

IEEE Standard for Smart Energy Profile Application Protocol

IEEE Communications Society

Sponsored by the
Power Line Communications Committee

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

IEEE Std 2030.5™-2018
(Revision of
IEEE Std 2030.5-2013)

IEEE Standard for Smart Energy Profile Application Protocol

Sponsor

Power Line Communications Committee
of the
IEEE Communications Society

Approved 7 May 2018

IEEE-SA Standards Board

Abstract: The application layer with TCP/IP providing functions in the transport and Internet layers to enable utility management of the end user energy environment, including demand response, load control, time of day pricing, management of distributed generation, electric vehicles, etc. is defined in this standard. Depending on the physical layer in use (e.g., IEEE 802.15.4™, IEEE 802.11™, IEEE 1901™, IEEE 1901.2™), a variety of lower layer protocols may be involved in providing a complete solution. Generally, lower layer protocols are not discussed in this standard except where there is direct interaction with the application protocol. The mechanisms for exchanging application messages, the exact messages exchanged including error messages, and the security features used to protect the application messages are defined in this standard. With respect to the Open Systems Interconnection (OSI) network model, this standard is built using the four layer Internet stack model. The defined application profile sources elements from many existing standards, including IEC 61968 and IEC 61850, and follows a RESTful architecture (Fielding [B3]) using IETF protocols such as HTTP.

Keywords: adoption, application, application protocol, demand response, distributed energy resources, energy usage information, IEEE 2030.5™, load control, metering, plugin electric vehicles, prepayment, pricing communication, RESTful, SEP 2, smart energy, smart energy profile, Smart Energy Profile 2

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

Copyright © 2018 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 21 December 2018. Printed in the United States of America.

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

2030 is a registered trademark of The Institute of Electrical and Electronics Engineers, Incorporated.

PDF: ISBN 978-1-5044-4907-6 STD23130
Print: ISBN 978-1-5044-4908-3 STDPD23130

IEEE prohibits discrimination, harassment, and bullying.

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

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

Important Notices and Disclaimers Concerning IEEE Standards Documents

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

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

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

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

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

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

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

Translations

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

Official statements

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

Comments on standards

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

Comments on standards should be submitted to the following address:

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

Laws and regulations

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

Copyrights

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

Photocopies

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

Updating of IEEE Standards documents

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

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

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

Errata

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

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken by the IEEE with respect to the existence or validity of any patent rights in connection therewith. If a patent holder or patent applicant has filed a statement of assurance via an Accepted Letter of Assurance, then the statement is listed on the IEEE-SA Website at <http://standards.ieee.org/about/sasb/patcom/patents.html>. Letters of Assurance may indicate whether the Submitter is willing or unwilling to grant licenses under patent rights without compensation or under reasonable rates, with reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to obtain such licenses.

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

Participants

At the time this IEEE standard was completed, the Smart Energy Profile 2.0 Working Group had the following membership:

Robert Heile, *Chair*
Robby Simpson, *Vice Chair*

Electric Power Research Institute, Inc.	Landis+Gyr
General Electric	Sensus
Itron, Inc.	SunSpec Alliance
KITU Systems	Wi-SUN Alliance

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

Electric Power Research Institute, Inc.	Tim Godfrey
General Electric	Robby Simpson
Itron, Inc.	Gene Falendysz
KITU Systems	Gordon Lum
Panasonic Corporation of North America	Michael Stelts
Power Plus Communications AG	Scott Willy
Sensus	Michael Cowan
Silver Spring Networks, Inc.	Don Sturek
Wi-SUN Alliance	Robert Heile

When the IEEE-SA Standards Board approved this standard on 7 May 2018, it had the following membership:

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

Chuck Adams	Michael Janezic	Robby Robson
Masayuki Ariyoshi	Thomas Koshy	Dorothy Stanley
Ted Burse	Joseph L. Koepfinger*	Adrian Stephens
Stephen Dukes	Kevin Lu	Mehmet Ulema
Doug Edwards	Daleep Mohla	Phil Wennblom
J. Travis Griffith	Damir Novosel	Howard Wolfman
Gary Hoffman	Ronald C. Petersen	Yu Yuan
	Annette D. Reilly	

*Member Emeritus

Introduction

This introduction is not part of IEEE Std 2030.5-2018™, IEEE Standard for Smart Energy Profile Application Protocol.

The empowerment of consumers to manage their usage and generation of energy is a critical feature of the Smart Grid and is a basis of innovation for new products and services in energy management. To enable this capability, information flow between devices such as meters, smart appliances, plug-in electric vehicles, energy management systems, and distributed energy resources (including renewable energy and storage elements) must occur in an open, standardized, secure, and interoperable fashion. The following standard is intended to fulfill those needs.

The first publication of this standard (IEEE Std 2030.5-2013) was driven by, and sought to address the requirements of, many activities across the globe. Of note were the efforts within the United States by the National Institute of Standards and Technology (NIST) and the Smart Grid Interoperability Panel (SGIP) (in particular, Priority Action Plans 3, 9, 10, 11, and 18, with influence from many of the others) in fulfillment of the EISA 2007 legislation, the European Mandate on Smart Metering (M/441) (in particular, efforts within CEN/CENELEC and ETSI, and the Smart Meter Working Group), as well as similar efforts in Australia, the United Kingdom, Japan, and China, and electric vehicle standardization efforts (in particular, ISO/IEC JWG automotive EV standards and SAE EV standards), to name only a few.

This revision of IEEE Std 2030.5 was made with particular attention to the activities underway in California as part of the Rule 21 revision and the associated Smart Inverter Working Group as well as the revision of IEEE Std 1547™. This revision also seeks to address any errors and ambiguities discovered in the testing and deployment of the first publication.

This standard is also intended to enable communications that are link-layer agnostic and run over the Internet Protocol. Careful consideration was given to premises networks with various architectures, numbers of devices, and constraints, while maintaining flexibility, extensibility, and security.

Contents

1. Overview	11
1.1 Scope	11
1.2 Purpose	11
1.3 Document organization.....	11
1.4 Requirement language	12
1.5 Typography conventions used	12
1.6 Design principles	12
2. Normative references.....	12
3. Definitions, acronyms, and abbreviations	14
3.1 Definitions	14
3.2 Acronyms and abbreviations	16
4. Design pattern.....	18
4.1 Protocol flexibility	18
4.2 General rules/best practices	18
4.3 WADL	19
4.4 Schema	19
4.5 Uniform resource identifiers.....	20
4.6 List resources	20
4.7 Resource design rules	24
5. Application support	24
5.1 Overview	24
5.2 Use of TCP	24
5.3 URI encoding.....	25
5.4 HTTP headers	25
5.5 HTTP response codes	27
5.6 Application payload syntax	30
5.7 Content negotiation.....	31
6. Security.....	32
6.1 Introduction	32
6.2 Security attributes	32
6.3 Device credentials.....	39
6.4 Resource access authentication and authorization context	40
6.5 Resource access authentication.....	42
6.6 Resource access authorization	42
6.7 Cipher suites	42
6.8 Default security policy.....	43
6.9 Registration.....	44
6.10 Security LogEvents.....	46
6.11 Certificate management.....	46
7. Discovery.....	59
7.1 Introduction	59
7.2 Service instance	60
7.3 Service name.....	61
7.4 TXT record	61
7.5 Subtype queries.....	62

7.6 Discovery procedure.....	64
8. Support resources	64
8.1 Introduction	64
8.2 Resource section outlines	65
8.3 Device Capabilities function set	66
8.4 Self Device function set.....	67
8.5 End Device function set.....	67
8.6 Function Set Assignments function set.....	69
8.7 Subscription/Notification function set	70
8.8 Response function set	73
9. Common resources	77
9.1 Introduction	77
9.2 Time function set.....	77
9.3 Device Information function set	79
9.4 Power Status function set.....	79
9.5 Network Status function set.....	80
9.6 Log Event function set.....	81
9.7 Configuration function set	83
9.8 File Download function set.....	83
10. Smart energy resources.....	88
10.1 Introduction	88
10.2 Common application functionality	88
10.3 Demand Response and Load Control function set.....	95
10.4 Metering function set.....	99
10.5 Pricing function set.....	105
10.6 Messaging function set	109
10.7 Billing function set	111
10.8 Prepayment function set	114
10.9 Flow Reservation function set	116
10.10 Distributed Energy Resources function set.....	117
10.11 Metering Mirror function set	128
11. Manufacturer-specific proprietary extensions	130
11.1 Overview	130
11.2 xmDNS/DNS-SD.....	130
11.3 URIs.....	130
11.4 Resources.....	130
11.5 DeviceCapabilities resource	131
Annex A (informative) Web-application description language (WADL)	132
A.1 Introduction	132
A.2 Support resources section	132
A.3 Common resources section	136
A.4 Smart energy resources section.....	139
Annex B (informative) IEEE 2030.5 model	153
B.1 Introduction.....	153
B.2 IEEE 2030.5 package.....	153
Annex C (informative) Examples and guideline	267
C.1 Introduction.....	267
C.2 Registration: Remote	267
C.3 Registration: Local.....	270

C.4 Discovery: Function Set Assignment.....	272
C.5 Discovery: Without Function Set Assignment.....	274
C.6 Discovery: Undirected without Function Set Assignment	276
C.7 Subscription/Notification	277
C.8 Demand response: General	280
C.9 Demand response: Cancel.....	284
C.10 Distributed energy resource: General.....	286
C.11 Metering: Reading	290
C.12 Metering: Interval	295
C.13 Metering: Instantaneous	302
C.14 Metering: Mirroring	305
C.15 Pricing: Time of use.....	317
C.16 Billing: Billing period.....	322
C.17 Billing: Historical.....	324
C.18 Billing: Projection.....	328
C.19 File loading	330
C.20 Flow Reservation: General.....	334
C.21 Flow Reservation: Cancel	339
C.22 Event randomization	343
Annex D (informative) Guidelines	348
D.1 Pricing implementation guidelines	348
D.2 PEV implementation guidelines (subject to work with SAE and ISO/IEC)	353
Annex E (informative) Mapping to IEEE Std 1547-2018	354
Annex F (informative) Bibliography	359

IEEE Standard for Smart Energy Profile Application Protocol

1. Overview

1.1 Scope

This standard defines the application layer with TCP/IP providing functions in the transport and Internet layers to enable utility management of the end user energy environment, including demand response, load control, time of day pricing, management of distributed generation, electric vehicles, etc. Depending on the physical layer in use (e.g., IEEE 802.15.4™, IEEE 802.11™, IEEE 1901™, IEEE 1901.2™), a variety of lower layer protocols may be involved in providing a complete solution. Generally, lower layer protocols are not discussed in this standard except where there is direct interaction with the application protocol. This standard defines the mechanisms for exchanging application messages, the exact messages exchanged including error messages, and the security features used to protect the application messages. With respect to the Open Systems Interconnection (OSI) network model, this standard is built using the four layer Internet stack model. The defined application profile sources elements from many existing standards, including IEC 61968 and IEC 61850, and follows a RESTful architecture (Fielding [B3]) using IETF protocols such as HTTP.

