

IEEE Standard for eXtensible Event Stream (XES) for Achieving Interoperability in Event Logs and Event Streams

IEEE Computational Intelligence Society

Developed by the
Standards Committee

IEEE Std 1849™-2023
(Revision of IEEE Std 1849-2016)

STANDARDS

IEEE Standard for eXtensible Event Stream (XES) for Achieving Interoperability in Event Logs and Event Streams

Developed by the

Standards Committee
of the
IEEE Computational Intelligence Society

Approved 5 June 2023

IEEE SA Standards Board

Abstract: A grammar for a tag-based language whose aim is to provide designers of information systems with a unified and extensible methodology for capturing systems' behaviors by means of event logs and event streams is defined in the XES standard. An "XML Schema" describing the structure of an XES event log/stream and an "XML Schema" describing the structure of an extension of such a log/stream are included in this standard. Moreover, a basic collection of so-called "XES extension" prototypes that provide semantics to certain attributes as recorded in the event log/stream is included in this standard.

Keywords: event log, event stream, extensions, IEEE 1849™, system behavior, XML

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

Copyright © 2023 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 8 September 2023. 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.

PDF: ISBN 978-1-5044-9805-0 STD26231
Print: ISBN 978-1-5044-9806-7 STDPD26231

IEEE prohibits discrimination, harassment, and bullying.

For more information, visit <https://www.ieee.org/about/corporate/governance/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 Standards documents are made available for use subject to important notices and legal disclaimers. These notices and disclaimers, or a reference to this page (<https://standards.ieee.org/ipr/disclaimers.html>), appear in all standards and may be found under the heading “Important Notices and Disclaimers Concerning IEEE Standards Documents.”

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

IEEE Standards documents are developed within IEEE Societies and subcommittees of IEEE Standards Association (IEEE SA) Board of Governors. IEEE develops its standards through an accredited consensus development process, which brings together volunteers representing varied viewpoints and interests to achieve the final product. IEEE Standards are documents developed by volunteers with scientific, academic, and industry-based expertise in technical working groups. Volunteers are not necessarily members of IEEE or IEEE SA 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 makes no warranties or representations concerning its standards, and expressly disclaims all warranties, express or implied, concerning this standard, including but not limited to the warranties of merchantability, fitness for a particular purpose and non-infringement. In addition, IEEE does not warrant or represent that the use of the material contained in its standards is free from patent infringement. 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: THE NEED TO PROCURE 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 is 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, nor be relied upon as, a formal position of IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that the presenter's views should be considered the personal views of that individual rather than the formal position of IEEE, IEEE SA, the Standards Committee, or the Working Group. Statements made by volunteers may not represent the formal position of their employer(s) or affiliation(s).

Comments on standards

Comments for revision of IEEE Standards documents are welcome from any interested party, regardless of membership affiliation with IEEE or IEEE SA. However, **IEEE does not provide interpretations, 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 subcommittees of the IEEE SA Board of Governors 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 evaluating comments or in revisions to an IEEE standard is welcome to join the relevant IEEE working group. You can indicate interest in a working group using the Interests tab in the Manage Profile and Interests area of the [IEEE SA myProject system](#).¹ An IEEE Account is needed to access the application.

Comments on standards should be submitted using the [Contact Us](#) form.²

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 constitute 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.

Data privacy

Users of IEEE Standards documents should evaluate the standards for considerations of data privacy and data ownership in the context of assessing and using the standards in compliance with applicable laws and regulations.

Copyrights

IEEE draft and approved standards are copyrighted by IEEE under US 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

¹Available at: <https://development.standards.ieee.org/myproject-web/public/view.html#landing>.

²Available at: <https://standards.ieee.org/content/ieee-standards/en/about/contact/index.html>.

by public authorities and private users, neither IEEE nor its licensors waive any rights in copyright to the documents.

Photocopies

Subject to payment of the appropriate licensing fees, 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; <https://www.copyright.com/>. 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 10 years. When a document is more than 10 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 [IEEE Xplore](#) or [contact IEEE](#).³ For more information about the IEEE SA or IEEE's standards development process, visit the IEEE SA Website.

Errata

Errata, if any, for all IEEE standards can be accessed on the [IEEE SA Website](#).⁴ Search for standard number and year of approval to access the web page of the published standard. Errata links are located under the Additional Resources Details section. Errata are also available in [IEEE Xplore](#). Users are encouraged to periodically check for errata.

Patents

IEEE Standards are developed in compliance with the [IEEE SA Patent Policy](#).⁵

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 <https://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.

