also containing their transmit time relative to the start time of current superframe.

A node requiring a time reference in the BAN needs to derive and recalibrate the boundaries of beacon periods (superframes) and allocation slots from reception of beacons or/and timed frames (T-Poll frames).

A frame transmission may span more than one allocation slot, starting or ending not necessarily on an allocation slot boundary.

4.5 MAC and security state diagrams

All nodes and hubs are to go through certain stages, i.e., states, at the MAC level before they exchange user (MAC client) data, as shown in Figure 4, where frames permitted or required to exchange between a node and a hub at each state are also indicated. State classification and transition are defined with respect to a pair of a node and a hub, but is often referenced in the name of the node only. Broadcast or multicast frames, such as beacons, are frame exchanges between a hub and a group of nodes but not between a hub and a node; their transmissions are not dictated by the pairwise state diagrams specified in this subclause.

4.5.1 Secured communication

A node and a hub are to follow the MAC state diagram Figure 4(a) for secured communication if either of them requires secured frame exchanges with the other.




Figure 4—MAC and security state diagrams

4.5.1.1 Orphan state

At this state, the node does not have any relationship with the hub for secured communication. It is the state that the node initially enters in relation with the hub. The node and the hub are not allowed to transmit any frames to each other—other than Security Association and control unsecured frames. The node is allowed to exchange Security Association frames with the hub to establish a security association, i.e., to activate a pre-shared master key (MK) or generate a new MK, and to authenticate with each other if so required, thereby transitioning to next state, *Associated* state. However, if the node and the hub fail to activate or establish a shared MK, they are not allowed to advance to *Associated* state.

4.5.1.2 Associated state

At this state, the node is associated, i.e., holds a shared MK, with the hub for their pairwise temporal key (PTK) creation. The node and the hub are not allowed to transmit any frames to each other—other than Security Disassociation and PTK unsecured frames, as well as control unsecured frames. The node is allowed to exchange PTK frames with the hub to establish a secured relationship, i.e., to confirm possession of a shared MK and to create a PTK, thereby transitioning to next state, *Secured* state. However, if the node and the hub fail to create a PTK, they are not allowed to advance to *Secured* state. If their MK is missing or invalid during their PTK creation, they are required to move back to *Orphan* state. To repeal the security association and hence the current MK, either the node or the hub is allowed to send a Security Disassociation frame to the other, thus moving back to *Orphan* state.

4.5.1.3 Secured state

At this state, the node is secured, i.e., holds a PTK, with the hub for message security, i.e., for secured frame exchanges. The node and the hub are not allowed to transmit any frames to each other—other than Security Disassociation, Connection Request, and Connection Assignment secured frames, as well as control secured or unsecured frames depending on whether authentication of control type frames was selected during their association. The node is allowed to exchange Connection Request and Connection Assignment frames with the hub to establish a connection, thereby transitioning to the next and final state, *Connected* state. However, if the node and the hub fail to establish a connection, they are not allowed to advance to *Connected* state. If their PTK is missing or invalid or the node's *Connected_NID* is lost, they are required to move back to *Associated* state. To repeal the security association and hence the current MK, either the node or the hub is allowed to send a Security Disassociation frame to the other, thus moving back to *Orphan* state.

4.5.1.4 Connected state

At this state, the node is connected, i.e., holds an assigned *Connected_NID*, a wakeup arrangement, and optionally one or more scheduled and unscheduled allocations with the hub for abbreviated node addressing, desired wakeup, and optionally scheduled and unscheduled access. The node and the hub are allowed to transmit any secured frames to each other—other than Security Association secured frames, but are not allowed to transmit any unsecured frames to each other—other than control unsecured frames if authentication of control type frames was not selected during their association. If their MK is missing or invalid during a PTK recreation, the node and the hub are required to move back to *Orphan* state. If their PTK is missing or invalid, the node and the hub are required to move back to *Associated* state. To repeal the security association and hence the current MK, either the node or the hub is allowed to send a Security Disassociation frame to the other, thus moving back to *Orphan* state. To repeal the connection and hence the node's *Connected_NID*, wakeup arrangement, and scheduled and unscheduled allocations if any, either the node or the hub is allowed to send a Disconnection frame to the other, thereby moving back to *Associated* state.

4.5.2 Unsecured communication

A node and a hub are required to follow the MAC state diagram Figure 4(b) for unsecured communication if neither of them requires secured frame exchanges with the other.

4.5.2.1 Orphan state

At this state, the node does not have any relationship with the hub for unsecured communication. It is the state that the node initially enters in relation with the hub. The node and the hub are not allowed to transmit any frames to each other—other than Connection Request, Connection Assignment, and control unsecured frames. The node is allowed to exchange Connection Request and Connection Assignment frames with the hub to establish a connection, thereby transitioning to next and final state, *Connected* state. However, if the node and the hub fail to establish a connection, they are not allowed to advance to *Connected* state.

4.5.2.2 Connected state

At this state, the node is connected, i.e., holds an assigned *Connected_NID*, a wakeup arrangement, and optionally one or more scheduled and unscheduled allocations with the hub for abbreviated node addressing, desired wakeup, and optionally scheduled and unscheduled access. The node and the hub are allowed to transmit any unsecured frames to each other—other than Security Association, Security Disassociation, and PTK frames, but are not allowed to transmit secured frames to each other. To repeal the

connection and hence the node's *Connected_NID*, wakeup arrangement, and scheduled and unscheduled allocations if any, either the node or the hub is allowed to send a Disconnection frame to the other, thereby moving back to *Orphan* state.

To change to secured communications between them, the node and the hub are required to disconnect from each other, thus moving back to *Orphan* state and then following the state diagram Figure 4(a) for secured communication.

