

A Universally Unique Identifier (UUID) URN Namespace

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the “Internet Official Protocol Standards” (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright © The Internet Society (2005). All Rights Reserved.

Abstract

This specification defines a Uniform Resource Name namespace for UUIDs (Universally Unique Identifier), also known as GUIDs (Globally Unique Identifier). A UUID is 128 bits long, and can guarantee uniqueness across space and time. UUIDs were originally used in the Apollo Network Computing System and later in the Open Software Foundation's (OSF) Distributed Computing Environment (DCE), and then in Microsoft Windows platforms.

This specification is derived from the DCE specification with the kind permission of the OSF (now known as The Open Group). Information from earlier versions of the DCE specification have been incorporated into this document.

Table of Contents

1. Introduction	3
2. Motivation	4
3. Namespace Registration Template	5
4. Specification	7
4.1. Format.....	7
4.1.1. Variant.....	7
4.1.2. Layout and Byte Order.....	7
4.1.3. Version.....	8
4.1.4. Timestamp.....	9
4.1.5. Clock Sequence.....	9
4.1.6. Node.....	10
4.1.7. Nil UUID.....	10
4.2. Algorithms for Creating a Time-Based UUID.....	10
4.2.1. Basic Algorithm.....	10
4.2.1.1. Reading Stable Storage.....	10
4.2.1.2. System Clock Resolution.....	11
4.2.1.3. Writing Stable Storage.....	11
4.2.1.4. Sharing State Across Processes.....	11
4.2.2. Generation Details.....	11
4.3. Algorithm for Creating a Name-Based UUID.....	12
4.4. Algorithms for Creating a UUID from Truly Random or Pseudo-Random Numbers.....	12
4.5. Node IDs that Do Not Identify the Host.....	13
5. Community Considerations	14
6. Security Considerations	15
7. Acknowledgments	16
8. Normative References	17
A. Appendix A - Sample Implementation	18
B. Appendix B - Sample Output of utest	20
C. Appendix C - Some Name Space IDs	21
Authors' Addresses	22
Intellectual Property and Copyright Statements	23

1. Introduction

This specification defines a Uniform Resource Name namespace for UUIDs (Universally Unique Identifier), also known as GUIDs (Globally Unique Identifier). A UUID is 128 bits long, and requires no central registration process.

The information here is meant to be a concise guide for those wishing to implement services using UUIDs as URNs. Nothing in this document should be construed to override the DCE standards that defined UUIDs.

There is an [ITU-T Recommendation and ISO/IEC Standard](#) [3] that are derived from earlier versions of this document. Both sets of specifications have been aligned, and are fully technically compatible. In addition, a global registration function is being provided by the Telecommunications Standardisation Bureau of ITU-T; for details see <http://www.itu.int/ITU-T/asn1/uuid.html>.

2. Motivation

One of the main reasons for using UUIDs is that no centralized authority is required to administer them (although one format uses IEEE 802 node identifiers, others do not). As a result, generation on demand can be completely automated, and used for a variety of purposes. The UUID generation algorithm described here supports very high allocation rates of up to 10 million per second per machine if necessary, so that they could even be used as transaction IDs.

UUIDs are of a fixed size (128 bits) which is reasonably small compared to other alternatives. This lends itself well to sorting, ordering, and hashing of all sorts, storing in databases, simple allocation, and ease of programming in general.

Since UUIDs are unique and persistent, they make excellent Uniform Resource Names. The unique ability to generate a new UUID without a registration process allows for UUIDs to be one of the URNs with the lowest minting cost.

3. Namespace Registration Template

Namespace ID: UUID

Registration Information:

Registration date: 2003-10-01

Declared registrant of the namespace:

JTC 1/SC6 (ASN.1 Rapporteur Group)

Declaration of syntactic structure:

A UUID is an identifier that is unique across both space and time, with respect to the space of all UUIDs. Since a UUID is a fixed size and contains a time field, it is possible for values to rollover (around A.D. 3400, depending on the specific algorithm used). 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.

The internal representation of a UUID is a specific sequence of bits in memory, as described in [Section 4](#). To accurately represent a UUID as a URN, it is necessary to convert the bit sequence to a string representation.

Each field is treated as an integer and has its value printed as a zero-filled hexadecimal digit string with the most significant digit first. The hexadecimal values "a" through "f" are output as lower case characters and are case insensitive on input.