³Available at: <https://ieeexplore.ieee.org/browse/standards/collection/ieee>.

⁴Available at: <https://standards.ieee.org/standard/index.html>.

⁵Available at: <https://standards.ieee.org/about/sasb/patcom/materials.html>.

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.

IMPORTANT NOTICE

IEEE Standards do not guarantee or ensure safety, security, health, or environmental protection, or ensure against interference with or from other devices or networks. IEEE Standards development activities consider research and information presented to the standards development group in developing any safety recommendations. Other information about safety practices, changes in technology or technology implementation, or impact by peripheral systems also may be pertinent to safety considerations during implementation of the standard. Implementers and users of IEEE Standards documents are responsible for determining and complying with all appropriate safety, security, environmental, health, and interference protection practices and all applicable laws and regulations.

Participants

At the time this draft standard was completed, the XES Working Group had the following membership:

Moe Wynn, *Chair*
Wil van der Aalst, *Vice Chair*
Eric Verbeek, *Secretary*

Estefania Serral Asensio
Peter Blank
Marcus Brenscheidt
Sergio Correia

Claudio Di Ciccio
Philipp Herrmann
Nancy Landreville
Julian Lebherz

Rajesh Murthy
Arif Selcuk Ogreni
Andreea Ungureanu

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

Estefania Serral Asensio
Bruno DeStefano
Werner Hoelzl
Ayodele Ishola-Salawu
Piotr Karocki

Greg Luri
Rajesh Murthy
Cole Scott
Wil van der Aalst
Eric Verbeek

Moe Wynn
Yu Yuan
Oren Yuen

When the IEEE SA Standards Board approved this standard on 5 June 2023, it had the following membership:

David J. Law, *Chair*
Ted Burse, *Vice Chair*
Gary Hoffman, *Past Chair*
Konstantinos Karachalios, *Secretary*

Sara R. Biyabani
Doug Edwards
Ramy Ahmed Fathy
Guido R. Hiertz
Yousef Kimiagar
Joseph L. Koepfinger*
Thomas Koshy
John D. Kulick

Joseph S. Levy
Howard Li
Gui Lin
Johnny Daozhuang Lin
Kevin W. Lu
Daleep C. Mohla
Andrew Myles
Paul Nikolich

Annette D. Reilly
Robby Robson
Lei Wang
F.Keith Waters
Karl Weber
Philip B. Winston
Don Wright

*Member Emeritus

Introduction

This introduction is not part of IEEE Std 1849-2023, IEEE Standard for eXtensible Event Stream (XES) for Achieving Interoperability in Event Logs and Event Streams.

Event logs contain information on how processes have evolved in running systems. As more and more systems capture such information, there is a need to be able to transfer this information from these running systems to a site where the information can be analyzed, either automatically by software from the computational intelligence field, or manually (at least in part) using such software.

This standard addresses this need by defining an eXtensible Event Stream (XES) structure for such event logs.

Furthermore, this standard defines the World Wide Web Consortium (W3C) Extensible Markup Language (XML) structure and constraints on the contents of XML 1.1 documents that can be used to represent XES instances, and a likewise structure (called XEEXT) that can be used to represent so-called extensions to this structure.

The purpose of this standard is to allow the generation and analysis of event logs using XML. This standard uses the W3C XML Schema definition language as the encoding, which allows for interoperability and the exchange of XES XML instances between various systems.

Contents

1. Overview	11
1.1 Scope	11
1.2 Purpose	11
1.3 Word usage	12
2. Normative references	12
3. Definitions, acronyms, and abbreviations	12
3.1 Definitions	12
3.2 Acronyms and abbreviations	13
4. XES metadata structure	13
4.1 Hierarchical components	13
4.2 Attribute component	13
4.3 Global attributes	14
4.4 Classifiers	15
4.5 Extensions	16
5. XES XML serialization	17
5.1 Log element	18
5.2 Trace element	19
5.3 Event element	19
5.4 Attribute element	20
5.5 Extension element	21
5.6 Global element	21
5.7 Classifier element	22
6. XESEXT XML serialization	22
6.1 XES Extension element	22
6.2 Log element	23
6.3 Trace element	25
6.4 Event element	25
6.5 Meta element	25
6.6 Attribute element	26
6.7 Alias	26
7. XES standard extensions	26
7.1 Concept extension	27
7.2 Lifecycle extension	27
7.3 Organizational extension	29
7.4 Time extension	29
7.5 Semantic extension	29
7.6 ID extension	30
7.7 Cost extension	30
7.8 Micro extension	31
7.9 Software Event extension	31
7.10 Software Communication extension	33
7.11 Software Telemetry extension	34
7.12 Artifact Lifecycle extension	35
8. XES Conformance	35
8.1 Strictly conforming XES instances	35
8.2 Conforming XES instances	36