4.6 Security paradigm

All nodes and hubs are to choose three security levels in this standard, as follows:

- Level 0—unsecured communication. At this level, messages are transmitted in unsecured frames, which provide no measures for message authenticity and integrity validation, confidentiality and privacy protection, and replay defense.
- Level 1—authentication but not encryption. At this level, messages are transmitted in secured authenticated but not encrypted frames, which provide measures for message authenticity and integrity validation and replay defense but not confidentiality and privacy protection.
- Level 2—authentication and encryption. At this level, messages are transmitted in secured authenticated and encrypted frames, which provide measures for message authenticity and integrity validation, confidentiality and privacy protection, and replay defense.

During association, a node and a hub need to jointly select a security level suitable for their subsequent secured frame exchanges, and whether to require authentication of control type frames, based on their respective security requirements and certain information specific to each other.

For unicast secured communication, the node and the hub need to activate a pre-shared MK or establish a new MK via an unauthenticated or authenticated association, and create a PTK via a PTK creation procedure. For multicast secured communication, the hub needs to distribute a GTK to the corresponding multicast group in a unicast manner.

The node and the hub are to follow the security structures shown in Figure 5 to perform security key generations and provide message security services. A “session” indicated in this figure refers to a time span in which a temporal key (TK) remains valid. The length of the session is determined by the security policy governing data transfers between the two communicating parties. It is further limited by the technical restrictions on the reuse of the same TK over successive messages (i.e., MAC frames in this specification).




Figure 5—Security hierarchy

5. MAC frame formats

5.1 Conventions

A MAC frame is an ordered sequence of fields delivered to or from the physical layer service access point (PHY SAP). Each figure in Clause 5 depicts the fields contained in a MAC frame, or a part thereof, from left to right in the transmit order, with fields that are optional or selectively absent drawn in dashes where possible. The transmit order starts from top to bottom if the fields are depicted in multiple rows, symbolically linked by dashes extended to the right of one row and to the left of next row, respectively. Also indicated is the number of octets contained in each field along with the corresponding octet transmit order, on top of the field, as illustrated in Figure 6.




Figure 6—Example of fields aligned with octet boundaries

Unless otherwise noted, an atomic field, i.e., a field that is not in turn comprised of other fields, denotes a numerical value encoded in unsigned binary. If such a field contains F octets ($F > 1$), octet 0 is the octet containing the least significant bits (LSBs) of that field and is the first octet transmitted of the field, whereas octet $F-1$ is the octet containing the most significant bits (MSBs) of that field and is the last octet transmitted of the field. The octet order is indicated as “L-R,” i.e., from left to right, above a multi-octet non-atomic field.

In a field encoding an EUI-48, it is set to a string of six octets per Clause 9 of IEEE Std 802-2001.⁹ Octets 0–2 of the field are set sequentially to the first three octets of the string that specify an organizationally unique identifier (OUI), with bit 0 of octet 0 being the individual/group (I/G) address bit. Octets 3–5 are set to the last three octets selected as an extension identifier (ID) by the organization identified by the OUI.

In a figure that depicts certain fields not aligned with octet boundaries, the number of bits and the corresponding bit order of encoding are shown instead for each field in the figure, as illustrated in Figure 7. Bits are ordered continually across the fields that are not aligned with octet boundaries, from left to right, starting from bit 0, i.e., the LSB of the bits comprising these fields. Bit numbering restarts from zero in fields located on octet boundaries.




Figure 7—Example of fields not aligned with octet boundaries

Each field is defined, or set, based on the perspective of the node or the hub that is sending the frame containing that field, referred to as the *sender*. It is parsed based on its definition by the hub or the node intended to receive the frame containing it, referred to as the *recipient* or the *intended recipient*.

⁹ Information on references can be found in Clause 2.

Reserved fields are set to zero on transmission and ignored on reception. If some values in a field are reserved, the field is not set to any of those reserved values on transmission. Unless otherwise noted, fields that are set to reserved values or defined based on other fields that are set to reserved values are ignored on reception.

MAC constants are referenced through parameters as listed in 6.15. These parameters are denoted with a preceding “m” if they are PHY independent or with a preceding “p” if they are PHY dependent.

5.2 General format

A MAC frame consists of a fixed-length MAC header, a variable-length MAC frame body, and a fixed-length Frame Check Sequence (FCS) field as shown in Figure 8. The MAC frame body has an octet length L_FB such that $0 \leq L_FB \leq pMaxFrameBodyLength$, and is present only if it has a nonzero length.




Figure 8—MAC frame format

5.2.1 MAC Header

The MAC Header is formatted as shown in Figure 9.




Figure 9—MAC Header format

5.2.1.1 Frame Control

The Frame Control is formatted as shown in Figure 10.




Figure 10—Frame Control format

5.2.1.1.1 Protocol Version

The Protocol Version field is set to zero for this revision of the standard. All other values are reserved. This field is invariant in size and place across all revisions of this standard.

5.2.1.1.2 Acknowledgment (Ack) Policy

The Ack Policy field is set according to Table 1 to indicate the acknowledgement requirement of the current frame.

Table 1—Acknowledgement Policy field encoding

Field value b2 b1	Acknowledgment requirement
00	No acknowledgment (N-Ack) or group acknowledgment (G-Ack)
01	Immediate acknowledgment (I-Ack)
11	Block acknowledgment later (L-Ack)
10	Block acknowledgment (B-Ack)

The group acknowledgment (G-Ack) value is applicable to frames sent to a hub and of frame type set to data and frame subtype set to mG-AckDataSubtype.

5.2.1.1.3 Security Level

The Security Level field is set according to Table 2 such that it indicates the security level of the current frame.