1.2 Purpose

The purpose of this document is to define the application protocol to enable utility management of the end user energy environment, including demand response, load control, time of day pricing, management of distributed generation, electric vehicles, etc. The defined application profile sources elements from many existing standards, including IEC 61968 and IEC 61850, and follows a RESTful architecture (Fielding [B3]) using IETF protocols such as HTTP.

1.3 Document organization

The following documents comprise the definition of IEEE 2030.5 and all IEEE 2030.5 devices will be required to maintain compliance to these documents:

- IEEE Std 2030.5 (this document)
- IEEE 2030.5 XML Schema Definition (XSD) (sep.xsd in the supplemental material of IEEE Std 2030.5)
- IEEE 2030.5 WADL (sep_wadl.xml in the supplemental material of IEEE Std 2030.5)

The IEEE 2030.5 XSD contains the definitions of the IEEE 2030.5 resources, attributes, and elements as well as their textual descriptions. The IEEE 2030.5 WADL contains the recommended URI structures and use of HTTP methods associated with these objects. In addition, an IEEE 2030.5 UML model (also contained in the supplemental material of IEEE Std 2030.5) has been utilized in the creation of the IEEE 2030.5 XSD. The IEEE 2030.5 UML model is an informative document. Informative textual extracts of the IEEE 2030.5 XSD and the IEEE 2030.5 WADL are contained in this document for convenience.

1.4 Requirement language

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document, when capitalized and in a normative section, constitute normative text and are to be interpreted as described in IETF RFC 2119.

1.5 Typography conventions used

Example URIs, protocol requests, protocol responses, and XML are presented in `fixed-width Courier New font, with a size of 8.`

1.6 Design principles

As IEEE Std 2030.5 follows a RESTful architecture (Fielding [B3]), several design principles are important for the standard, including:

- While devices MAY maintain state, interfaces SHOULD be stateless.
- URI structure SHOULD be clear but as efficient as possible.
- Minimize the number of transactions required to achieve a given function.

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.

IEC 61968, Application integration at electric utilities—System interfaces for distribution management—All parts.¹

IEEE Std 2030.5™, Supplemental material (<http://standards.ieee.org/downloads/2030.5/>).²

IEEE Std 802.1AR™, IEEE Standard for Local and Metropolitan Area Networks—Secure Device Identity (<http://www.ieee802.org/1/pages/802.1ar.html>).

¹ IEC publications are available from the International Electrotechnical Commission (<http://www.iec.ch>) and the American National Standards Institute (<http://www.ansi.org/>).

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

IETF, Extended Multicast DNS (<https://tools.ietf.org/html/draft-lynn-homenet-site-mdns-01>).

IETF RFC 793, Transmission Control Protocol (<http://tools.ietf.org/html/rfc793>).

IETF RFC 1630, Universal Resource Identifiers in WWW (<http://tools.ietf.org/html/rfc1630>).

IETF RFC 2119, Key Words for Use in RFCs to Indicate Requirement Levels (<http://tools.ietf.org/html/rfc2119>).

IETF RFC 2818, HTTP Over TLS (<http://tools.ietf.org/html/rfc2818>).

IETF RFC 2863, The Interface MIB (<https://tools.ietf.org/html/rfc2863>).

IETF RFC 3986, Uniform Resource Identifier (URI): Generic Syntax (<https://tools.ietf.org/html/rfc3986>).

IETF RFC 4108, Using Cryptographic Message Syntax (CMS) to Protect Firmware (<http://tools.ietf.org/html/rfc4108>).

IETF RFC 4193, Unique IPv6 Unicast Addresses (<http://tools.ietf.org/html/rfc4193>).

IETF RFC 4291, IP Version 6 Addressing Architecture (<http://tools.ietf.org/html/rfc4291>).

IETF RFC 5246, Transport Layer Security (TLS) Protocol Version 1.2 (<http://tools.ietf.org/html/rfc5246>).

IETF RFC 5280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile (<http://tools.ietf.org/html/rfc5280>).

IETF RFC 5480, Elliptic Curve Cryptography Subject Public Key Information (<http://tools.ietf.org/html/rfc5480>).

IETF RFC 6550, RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. (<http://tools.ietf.org/html/rfc6550>).

IETF RFC 6554, An IPv6 Routing Header for Source Routes with the Routing Protocol for Low-Power and Lossy Networks (RPL) (<http://tools.ietf.org/html/rfc6554>).

IETF RFC 6763, DNS-Based Service Discovery (<http://tools.ietf.org/html/rfc6763>).

IETF RFC 6960, X.509 Internet Public Key Infrastructure Online Certificate Status Protocol—OCSP (<https://tools.ietf.org/html/rfc6960>).

IETF RFC 7230, Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing (<https://tools.ietf.org/html/rfc7230>).

IETF RFC 7231, Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content (<https://tools.ietf.org/html/rfc7231>).

IETF RFC 7232, Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests (<https://tools.ietf.org/html/rfc7232>).

IETF RFC 7233, Hypertext Transfer Protocol (HTTP/1.1): Range Requests (<https://tools.ietf.org/html/rfc7233>).

IETF RFC 7234, Hypertext Transfer Protocol (HTTP/1.1): Caching (<https://tools.ietf.org/html/rfc7234>).

IETF RFC 7235, Hypertext Transfer Protocol (HTTP/1.1): Authentication (<https://tools.ietf.org/html/rfc7235>).

IETF RFC 7251, AES-CCM Elliptic Curve Cryptography (ECC) Cipher Suites for TLS (<http://tools.ietf.org/html/rfc7251>).

Internet Assigned Numbers Authority (IANA), Private Enterprise Number (PEN) Request.³

ISO 4217, Codes for the representation of currencies.⁴

ITU-T Recommendation X.680 (08/2015), Information technology—Abstract Syntax Notation One (ASN.1): Specification of basic notation.⁵

NIST Special Publication (SP) 800-57, Recommendation for Key Management.⁶

NIST Special Publication (SP) 800-131A, Transition: Recommendation for Transitioning the Use of Cryptographic Algorithms and Key Lengths.

W3C, Efficient XML Interchange (EXI) Format 1.0 (<http://www.w3.org/TR/exi/>).⁷

W3C, Extensible Markup Language (XML) (<http://www.w3.org/TR/REC-xml/>).

W3C, Web Application Description Language (WADL) (<http://www.w3.org/Submission/wadl/>).

3. Definitions, acronyms, and abbreviations

3.1 Definitions

For the purposes of this document, the following terms and definitions apply. The *IEEE Standards Dictionary Online* should be consulted for terms not defined in this clause.⁸

access control list: A security mechanism in which entities and authorizations (e.g., read, write, create, delete) are related to resources to determine the entities' allowed operations on the resources.

certificate authority: An entity that issues digital certificates for use by other entities.

certificate chain: A chain of certificates, with each certificate's signature verified using the key from the next certificate in the chain. The single exception is the certificate at the end of the chain (the trust anchor), known as the root certificate authority (CA) certificate, which is self-signed.

client: The device or host that interacts with a server to obtain information related to a resource hosted by the server.

³ Available online: <http://pen.iana.org/pen/PenApplication.page>.

⁴ ISO publications are available from the International Organization for Standardization (<http://www.iso.org/>) and the American National Standards Institute (<http://www.ansi.org/>).

⁵ ITU-T publications are available from the International Telecommunications Union (<http://www.itu.int/>).

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

⁷ W3C publications are available from the World Wide Web Consortium (<https://www.w3.org/>).

⁸ *IEEE Standards Dictionary Online* subscription is available at: <http://dictionary.ieee.org>.

device certificate: A digital certificate installed within a device that binds the device identity to the device. Device certificates are exchanged by network access control and application protocols to authenticate devices as genuine IEEE 2030.5 and further to prove specific device identity.

energy services interface (ESI): A device, with multiple network interfaces, which is a member of both the home smart energy network and a service provider's private network. This is the primary mechanism for the service provider to contribute data and directives into the smart energy network and to receive responses from smart energy devices.

fingerprint: This is the result of summarizing a certificate with a secure hash function. The fingerprint is generally expressed as a hex string. It is used to confirm the integrity of a certificate obtained over an untrusted channel.

function set: A logical grouping of resources that cooperate to implement IEEE 2030.5 features (e.g., metering, demand response, and load control).

Function Set Assignments: A logical addressing mechanism in IEEE Std 2030.5 that allows devices to be directed to use specific resources (e.g., to facilitate a device's participation in a program). Please see 8.6 for details.

function set instance: A single, top-level, instance of a function set (e.g., a single UsagePoint for the Metering function set, a single DemandResponseProgram for the Demand Response and Load Control function set).

host: This is the representation of a device in its application context. Typically represented by an Internet Protocol (IP) address or domain name.

intermediate CA: A certificate authority (CA) below the root CA that issues certificates to subordinate CAs.

issuing CA: A certificate authority (CA) that issues certificates to devices or code-signers.

manufacturer's CA (MCA): An intermediate certificate authority (CA) operated by a specific manufacturer for the purpose of issuing manufacturing issuing CAs for that manufacturer.

manufacturing issuing CA (MICA): An issuing certificate authority (CA) that issues certificates to devices during the manufacturing process.

manufacturing public key infrastructure (PKI): The set of certificate authorities (CAs) that issue certificates to devices during the manufacturing process. The set includes the Smart Energy Root CA, Manufacturer's CAs, and Manufacturing Issuing CAs.

master resource identifier (mRID): An mRID is the global identifier attribute of an object. For further definition and formatting, see the schema (IEEE Std 2030.5 supplemental material).

node: This is the representation of a device in its network context, typically represented as an Internet Protocol (IP) address.

object identifier (OID): An OID consists of a node in a hierarchically-assigned namespace, formally defined using ITU-T's ASN.1 standard (ITU-T Recommendation X.680).

privacy-enhanced electronic mail (PEM) format certificate: An X.509 certificate that has been Base64 encoded and wrapped in "-----BEGIN CERTIFICATE-----", "-----END CERTIFICATE-----" sentinels for transport as a text file or block.

public key infrastructure (PKI): A set of hardware, software, people, policies, and procedures needed to create, manage, distribute, use, store, and revoke digital certificates.

registered: The state of a device with regards to a particular server wherein an EndDevice record for the device has been populated on the server and that record contains a valid registration record. These records are typically populated with device information transmitted out-of-band to the server's owner.

resource: Uniform resource identifier (URI) addressable object that is manipulated via the RESTful uniform interface.

resource discovery: The process whereby clients identify resources being served on the network. Clients issue a request to all devices on the network requesting resource(s) of interest. Servers hosting the requested resource(s) respond with information necessary to access the server and its resource(s).

root CA: A certificate authority (CA) whose certificate or public key is a trust anchor for any other certificates in a chain of trust.

root certificate: The root certificate authority's self-signed certificate. Generally also a trust anchor.

self-signed certificate: A certificate whose issuer and subject are identical, and whose public key verifies its signature.

server: The device or host that holds a resource and exposes representations of that resource.