The formal definition of the UUID string representation is provided by the following ABNF [7]:

```

UUID          = time-low "-" time-mid "-"
                time-high-and-version "-"
                clock-seq-and-reserved
                clock-seq-low "-" node

time-low      = 4hexOctet
time-mid      = 2hexOctet
time-high-and-version = 2hexOctet
clock-seq-and-reserved = hexOctet
clock-seq-low  = hexOctet
node          = 6hexOctet
hexOctet      = hexDigit hexDigit
hexDigit =
    "0" / "1" / "2" / "3" / "4" / "5" / "6" / "7" / "8" / "9" /
    "a" / "b" / "c" / "d" / "e" / "f" /
    "A" / "B" / "C" / "D" / "E" / "F"

```

The following is an example of the string representation of a UUID as a URN:

```
urn:uuid:f81d4fae-7dec-11d0-a765-00a0c91e6bf6
```

Relevant ancillary documentation:

[1][2]

Identifier uniqueness considerations:

This document specifies three algorithms to generate UUIDs: the first leverages the unique values of 802 MAC addresses to guarantee uniqueness, the second uses pseudo-random number generators, and the third uses cryptographic hashing and application-provided text strings. As a result, the UUIDs generated according to the mechanisms here will be unique from all other UUIDs that have been or will be assigned.

Identifier persistence considerations:

UUIDs are inherently very difficult to resolve in a global sense. This, coupled with the fact that UUIDs are temporally unique within their spatial context, ensures that UUIDs will remain as persistent as possible.

Process of identifier assignment:

Generating a UUID does not require that a registration authority be contacted. One algorithm requires a unique value over space for each generator. This value is typically an IEEE 802 MAC address, usually already available on network-connected hosts. The address can be assigned from an address block obtained from the IEEE registration authority. If no such address is available, or privacy concerns make its use undesirable, [Section 4.5](#) specifies two alternatives. Another approach is to use version 3 or version 4 UUIDs as defined below.

Process for identifier resolution:

Since UUIDs are not globally resolvable, this is not applicable.

Rules for Lexical Equivalence:

Consider each field of the UUID to be an unsigned integer as shown in the table in section [Section 4.1.2](#). Then, to compare a pair of UUIDs, arithmetically compare the corresponding fields from each UUID in order of significance and according to their data type. Two UUIDs are equal if and only if all the corresponding fields are equal.

As an implementation note, equality comparison can be performed on many systems by doing the appropriate byte-order canonicalization, and then treating the two UUIDs as 128-bit unsigned integers.

UUIDs, as defined in this document, can also be ordered lexicographically. For a pair of UUIDs, the first one follows the second if the most significant field in which the UUIDs differ is greater for the first UUID. The second precedes the first if the most significant field in which the UUIDs differ is greater for the second UUID.

Conformance with URN Syntax:

The string representation of a UUID is fully compatible with the URN syntax. When converting from a bit-oriented, in-memory representation of a UUID into a URN, care must be taken to strictly adhere to the byte order issues mentioned in the string representation section.

Validation mechanism:

Apart from determining whether the timestamp portion of the UUID is in the future and therefore not yet assignable, there is no mechanism for determining whether a UUID is 'valid'.

Scope:

UUIDs are global in scope.

4. Specification

4.1. Format

The UUID format is 16 octets; some bits of the eight octet variant field specified below determine finer structure.

4.1.1. Variant

The variant field determines the layout of the UUID. That is, the interpretation of all other bits in the UUID depends on the setting of the bits in the variant field. As such, it could more accurately be called a type field; we retain the original term for compatibility. The variant field consists of a variable number of the most significant bits of octet 8 of the UUID.

The following table lists the contents of the variant field, where the letter "x" indicates a "don't-care" value.

Msb0	Msb1	Msb2	Description
0	x	x	Reserved, NCS backward compatibility.
1	0	x	The variant specified in this document.
1	1	0	Reserved, Microsoft Corporation backward compatibility
1	1	1	Reserved for future definition.

Interoperability, in any form, with variants other than the one defined here is not guaranteed, and is not likely to be an issue in practice.