Table 2—Security Level field encoding

Field value b4 b3	Security level of current frame
00	Level 0—frame not secured
01	Level 1—frame authenticated but not encrypted
10	Level 2—frame authenticated and encrypted
11	Reserved

5.2.1.1.4 Temporal key (TK) Index

The TK Index field is set as follows to indicate the PTK or GTK being used to secure the current frame:

- In frames secured by a PTK, it is set to the value of the PTK Index field in the PTK frames that were exchanged in creating the PTK.
- In frames secured by a GTK, it is set to the value of the GTK Index field in the GTK frame that was exchanged in distributing the GTK.

The TK Index field is reserved in unsecured frames.

5.2.1.1.5 BAN Security/Relay

The BAN Security/Relay field is set as follows:

- a) In beacon, Poll, and T-Poll frames sent by a hub, it is used as a BAN Security field, which is set to one if this hub accepts only secured communication with it as described in 4.5.1, or is set to zero if this hub accepts either secured or unsecured communication with it as described in 4.5.1 or 4.5.2.
- b) In frames sent to or from a relaying node in a two-hop extended star network communication, it is used as a Relay field, which is set to one.
- c) In all other frames, it is reserved.

5.2.1.1.6 Ack Timing/EAP Indicator/First Frame On Time

The Ack Timing/EAP Indicator/First Frame On Time field is set as follows:

- a) In beacon frames, it is used as an EAP Indicator field, which is set to one if exclusive access phase 1 (EAP 1) in the current (with beacon shifting not enabled) or next (with beacon shifting enabled) beacon period has a nonzero length, or is set to zero otherwise.
- b) In non-beacon management or data type frames or Poll frames sent by a hub to a node, it is used as a First Frame On Time field, which is set to one if this is the first frame sent by the hub to the node at the start of an allocation interval of a scheduled allocation assigned to the node, or is set to zero otherwise.
- c) In non-beacon management or data type frames with the Ack Policy field of the MAC header set to I-Ack or B-Ack sent by a node to a hub, it is used as an Ack Timing field, which is set to one if the following acknowledgment (I-Ack, B-Ack, I-Ack+Poll, or B-Ack+Poll) frame is to include a timestamp in its frame payload, or is set to zero otherwise. The timestamp encodes the start time of the acknowledgment frame transmission based on the hub's clock.
- d) In I-Ack, B-Ack, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as an Ack Timing field, which is set to one if the frame includes a timestamp in the frame payload, or is set to zero otherwise.
- e) In all other frames, it is reserved.

5.2.1.1.7 Frame Subtype

The Frame Subtype field is set to indicate the subtype of the current frame of a given type according to Table 3. The name of the frame subtype of a frame is used as the name of the frame throughout Clause 2 through Clause 7.

5.2.1.1.8 Frame Type

The Frame Type field is set to indicate the type of the current frame according to Table 3.

Table 3—Frame Type and Frame Subtype field encoding

Frame Type value b5 b4	Frame Type name	Frame Subtype value b3 b2 b1 b0	Frame Subtype name
00	Management	0000	Beacon
00	Management	0001	Reserved
00	Management	0010	Security Association
00	Management	0011	Security Disassociation
00	Management	0100	PTK
00	Management	0101	GTK
00	Management	0110–0111	Reserved
00	Management	1000	Connection Request
00	Management	1001	Connection Assignment
00	Management	1010	Disconnection
00	Management	1011–1110	Reserved
00	Management	1111	Command
01	Control	0000	I-Ack
01	Control	0001	B-Ack
01	Control	0010–0011	Reserved
01	Control	0100	I-Ack+Poll
01	Control	0101	B-Ack+Poll
01	Control	0110	Poll
01	Control	0111	T-Poll
01	Control	1000–1101	Reserved
01	Control	1110	Wakeup
01	Control	1111	B2
10	Data	0000	User Priority 0 or Allocation Mapped Data Subtype
10	Data	0001	User Priority 1 or Allocation Mapped Data Subtype
10	Data	0010	User Priority 2 or Allocation Mapped Data Subtype
10	Data	0011	User Priority 3 or Allocation Mapped Data Subtype
10	Data	0100	User Priority 4 or Allocation Mapped Data Subtype
10	Data	0101	User Priority 5 or Allocation Mapped Data Subtype
10	Data	0110	User Priority 6 or Allocation Mapped Data Subtype
10	Data	0111	Emergency
10	Data	1000–1111	Allocation Mapped Data Subtype
11	Reserved	0000–1111	Reserved

5.2.1.1.9 More Data

The More Data field is set as follows:

- a) In management or data type frames sent by a node to a hub,
 - 1) it is set to zero if the node has no management or data type frame pending for transmission to the hub, except for a possible retransmission of the current frame, or
 - 2) it is set to one if the node has at least a management or data type frame pending for transmission or retransmission to the hub.
- b) In I-Ack and B-Ack frames sent by a node to a hub,
 - 1) it is set to zero if the node has no management or data type frame pending for transmission or retransmission to the hub, or
 - 2) it is set to one if the node has at least one management or data type frames pending for transmission or retransmission to the hub.
- c) In non-beacon management or data type frames sent by a hub to a node,
 - 1) it is set to zero if the hub has no poll or post pending for transmission to the node, except for a possible retransmission of the current frame, or
 - 2) it is set to one if the hub has at least a poll or post pending for transmission to the node.
- d) In I-Ack and B-Ack frames sent by a hub to a node,
 - 1) it is set to zero if the hub is not to send a poll or post to the node immediately after the current allocation interval is reclaimed or ended, or
 - 2) it is set to one if the hub is to send a poll or post to the node immediately after the current allocation interval is reclaimed or ended.
- e) In Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node,
 - 1) if Access Mode = 0 indicating that the hub is operating in beacon or non-beacon mode with superframes,
 - i) it is set to zero if via the current frame the hub grants a type-I (immediate) polled allocation that starts pSIFS thereafter and ends at a time as encoded in the Poll-Post Window and Next fields therein, or
 - ii) it is set to one if via the current frame the hub grants no polled allocation but is to send a (future) poll or post to the node at a time as encoded in the Poll-Post Window and Next fields thereof;
 - 2) if Access Mode = 1 indicating that the hub is operating in non-beacon mode without superframes, and that via the current frame the hub grants a type-II (immediate) polled allocation that starts pSIFS thereafter and allows for the node to send up to a specified number of frames as encoded in the Poll-Post Window field thereof, it is reserved.
- f) In all other frames, it is reserved.