9. XEEXT Conformance.....	36
9.1 Strictly conforming XEEXT instances.....	36
9.2 Conforming XEEXT instances.....	36
Annex A (informative) History of XES	37
Annex B (informative) Current status of XES	38
Annex C (informative) XES support	39
Annex D (normative) XES Schema definition (XSD)	42
Annex E (normative) XEEXT Schema definition (XSD)	46
Annex F (informative) XEEXT Example.....	48
Annex G (informative) Artifact lifecycle extension example	50
Annex H (informative) Bibliography	52

IEEE Standard for eXtensible Event Stream (XES) for Achieving Interoperability in Event Logs and Event Streams

1. Overview

1.1 Scope

This standard defines World Wide Web Consortium (W3C) Extensible Markup Language (XML) structure and constraints on the contents of XML 1.1 documents that can be used to represent extensible event stream (XES) instances.⁶ A XES instance corresponds to a file-based event log or a formatted event stream that can be used to transfer event-driven data in a unified and extensible manner from a first site to a second site. Typically, the first site will be the site generating this event-driven data (e.g., workflow systems, case handling systems, procurement systems, devices like wafer steppers and X-ray machines, and hospitals) while the second site will be the site analyzing this data (e.g., by data scientists and/or advanced software systems).

To transfer event-driven data in a unified manner, this standard includes a W3C XML Schema describing the structure of an XES instance. To transfer this data in an extensible manner, this standard also includes a W3C XML Schema describing the structure of an extension to such an XES instance. Basically, such an extension provides semantics to the structure as prescribed by the XES instance. Finally, this standard includes a basic collection of such extensions.

1.2 Purpose

The purpose of this standard is to provide a generally acknowledged XML format for the interchange of event data between information systems in many applications domains on the one hand and analysis tools for such data on the other hand. As such, this standard aims to fix the syntax and the semantics of the event data which, for example, is being transferred from the site generating this data to the site analyzing this data. As a result of this standard, if the event data is transferred using the syntax as described by this standard, its semantics will be well understood and clear at both sites.

⁶Although XES is now an IEEE standard, XES was already used before it became an IEEE standard. [Annex A](#) shows the history of XES leading up to this revised IEEE standard, whereas [Annex B](#) contains a list of data sets already published using XES and a list of publications referring to XES.

1.3 Word usage

The word *shall* indicates mandatory requirements strictly to be followed in order to conform to the standard and from which no deviation is permitted (*shall* equals *is required to*).^{7,8}

The word *should* indicates that among several possibilities one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required (*should* equals *is recommended that*).

The word *may* is used to indicate a course of action permissible within the limits of the standard (*may* equals *is permitted to*).

The word *can* is used for statements of possibility and capability, whether material, physical, or causal (*can* equals *is able to*).

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.

ISO 639-1, Codes for the Representation of Names of Languages—Part 1: Alpha-2 Code.⁹

ISO 639-2, Codes for the Representation of Names of Languages—Part 2: Alpha-3 Code.

ISO 4217:2008, Codes for the Representation of Currencies and Funds.

ISO 8601, Data elements and interchange formats—Information interchange—Representation of dates and times.

XML Schema Part 1: Structures, Second Edition, W3C Recommendation (28 October 2004).¹⁰

XML Schema Part 2: Datatypes, Second Edition, W3C Recommendation (28 October 2004).¹¹

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.¹²

component: An extensible event stream (XES) element that may contain XES attributes, that is, a log, a trace, an event, or an attribute.

⁷The use of the word *must* is deprecated and cannot be used when stating mandatory requirements; *must* is used only to describe unavoidable situations.

⁸The use of *will* is deprecated and cannot be used when stating mandatory requirements; *will* is only used in statements of fact.

⁹ISO publications are available from the ISO Central Secretariat, <https://www.iso.org/>. ISO publications are also available in the United States from the American National Standards Institute, <https://www.ansi.org/>.

¹⁰Available at: <https://www.w3.org/TR/2004/REC-xmlschema-1-20041028/>.

¹¹Available at <https://www.w3.org/TR/xmlschema-2/>.

¹²*IEEE Standards Dictionary Online* is available at: <https://dictionary.ieee.org>. An IEEE account is required for access to the dictionary, and one can be created at no charge on the dictionary sign-in page.