Smart Energy Root CA (SERCA): The top-level certificate authority (CA) for the IEEE 2030.5 manufacturing public key infrastructure (PKI).

trust anchor: The root of trust for a certificate chain. This is an authoritative entity represented by a public key and associated data and is generally provided in the IEEE 2030.5 hierarchy in the form of a self-signed certificate.

trusted root store: An integrity-protected location for storing root certificates.

3.2 Acronyms and abbreviations

This subclause provides a list of acronyms and abbreviations introduced in this document.

ACL	access control list
CA	certificate authority
CRL	certificate revocation list
DER	distributed energy resource
DNS	Domain Name System
DNS-SD	DNS-based Service Discovery
DRLC	demand response and load control
ECDSA	Elliptic Curve Digital Signature Algorithm

EMS	energy management system
ESI	energy services interface
EVSE	electric vehicle supply equipment
EXI	Efficient XML Interchange
HEMS	home energy management system
HTTP	Hypertext Transfer Protocol
IETF	Internet Engineering Task Force
IP	Internet Protocol
LFDI	long form device identifier
MCA	manufacturer's certificate authority
MICA	manufacturing issuing certificate authority
OCSP	Online Certificate Status Protocol
OID	object identifier
PKI	public key infrastructure
REST	representational state transfer
SERCA	smart energy root certificate authority
SFDI	short form device identifier
TCP	Transmission Control Protocol
UOM	units of measure
URI	uniform resource identifier
WADL	Web Application Description Language
XML	Extensible Markup Language

4. Design pattern

4.1 Protocol flexibility

IEEE Std 2030.5 is designed to implement a REST architecture. It is built around the core actions of GET, HEAD, PUT, POST, and DELETE (as used in (Fielding [B3])), with the addition of a lightweight subscription mechanism as discussed in 8.7. Any application protocol that can implement a RESTful command set could likely be used with IEEE Std 2030.5, but Hypertext Transfer Protocol (HTTP) is a required baseline for interoperable IEEE 2030.5 implementations. HTTP utilizes Transmission Control Protocol (TCP) as its transport protocol. As a result, TCP manages the session providing delivery assurance and windowing.

IEEE 2030.5 servers and clients SHALL be compliant with IETF RFC 7230, IETF RFC 7231, IETF RFC 7232, IETF RFC 7233, IETF RFC 7234, and IETF RFC 7235.

4.2 General rules/best practices

This standard shall not make distinctions between servers, clients, or devices when defining interfaces and uniform resource identifiers (URIs). The goal is to avoid having a resource that has one behavior on a server and a different behavior on a client or different type of server. The distinction between the server and client role depends on whether a device exposes a resource (server) or interacts with the resource (client).

The default mechanism for obtaining a resource representation is a pull mechanism, implemented with GET. A client requests and retrieves data from a server or creates, modifies, or deletes data on a server.

The default polling rate for a given function set is specified at the top-level resource of that function set in the pollRate attribute.

The use of a subscription mechanism for retrieving a resource representation is also optionally supported, where convenient and appropriate. Resources that support subscription are denoted by their subscribable attribute.

Objects have a defined granularity and whole objects (not partial objects) are to be updated with the granularity defined in the schema.

Clients that expect to have intermittent connections to the network (e.g., battery-powered sleepy devices, mobile devices, etc.) use a pull mechanism as their default behavior for resource retrieval, as a Subscription/Notification mechanism may not be reliable. It should be noted that clients that expect to have intermittent connections to the network may still POST, PUT, and DELETE resources, provided they have the appropriate security permissions.

The TCP ports used for HTTP or Hypertext Transfer Protocol Secure (HTTPS) SHALL be specified in the xmdns service advertisement for the service.

Content SHALL be transferred with either one of the content types: “application/sep+xml” or “application/sep-exi.”

Devices do not assume the use of the URIs and their structures given throughout this standard. All resources are self-describing as it is acknowledged that URIs, schemas, and resources might change in the future. All resources SHALL contain links to their subordinate resources to support flexibility in URIs and future extensibility. Thus, to allow for extensibility and granularity, all objects are described in schemas and referenced via URIs. The URIs presented throughout this standard are recommendations. Thus, clients

do not assume that URIs for resources are fixed on all servers or even on a given server (over time), but rather retrieve the appropriate URIs through resource discovery and links within resources. For network efficiency, devices MAY assume URIs are fixed on a particular server over time. If a URI returns an unexpected result, the client SHOULD execute resource discovery to determine the new URI value.

Version information should not be presented in the URI unless that version information is inherent to the name of that resource. If necessary and for reasons of extensibility, version information is provided within the associated resources and/or schemas.

All values in XML and EXI SHOULD be represented as compactly as possible. Decimal values SHOULD be represented without leading zeros. Hexadecimal values SHALL be represented with one leading zero, if needed, to ensure an even number of digits.

4.3 WADL

The IEEE 2030.5 RESTful interface is defined using the Web Application Description Language (WADL) (IEEE Std 2030.5 supplemental material).

IEEE 2030.5 devices SHALL conform to the interface specifications contained in the WADL as follows:

- a) Devices SHALL conform to the WADL specification as per W3C WADL.
- b) Devices SHALL conform to the WADL definition in the supplemental material of IEEE Std 2030.5. By implication, all resource representations SHALL validate per the schema (IEEE Std 2030.5 supplemental material) within the standardized IEEE 2030.5 XML namespace (`urn:ieee:std:2030.5:ns`).
- c) Compliance (MODE) designations are interpreted as follows:
 - 1) Mandatory ("M"): Devices MUST implement and conform.
 - 2) Optional ("O"): Devices MAY implement, and if implemented, MUST conform.
 - 3) Discouraged ("D"): Devices SHOULD NOT implement, but if implemented, MUST conform.
 - 4) Error ("E"): Devices MUST return one of the specified response status codes (e.g., 400 Bad Request or 405 – Method Not Allowed).

4.4 Schema

Resources located at URIs returned in the href attribute of "Link" specializations (e.g., EndDeviceListLink, SelfDeviceLink) SHALL conform to the schema definition for that object, which is the tag name with the "Link" suffix removed. For example, the resource at SelfDeviceLink href follows the definition for the SelfDevice resource.

If a client PUTs or POSTs a resource to a server containing attributes or elements that instead are to be populated by the server (e.g., href), the server SHALL return an HTTP 400 error.

If a function set is not implemented, Link elements to resources in that function set SHALL NOT be included.

4.5 Uniform resource identifiers

HTTP uses ASCII text for transferring URIs between clients and servers, as well as for including options and details regarding the message content. These transfers occur with every transaction. If the naming scheme used for URIs is overly verbose, these transactions become needlessly inefficient on constrained networks. Of course, if the naming scheme used is overly cryptic, the advantages of a text-based protocol are lost, and it becomes difficult for someone troubleshooting to decipher a transaction.

The following conventions are used for URI naming:

- a) URI elements are at most four characters, but recognizable to a knowledgeable engineer. Element names as short as one character are acceptable provided their meaning is clear.
- b) URI elements are constructed of consonants only, unless inclusion of a vowel adds clarity, such as a leading vowel or well-known abbreviation.
- c) URI elements are in all lower case.
- d) URIs SHALL NOT be greater than 255 bytes in length. In practice, URIs SHOULD be much smaller than 80 bytes.

4.6 List resources

4.6.1 Introduction

Many resources within this standard are derived from the <List> object. Throughout this standard, these resources will collectively be referred to as *list resources*.

The following attributes are defined for list resources:

- `all` is used to indicate the total number of items (subordinate resources) that exist in the list resource. This number may vary according to the client's access privileges.
- `results` is used to indicate the number of items (subordinate resources) included in a specific subset of the list (result from a paged GET query to the list, etc.). This value will always be less than or equal to "all."

Clients and servers use these attributes, combined with the query string parameters described below, to implement paged access to lists. Client control of list paging is important for resource-constrained devices.

List items (subordinate resources) are read using one of the two idioms described below:

- 1) Ordinal access to the first, second, nth, etc. item in the list is supported via query string parameters included with a GET to the list resource URI.
- 2) Random access to specific list items is supported via a GET directly to the URI of the list item.

Some list items may be created by POSTing a list item representation to the URI of the list (e.g., Notifications), while others may be created using private interfaces over the provider network. Ordinal placement of the new resource within the list is determined by the list sort order (defined by each list).

Each function set defined in this standard that contains list resources includes a list ordering table. Each list resource has an entry in the corresponding table that describes the details of one or more unique sort keys and the precedence of those keys. A list resource's elements SHALL be ordered according to this specified list ordering.

4.6.2 Query string parameters

Query string parameters are parameters added to a URI to provide filtering/paging of list items returned in query results.

The list paging mechanism allows GET requests to specify the range of list items to be returned in a query result set. The general syntax of a paged query is as follows:

`{URI}?s={x}&a={y}&l={z}`

Where {URI} represents a URI used to address a list resource, (s | a | l) represent query string parameters (further defined below), and {x}, {y}, and {z} represent the respective query string parameter values.

The query string parameters are defined as:

- s (“start”) is used to indicate the first ordinal position in the list to be returned in the query result list as determined by the list’s ordering. The value is specified in decimal. The first ordinal position of the list SHALL be designated with a value of “0” and the maximum possible value is “4294967295.” If this query string parameter is not specified, the default start value SHALL be “0.”
- a (“after”) is used to indicate that only items whose primary key occurs after the given date/time parameter should be included in the query result list. This query string parameter is only applied to list resources that are ordered using a time-based primary key. The parameter SHALL be ignored if the primary key is not time-based. The format of the parameter SHALL be a 64-bit decimal number with identical semantics as that of the TimeType (see subclause 12.2 and the XML XSD in the supplemental material of IEEE Std 2030.5).
- l (“limit”) is used to set the maximum number of list items (up to 4 294 967 295) to be included in the query result list. The value is specified in decimal. If this query string parameter is not specified, the default limit SHALL be “1.” Servers MAY return a result list smaller than that specified by the client-provided limit.

If both a “start” and “after” query string parameter are used simultaneously, the “after” query string parameter SHALL have precedence. The “start” position 0 SHALL be relative to the position specified by the “after” parameter.

If a query string requests a List element that does not exist (e.g., s=3 when there are two items in the list), servers SHALL return an empty list representation.

Readers should note that the “after” query string parameter SHOULD NOT be used alone for paging through a list. As some list resources MAY contain multiple subordinate resources with the same time-based primary key, it is RECOMMENDED that clients wishing to paginate a list resource while using the “after” query string parameter should keep the value for the “after” query string parameter constant while changing the “start” query string parameter.

If a particular query string parameter appears more than once, then the first occurrence of the query string parameter SHALL be used (in left-to-right order) and subsequent occurrences MUST be ignored.

Server receipt of a query parameter unknown to the server MUST be ignored by the server and MUST NOT generate an HTTP error. Servers MUST NOT generate resource representations containing href attributes that contain query parameters. Clients MUST ignore query parameters contained in resource hrefs, but SHOULD NOT remove them if the URI is used for subsequent RESTful exchanges.

Should an empty list representation be requested (either through the use of query string parameters such as `l=0` or when the list itself is empty), the server SHALL return no subordinate representations, but SHALL return any other elements that may be defined for the list.