5.2.1.1.10 Last Frame/Access Mode/B2

The Last Frame/Access Mode/B2 field is set as follows:

- a) In beacon frames, it is used as a B2 field, which is set to one if a B2 frame is to be transmitted in the current beacon period, or is set to zero otherwise.

- b) In non-beacon management or data type frames, it is used as a Last Frame field, which is
 - 1) set to zero if the sender is likely to send another frame in the current allocation interval after the current frame, or
 - 2) set to one if the sender is definitely not to send another frame in the current allocation interval after the current frame.
- c) In I-Ack, B-Ack, I-Ack+Poll, B-Ack+Poll, Poll, and T-Poll frames sent by a hub to a node, it is used as an Access Mode field, which is
 - 1) set to zero if the hub that is the sender or recipient of the current frame is operating in beacon or non-beacon mode with superframes, or
 - 2) set to one if the hub is operating in non-beacon mode without superframes.
- d) In all other frames, it is reserved.

5.2.1.1.11 Sequence Number/Poll-Post Window

The Sequence Number/Poll-Post Window field is set as follows:

- a) In beacon frames, it is used as a Sequence Number field, which is
 - 1) incremented by one, modulo 256, from its value applicable to the last beacon period (superframe), active or inactive;
 - 2) also referred to as the sequence number of the current beacon (superframe) or of the current beacon period (superframe).
- b) In non-beacon management type frames, it is used as a Sequence Number field, which is
 - 1) set to zero if the frame is the first non-beacon management type frame sent to a recipient or a group of recipients;
 - 2) incremented by one, modulo 256, from its value in the frame that was of the same frame type and addressed to the same recipient(s); or
 - 3) kept with the same value if the frame is being retransmitted to the same recipient(s).
- c) In Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as a Poll-Post Window field, which is set as follows:
 - 1) If Access Mode = 0 indicating that the hub is operating in beacon or non-beacon mode with superframes, and
 - i) if More Data = 0 indicating that via the current frame the hub grants a type-I (immediate) polled allocation that starts pSIFS thereafter, it is set to E such that the polled allocation ends at the end of the allocation slot that is numbered E and located in a beacon period (superframe) encoded in the Next field thereof; or
 - ii) if More Data = 1 indicating that via the current frame the hub grants no polled allocation, it is set to S such that the hub is to send a (future) poll or post to the node at the start of the allocation slot that is numbered S and located in a beacon period (superframe) encoded in the Next field thereof.
 - 2) If Access Mode = 1 indicating that the hub is operating in non-beacon mode without superframes, and that via the current frame the hub grants a type-II (immediate) polled allocation that starts pSIFS thereafter, it is set to M such that the node is allowed to send up to M frames in the polled allocation.
- d) In Wakeup frames, it is used as a Poll-Post Window field, which is set to $Poll_Post_Window$ such that this hub is to send a (future) poll after $2 \times (1 + Scale) \times (256 \times Next + Poll_Post_Window)$

milliseconds of the current frame, where *Next* and *Scale* are the values of the Next and Scale fields, respectively.

- e) In B2 frames, it is used as a Poll-Post Window field, which is set to *C* such that the current frame starts a contention access phase (CAP) that ends at the end of the allocation slot numbered *C* in the current beacon period (superframe).
- f) In data type frames, it is used as a Sequence Number field, which is
 - 1) set to zero if the frame is the first data type frame of this frame subtype sent to a recipient or a group of recipients;
 - 2) incremented by one, modulo 256, from its value in the data type frame that was of the same frame subtype and addressed to the same recipient(s) and that contained the previous MSDU or final fragment thereof;
 - 3) kept with the same value in frames containing fragments of the same MSDU; or
 - 4) kept with the same value if the frame is being retransmitted to the same recipient(s).
- g) In all other frames, it is reserved.

5.2.1.1.12 Fragment Number/Next/Coexistence

The Fragment Number/Next/Coexistence field is set as follows:

- a) In beacon and B2 frames, it is used as a Coexistence field, which is formatted as shown in Figure 11.
 - 1) The Beacon Shifting field is set to one if beacon shifting is currently enabled, or is set to zero otherwise.
 - 2) The Channel Hopping field is set to one if channel hopping is currently enabled, or is set to zero otherwise.
 - 3) The Superframe Interleaving field is set to one if the sender (a hub) supports active superframe interleaving and Command frames, or is set to zero otherwise.