3.2 Acronyms and abbreviations

BPAF	Business Process Analytics Format
UUID	universally unique identifier
URI	uniform resource identifier
XES	extensible event stream
XEEXT	XES extension
XML	Extensible Markup Language

4. XES metadata structure

4.1 Hierarchical components

4.1.1 Log component

A *log component* represents information that is related to a specific process. Examples of processes include handling insurance claims, using a complex X-ray machine, and browsing a website. A log shall contain a (possibly empty) collection of traces followed by a (possibly empty) list of events. The order of the events in this list shall be important, as it signifies the order in which the events have been observed.

If the log contains only events and no traces, then the log is also called a *stream*.

4.1.2 Trace component

A *trace component* represents the execution of a single case, that is, of a single execution (or enactment) of the specific process. A trace shall contain a (possibly empty) list of events that are related to a single case. The order of the events in this list shall be important, as it signifies the order in which the events have been observed.

4.1.3 Event component

An *event component* represents an atomic granule of activity that has been observed. If the event occurs in some trace, then it is clear to which case the event belongs. If the event does not occur in some trace, that is, if it occurs in the log, then ways to relate events to cases are needed. For this, a combination of a trace classifier and an event classifier will be used (see 4.4).

4.2 Attribute component

Information on any component (log, trace, or event) is stored in *attribute components*. Attributes describe the enclosing component, which may contain an arbitrary number of attributes. However, no two attributes of the same component may share the same key; that is, every key shall occur only once in a single component.

For providing maximum flexibility, this standard allows for *nested* attributes, that is, attributes that themselves have child attributes. While this feature is necessary for efficient encoding of certain information types, it is *optional for tools to implement nested attributes*; that is, the feature is not strictly required in order to be compliant to this standard. Nevertheless, a tool that does not support nested attributes shall be able to read documents that feature nested attributes. These tools shall transparently ignore and discard any nested attributes, and, where feasible, alert the user to the fact that some information may not be available.

An attribute shall be either *elementary* or *composite*.

4.2.1 Elementary attributes

An *elementary attribute* is an attribute that shall contain an elementary (single, basic) value. In this standard, an elementary attribute shall be a *string attribute*, a *date and time attribute*, an *integer number attribute*, a *real number attribute*, a *Boolean attribute*, or an *ID attribute*.

4.2.1.1 String attributes

Valid values for a string attribute are values that conform to the *xs:string* datatype.

4.2.1.2 Date and time attributes

Values for a date and time attribute should be specified in UTC time (see ISO 8601¹³), also known as Zulu time, and represented as *xs:dateTime* datatype.

4.2.1.3 Integer number attributes

Valid values for an integer number attribute are values that conform to the *xs:long* datatype.

4.2.1.4 Real number attributes

Valid values for a real number attribute are values that conform to the *xs:double* datatype.

4.2.1.5 Boolean attributes

Valid values for a Boolean attribute are values that conform to the *xs:boolean* datatype.

4.2.2 ID attributes

Valid values for an ID attribute are values that conform to the ID datatype, that is, all string representations of universally unique identifiers (UUIDs).

4.2.3 Composite attributes

A *composite attribute* is an attribute that may contain multiple values. In this standard, a composite attribute shall be a *list attribute*.

4.2.3.1 List attribute

Valid values for the list datatype are all lists (series) of attribute values. The order between the child attributes in this list shall be important. In contrast to attributes enclosed in a component, attributes enclosed in such an attribute value list may share the same key. Note that the attribute value list is not the same as the list attribute: The former is contained in the latter, and the latter is a component, but the former is not.

4.3 Global attributes

A log shall hold a (possibly empty) list of global attribute declarations. A global attribute declaration shall have a valid *key*, a valid *datatype*, and a valid *value* for the datatype. A global attribute declaration shall either be a global event attribute or a global trace attribute.

Global attributes are a required feature for compliance to this standard. Nevertheless, a defensive approach is recommended with respect to global attributes, as there is no way to undo a global declaration.

¹³Information on references can be found in [Clause 2](#).

4.3.1 Global event attributes

Global event attributes are event attributes that are understood to be *available* and *properly defined* for each event in the log (be it in a trace or not). As a result, every event in the log shall contain an attribute with the given key and the given datatype, but possibly with a different valid value. The value provided for a global event attribute declaration is *only significant* in case an event *needs to be created* (for some reason) for which *no value is provided* for that attribute. In that case, the value of the declaration shall be used as the value for the attribute. In all other cases, the value of the declaration is insignificant and shall not be used.