Clients MUST NOT assume any index semantics for list URIs. For example, a client desiring to read the third item from the list at `http://some-host/somelist` MUST NOT assume a GET to `http://some-host/somelist/2` will return the third item. The correct access is supported by the client issuing a GET to `http://some-host/somelist?s=2&l=1`.

The following examples demonstrate the use of query string parameters with a list resource. Consider the `MyTypeList` resource, depicted below:

```
<MyTypeList href="http://host1/the/list" all="7" results="7">
  <MyType href="http://host1/instance/of/type/red">
    <timeStamp>100</timeStamp>
  </MyType>
  <MyType href="http://host2/instance/of/type/green">
    <timeStamp>200</timeStamp>
  </MyType>
  <MyType href="http://host3/instance/of/type/blue">
    <timeStamp>300</timeStamp>
  </MyType>
  <MyType href="http://host4/instance/of/type/yellow">
    <timeStamp>400</timeStamp>
  </MyType>
  <MyType href="http://host5/instance/of/type/black">
    <timeStamp>500</timeStamp>
  </MyType>
  <MyType href="http://host6/instance/of/type/white">
    <timeStamp>600</timeStamp>
  </MyType>
  <MyType href="http://host7/instance/of/type/orange">
    <timeStamp>700</timeStamp>
  </MyType>
</MyTypeList>
```

This list is sorted in ascending order per the `<timeStamp/>` element of `MyType`.

A GET to `http://host1/the/list?s=0&l=1` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="1">
  <MyType href="http://host1/instance/of/type/red">
    <timeStamp>100</timeStamp>
  </MyType>
</MyTypeList>
```

A GET to `http://host1/the/list?s=0&l=5` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="5">
  <MyType href="http://host1/instance/of/type/red">
    <timeStamp>100</timeStamp>
  </MyType>
  <MyType href="http://host2/instance/of/type/green">
    <timeStamp>200</timeStamp>
  </MyType>
  <MyType href="http://host3/instance/of/type/blue">
    <timeStamp>300</timeStamp>
  </MyType>
  <MyType href="http://host4/instance/of/type/yellow">
    <timeStamp>400</timeStamp>
  </MyType>
  <MyType href="http://host5/instance/of/type/black">
    <timeStamp>500</timeStamp>
  </MyType>
</MyTypeList>
```

A GET to `http://host1/the/list?s=5&l=1` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="1">
  <MyType href="http://host6/instance/of/type/white">
    <timeStamp>600</timeStamp>
  </MyType>
</MyTypeList>
```

A GET to `http://host1/the/list?s=5&l=5` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="2">
  <MyType href="http://host6/instance/of/type/white">
    <timeStamp>600</timeStamp>
  </MyType>
  <MyType href="http://host7/instance/of/type/orange">
    <timeStamp>700</timeStamp>
  </MyType>
</MyTypeList>
```

A GET to `http://host1/the/list?s=12&l=2` will return a list representation containing zero items:

```
<MyTypeList href="http://host1/the/list" all="7" results="0">
</MyTypeList>
```

or `<MyTypeList href="http://host1/the/list" all="7" results="0" />`

A GET to `http://host6/instance/of/type/white` will return:

```
<MyType href="http://host6/instance/of/type/white">
  <timeStamp>600</timeStamp>
</MyType>
```

A GET to `http://host1/the/list?a=400&l=4` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="3">
  <MyType href="http://host5/instance/of/type/black">
    <timeStamp>500</timeStamp>
  </MyType>
  <MyType href="http://host6/instance/of/type/white">
    <timeStamp>600</timeStamp>
  </MyType>
  <MyType href="http://host7/instance/of/type/orange">
    <timeStamp>700</timeStamp>
  </MyType>
</MyTypeList>
```

A GET to `http://host1/the/list?a=400&s=0&l=2` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="2">
  <MyType href="http://host5/instance/of/type/black">
    <timeStamp>500</timeStamp>
  </MyType>
  <MyType href="http://host6/instance/of/type/white">
    <timeStamp>600</timeStamp>
  </MyType>
</MyTypeList>
```

A GET to `http://host1/the/list?a=400&s=2&l=2` will return:

```
<MyTypeList href="http://host1/the/list" all="7" results="1">
  <MyType href="http://host7/instance/of/type/orange">
    <timeStamp>700</timeStamp>
  </MyType>
</MyTypeList>
```

4.7 Resource design rules

The following rules apply to the design and use of list resources:

- List resources SHALL support “start” and “limit” query string parameters, thus always supporting paging.
- List resources that have a time-based primary key SHALL support the “after” query string parameters.
- All subordinate resources of list resources SHALL include an `href` attribute containing the URI of the subordinate resource.
- Each list resource described in this standard includes a list ordering that defines how the list is ordered, including the details of one or more keys used to order the list and the precedence of these keys.
- When queried, list resources SHALL return subordinate resources in the order defined by the list ordering.
- All subordinate resources of list resources that support multiple types, for example, NotificationList, SHALL include an `xsi:type` attribute. In this case, the XML Schema Instance Namespace must also be declared. See Clause 16 for examples.

The following rules apply to the design and use of non-list resources:

- Non-list resources SHALL NOT support the defined query string parameters. Query string parameters applied to non-list resource URIs SHOULD be ignored.

5. Application support

5.1 Overview

This clause details the standards-based application transport protocol and other lower layer requirements required for interoperability of IEEE 2030.5 devices.

The application support layer provides the following services:

- RESTful HTTP/1.1 as the application data exchange semantics, as described in IETF RFC 7231 and (Fielding [B3])
- XML (W3C XML) and/or EXI (W3C EXI) encoding as the data payload of the RESTful operations
- Transport authentication and encryption using HTTP over TLS (IETF RFC 2818 and IETF RFC 5246)

5.2 Use of TCP

The choice of HTTP/1.1 as the application data exchange protocol leads directly to the use of TCP as the transport protocol. IETF RFC 7230 states that:

“HTTP messaging is independent of the underlying transport- or session-layer connection protocol(s). HTTP only presumes a reliable transport with in-order delivery of requests and the corresponding in-order delivery of responses. The mapping of HTTP request and response structures onto the data units of an underlying transport protocol is outside the scope of this specification.”

Hence, albeit that other transport protocols may be used to carry HTTP, they are required to provide reliable transport. Clearly, the de facto choice is TCP (IETF RFC 793).

5.3 URI encoding

The address of the resources presented by an IEEE 2030.5 host will use the standard URI syntax specific to HTTP/1.1 (i.e., `http://`) as per IETF RFC 1630 and IETF RFC 3986.

Given the constrained nature of many IEEE 2030.5 hosts, it is critical to devise the URL namespace scheme (hierarchy) that is both descriptive and compact. For more information, see 4.5.

URIs generated by IEEE 2030.5 hosts SHOULD use only xDNS domain names or Internet Protocol (IP) addresses in fully qualified URIs, as IEEE 2030.5 devices may not have access to name servers.

5.4 HTTP headers

5.4.1 Introduction

HTTP/1.1 defines a variety of header types that have the potential to be verbose. As an example, the Accept header is illustrated below showing lengths varying from 35 octets to 106 octets:

- Accept: audio/*; q=0.2, audio/basic (35 octets)
- Accept: text/*; q=0.3, text/html; q=0.7, text/html; level=1, text/html; level=2; q=0.4, */*; q=0.5 (92 octets)
- Accept: text/xml, application/xml, application/xhtml+xml, text/html; q=0.9, text/plain; q=0.8, image/png, */*; q=0.5 (106 octets)

Given the packet size constraints for many IEEE 2030.5 networks, the set of HTTP headers supported must be curtailed through best practices recommendations realizing that a host cannot prevent a verbose message from being sent to it.

IEEE 2030.5 provides a set of *mandatory*, *optional*, and *discouraged* recommendations for each HTTP header. IEEE 2030.5 implementations SHOULD employ only mandatory HTTP headers, minimize the use of optional HTTP headers, and avoid the use of discouraged headers.

5.4.2 HTTP header field recommended usage

In the following table, the HTTP/1.1 header fields have been annotated with the following labels:

- MANDATORY: Support for the field is REQUIRED.
- OPTIONAL: Support for this field is left to the implementer's discretion.
- DISCOURAGED: To conserve code space and/or bandwidth, support for this field, while not explicitly forbidden, is not recommended.

Table 1 summarizes the recommended use of HTTP/1.1 headers in IEEE Std 2030.5. This table is provided as informative guidance, particularly to suit constrained networks and devices, and IETF RFC 7230, IETF RFC 7231, IETF RFC 7232, IETF RFC 7233, IETF RFC 7234, and IETF RFC 7235 are the authoritative references for compliance.

Table 1—HTTP headers

Header	Used in message type		IETF RFC required/optional	IEEE 2030.5 use
Accept	Request		Optional	Mandatory
Accept-Charset	Request		Optional	Discouraged
Accept-Encoding	Request		Optional	Discouraged
Accept-Language	Request		Optional	Discouraged
Accept-Ranges	Request	Response	Optional	Discouraged
Age		Response	Optional (required for a cache)	Discouraged
Allow		Response	Required	Mandatory
Authorization	Request	Response	Optional	Discouraged
Cache-Control	Request	Response	Optional	Discouraged
Connection	Request		Optional (required in some situations [e.g., HTTP/1.1 applications that do not support persistent connections])	Optional (mandatory in some situations)
Content-Encoding		Response	Optional (required when an encoding is applied)	Discouraged
Content-Language		Response	Optional	Discouraged
Content-Length	Request	Response	Optional (required in many situations, see Section 3.3.2 of IETF RFC 7230)	Optional (required in many situations, see Section 3.3.2 of IETF RFC 7230)
Content-Location		Response	Optional	Discouraged
Content-MD5		Response	Optional	Discouraged
Content-Range		Response	Optional	Optional (see 9.8.2.3)
Content-Type	Request	Response	Required	Mandatory
Cookies	Request		Optional	Discouraged
Date	Request	Response	Mandatory	Mandatory
ETag		Response	Optional	Optional (see 9.8.2.3)
Expect	Request		Optional	Discouraged
Expires		Response	Optional	Discouraged
From	Request		Optional	Discouraged
Host	Request		Required	Mandatory
If-Match	Request		Optional	Discouraged
If-Modified-Since	Request		Optional	Discouraged
If-None-Match	Request		Optional	Discouraged
If-Range	Request		Optional	Discouraged
If-Unmodified-Since	Request		Optional	Discouraged
Last-Modified		Response	Optional	Discouraged
Location		Response	Optional	Mandatory in many situations (e.g., POST responses)
Max-Forwards	Request		Optional	Discouraged
Pragma	Request	Response	Optional	Discouraged
Proxy-Authenticate		Response	Optional	Discouraged
Proxy-Authorization	Request		Optional	Discouraged
Range	Request		Optional	Optional (see 9.8.2.3)
Referer	Request		Optional	Discouraged
Retry-After		Response	Optional	Discouraged
Server		Response	Optional	Discouraged
Set-Cookie		Response	Optional	Discouraged
TE	Request		Discouraged	Discouraged
Trailer		Response	Discouraged	Discouraged
Transfer-Encoding		Response	Optional	Discouraged
Upgrade	Request		Optional	Discouraged
User-Agent	Request		Optional	Discouraged
Vary		Response	Discouraged	Discouraged
Via	Request	Response		Discouraged
Warning	Request	Response	Discouraged	Discouraged
WWW-Authenticate		Response	Optional	Discouraged

5.5 HTTP response codes

5.5.1 Introduction

Response codes are expected to be generalized across RESTful platforms. The specific uses detailed below are likely to be generalized. In the interest of clarity and completeness, they are included here. Please note that these response codes follow general best practices for RESTful interfaces, though they are tuned to address some of the limitations of the embedded space.