Figure 11 —Coexistence format

- b) In non-beacon management type frames, it is used as a Fragment Number field, which is
 - 1) set to zero if the current frame contains a frame payload not fragmented or the first fragment of a fragmented frame payload;
 - 2) incremented by one from its value in the frame containing the previous fragment of the frame payload if the current frame contains a non-first fragment of a fragmented frame payload; or
 - 3) kept with the same value if the frame is being retransmitted to the same recipient(s).
- c) In data type frames, it is used as a Fragment Number field, which is
 - 1) set to zero if the current frame contains no MSDU, an MSDU not fragmented, or the first fragment of a fragmented MSDU;
 - 2) incremented by one from its value in the frame containing the previous fragment of the MSDU if the current frame contains a non-first fragment of a fragmented MSDU; or

- 3) kept with the same value if the frame is being retransmitted to the same recipient(s).
- d) In Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as a Next field, which is set as follows:
 - 1) If Access Mode = 0 indicating that the hub is operating in beacon or non-beacon mode with superframes, and
 - i) if More Data = 0 indicating that via the current frame the hub grants a type-I (immediate) polled allocation that starts pSIFS thereafter and ends at the end of an allocation slot as encoded in the Poll-Post Window field thereof, it is set to N such that the allocation slot is the one located in the current beacon period (superframe) for $N = 0$ or in the next N th beacon period (superframe) not counting the current one for $N > 0$ (a case possible if no beacon is to be sent in the polled allocation); or
 - ii) if More Data = 1 indicating that via the current frame the hub grants no polled allocation but is to send a (future) poll or post to the node at the start of an allocation slot as encoded in the Poll-Post Window field thereof, it is set to F such that the allocation slot is the one located in the current beacon period (superframe) for $F = 0$ or in the next F th beacon period (superframe) not counting the current one for $F > 0$;
 - 2) If Access Mode = 1 indicating that the hub is operating in non-beacon mode without superframes, and that via the current frame the hub grants a type-II (immediate) polled allocation that starts pSIFS thereafter and allows for the node to send up to a specified number of frames as encoded in the Poll-Post Window field thereof, it is reserved.
- e) In Wakeup frames, it is used as a Next field, which is set to $Next$ such that this hub is to send a (future) poll after $2 \times (1 + Scale) \times (256 \times Next + Poll_Post_Window)$ milliseconds of the current frame, where $Poll_Post_Window$ and $Scale$ are the values of the Poll-Post Window and Scale fields, respectively.
- f) In all other frames, it is reserved.

5.2.1.1.13 Non-final Fragment/Cancel/Scale/Inactive

The Non-final Fragment/Cancel/Scale/Inactive field is set as follows:

- a) In beacon and B2 frames, it is used as an Inactive field, which is set to one if one or more inactive superframes are enabled (starting) at the end of the current beacon period (superframe), or is set to zero otherwise.
- b) In non-beacon management type frames, it is used as a Non-final Fragment field, which is
 - 1) set to zero if the current frame contains no frame payload, a frame payload not fragmented, or the final fragment of a fragmented frame payload; or
 - 2) set to one if the frame contains a non-final fragment of a fragmented frame payload.
- c) In data type frames, it is used as a Non-final Fragment field, which is
 - 1) set to zero if the current frame contains no MSDU, an MSDU not fragmented, or the final fragment of a fragmented MSDU; or
 - 2) set to one if the frame contains a non-final fragment of a fragmented MSDU.
- d) In I-Ack, B-Ack, Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as a Cancel field, which is
 - 1) set to zero if the current frame does not cancel any future polls/posts previously improvised by this hub; or
 - 2) set to one if the current frame cancels all future polls/posts previously improvised by this hub.

- e) In Wakeup frames, it is used as a Scale field, which is set to *Scale* such that this hub is to send a (future) poll after $2 \times (1 + \text{Scale}) \times (256 \times \text{Next} + \text{Poll_Post_Window})$ milliseconds of the current frame, where *Poll_Post_Window* and *Next* are the values of the Poll-Post Window and Next fields, respectively.
- f) In all other frames, it is reserved.

5.2.1.2 Recipient ID

The Recipient ID field is set to the abbreviated address (i.e., NID or HID) of the recipient of the current frame.

5.2.1.3 Sender ID

The Sender ID field is set to the abbreviated address (i.e., NID or HID) of the sender of the current frame.

5.2.1.4 BAN ID

The BAN ID field is set to the abbreviated address of the BAN in which the current frame is transferred.

5.2.2 MAC Frame Body

The MAC Frame Body, when it has a nonzero length, is formatted as shown in Figure 12. The length of the MAC Frame Body *L_{FB}* is not to exceed pMaxFrameBodyLength.




Figure 12—MAC Frame Body format

The Low-Order Security Sequence Number and Message Integrity Code (MIC) fields are not present in unsecured frames, as indicated by the Security Level field of the MAC header of the current frame.

5.2.2.1 Low-Order Security Sequence Number

The Low-Order Security Sequence Number field is set as follows for message freshness as needed for nonce construction and replay detection:

- a) It is set to zero, if the current frame is secured with a PTK or GTK that has never been used.
- b) It is incremented by one from its value in the last frame secured with the same PTK or GTK, containing a valid MIC value, and transmitted by the same sender, even if the current frame transmission is a retransmission of the last frame or an earlier frame. Incrementing by one from the maximum value of the field wraps around to zero, causing the High-Order Security Sequence Number maintained internally to increment by one.

The value of the Low-Order Security Sequence Number field increments in frames secured with the same PTK or GTK, rather than in frames of the same frame type or frame subtype. It increments even if the current frame transmission is a retransmission of the last frame or an earlier frame. It is set independently between the frames sent from a hub to a node and the frames sent from the node to the hub, although the same PTK applies to secured frames sent in either direction.

5.2.2.2 Frame Payload

The Frame Payload field is set as follows:

- a) In management type frames, prior to encryption (if any) it is set to a sequence of fields to be communicated to the recipient(s), with the fields defined in 5.3.
- b) In control type frames, it is set to a sequence of fields to be communicated to the recipient(s), with the fields defined in 5.4.
- c) In data type frames, prior to encryption (if any) it is set to a sequence of octets passed as an MSDU through the MAC SAP to the MAC, without altering the order of the sequence.
- d) In data type frames with the Relay field in their MAC header set to one, prior to encryption (if any) it is set to the relayed MAC frame that would otherwise be sent directly between a node and a hub without relay.

