INTERNATIONAL STANDARD

ISO/IEC 9834-8

Third edition
2014-08-15

Information technology — Procedures for the operation of object identifier registration authorities —

Part 8: Generation of universally unique identifiers (UUIDs) and their use in object identifiers

Technologies de l'information — Procédures opérationnelles pour les organismes d'enregistrement d'identificateur d'objet —

Partie 8: Génération des identificateurs uniques universels (UUID) et utilisation de ces identificateurs dans les composants d'identificateurs d'objets

Reference number
ISO/IEC 9834-8:2014(E)

© ISO/IEC 2014



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2014

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

 © ISO/IEC 2014 – All rights reserved

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 9834-8 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*, in collaboration with ITU-T. The identical text is published as Rec. ITU-T X.667 (10/2012).

This third edition cancels and replaces the second edition (ISO/IEC 9834-8:2008), which has been technically revised.

ISO/IEC 9834 consists of the following parts, under the general title *Information technology — Procedures for the operation of object identifier registration authorities*:

— *Part 1: General procedures and top arcs of the international object identifier tree*

— *Part 2: Registration procedures for OSI document types*

— *Part 3: Registration of Object Identifier arcs beneath the top-level arc jointly administered by ISO and ITU-T*

— *Part 4: Register of VTE Profiles*

— *Part 5: Register of VT Control Object Definitions*

— *Part 6: Registration of application processes and application entities*

— *Part 7: Joint ISO and ITU-T Registration of International Organizations*

— *Part 8: Generation of universally unique identifiers (UUIDs) and their use in object identifiers*

— *Part 9: Registration of object identifier arcs for applications and services using tag-based identification*

© ISO/IEC 2014 – All rights reserved

CONTENTS

Page

1 Scope 1
2 Normative references 1
2.1 Identical Recommendations | International Standards 1
2.2 Other normative references 1
3 Terms and definitions 2
3.1 ASN.1 notation 2
3.2 Registration authorities 2
3.3 Network terms 2
3.4 Additional definitions 2
4 Abbreviations 3
5 Notation 3
6 UUID structure and representation 3
6.1 UUID field structure 3
6.2 Binary representation 4
6.3 Representation as a single integer value 4
6.4 Hexadecimal representation 4
6.5 Formal syntax of the hexadecimal representation 4
7 Use of a UUID as the primary integer value and Unicode label of a Joint UUID arc 5
8 Use of a UUID to form a URN 5
9 Rules for comparison and ordering of UUIDs 5
10 Validation 6
11 Variant bits 6
12 Use of UUID fields and transmission byte order 6
12.1 General 6
12.2 Version 7
12.3 Time 7
12.4 Clock sequence 8
12.5 Node 8
13 Setting the fields of a time-based UUID 9
14 Setting the fields of a name-based UUID 9
15 Setting the fields of a random-number-based UUID 10
Annex A – Algorithms for the efficient generation of time-based UUIDs 11
A.1 Basic algorithm 11
A.2 Reading stable storage 11
A.3 System clock resolution 11
A.4 Writing stable storage 12
A.5 Sharing state across processes 12
Annex B – Properties of name-based UUIDs 13
Annex C – Generation of random numbers in a system 14
Annex D – Sample implementation 15
D.1 Files provided 15
D.2 The `copyrt.h` file 15
D.3 The `uuid.h` file 15
D.4 The `uuid.c` file 16
D.5 The `sysdep.h` file 19

Page
D.6 The `sysdep.c` file 19
D.7 The `utest.c` file 21
D.8 Sample output of `utest` 21
D.9 Some name space IDs 22
Bibliography 23

Introduction

This Recommendation | International Standard standardizes the generation of universally unique identifiers (UUIDs).

UUIDs are an octet string of 16 octets (128 bits). The 16 octets can be interpreted as an unsigned integer encoding, and the resulting integer value can be used as the primary integer value (defining an integer-valued Unicode label) for an arc of the International Object Identifier tree under the Joint UUID arc. This enables users to generate object identifier and OID internationalized resource identifier names without any registration procedure.

UUIDs are also known as globally unique identifiers (GUIDs), but this term is not used in this Recommendation | International Standard. UUIDs were originally used in the network computing system (NCS) [7] and later in the Open Software Foundation's Distributed Computing Environment (DCE) [6]. ISO/IEC 11578 [5] contains a short definition of some (but not all) of the UUID formats specified in this Recommendation | International Standard. The specification in this Recommendation | International Standard is consistent with all these earlier specifications.