This subclause attempts to highlight HTTP response codes that are felt to be more important or that need special attention from developers. However, IEEE 2030.5 clients may encounter various HTTP response codes and, all use of, and response to HTTP response codes SHALL be standard and RFC compliant.

5.5.2 Common responses

5.5.2.1 Introduction

The following HTTP response codes are those considered to be of utmost importance for this standard.

5.5.2.2 1xx (Informational)

These response codes are informational in purpose and are used to indicate that the server is continuing to process in some fashion.

IETF RFC 7231 states, “A client MUST be able to parse one or more 1xx responses received prior to a final response, even if the client does not expect one. A user agent MAY ignore unexpected 1xx responses.”

5.5.2.3 200 ("OK")

This response code is sent to indicate a successful transaction.

This response code is often used in response to a successful GET request, with the entity-body containing a representation of the requested resource. Use of this response code in response to PUT, POST, or DELETE requests is discouraged, to avoid the potentially unnecessary traffic generated by returning the resource representation in the entity-body (see 201 ["Created"] and 204 ["No Content"]).

5.5.2.4 201 ("Created")

This response code is sent to indicate a new resource has been created, at the client's request with a PUT or POST.

The Location header SHALL be used in conjunction with this response code to indicate the URI of the newly created resource. The inclusion of a representation of the newly created resource in the entity-body of the response is discouraged, to conserve bandwidth.

IETF RFC 7231 states, “If the target resource does not have a current representation and the PUT successfully creates one, then the origin server MUST inform the user agent by sending a 201 (Created) response.”

5.5.2.5 204 ("No Content")

This response code is sent to indicate a successful transaction, but one where the response does not include an entity-body.

This response code is often used in response to a successful PUT or POST request, where the resource is modified, not created. This response code is also sent in response to a successful DELETE request. This response code is also sent in response to a successful GET request, where the resource exists but has an empty representation.

Further, when there are URIs that point to a resource that does not yet have content (an “empty representation”), this response code SHOULD be returned. For instance, if a client created a new resource with a POST and that new resource contains URIs pointing to resources that were not yet created and then a client were to request those linked resources, this response code would be the best response. When those resources are created (via a PUT, for instance), this response code (204) SHOULD be returned (in response to the PUT, for instance).

5.5.2.6 206 ("Partial Content")

This response code is sent to indicate the server has fulfilled the partial GET request (as specified by the Range header) for a resource. Note that IETF RFC 7233 requires the Content-Range and Date headers MUST be present in the response.

5.5.2.7 301 ("Moved Permanently")

This response code is sent to indicate that the requested resource has a new URI. The Location header SHOULD be used in conjunction with this response code to indicate the new URI of the requested resource. The entity-body of the response SHOULD be empty. Upon unexpected receipt of this response code, clients SHOULD perform resource discovery to determine which resources have changed location.

5.5.2.8 302 ("Redirect")

The Location header SHALL be used in conjunction with this response code to indicate the new URI of the requested resource. The entity-body of the response SHOULD be empty.

This response code is often used to redirect URIs requested as HTTP to HTTPS.

5.5.2.9 400 ("Bad Request")

This response code is used to indicate a client-side error and is used when no other 4xx response code is appropriate. Often, this response code indicates that the representation sent by a client with a PUT or POST is not appropriate or is malformed.

IETF RFC 7230 states, “A server MUST respond with a 400 (Bad Request) status code to any HTTP/1.1 request message that lacks a Host header field and to any request message that contains more than one Host header field or a Host header field with an invalid field-value.”

5.5.2.10 401 ("Unauthorized")

This response code is used when a client does not have proper authorization to perform the requested action on a resource.

Note, if a server did not wish a client to know of the existence of the resource, it should instead send a 404 ("Not Found") response code.

IETF RFC 7235 states, "The server generating a 401 response MUST send a WWW-Authenticate header field (Section 4.1) containing at least one challenge applicable to the target resource."

5.5.2.11 404 ("Not Found")

This response code is used to indicate that no resource can be found at the specified URI.

This response code MAY also be used in lieu of a 401 response code.

5.5.2.12 405 ("Method Not Allowed")

This response code is used to indicate that the resource does not allow the HTTP method used by the client.

IETF RFC 7231 states, "The origin server MUST generate an Allow header field in a 405 response containing a list of the target resource's currently supported methods."

5.5.2.13 406 ("Not Acceptable")

This response code is used to indicate that a server is unable to generate a response that is acceptable according to the Accept headers sent in the request.

5.5.2.14 413 ("Request Entity Too Large")

This response code is used to indicate that a server is refusing to process a request, as the request is larger than the server is willing or able to process.

5.5.2.15 416 ("Requested Range Not Satisfiable")

This response code is used to indicate that a server has received a Range request that does not overlap any of the resource content.

5.5.2.16 417 ("Expectation Failed")

This response code is used to indicate that an expectation given in an Expect request-header field cannot be met by the server or that the server does not support the given expectation.

5.5.2.17 500 ("Internal Server Error")

This response code is used to indicate that the server has an internal problem and is a generic server error response.

5.5.2.18 501 ("Not Implemented")

This response code is used when a client attempts to use a feature of HTTP (such as a method) that the server does not support.

5.5.2.19 503 ("Service Unavailable")

This response code is used when a server, due to a temporary overload condition, is unable to service a request.

5.5.3 Minimal understanding

Should a client wish to operate with minimal understanding of HTTP response codes, it need only examine the first digit of the response code to understand the general category of the response and “treat an unrecognized status code as being equivalent to the x00 status code of that class, with the exception that a recipient MUST NOT cache a response with an unrecognized status code” (IETF RFC 7231).

5.6 Application payload syntax

5.6.1 Introduction

Application payload message encoding using both XML (W3C XML) and EXI (W3C EXI) SHALL be supported by all servers. Application payload message encoding using either XML (W3C XML) or EXI (W3C EXI) SHALL be supported by all clients. Certain function sets may require additional encodings and, if so, are specified in their respective sections.

5.6.2 XML encoding

The XML declaration is optional as per W3C XML and SHOULD NOT be included in IEEE 2030.5 transactions, to reduce packet sizes. The XML version used SHALL be 1.0. For XML payloads, the encoding SHALL be UTF-8.

5.6.3 EXI encoding

The options for encoding EXI documents SHALL be as follows, and transactions will likely fail if different options are declared in the EXI option header. Options marked as (default) are EXI specification (W3C EXI) defaults and SHOULD NOT be specified explicitly.

- Non-strict schema-informed grammar with the schema (IEEE Std 2030.5 supplemental material)
- Alignment is bit-packed (default)
- Compression is false (default)
- Strict is false (default)
- Fragment is false (default)
- Preserve options are all false (default)
- selfContained is false (default)

- schemaId is “S1” (two bytes: 0x53, 0x31, without quotes). This schemaId corresponds to the normative schema of IEEE Std 2030.5-2018 (supplemental material).
- datatypeRepresentationMap is not used (default)
- valueMaxLength is unbounded (default)
- valuePartitionCapacity is unbounded (default)
- No user defined meta-data

The following XML document describes the EXI option header that SHALL be used to encode the messages:

```
<header xmlns="http://www.w3.org/2009/exi">  
  <common><schemaId>S1</schemaId></common>  
</header>
```

5.7 Content negotiation

5.7.1 General

A client SHALL declare acceptable media types using the HTTP Accept header.

5.7.2 Schema version negotiation

When specifying an “application/sep-exi” media type, an extensibility level (“level”) SHALL be specified in the HTTP Accept header for schema version negotiation. “level” is an Accept-extension defined in IETF RFC 7231. The extensibility level specified using an HTTP Accept header supersedes an extensibility declaration discovered during resource discovery (see 7.4).

The extensibility level defines the base schema and its capability for arbitrary extension. The extensibility level is one of “-S1” or “+S1”. The S1 indicates the base schema version: IEEE Std 2030.5-2018. “-S1” indicates the node does not accept arbitrary tags that are not defined in the base schema, and “+S1” indicates it accepts arbitrary tags. A node with “-S1” will likely fail on an EXI document using arbitrary types, elements, and attributes that are not defined in the schema used for encoding. Devices SHALL NOT send messages to nodes that declare “-S1” using arbitrary types, elements, and attributes.

The grammar used for EXI SHALL be generated as a non-strict grammar only, as having both strict and non-strict grammars would put a large burden on storage requirements for certain devices. The use of a non-strict grammar allows for extensions without schema modification. An invalid (i.e., not defined in the schema) part of an EXI document is allowed in a non-strict grammar and can carry arbitrary tags, attributes, and text encoded using the built-in grammar.

Due to strict memory constraints, some nodes may not be able to parse invalid parts of an EXI document encoded using the built-in grammar. To avoid such errors, a node may declare its inability to receive arbitrary extensions using the “-” (minus) prefix in the extensibility level. Alternatively, nodes that declare the “+” (plus) prefix in the extensibility level will be able to parse extended parts of an EXI document.

Note that the extensibility level does not indicate whether the node can process the data, but only whether it can parse the data.

The format of the extensibility level is “(-/+)Sn” where n is a character to describe base schema version (currently “1”). As extensions of IEEE 2030.5 schemas are intended to be backward compatible, a node

that declares schemaId “S[i]” is intended to be compatible with all versions between “S1” and “S[i]”. However, an “S1” implementation is not backward compatible with an “S0” (IEEE 2030.5-2013) implementation.

For example:

```
Accept: application/sep-exi; level=-S1
```

indicates that the client wishes to receive content encoded using EXI where the base schema of the client is IEEE 2030.5-2018 and that the client does not accept arbitrary tags not defined in the schema.

A client SHOULD use the extensibility level discovered during resource discovery to determine if a server accepts non-strict parts of an EXI document prior to initiating PUT/POST operations where the content contains extended attributes/elements. A server SHOULD use the extensibility level specified in the Accept header to determine if a client accepts non-strict parts of an EXI document prior to responding to GET operations where the content contains extended attributes/elements.

6. Security

6.1 Introduction

Depending on the underlying physical network, messages may be encrypted at lower layers, in addition to the security features provided specifically for the application layer. This clause describes the security features that are provided at the application layer and that are REQUIRED for use over all networks.