If a frame payload is fragmented and carried in multiple frames, the Frame Payload field is set to a fragment of the otherwise unfragmented frame payload.

The length of the Frame Payload field, denoted as L_{FP} in Figure 12, needs to be such that the length of the MAC frame body does not exceed $pMaxFrameBodyLength$. If the Frame Payload has a zero length, i.e., if a frame does not have a frame payload, then the frame does not have a MAC frame body if it is not secured, but still has a MAC frame body containing the Low-Order Security Sequence Number and MIC fields if it is secured.

5.2.2.3 Message Integrity Code (MIC)

The MIC field is set to a message authentication code for preserving the authenticity and integrity of the MAC header and the MAC frame body of the current secured frame, as further specified in 7.3.1.5.

5.2.3 Frame Check Sequence (FCS)

The FCS field is formatted as shown in Figure 13, where its transmit order is defined such that a_{15} is the LSB of the field, and a_0 is the MSB. The bits $a_{15}, a_{14}, \dots, a_0$ are the binary coefficients of a cyclic redundancy check (CRC) polynomial of degree 15 denoted, as shown in Equation (1):

$$R(x) = a_{15}x^{15} + a_{14}x^{14} + \dots + a_1x + a_0 \quad (1)$$




Figure 13—FCS format

The CRC polynomial is calculated over a calculation field using the following CRC-16-CCITT standard generator polynomial of degree 16, shown in Equation (2):

$$G(x) = x^{16} + x^{12} + x^5 + 1 \quad (2)$$

The calculation field is the transmitted MAC frame except the FCS field for this specification. It is mapped to a message polynomial $M(x)$ of degree $k-1$, where k is the number of bits in the calculation field. The LSB of the first octet presented to the PHY SAP is the coefficient of the x^{k-1} term, the next LSB is the coefficient of the x^{k-2} term, and finally the MSB of the last octet transmitted is the coefficient of the x^0 term.

The CRC polynomial is the remainder $R(x)$ resulting from the (modulo 2) division of $[x^{16} \times M(x)]$ by $G(x)$, shown in Equation (3):

$$R(x) = x^{16} \times M(x) \bmod G(x) \quad (3)$$

The initial remainder of the division is set to zero, and the final remainder is not inverted as is the case in some other standards.

5.3 Management type frames

A management type frame contains certain mandatory fixed-length fields and some optional variable length components referred to as *information elements* (IEs).

5.3.1 Beacon

A beacon frame contains a Frame Payload that is formatted as shown in Figure 14. It is locally broadcast by a hub in every beacon period (superframe).




Figure 14—Frame Payload format for Beacon frames

5.3.1.1 Sender Address

The Sender Address field is set to the EUI-48 of the hub sending the current beacon.

5.3.1.2 Beacon Period Length

The Beacon Period Length field is set to the length of the current beacon period (superframe), in units of allocation slots. It is set to zero to encode a value of 256 allocation slots.

5.3.1.2.1 Allocation slot numbering

If beacon shifting is not enabled, the allocation slots in a beacon period (superframe) are numbered $0, 1, \dots$, starting from the allocation slot that starts at the beacon transmission time of the beacon period (superframe) consecutively to the allocation slot that ends at the end of the beacon period (superframe), as shown in Figure 15.




Figure 15—Allocation slot ordering and numbering in beacon periods (superframes)

5.3.1.2.2 Allocation slot numbering with beacon shifting

If beacon shifting is enabled, a beacon period (superframe) needs to have $4N$ allocation slots in length, where N is an integer. The allocation slots in a beacon period (superframe) are numbered $0, 1, \dots, 4N-1$, starting from the allocation slot that starts at the beacon transmission time of the beacon period (superframe) to the allocation slot that ends at the end of the beacon period (superframe) and, if the beacon transmission time is not at the start of the beacon period (superframe), wrapping back to the allocation slot that starts at the start of the beacon period (superframe) and finally to the allocation slot that ends at the beacon transmission time, as shown in Figure 16. A beacon period (superframe) has four quarters, which are comprised of allocation slots 0 to $N-1$, N to $2N-1$, $2N$ to $3N-1$, and $3N$ to $4N-1$, respectively.




Figure 16—Allocation slot ordering and numbering in beacon periods (superframes) with exemplary beacon shifting

5.3.1.3 Allocation Slot Length

The Allocation Slot Length field is set to L such that the length of an allocation slot is equal to $p\text{AllocationSlotMin} + L \times p\text{AllocationSlotResolution}$.

5.3.1.4 RAP1 Start

The RAP1 Start field is present only if EAP1 has a nonzero length as indicated by the EAP Indicator field of the MAC header of the current beacon frame. When present, it is set to $S1$ such that random access phase 1 (RAP1) starts at the beginning of the allocation slot that is numbered $S1$ and located in the current beacon period ends EAP1, and it occurs after the PHY Capability field.

Both EAP1 (if of nonzero length) and RAP1 (if of nonzero length) occur in the current beacon period if beacon shifting is not enabled or in the next beacon period if beacon shifting is enabled. If EAP1 has a nonzero length, it immediately follows the preceding beacon, and is immediately followed by RAP1. If EAP1 has a zero length, RAP1 immediately follows the preceding beacon.

5.3.1.5 RAP1 End

The RAP1 End field is set to $E1$ such that RAP1 ends at the end of the allocation slot that is numbered $E1$ and located in the current beacon period. The value of this field is not to be smaller than the value of the Earliest RAP1 End field in any Connection Assignment frame sent by the hub transmitting this beacon.