UUIDs forming a component of an OID are represented in ASN.1 value notation as the decimal representation of their integer value, but for all other display purposes it is more usual to represent them with hexadecimal digits with a hyphen separating the different fields within the 16-octet UUID. This representation is defined in this Recommendation | International Standard.

If generated according to one of the mechanisms defined in this Recommendation | International Standard, a UUID is either guaranteed to be different from all other UUIDs generated before 3603 A.D., or is extremely likely to be different (depending on the mechanism chosen).

No centralized authority is required to administer UUIDs. Centrally generated UUIDs are guaranteed to be different from all other UUIDs centrally generated.

A UUID can be used for multiple purposes, from tagging objects with an extremely short lifetime, to reliably identifying very persistent objects across a network, particularly (but not necessarily) as part of an object identifier or OID internationalized resource identifier value, or in a uniform resource name (URN).

The UUID generation algorithm specified in this Recommendation | International Standard supports very high allocation rates: 10 million per second per machine if necessary, so UUIDs can also be used as transaction IDs. An informative annex provides a program in the C language that will generate UUIDs in accordance with this Recommendation | International Standard.

Three algorithms are specified for the generation of unique UUIDs, using different mechanisms to ensure uniqueness. These produce different versions of a UUID.

The first (and most common) mechanism produces the so-called time-based version. These UUIDs can be generated at the rate of 10 million per second. For UUIDs generated within a single computer system, a 60-bit time-stamp (used as a Clock value) with a granularity of 100 nanoseconds, based on coordinated universal time (UTC) is used to guarantee uniqueness over a period of approximately 1600 years. For UUIDs generated with the same time-stamp by different systems, uniqueness is obtained by use of 48-bit media access control (MAC) addresses, specified in ISO/IEC 8802-3 (this is used as a Node value). (These addresses are usually already available on most networked systems, but are otherwise obtainable from the IEEE Registration Authority for MAC addresses – see [4].) Alternative ways of generating Clock and Node values are specified for the time-based version if UTC time is not available on a system, or if there is no MAC address available.

The second mechanism produces a single UUID that is a name-based version, where cryptographic hashing is used to produce the 128-bit UUID value from a globally unambiguous (text) name.

The third mechanism uses pseudo-random or truly random number generation to produce most of the bits in the 128-bit value.

Clause 5 specifies the notation used for octet-order and bit-order naming, and for specification of transmission order.

Clause 6 specifies the structure of a UUID and the representation of it in binary, hexadecimal, or as a single integer value.

Clauses 7 and 8 specify the use of a UUID in an OID or a URN respectively.

Clause 9 specifies rules for comparing UUIDs to test for equality or to provide an ordering relation between two UUIDs.

Clause 10 discusses the possibility of checking the validity of a UUID. In general, UUIDs have little redundancy, and there is little scope for checking their validity.

Clause 11 describes the historical use of some bits in the UUID to define different variants of the UUID format, and specifies the value of these bits for UUIDs defined in accordance with this Recommendation | International Standard.

Clause 12 specifies the use of the fields of a UUID in the different versions that are defined (time-based, name-based, and random-number based versions). It also defines the transmission byte order.

Clause 13 specifies the setting of the fields of a time-based UUID.

Clause 14 specifies the setting of the fields of a name-based UUID.

Clause 15 specifies the setting of the fields of a random-number-based UUID.

All annexes are informative.

Annex A describes various algorithms for the efficient generation of time-based UUIDs.

Annex B discusses the properties that a name-based UUID should have, affecting the selection of name spaces for use in generating such UUIDs.

Annex C provides guidance on mechanisms that can be used to generate random numbers in a computer system.

Annex D contains a complete program in the C programming language that can be used to generate UUIDs.

INTERNATIONAL STANDARD
RECOMMENDATION ITU-T

Information technology – Procedures for the operation of object identifier registration authorities: Generation of universally unique identifiers and their use in object identifiers

1 Scope

This Recommendation | International Standard specifies the format and generation rules that enable users to produce 128-bit identifiers that are either guaranteed to be globally unique, or are globally unique with a high probability.