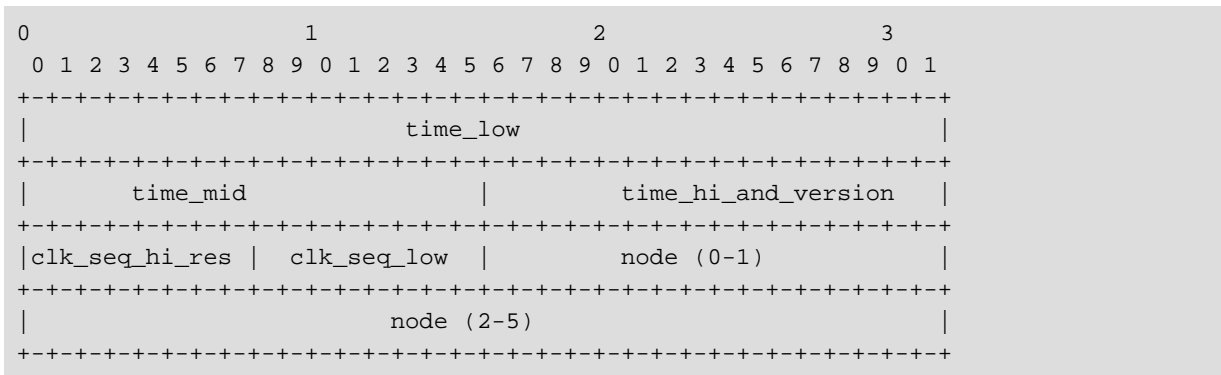
4.1.2. Layout and Byte Order

To minimize confusion about bit assignments within octets, the UUID record definition is defined only in terms of fields that are integral numbers of octets. The fields are presented with the most significant one first.

Field	Data Type	Octet #	Note
time_low	unsigned 32 bit integer	0-3	The low field of the timestamp
time_mid	unsigned 16 bit integer	4-5	The middle field of the timestamp
time_hi_and_version	unsigned 16 bit integer	6-7	The high field of the timestamp multiplexed with the version number
clock_seq_hi_and_reserved	unsigned 8 bit integer	8	The high field of the clock sequence multiplexed with the variant
clock_seq_low	unsigned 8 bit integer	9	The low field of the clock sequence
node	unsigned 48 bit integer	10-15	The spatially unique node identifier

In the absence of explicit application or presentation protocol specification to the contrary, a UUID is encoded as a 128-bit object, as follows:

The fields are encoded as 16 octets, with the sizes and order of the fields defined above, and with each field encoded with the Most Significant Byte first (known as network byte order). Note that the field names, particularly for multiplexed fields, follow historical practice.



4.1.3. Version

The version number is in the most significant 4 bits of the time stamp (bits 4 through 7 of the time_hi_and_version field).

The following table lists the currently-defined versions for this UUID variant.

Msb0	Msb1	Msb2	Msb3	Version	Description
0	0	0	1	1	The time-based version specified in this document.
0	0	1	0	2	DCE Security version, with embedded POSIX UIDs.
0	0	1	1	3	The name-based version specified in this document that uses MD5 hashing.
0	1	0	0	4	The randomly or pseudo-randomly generated version specified in this document.
0	1	0	1	5	The name-based version specified in this document that uses SHA-1 hashing.

The version is more accurately a sub-type; again, we retain the term for compatibility.

4.1.4. Timestamp

The timestamp is a 60-bit value. For UUID version 1, this is represented by Coordinated Universal Time (UTC) as a count of 100-nanosecond intervals since 00:00:00.00, 15 October 1582 (the date of Gregorian reform to the Christian calendar).

For systems that do not have UTC available, but do have the local time, they may use that instead of UTC, as long as they do so consistently throughout the system. However, this is not recommended since generating the UTC from local time only needs a time zone offset.

For UUID version 3 or 5, the timestamp is a 60-bit value constructed from a name as described in [Section 4.3](#).

For UUID version 4, the timestamp is a randomly or pseudo-randomly generated 60-bit value, as described in [Section 4.4](#).

4.1.5. Clock Sequence

For UUID version 1, the clock sequence is used to help avoid duplicates that could arise when the clock is set backwards in time or if the node ID changes.

If the clock is set backwards, or might have been set backwards (e.g., while the system was powered off), and the UUID generator can not be sure that no UUIDs were generated with timestamps larger than the value to which the clock was set, then the clock sequence has to be changed. If the previous value of the clock sequence is known, it can just be incremented; otherwise it should be set to a random or high-quality pseudo-random value.

Similarly, if the node ID changes (e.g., because a network card has been moved between machines), setting the clock sequence to a random number minimizes the probability of a duplicate due to slight differences in the clock settings of the machines. If the value of clock sequence associated with the changed node ID were known, then the clock sequence could just be incremented, but that is unlikely.

The clock sequence **MUST** be originally (i.e., once in the lifetime of a system) initialized to a random number to minimize the correlation across systems. This provides maximum protection against node identifiers that may move or switch from system to system rapidly. The initial value **MUST NOT** be correlated to the node identifier.