5.3.1.6 RAP2 Start

The RAP2 Start field is set to $S2$ such that random access phase 2 (RAP2) starts at the beginning of the allocation slot that is numbered $S2$ and located in the current beacon period if either exclusive access phase 2 (EAP2) or RAP2 is of nonzero length, or is set to zero otherwise. If EAP2 is of nonzero length, it ends at the time indicated by this field. The start time of EAP2 is encoded in Connection Assignment frames defined in 5.3.6.15.

Both EAP2 (if of nonzero length) and RAP2 (if of nonzero length) occur in the current beacon period if beacon shifting is not enabled or in the next beacon period if beacon shifting is enabled.

5.3.1.7 RAP2 End

The RAP2 End field is set to $E2$ such that RAP2 ends at the end of the allocation slot that is numbered $E2$ and located in the current beacon period if RAP2 is of nonzero length, or is set to zero otherwise.

5.3.1.8 MAC Capability

The MAC Capability is as defined in 5.6.1.

5.3.1.9 PHY Capability

The PHY Capability is as defined in 5.6.2.

5.3.1.10 Beacon Shifting Sequence

The Beacon Shifting Sequence is present only if beacon shifting is currently enabled as encoded according to Figure 11. When present, it is formatted as shown in Figure 17 to indicate the beacon transmission time in the current beacon period (superframe).




Figure 17—Beacon Shifting Sequence format

5.3.1.10.1 Beacon Shifting Sequence Index

The Beacon Shifting Sequence Index field is set according to Table 4 to the index m of the PN sequence $PN_m(n)$ governing the beacon transmission time pattern.

Table 4—Beacon Shifting Sequence field encoding

Beacon Shifting Sequence Index m in decimal value	Beacon Shifting Sequence as function of Beacon Shifting Sequence Phase $n = 0, 1, 2, \dots, 15$	Beacon Shifting Sequence pattern (“...” denotes pattern repeat)
0	$PN_0(n) = n \bmod 2$	$PN_0(n) = 0, 1, 0, 1, \dots$
1	$PN_1(n) = 2 \times PN_0(n)$	$PN_1(n) = 0, 2, 0, 2, \dots$
2	$PN_2(n) = n \bmod 4$	$PN_2(n) = 0, 1, 2, 3, \dots$
3	$PN_3(n) = [PN_0(n) + PN_2(n)]/2 \bmod 2 + [PN_0(n) + PN_1(n) + PN_2(n)] \bmod 4$	$PN_3(n) = 0, 1, 3, 2, \dots$
4	$PN_4(n) = [PN_0(n) + PN_1(n) + PN_2(n)]/2$	$PN_4(n) = 0, 2, 1, 3, \dots$
5	$PN_5(n) = \{PN_2(n) + [PN_0(n) + PN_2(n)]/2\} \bmod 4$	$PN_5(n) = 0, 2, 3, 1, \dots$
6	$PN_6(n) = PN_1(n) + \{[PN_0(n) + PN_2(n)]/2 \bmod 2\}$	$PN_6(n) = 0, 3, 1, 2, \dots$
7	$PN_7(n) = [PN_1(n) + PN_2(n)] \bmod 4$	$PN_7(n) = 0, 3, 2, 1, \dots$
8–15	Reserved	Reserved

5.3.1.10.2 Beacon Shifting Sequence Phase

The Beacon Shifting Sequence Phase field is incremented by one, modulo 16, from its value applicable to the last beacon period (superframe), active or inactive.

5.3.1.11 Channel Hopping State

The Channel Hopping State field is present only if channel hopping is currently enabled as encoded according to Figure 11. When present, it is set to the current state of a 16-bit maximum-length linear feedback shift register (LFSR) used to generate the channel hopping sequence by the hub sending this beacon, as further specified in 6.13.2.

5.3.1.12 Next Channel Hop

The Next Channel Hop field is present only if channel hopping is currently enabled as encoded according to Figure 11. When present, it is set to the sequence number of the beacon period (superframe) in which the hub sending the current beacon is to hop to another channel according to its channel hopping sequence. Channel hopping is independent of inactive duration as defined in 5.3.1.13.

5.3.1.13 Inactive Duration

The Inactive Duration field is present only if one or more inactive superframes are starting at the end of the current beacon period (superframe) as encoded by the Non-final Fragment/Cancel/Scale/Inactive field according to 5.2.1.1.13. When present, it is set to the number of inactive superframes after each active superframe.

5.3.2 Security Association

A Security Association frame contains a Frame Payload that is formatted as shown in Figure 18. It is exchanged between a node and a hub during the execution of a security association protocol to activate a pre-shared MK or generate a new shared MK.




Figure 18—Frame Payload format for Security Association frames

5.3.2.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame, or is set to zero if such an EUI-48 is yet unknown.

5.3.2.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.2.3 Security Suite Selector

The Security Suite Selector is formatted as shown in Figure 19.




Figure 19—Security Suite Selector format

5.3.2.3.1 Security Association Protocol

The Security Association Protocol field is set according to Table 5 to select a desired security association protocol.

Table 5—Security Association Protocol field encoding

Field value in decimal	Security association protocol
0	Master key pre-shared association
1	Unauthenticated association
2	Public key hidden association
3	Password authenticated association
4	Display authenticated association
5–7	Reserved

5.3.2.3.2 Security Level

The Security Level field is set to the security level required by this sender according to Table 6.

Table 6—Security Level field encoding

Field value in decimal	Security level required
0	Level 0—unsecured communication
1	Level 1—authentication but not encryption
2	Level 2—authentication and encryption
3	Reserved

5.3.2.3.3 Control Frame Authentication

The Control Frame Authentication field is set to one if control type frames sent from or to this sender needs to be authenticated but not encrypted when they are required to have security level 1 or 2. It is set to zero if control type frames sent from or to this sender needs to be neither authenticated nor encrypted even when they are otherwise required to have security level 1 or 2.