4.3.2 Global trace attributes

Global trace attributes are trace attributes that are understood to be *available* and *properly defined* for each trace in the log. As a result, every trace shall contain an attribute with the given key and the given datatype, but possibly with a different valid value. The value provided for a global attribute declaration is *only significant* in case a trace *needs to be created* (for some reason) for which *no value is provided* for that attribute. In that case, the value of the declaration shall be used as the value for the attribute. In all other cases, the value of the declaration is insignificant and shall not be used.

4.4 Classifiers

In this standard, there are no predefined attributes with any well-understood meaning per se. Instead, a log shall hold a (possibly empty) list of classifiers.

A classifier assigns an *identity* to each event that makes it comparable to others (via their assigned identity). Examples of such identities include the descriptive name of the event, the descriptive name of the case the event relates to, the descriptive name of the cause of the event, and the descriptive name of the case related to the event.

A classifier shall either be an *event classifier* or a *trace classifier*.

In case the log contains events that do not occur in a trace, it is necessary to be able to relate these events to cases. For this reason, it is assumed that one of the existing event classifiers provides the descriptive name of its case. Two events for which this classifier results in the same identity belong to the same case. Furthermore, it is assumed that one of the trace classifiers provides the descriptive name for the case. If this classifier and the event classifier mentioned earlier return the same identity, the corresponding events belong to the same case as the corresponding trace. As such, the event shall be appended to this trace in the same order as they appear in the log. If no matching trace exists, a new trace shall be constructed from these events again in the same order as they appear in the log.

4.4.1 Event classifiers

An event classifier shall be defined via an ordered list of attribute keys. The identity of the event shall be derived from the actual values of the attributes with these keys. An attribute whose key appears in an event classifier list shall be declared as a global event attribute before the event classifier is defined, as the actual value for the attribute is required by the event classifier.

4.4.2 Trace classifiers

A trace classifier shall be defined via an ordered list of attribute keys. The identity of the trace shall be derived from the actual values of the attributes with these keys. An attribute whose key appears in a trace classifier list shall be declared as a global trace attribute before the trace classifier is defined, as the actual value for the attribute is required by the trace classifier.

4.4.3 Event ordering

Within the context of a single trace, the ordering of events shall be important: An event that occurs in a log (be it in a trace or in the log itself) before another event that is related to the same trace shall be assumed to have occurred before that other event. However, the notion of a trace also depends on the trace and event classifiers selected by the user. As such, the notion of a trace is not necessarily predefined for a log. As a result, the notion of the event ordering may be affected by the choice of classifiers.

Whatever classifier is selected by the user, the ordering in the log shall be maintained: The first event that belongs to some trace shall be the first event encountered in the log (be it in a trace or in the log itself). For example, consider the event log that contains the following:

- a) A trace containing events e_{11} and e_{12}
- b) A trace containing an event e_{21}
- c) A trace containing events e_{31} , e_{32} , and e_{33}
- d) An event e_4

Now, assume that the user has selected an event classifier and a trace classifier that causes the events e_{11} , e_{31} , e_{33} , and e_4 to belong to the same trace. Then the ordering in this (classified) trace shall be e_{11} , e_{31} , e_{33} , and e_4 .

4.5 Extensions

This standard does not define a specific set of attributes per component. As such, the semantics of the data attributes these elements do contain may necessarily be ambiguous, hampering the interpretation of that data.

This ambiguity is resolved by the concept of *extensions* in this standard. An extension defines a (non-empty) set of attributes by defining a (possibly empty) set of attributes for every type of component. The extension provides points of reference for interpreting these attributes, and, thus, their components. Extensions, therefore, are *primarily a vehicle for attaching semantics to a set of defined attributes per component*.

Extensions have many possible uses. One important use is to introduce a set of commonly understood attributes that are vital for a specific perspective or dimension of event log analysis (and which may even not have been foreseen at the time of developing this standard). See [Clause 7](#) for the current set of standard extensions.

Other uses include the definition of generally understood attributes for a specific application domain (e.g., medical attributes for hospital processes), or for supporting special features or requirements of a specific application.

An extension shall have a descriptive *name*, a *prefix*, and a *uniform resource identifier (URI)*. The prefix is the prefix of all attributes defined by the extension. This means the keys of all attributes defined by the extension shall be prepended with this prefix and colon separation character (like a namespace in XML). The URI is a unique URI that points to the definition of the extension.

The definition of the extension shall contain a (possibly empty) list of attribute declarations for every component. An attribute declaration shall contain the *key* of the attribute, the *datatype* of the attribute, and a (possible empty) list of *aliases*. An alias shall contain the *descriptive text* for the attribute (that is, the commonly understood semantics for the attribute) and the *language code* of the language of this descriptive text.