For UUID version 3 or 5, the clock sequence is a 14-bit value constructed from a name as described in [Section 4.3](#).

For UUID version 4, clock sequence is a randomly or pseudo-randomly generated 14-bit value as described in [Section 4.4](#).

4.1.6. Node

For UUID version 1, the node field consists of an IEEE 802 MAC address, usually the host address. For systems with multiple IEEE 802 addresses, any available one can be used. The lowest addressed octet (octet number 10) contains the global/local bit and the unicast/multicast bit, and is the first octet of the address transmitted on an 802.3 LAN.

For systems with no IEEE address, a randomly or pseudo-randomly generated value may be used; see [Section 4.5](#). The multicast bit must be set in such addresses, in order that they will never conflict with addresses obtained from network cards.

For UUID version 3 or 5, the node field is a 48-bit value constructed from a name as described in [Section 4.3](#).

For UUID version 4, the node field is a randomly or pseudo-randomly generated 48-bit value as described in [Section 4.4](#).

4.1.7. Nil UUID

The nil UUID is special form of UUID that is specified to have all 128 bits set to zero.

4.2. Algorithms for Creating a Time-Based UUID

Various aspects of the algorithm for creating a version 1 UUID are discussed in the following sections.

4.2.1. Basic Algorithm

The following algorithm is simple, correct, and inefficient:

- Obtain a system-wide global lock
- From a system-wide shared stable store (e.g., a file), read the UUID generator state: the values of the timestamp, clock sequence, and node ID used to generate the last UUID.
- Get the current time as a 60-bit count of 100-nanosecond intervals since 00:00:00.00, 15 October 1582.
- Get the current node ID.
- If the state was unavailable (e.g., non-existent or corrupted), or the saved node ID is different than the current node ID, generate a random clock sequence value.
- If the state was available, but the saved time stamp is later than the current timestamp, increment the clock sequence value.
- Save the state (current time stamp, clock sequence, and node ID) back to the stable store.
- Release the global lock.
- Format a UUID from the current timestamp, clock sequence, and node ID values according to the steps in [Section 4.2.2](#).

If UUIDs do not need to be frequently generated, the above algorithm may be perfectly adequate. For higher performance requirements, however, issues with the basic algorithm include:

- Reading the state from stable storage each time is inefficient.
- The resolution of the system clock may not be 100-nanoseconds.
- Writing the state to stable storage each time is inefficient.
- Sharing the state across process boundaries may be inefficient.

Each of these issues can be addressed in a modular fashion by local improvements in the functions that read and write the state and read the clock. We address each of them in turn in the following sections.

4.2.1.1. Reading Stable Storage

The state only needs to be read from stable storage once at boot time, if it is read into a system-wide shared volatile store (and updated whenever the stable store is updated).

If an implementation does not have any stable store available, then it can always say that the values were unavailable. This is the least desirable implementation because it will increase the frequency of creation of new clock sequence numbers, which increases the probability of duplicates.

If the node ID can never change (e.g., the net card is inseparable from the system), or if any change also reinitializes the clock sequence to a random value, then instead of keeping it in stable store, the current node ID may be returned.

4.2.1.2. System Clock Resolution

The timestamp is generated from the system time, whose resolution may be less than the resolution of the UUID timestamp.

If UUIDs do not need to be frequently generated, the time stamp can simply be the system time multiplied by the number of 100-nanosecond intervals per system time interval.

If a system overruns the generator by requesting too many UUIDs within a single system time interval, the UUID service **MUST** either return an error, or stall the UUID generator until the system clock catches up.

A high resolution timestamp can be simulated by keeping a count of the number of UUIDs that have been generated with the same value of the system time, and using it to construct the low order bits of the timestamp. The count will range between zero and the number of 100-nanosecond intervals per system time interval.

Note: If the processors overrun the UUID generation frequently, additional node identifiers can be allocated to the system, which will permit higher speed allocation by making multiple UUIDs potentially available for each time stamp value.

4.2.1.3. Writing Stable Storage

The state does not always need to be written to stable store every time a UUID is generated. The timestamp in the stable store can be periodically set to a value larger than any yet used in a UUID. As long as the generated UUIDs have timestamps less than that value, and the clock sequence and node ID remain unchanged, only the shared volatile copy of the state needs to be updated. Furthermore, if the time stamp value in stable store is in the future by less than the typical time it takes the system to reboot, a crash will not cause a reinitialization of the clock sequence.