5.3.2.3.4 Cipher Function

The Cipher Function field is set according to Table 7 to indicate the underlying cipher function selected by this sender for performing security services as specified in Clause 7.

Table 7—Cipher Function field encoding

Field value in decimal	Cipher function
0	AES-128 forward cipher function
1	Camellia-128 forward cipher function
2–15	Reserved

5.3.2.4 Association Control

The Association Control is formatted as shown in Figure 20.




Figure 20—Association Control format

5.3.2.4.1 Association Sequence Number

The Association Sequence Number field is set to the number (i.e., position) of the current Security Association frame in the execution of the selected security association protocol. In particular, it is set to one in the first Security Association frame of the protocol, two in the second Security Association frame, etc. The first Security Association frame is the Security Association frame transmitted or retransmitted by the node initializing the security association, the second Security Association frame is the next Security Association frame transmitted or retransmitted by the responding hub, and the like.

5.3.2.4.2 Association Status

The Association Status field in the second Security Association frame of a security association protocol (procedure) is set according to Table 8 to indicate the status of the current security association procedure. It is reserved in other Security Association frames.

Table 8—Association Status field encoding

Field value in decimal	Association status
0	Joining the security association procedure
1	Aborting the security association procedure with a different security suite selector
2	Aborting the security association procedure due to lack of needed security credential (no MK pre-shared with the initiator of the security association procedure for MK pre-shared association; no public key of the initiator of the security association procedure for public key hidden association; no password shared with the initiator of the security association procedure for password authenticated association)
3	Aborting the security association procedure due to temporary lack of resources
4	Aborting the security association procedure due to security policy restrictions as imposed by the administrator/owner of this hub on the communications in its BAN
5–15	Reserved

5.3.2.5 Security Association Data

The Security Association Data field is specific to the selected security association protocol.

For MK pre-shared association, the Security Association Data field is not present.

For unauthenticated association, public key hidden association, password authenticated association, and display authenticated association, the Security Association Data field is formatted as shown in Figure 21.




Figure 21 —Security Association Data format for security association protocols 1 to 4

5.3.2.5.1 Sender Nonce

The Sender Nonce field is set to a statistically unique number per sender and per security association procedure, except otherwise indicated, as follows:

- a) For unauthenticated association, public key hidden association, and password authenticated association
 - 1) In the first Security Association frame of the security association procedure, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders.
 - 2) In the second Security Association frame of the procedure, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders if the sender of the frame is to join the security association procedure, or is set to zero otherwise.
 - 3) In the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
 - 4) In the fourth Security Association frame, the field is set to its value contained in the first Security Association frame of the procedure.
- b) For display authenticated association
 - 1) In the first Security Association frame of the security association procedure, the field is set to a witness (message authentication code) of its value contained in the fourth Security Association frame of the procedure.
 - 2) In the second Security Association frame of the procedure, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders if the sender of the frame is to join the security association procedure, or is set to zero otherwise.
 - 3) In the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
 - 4) In the fourth Security Association frame, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders.

5.3.2.5.2 Sender PK_x

The Sender PK_x field is set to the *X*-coordinate of the sender's elliptic curve public key, except otherwise indicated, as follows:

- a) For unauthenticated association and display authenticated association,
 - 1) in the first Security Association frame of the current security association procedure, the field is set to the *X*-coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$;
 - 2) in the second Security Association frame of the procedure, the field is set to the *X*-coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;

- 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure;
 - 4) in the fourth Security Association frame, the field is set to its value contained in the first Security Association frame of the procedure.
- b) For public key hidden association,
- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the second Security Association frame, the field is set to the X -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
- c) For password authenticated association,
- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to the X -coordinate of the sender's password-scrambled elliptic curve public key $PK' = (PK'_x, PK'_y)$;
 - 2) in the second Security Association frame, the field is set to the X -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.

5.3.2.5.3 Sender PK_y

The Sender PK_y field is set to the Y -coordinate of the sender's elliptic curve public key, except otherwise indicated, as follows:

- a) For unauthenticated association and display authenticated association,
- 1) in the first Security Association frame of the current security association procedure, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$;
 - 2) in the second Security Association frame of the procedure, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure;
 - 4) in the fourth Security Association frame, the field is set to its value contained in the first Security Association frame of the procedure.
- b) For public key hidden association,
- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the second Security Association frame, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
- c) For password authenticated association,

- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to the Y -coordinate of the sender's password-scrambled elliptic curve public key $PK' = (PK'_x, PK'_y)$;
- 2) in the second Security Association frame, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
- 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.

5.3.2.5.4 MK_KMAC

The MK_KMAC field is set to a key message authentication code (KMAC) for certain fields of the frame payloads of the Security Association frames of the current security association procedure, except otherwise indicated, as follows:

- a) For unauthenticated association and display authenticated association,
 - 1) in the first and second Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the third Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the first and second Security Association frames of the procedure;
 - 3) in the fourth Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the second and fourth Security Association frames.
- b) For public key hidden association,
 - 1) in the first and second Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the third Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the first and second Security Association frames of the procedure;
 - 3) in the fourth Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the second and fourth Security Association frames.
- c) For password authenticated association,
 - 1) in the first and Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the third Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the first and second Security Association frames of the procedure;
 - 3) in the fourth Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the second and fourth Security Association frames.

5.3.3 Security Disassociation

A Security Disassociation frame contains a Frame Payload that is formatted as shown in Figure 22. It is transmitted by either an associated node or a hub to repeal an existing security association, i.e., the shared MK.