Securing transactions between clients and servers is based on using HTTP over TLS (IETF RFC 2818) (also known as HTTPS) using TLS version 1.2 (IETF RFC 5246). The TLS records are then transported using TCP. The TLS handshake mechanism provides mutual authentication based on device certificates or self-signed certificates and TLS records provide encryption and message authentication using the AES-CCM mode of operation. Access control lists allow or deny use of resources based on authentication level and address information. A registration list is used for authorizing clients.

6.2 Security attributes

6.2.1 Introduction

In this subclause, we define some abstract data structures for managing registration (see 6.9) and access control. How this functionality is accomplished is left to the implementer. No access to these data structures is defined in this standard.

6.2.2 Local registration attributes

Local registration attributes represent the data which would be used to hold information passed out-of-band as part of the registration process prior to resources being established.

Table 2—Local registration attributes

Attribute	Identifier	Type	Range	Description	Default
<i>aclLocalRegistration</i>	0x00	List	-	A table of Registration Descriptors each with information for a specific registration	(empty)
<i>aclLocalRegistrationEntries</i>	0x01	Integer	Implementation specific	Number of entries in <i>aclRegistration</i>	0

Table 3—Local registration descriptor entry

Name	Type	Range	Description	Default
PIN	Integer	0 to 999999	6-digit personal identification number (PIN) (5 plus check digit) for basic server validation. The PIN is also reflected in the Registration resource linked to the EndDevice resource.	-
SFDI	SFDI	-	SFDI of registering device. The SFDI is also reflected in the EndDevice resource.	-
DeviceType	Integer	0 to 3	Minimum required device type	0
HardwareModuleName	String	-	Optional additional hardware module information	-

6.2.3 Access control list (ACL) attributes

6.2.3.1 Introduction

Access control list (ACL) attributes represent the data that would be used to hold information to determine whether access to a particular resource by a particular client is allowed or denied. An ACL can enforce more granular access control based on various criteria (e.g., client identity). Conceptually, an ACL exists for every single accessible resource, however in practice it is likely only certain resources with more complex access policies would require a representation of all the data specified in this subclause.

Table 4—ACL attributes

Attribute	Identifier	Type	Range	Description	Default
<i>aclDefaultAccess</i>	0x02	Access descriptor		Default access to resource	-
<i>aclSpecificID</i>	0x03	List		A list of specificIDDescriptors for each specific client access to resource	-
<i>aclSpecificIDEntries</i>	0x04	Integer	Implementation specific	Number of entries in <i>aclSpecificID</i>	0

Table 5—AccessDescriptorEntry

Name	Type	Range	Description	Default
Method	Bitmap	0x0 to 0xf	Bitmap of which methods are supported: 0x1: GET 0x2: PUT 0x4: POST 0x8: DELETE 0x10: HEAD	0x0
AuthType	Integer	0x0 to 0x0f	Bitmap of which authentication types are allowed: 0x1: No authentication 0x2: User authentication 0x4: Self-signed certificate 0x8: Device certificate Remaining bits are reserved for authentication types not defined by this standard but by an additional security policy.	0x0
DeviceType	Integer	0 to 3	DeviceType: 0: Any device type	0

Table 6—SpecificIDDescriptor entry

Name	Type	Range	Description	Default
Access	Access descriptor	-	Access levels required for client	-
IPAddr	IPAddr	-	IP address of client	-
Port	Integer	0x0000 to 0xffff	Port of client 0: Wildcard (any port)	0

Access control list (ACL) attributes provide a mechanism for granting and revoking privileges to use specified methods with a particular resource, applicable to all resources described in this standard. Access control that is more granular than a resource is out of scope for this standard. The mechanisms for authentication and the binding of that authentication to a specified identity (such as an IP address) are specified in 6.5 and 6.6.

ACLs are used to set default privileges and grant additional privileges, not to deny privileges. Thus, if a given client's request does not explicitly meet the required privilege in the associated ACL, then that client does not have that privilege. The default configuration of an ACL for a given resource means that resource

is not accessible to clients. In practice, this state only exists ephemeral and all ACLs will be initialized appropriately at startup according to the security policy and will be subsequently modified according to registration and authentication.

Subordinate resources created dynamically will initially inherit their ACL from their parent resource.

ACLs may be statically fixed with a default operation for some resources and may be dynamic and extensible for others. If a resource does not have an ACL, access is granted to the resource unconditionally.

Initialization of ACLs, beyond minimal requirements, is out of scope for this standard and is governed by overall security policy. Default settings for ACLs for particular function sets are described in 6.8.

6.2.3.2 *aclDefaultAccess*

aclDefaultAccess in the ACL is used for settings irrespective of the client identity of an incoming request and is used initially to authorize an incoming HTTP request.

6.2.3.3 *aclSpecificIDList*

A *SpecificIDDescriptor* entry in the *aclSpecificIDList* part of the ACL is an additional entry used to allow specific additional checks to be done to authorize based on client identity (IP address and port).

6.2.3.4 Access authorization

6.2.3.4.1 Introduction

These are controls that are independent of the source identity (IP address and port) and thus can be configured in both *aclDefaultAccess* and a *SpecificIDDescriptor* entry in *aclSpecificIDList* as shown in Table 5 and Table 6.

In the following, *corresponding ACL entry* means:

- The first *SpecificIDDescriptor* entry in *aclSpecificIDList* which matches the incoming client HTTP request's source IP address and port
- *aclDefaultAccess* otherwise

6.2.3.4.2 Method attribute

The Method attribute is used to control which HTTP request method is allowed for client access.

- GET
- PUT
- POST
- DELETE
- HEAD

The HTTP method of an incoming HTTP request is checked and Method authorization will be TRUE if the method is allowed in the corresponding ACL entry.

6.2.3.4.3 AuthType attribute

The AuthType attribute is used to control the required authentication types the client can use in its incoming HTTP/HTTPS request.

- 0x1: No authentication
- 0x2: User authentication
- 0x4: Self-signed certificate
- 0x8: Device certificate

If an incoming HTTP request is destined for the port associated with HTTPS (typically port 443), the incoming authentication type will be set to the corresponding TLS session authentication type. If an incoming HTTP request is destined for the port associated with HTTP (typically port 80), the authentication type will be set to 0x1 (no authentication). The user authentication AuthType (0x02) can be used if additional user authentication outside of the scope of this standard takes place.

The incoming authentication type is compared with the bitmap in AuthType in the corresponding ACL entry and AuthType authorization will be TRUE if the corresponding authentication type is set in AuthType in the corresponding ACL entry. Representing this logically:

```
if (authentication type & AuthType) != 0:  
    AuthType authorization = TRUE  
else  
    AuthType authorization = FALSE
```

6.2.3.4.4 DeviceType attribute

The DeviceType attribute is used to control the device type required of the client's incoming HTTP request. It is based on the deviceType OID in the certificate (see 6.11.7.2).

If an incoming HTTP request is destined for the port associated with HTTPS, the incoming device type will be set to the corresponding TLS session device type based on the deviceType OID in the certificate. If an incoming HTTP request is destined for the port associated with HTTP, the device type will be set to 0 (any device type).

If the DeviceType in the ACL is set to 0, device type authorization will be TRUE unconditionally. Otherwise the incoming device type is compared with the DeviceType in the corresponding ACL entry and DeviceType authorization will be TRUE if the device type is equal to the DeviceType in the corresponding ACL entry.

6.2.3.5 Authorization logic

Authorization is granted if Method authorization, AuthType authorization, and DeviceType authorization are all TRUE. Access to the resource can then take place.

If Method authorization is not granted, the server MAY respond with either:

```
HTTP/1.1 400 Bad Request
```

or

HTTP/1.1 405 Method Not Allowed

If AuthType or DeviceType authorization is not granted, the server SHOULD immediately respond:

HTTP/1.1 404 Not Found

The server MAY respond:

HTTP/1.1 401 Unauthorized

Otherwise, processing will continue.

6.2.3.6 ACL examples

6.2.3.6.1 Introduction

This subclause contains two informative examples to illustrate the use of ACLs, using EndDeviceList and EndDevice.

6.2.3.6.2 EndDeviceList

The EndDeviceList resource is usually accessible by any client to find its own entry in the list. Therefore, the ACL will be as follows:

Table 7—Example ACL for EndDeviceList

Attribute	Comment
<i>aclDefaultAccess</i>	Method: 0x01 (GET only) AuthType: 0x1 (no authentication) DeviceType: 0 (any device type)
<i>aclSpecificID</i>	Empty
<i>aclSpecificIDEntries</i>	0

In this example, there will never be any SpecificIDDescriptors required, as this resource needs to be accessible to any device.

6.2.3.6.3 EndDevice

In this example, an EndDevice resource for a client does not exist prior to registration.

The earliest point it can be created is at the point of registration and the ACL would be set as shown in Table 8.

Table 8—Example ACL entry for EndDevice at point of registration

Attribute	Comment
<i>aclDefaultAccess</i>	Method: 0x00 (no default access) AuthType: 0x1 (no authentication) DeviceType: 0 (any device type)
<i>aclSpecificID</i>	Empty
<i>aclSpecificIDEntries</i>	0

In this example, there is no default access as only the client device associated with the EndDevice is able to access the EndDevice. Also, at this point, there has been no client communication, therefore there would be no SpecificIDDescriptor entry in *aclSpecificIDList*.

There would also be an entry placed in the local registration list corresponding to the client, as shown in Table 9.

Table 9—Example local registration entry for registering device

Name	Type	Range	Description	Default
PIN	Integer	0 to 999 999	6-digit PIN (5 plus check digit) of registering device	-
SFDI	SFDI	0 to 68 719 476 735, $2^{36} - 1$	SFDI of registering device	-
DeviceType	Integer	0 to 3	Registering device type	0
HardwareModuleName	String	-	Optional additional hardware module information	-

In this state, it is primed to be populated with a SpecificIDDescriptor when the device actually accesses the EndDeviceList resource.

When the client attempts access to any resource on a function set server using TLS, it will start performing the TLS handshake, which involves transferring the client's certificate. At the point of access, if there is a pending registration, it will be checked against the client's device. If there is a match, and the validated certificate's SFDI matches, the ACL for the EndDevice will be populated with an additional SpecificIDDescriptor as shown in Table 10 and Table 11.

Table 10—Example ACL entry for EndDevice at point of client access

Attribute	Comment
<i>aclDefaultAccess</i>	Method: 0x00 (no default access) AuthType: 0x1 (no authentication) DeviceType: 0 (Any device type)
<i>aclSpecificID</i>	One SpecificIDDescriptor for client
<i>aclSpecificIDEntries</i>	1

Table 11 —Example SpecificIDDescriptor for EndDevice at point of client access

Name	Type	Range	Description	Default
Access	AccessDescriptor	-	Access levels required for client: Method: 0xF (GET, PUT, POST, DELETE) AuthType: 0x8 (device certificate) DeviceType: (as appropriate for device)	-
IPAddr	IPAddr	-	IP address of client	-
Port	Integer	0x0000 to 0xffff	Port of client	0

6.3 Device credentials

6.3.1 Introduction

There are three credentials per device:

- Short form device identifier (SFDI)
- Long form device identifier (LFDI)
- PIN

6.3.2 Certificate fingerprint

The certificate fingerprint is the result of performing a SHA256 operation over the whole DER-encoded certificate and is used to derive the SFDI and LFDI. A certificate fingerprint is not confidential and is never used to derive subsequent keying material.