The universally unique identifiers (UUIDs) generated in conformance with this Recommendation | International Standard are suitable either for transient use, with generation of a new UUID every 100 nanoseconds, or as persistent identifiers.

This Recommendation | International Standard is derived from earlier non-standard specifications of UUIDs and their generation, and is technically identical to those earlier specifications.

This Recommendation | International Standard also specifies and allows the use of UUIDs as primary values (which define Unicode labels) for arcs beneath the Joint UUID arc. This enables users to generate and use such arcs without any registration procedures.

This Recommendation | International Standard also specifies and allows the use of UUIDs to form a uniform resource name (URN).

2 Normative references

The following Recommendations and International Standards contain provisions which, through reference in this text, constitute provisions of this Recommendation | International Standard. At the time of publication, the editions indicated were valid. All Recommendations and Standards are subject to revision, and parties to agreements based on this Recommendation | International Standard are encouraged to investigate the possibility of applying the most recent edition of the Recommendations and Standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards. The Telecommunication Standardization Bureau of the ITU maintains a list of currently valid ITU-T Recommendations.

2.1 Identical Recommendations | International Standards

- Recommendation ITU-T X.660 (2011) | ISO/IEC 9834-1:2012, *Information technology – Procedures for the operation of object identifier registration authorities: General procedures and top arcs of the international object identifier tree*.
- Recommendation ITU-T X.680 (2008) | ISO/IEC 8824-1:2008, *Information technology – Abstract Syntax Notation One (ASN.1): Specification of basic notation*.

2.2 Other normative references

- ISO/IEC 8802-3:2000, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specification*s.
- ISO/IEC 10118-3:2004, *Information technology – Security techniques – Hash functions – Part 3: Dedicated hash-functions*.
- ISO/IEC 10646:2012, *Information technology – Universal Coded Character Set (UCS)*.
- IETF RFC 1321 (1992), *The MD5 Message-Digest Algorithm*.
- IETF RFC 2141 (1997), *URN Syntax*.
- FIPS PUB 180-3:2008, *Federal Information Processing Standards Publication, Secure Hash Standard (SHS)*.

3 Terms and definitions

For the purposes of this Recommendation | International Standard, the following definitions apply.

3.1 ASN.1 notation

This Recommendation | International Standard uses the following terms defined in Rec. ITU-T X.680 | ISO/IEC 8824-1:

a) Coordinated universal time (UTC);

b) object identifier type;

c) OID internationalized resource identifier type.

3.2 Registration authorities

This Recommendation | International Standard uses the following terms defined in Rec. ITU-T X.660 | ISO/IEC 9834-1:

a) International Object Identifier tree;

b) IRI/URI value;

c) OID internationalized resource identifier;

d) object identifier;

e) registration;[2]

f) registration authority;

g) registration procedures;

h) secondary identifier;

i) Unicode character;

j) Unicode label.

3.3 Network terms

This Recommendation | International Standard uses the following term defined in ISO/IEC 8802-3:

– MAC address.

3.4 Additional definitions

3.4.1 cryptographic-quality random-number: A random number or pseudo-random number generated by a mechanism, which ensures sufficient spread of repeatedly-generated values to be acceptable for use in cryptographic work (and is used in such work).

3.4.2 Joint UUID arc: An arc beneath the node of the International Object Identifier tree identified by the ASN.1 object identifier value **`{joint-iso-itu-t(2) uuid(25)}`** and the ASN.1 OID internationalized resource identifier value **`"/UUID"`**.

3.4.3 name-based version: A UUID that is generated using cryptographic hashing of a name space name and a name in that name space.

3.4.4 name space: A system for generating names of objects that ensures unambiguous identification within that name space.

NOTE – Examples of name spaces are the network domain name system, URNs, OIDs, Directory distinguished names (see [1]), and reserved words in a programming language.

3.4.5 random-number-based version: A UUID that is generated using a random or pseudo-random number.

3.4.6 standard UUID variant: The variant of the possible UUID formats that is specified by this Recommendation | International Standard.

NOTE – Historically, there have been other specifications of UUID formats that differ from the variant specified in this Recommendation | International Standard. UUIDs generated according to all these variant formats are all distinct.

3.4.7 time-based version: A UUID in which uniqueness is obtained by the use of a MAC address to identify a system, and a Clock value based on the current UTC time.