An example certificate fingerprint used for illustration in the following examples is:

3E4F-45AB-31ED-FE5B-67E3-43E5-E456-2E31-984E-23E5-349E-2AD7-4567-2ED1-45EE-213A

6.3.3 Short-form device identifier (SFDI)

The SFDI SHALL be the certificate fingerprint left-truncated to 36 bits. For display purposes, this SHALL be expressed as 11 decimal (base 10) digits, with an additional sum-of-digits checksum digit right-concatenated. Based on the example in 6.3.2, this would be 167-261-211-391.

Left truncation to 36 bits: 0x3E4F45AB3

Expressed as a decimal: 16726121139

Right-concatenation of check digit and hyphenation: 167-261-211-391

For input validation purposes, the sum of the digits of the fingerprint including the checksum digit, modulo 10, SHALL be zero. The SFDI has sufficient entropy (236) to uniquely identify the device in the context of its usage and is used to identify a device within a home area network (HAN) or site domain. It should not

be used in a truly global context (i.e., where the identity of the device cannot be qualified with the domain it is in).

For a device with a device certificate, the SFDI can be printed on the device packaging.

6.3.4 Long-form device identifier (LFDI)

The LFDI SHALL be the certificate fingerprint left-truncated to 160 bits (20 octets). For display purposes, this SHALL be expressed as 40 hexadecimal (base 16) digits in groups of four. Based on the example in 6.3.2, this would be “3E4F-45AB-31ED-FE5B-67E3-43E5-E456-2E31-984E-23E5.” The LFDI is used when a globally unique identity is required, for example in sending an event back to a service provider that is associated with a particular device.

6.3.5 6-digit PIN code

The SFDI and LFDI are derived from public information (i.e., a Certificate), therefore can potentially be recreated by an eavesdropper. Therefore, a device MAY also have an additional 6-digit PIN code, which can be shared out-of-band with a service provider in conjunction with the SFDI or LFDI. For display purposes, this SHALL be expressed as 5 decimal (base 10) digits, with an additional sum-of-digits checksum digit right-concatenated:

Original PIN: 12345

Right-concatenation of check digit and hyphenation: 123-455

For input validation purposes, the sum of the digits of the PIN including the checksum digit, modulo 10, SHALL be zero. The PIN MAY be obtainable from the EndDevice server through the Registration resource to validate that the client is in communication with the correct server. The PIN SHOULD be configurable on a device where possible for registration purposes, otherwise SHOULD be a random 5-digit value plus check digit pre-programmed into the device and printed on the device. The PIN is not overly secure and therefore SHALL NOT be used in any way to derive keys for actual data encryption.

6.3.6 Registration code

The SFDI and PIN are usually presented separately. However, in certain cases it may be convenient to provide a single registration code, which is simply the concatenation of the SFDI and the PIN expressed as a decimal (base 10) number:

SFDI || PIN

From the examples above, this would be 167-261-211-391-123-455.

6.4 Resource access authentication and authorization context

A node is able to perform network layer communication once it has been authenticated and authorized to join the network. However, application layer authentication and authorization MAY be required before clients and servers can exchange application layer messaging. Registration (see 6.9) with a utility or third party service provider MAY also be needed to provide explicit device and user authorization at the application layer.

Resource access requiring application layer authentication, data confidentiality, and integrity checking SHALL occur through requests from a client to the server using HTTP over TLS (IETF RFC 2818) (also known as HTTPS) using TLS version 1.2 (IETF RFC 5246). Resource access not requiring application layer authentication, data confidentiality, or integrity checking SHALL occur through requests from a client to the host server using HTTP (IETF RFC 7230).

If a request is made to the port associated with HTTPS, it is considered an HTTPS request and authentication SHALL have taken place. If authentication has not taken place, authentication SHALL be initiated as described in 6.5. When authenticated, the request is then passed to the ACL associated with the resource. Ancillary information about the request obtained from the secure session, notably the level of client authentication, will also be compared with the ACL.

If the request is made to the port associated with HTTP, it is considered an HTTP request and authentication SHALL NOT be REQUIRED and the request is then passed to the ACL associated with the resource. Ancillary information about the request stating that the client is unauthenticated will also be compared with the ACL (see 6.2.3.4.3).

Authorization on a request-by-request basis is determined by the ACL settings for the resource, which may be set up at the end of the authentication based on the level of client authentication. The Local Registration List (*aclLocalRegistrationList*) may be additionally used to authorize on a device-by-device basis.




Figure 1—Example device authentication procedure using HTTPS port 443

6.5 Resource access authentication

Resource access authentication only applies using HTTPS. It may be possible to authenticate at a higher level using authentication based on HTTP-only transactions, but this is out of scope for this standard.

The use of TLS (IETF RFC 5246) requires that all hosts implementing server functionality SHALL use a device certificate whereby the server presents its device certificate as part of the TLS handshake.

The application authentication process is as follows:

- 1) The resource's server listens on the TCP port associated with HTTPS.
- 2) The client initiates an HTTP request using a random unused source TCP port to the resource's server using the TCP port associated with HTTPS.
- 3) If no TLS session is in place, a TLS handshake SHALL occur between the client and server:
 - a) Authentication of the server SHALL be done as part of the TLS handshake by validating its device certificate as described in (IETF RFC 5246), Section 7 using the inherent PKI. If security policy dictates, additional certificate validation MAY be required.
 - b) If the client has a device certificate, authentication of the client SHALL be done as part of the TLS handshake by validating the client's device certificate as described in (IETF RFC 5246), Section 7 using the inherent PKI. If security policy dictates, additional certificate validation MAY be required. The authentication level to be compared with a resource's corresponding AuthType attribute will be 0x8 (device certificate).
 - c) If the client has a self-signed certificate, the self-signed certificate SHALL be validated for correctness. The authentication level to be compared with a resource's corresponding AuthType attribute will be 0x4 (self-signed certificate).

If the client does not have a certificate and the security policy allows, client authentication MAY NOT need to take place, or secondary client authentication MAY take place after the TLS handshake. If secondary client authentication has taken place, the authentication level to be compared with a resource's corresponding AuthType attribute will be 0x2 (user authentication). If no client authentication has taken place, the authentication level to be compared with a resource's corresponding AuthType attribute will be 0x1 (no authentication).

6.6 Resource access authorization

Pre-authorization for resources is normally set when the client registers with the host as described in 6.9. If the security policy allows, authorization MAY occur immediately after authentication based on implicit rules to allow a request to complete. This is to allow unregistered access to resources based on security policy. If the client uses a self-signed certificate, pre-authorization using the SFDI of the self-signed certificate MUST have taken place and authorization SHALL be granted if the SFDI of the presented self-signed certificate matches the SFDI presented as part of registration.

6.7 Cipher suites

All devices SHALL support the TLS_ECDHE_ECDSA_WITH_AES_128_CCM_8 cipher suite (IETF RFC 7251). The ECC cipher suite SHALL use elliptic curve secp256r1.

- All devices acting as a server SHALL support the ECC cipher suite. In particular, all devices acting as a server SHALL have an ECC certificate.
- All devices acting as a client SHALL support the ECC cipher suite for the purposes of validating an ECC certificate.
- All devices acting as a client SHOULD support the request for an ECC certificate.

Devices MAY support additional cipher suites; however, any additional cipher suites SHOULD provide a cryptographic strength at least equivalent to the mandatory cipher suite.

6.8 Default security policy

Service providers create security policies by balancing the requirements of their regulatory environment and the results of their risk assessments. Different regulatory environments may mandate requirements that trade ease of data access with information assurance. The use of TLS, ACLs, and other security controls give the service provider the flexibility to meet these needs. Security policies are a combination of ACL attribute values and additional security controls dictated by the service provider. Implementation of security policies is out of scope of this standard. For the purpose of certification testing, the following table represents the default security policy for each function set. Servers SHALL be configurable to support each default policy for all implemented function sets during certification testing.

The function set column in Table 12 reflects the functionsImplemented attribute in DeviceInformation.

Table 12—Attribute values for default security policy

Function set	<i>aclDefaultAccess</i> AuthType	Device certificate needed	Registered device
Device capability	0xf	No	No
Self device resource	0xc	No	Yes
End device resource	0xc	No	Yes
Function set assignments	0x8	Yes	Yes
Subscription/Notification mechanism	0x8	Yes	Yes
Response	0x8	Yes	Yes
Time	0x8	Yes	Yes
Device information	0x8	Yes	Yes
Power status	0x8	Yes	Yes
Network status	0x8	Yes	Yes
Log event	0x8	Yes	Yes
Configuration resource	0x8	Yes	Yes
Software download	0x8	Yes	Yes
DRLC	0x8	Yes	Yes
Metering	0x8	Yes	Yes
Pricing	0xc	No	Yes
Messaging	0xc	No	Yes
Billing	0x8	Yes	Yes
Prepayment	0x8	Yes	Yes
Flow reservation	0x8	Yes	Yes
DER control	0x8	Yes	Yes

The *aclDefaultAccess* attribute Method value SHOULD match the Allowed Methods for each resource enumerated in the IEEE 2030.5 WADL (IEEE Std 2030.5 supplemental material). The Method value MUST contain GET (0x01). The *aclDefaultAccess* attribute DeviceType value should be “any device type” (0). Servers SHALL support the default policies for certification testing. Servers MAY additionally support alternative policies. For example, to meet regulatory requirements, a utility may mandate a policy that provides unauthenticated pricing information from a pricing server over the port associated with HTTP to

any IEEE 2030.5 device. Based on risk assessments, service providers may have differing policies for devices enrolled in high-incentive Demand Response/Load Control programs than those enrolled in low-incentive programs, to include additional requirements such as DeviceType authorization. Servers SHOULD provide the functionality to support multiple security policies to meet the requirements of different service providers.

6.9 Registration

6.9.1 Introduction

Registration describes the procedure whereby an out-of-band procedure is used to convey client registration information a priori to the server that houses a resource that will subsequently be accessed by the client. The registration information is the client's SFDI and optionally, PIN, which uniquely identifies the client in the given context.

Registration may occur some time before the client attempts to access a resource, for example, using a web site or telephone to register the information with a service provider. The service provider will then provide the information to the EndDevice server using some out-of-band mechanism, (e.g., the AMI network) and the server will program its registration list accordingly.

Alternatively, there may be no actual registration before a client attempts to access a resource and, for example, the server may present the premises owner with the SFDI of the client attempting access via a user interface. The premises owner may then continue to authorize the client access, or deny access based on the information presented from the client's certificate.

This subclause describes a typical registration procedure for a client using a device certificate with an EndDevice server.

Registration for clients SHALL occur via an EndDevice resource corresponding to the client, which typically resides on an energy services interface (ESI) associated with the utility, premises owner, or third party service provider that is trusted to perform registration.




Figure 2—Device authentication with registration procedure examples using HTTPS

6.9.2 EndDeviceList

Clients SHALL locate HAN services by performing DNS service discovery (DNS-SD) queries to the HAN; see Clause 7 for details. The client can then resolve the URI of the EndDeviceList (given as /edev for illustration purposes) for registration and authentication purposes and know which port(s) the server for the EndDeviceList is listening on.

The EndDeviceList is the resource used by a client to complete the process initiated by registration of the client when the device owner wishes to register the device in a utility, premises owner, or service provider program. In some cases, registration MAY be required for access.

Upon registering a client, the EndDevice resource's server *aclLocalRegistrationList* will be configured with:

- The client SFDI, to be registered in the utility, premises owner, or third party service provider program.
- Optionally, the client PIN
- The required device types of the associated client.

Thus, at the point of registration, the EndDevice resource's server is able to perform authentication based on the device certificate and additional user authentication based on the client SFDI.

The EndDevice resource's server MAY allow access from clients that have not been pre-configured if the security policy allows.

The registration procedure is as follows:

- a) The EndDevice resource's server SHALL listen on the TCP port associated with HTTPS and follow the procedure described in (IETF RFC 5246) when a client attempts to access the EndDevice resource.
- b) Authorization SHALL then occur whereby the ACLs of the server resources corresponding to the registering client are set according to the security policy and the presence in *aclLocalRegistrationList*. This ensures that following registration, a client can typically proceed to access all the resources it is authorized to without having to perform any further procedures.
- c) If present, the client MUST verify that the EndDevice's associated Registration resource contains the correct PIN for the client. If the PIN does not match, the client SHOULD NOT attempt any further access to that server.
- d) The client SHALL subsequently re-use its device certificate to authenticate with any other server host. It does not need to re-authenticate with the EndDevice resource's server.

6.10 Security LogEvents

There are specific LogEvents attributable to security. These will use a function set enumeration of "security" as defined in the model. These are as shown in Table 13.

Table 13—Security LogEvents

LogEvent name	LogEvent code	LogEvent description
SEC_TLS_ALERT	0x00	SHOULD be issued when a TLS Alert is generated. The logEventID SHALL be set to the TLS Alert value (IETF RFC 5246).
SEC_REGISTRATION_MISS	0x01	SHOULD be issued when a received certificate does not have a corresponding SFDI entry in the registration list.
SEC_ACL_ACCESS_FAILED	0x02	SHOULD be issued when access to a resource fails due to failing access control criteria described in 6.2.3.

6.11 Certificate management

6.11.1 Introduction

This standard defines a public key infrastructure (PKI) in the IEEE 2030.5 certificate management system: the Manufacturing PKI. The Manufacturing PKI issues certificates to devices at the time of application installation (e.g., at manufacture). These certificates are intended for use during deployment (or redeployment) and on-going operation to authenticate the device to other IEEE 2030.5 devices implementing IEEE 2030.5 applications over TLS. It is expected that the market will implement one or more Manufacturing PKIs in accordance with the requirements set forth in this standard.

In addition, this document describes a few other types of certificates that IEEE 2030.5 applications may encounter depending on which services they support—see 6.11.8.4. While native IEEE 2030.5 devices are required to understand, support, and process Manufacturing PKI certificates, support of the additional

certificates is optional and generally targeted at specific classes of devices (e.g., energy services interface [ESI], web portal).

There are six classes of certificates that may be active in an IEEE 2030.5 system, depending on configuration and use, as follows:

- Device certificates: Issued under the Manufacturing PKI during manufacturing to purpose-built (aka “native”) IEEE 2030.5 certified devices for operational purposes.
- Device test certificates: Issued under the Manufacturing PKI during manufacturing to purpose-built (aka “native”) IEEE 2030.5 certified devices for test purposes.
- Additional certificates for IEEE 2030.5 devices: One or more OPTIONAL TLS server certificates issued by non-IEEE 2030.5 CAs to IEEE 2030.5 devices such as ESIs for use in addition to the device certificate.
- Generic client certificate for non-native entities: A TLS client certificate issued by a generic (non-IEEE 2030.5) certificate authority to a non-native entity.
- Generic server certificate for non-native entities: A TLS server certificate issued by a generic (non-IEEE 2030.5) certificate authority to a non-native entity.
- Self-signed client certificate for non-native entities: A TLS client certificate self-generated and self-signed by a customer or software.

6.11.2 Certificate usage—authentication versus authorization

Certificates provide a mechanism to authenticate an identity. Once authenticated (by proving possession of the associated private key, and by having the certificate chain to a known root of trust), that identity (or the service, person, or application associated with that identity) can be authorized access to resources, the ability to assume a role (e.g., system operator, general user), or perform various functions.

An authenticated certificate by itself does not generally grant authorization. Specific applications that accept the certificate MAY grant implicit authorization to access any resource under the purview of the application, but usually authorize access based on the identity represented by the certificate (e.g., an access control list entry). The following tables describe both the authentication and authorization uses of certificates described by this profile.

Table 14 describes the mechanisms for certificate validation—the authentication of the identity claimed in the certificate. Note that for a self-signed certificate, the authentication is limited only to validating the self-signature over the certificate and that provides only an integrity check.

The phrase “IEEE 2030.5 Cert - Indef” is meant to convey that Manufacturing PKI certificates are indefinitely valid and the check is limited solely to a check of the signatures on the certificate chain. The phrase “Optional OCSP” means that the server device (and optionally, the client device) may utilize Online Certificate Status Protocol (OCSP) (IETF RFC 6960) as an additional mechanism to determine if a certificate has been revoked. OCSP may only be used to verify non-IEEE 2030.5 certificates.

Table 15 describes the mechanisms for determining if the holder of a specific certificate (and its private key) may access specific resources. Generically, each resource has an ACL tagged to it. If the identity of the certificate holder has the appropriate rights to a specific resource, then the representation of the resource is returned via query. Note that there are certain resources that are accessible to all authenticated clients.

Table 14—TLS authentication matrix

Authentication		Server		
		Native IEEE 2030.5 application	Generic server	Self-signed
Client	Native IEEE 2030.5 application	IEEE 2030.5 Cert - Indef	Optional OCSP	Not allowed
	Generic client	Optional OCSP	Not specified here	Not specified here
	Self-signed	Signature validation	Not specified here	Not specified here

Table 14 should be read as describing the authentication on the offered client certificate by the server. For example, Generic client/Native IEEE 2030.5 application has an “Optional OCSP” entry meaning the IEEE 2030.5 Application can OPTIONALLY use OCSP to check the validity of the Generic Client certificate in addition to its normal certificate validation process.

“IEEE 2030.5 Cert - Indef” means that if the IEEE 2030.5 Manufacturing PKI certificate validly chains to the IEEE 2030.5 root it is considered valid—neither OCSP nor CRLs are used or issued by CAs within the Manufacturing PKI hierarchy.

Table 15—IEEE 2030.5 authorization matrix

	Native IEEE 2030.5 application	Generic server	Self-signed
Native IEEE 2030.5 application	ACL	ACL or public resources (server specific)	Not allowed
Generic client	ACL or public resources	Not specified here	Not specified here
Self-signed	ACL or public resources	Not specified here	Not specified here

The above table should be read as describing the authorization mechanism used by the server with respect to the offered client credential. For example, the Generic Client/Native IEEE 2030.5 application entry of “ACL” means that either there is a specific ACL for the generic client credential that permits access to specific data, or the absence of such ACL allows that generic client access only to public resources.

6.11.3 Manufacturing PKI

6.11.3.1 Introduction

This subclause covers only those certificates issued under the auspices of the Manufacturing PKI. It specifically excludes the certificates described in 6.11.8.4 as “Other Certificates.”

Any IEEE 2030.5 Manufacturing PKI SHOULD be a hierarchy with a depth of 2, 3, or 4 levels. At the top level, Manufacturing PKI hierarchy SHOULD have one SERCA. A commercial CA MAY operate one or more SERCAs as business needs dictate, and this is one possible model for the initial operation of the SERCAs. Private key material for a SERCA SHOULD be held in a form to allow secure transfer of the material to a new commercial CA if necessary. The SERCA issues MCA certificates and MICA certificates to authorized vendors.

A SERCA MAY issue device certificates on behalf of one or more manufacturers.

The Manufacturing PKI hierarchy MAY include one IEEE 2030.5 MCA. One or more MCAs SHOULD be out-sourced to an IEEE 2030.5 vendor, specifically for the issuance of vendor-specific MICAs. An IEEE

2030.5 vendor MAY contract with a commercial CA to host the IEEE 2030.5 vendor's MCA credentials (certificate and private key), but retains ownership of such credentials. MCAs may only issue MICA certificates.

The Manufacturing PKI hierarchy MAY include one IEEE 2030.5 MICA. One or more MICAs SHOULD be out-sourced to an IEEE 2030.5 vendor, specifically for the production of IEEE 2030.5 certified devices by that vendor.

All devices SHALL store exactly one SERCA in their certificate path. All devices SHALL include at least the public keys of all existing SERCAs and MAY include their certificates. In the course of any particular authentication, any device can therefore verify the chain of signatures leading up to any one of the roots.

The following certificate paths are the only valid instantiations of the Manufacturing PKI:

- SERCA > device certificate
- SERCA > MICA > device certificate
- SERCA > MCA > MICA > device certificate

6.11.3.2 Manufacturing certificate lifecycle

Certificates within the Manufacturing PKI have an indefinite lifetime—this includes, specifically, a SERCA. Nevertheless, CA certificates may be retired and subsequently replaced when circumstances dictate. In such cases, the retired certificate and its associated private key SHALL no longer be used for issuing certificates. However, parties MAY continue to rely on those certificates for validating subordinate certificates. If, however, the signature algorithm or parameter set used in a manufacturing certificate comes to be viewed as insufficiently secure for the purpose, parties MAY retire those certificates and associated public keys. A retired certificate and key may no longer be used for issuing new certificates, but is not considered revoked for the purpose of validation.

An MCA or MICA certificate SHALL NOT be re-issued (e.g., re-signed with the same SubjectName and public key, but with a new validity period). Instead, if it is desired to retire an existing key/certificate, a new key pair SHALL be generated and bound into a new MCA or MICA certificate generally with a new serialNumber component of the SubjectName. Operationally, the responsible CA SHOULD verifiably destroy⁹ the private key of the retired certificate. Retired certificates MUST remain available to verify subsidiary certificates. A replacement of a MCA or MICA SHOULD use a new name formed by incrementing or otherwise adjusting the serialNumber component of the subject name. See 6.11.8.3.1 and 6.11.8.3.2 for the format of the name.

A side effect of the indefinite lifetime requirement coupled with the permanent embedding of the device certificate and its certificate chain within a device is that the device, MCA, and MICA certificates MUST NOT and cannot be revoked once issued. Specifically, no MCA, MICA, or SERCA shall issue or be required to issue CRLs, and no MCA or MICA shall operate or have operated on their behalf any OCSP server for the purpose of providing validity information for any certificate under the Manufacturing PKI hierarchy. This does not preclude the use of a certificate and its corresponding identity to be used in a blacklist or whitelist for authorization purposes to allow or deny access to resources on a server.

⁹ The process for destroying private keys is a business issue that should be covered by any contract for SERCA services.