4.2.1.4. Sharing State Across Processes

If it is too expensive to access shared state each time a UUID is generated, then the system-wide generator can be implemented to allocate a block of time stamps each time it is called; a per-process generator can allocate from that block until it is exhausted.

4.2.2. Generation Details

Version 1 UUIDs are generated according to the following algorithm:

- Determine the values for the UTC-based timestamp and clock sequence to be used in the UUID, as described in [Section 4.2.1](#).
- For the purposes of this algorithm, consider the timestamp to be a 60-bit unsigned integer and the clock sequence to be a 14-bit unsigned integer. Sequentially number the bits in a field, starting with zero for the least significant bit.
- Set the `time_low` field equal to the least significant 32 bits (bits zero through 31) of the timestamp in the same order of significance.
- Set the `time_mid` field equal to bits 32 through 47 from the timestamp in the same order of significance.

- Set the 12 least significant bits (bits zero through 11) of the `time_hi_and_version` field equal to bits 48 through 59 from the timestamp in the same order of significance.
- Set the four most significant bits (bits 12 through 15) of the `time_hi_and_version` field to the 4-bit version number corresponding to the UUID version being created, as shown in the table above.
- Set the `clock_seq_low` field to the eight least significant bits (bits zero through 7) of the clock sequence in the same order of significance.
- Set the 6 least significant bits (bits zero through 5) of the `clock_seq_hi_and_reserved` field to the 6 most significant bits (bits 8 through 13) of the clock sequence in the same order of significance.
- Set the two most significant bits (bits 6 and 7) of the `clock_seq_hi_and_reserved` to zero and one, respectively.
- Set the node field to the 48-bit IEEE address in the same order of significance as the address.

4.3. Algorithm for Creating a Name-Based UUID

The version 3 or 5 UUID is meant for generating UUIDs from "names" that are drawn from, and unique within, some "name space". The concept of name and name space should be broadly construed, and not limited to textual names. For example, some name spaces are the domain name system, URLs, ISO Object IDs (OIDs), X.500 Distinguished Names (DNs), and reserved words in a programming language. The mechanisms or conventions used for allocating names and ensuring their uniqueness within their name spaces are beyond the scope of this specification.

The requirements for these types of UUIDs are as follows:

- The UUIDs generated at different times from the same name in the same namespace **MUST** be equal.
- The UUIDs generated from two different names in the same namespace should be different (with very high probability).
- The UUIDs generated from the same name in two different namespaces should be different with (very high probability).
- If two UUIDs that were generated from names are equal, then they were generated from the same name in the same namespace (with very high probability).

The algorithm for generating a UUID from a name and a name space are as follows:

- Allocate a UUID to use as a "name space ID" for all UUIDs generated from names in that name space; see [Appendix C](#) for some pre-defined values.
- Choose either [MD5](#) [4] or [SHA-1](#) [8] as the hash algorithm; If backward compatibility is not an issue, SHA-1 is preferred.
- Convert the name to a canonical sequence of octets (as defined by the standards or conventions of its name space); put the name space ID in network byte order.
- Compute the hash of the name space ID concatenated with the name.
- Set octets zero through 3 of the `time_low` field to octets zero through 3 of the hash.
- Set octets zero and one of the `time_mid` field to octets 4 and 5 of the hash.
- Set octets zero and one of the `time_hi_and_version` field to octets 6 and 7 of the hash.
- Set the four most significant bits (bits 12 through 15) of the `time_hi_and_version` field to the appropriate 4-bit version number from [Section 4.1.3](#).
- Set the `clock_seq_hi_and_reserved` field to octet 8 of the hash.
- Set the two most significant bits (bits 6 and 7) of the `clock_seq_hi_and_reserved` to zero and one, respectively.
- Set the `clock_seq_low` field to octet 9 of the hash.
- Set octets zero through five of the node field to octets 10 through 15 of the hash.
- Convert the resulting UUID to local byte order.

4.4. Algorithms for Creating a UUID from Truly Random or Pseudo-Random Numbers

The version 4 UUID is meant for generating UUIDs from truly-random or pseudo-random numbers.

The algorithm is as follows:

- Set the two most significant bits (bits 6 and 7) of the `clock_seq_hi_and_reserved` to zero and one, respectively.
- Set the four most significant bits (bits 12 through 15) of the `time_hi_and_version` field to the 4-bit version number from [Section 4.1.3](#).
- Set all the other bits to randomly (or pseudo-randomly) chosen values.

See [Section 4.5](#) for a discussion on random numbers.

4.5. Node IDs that Do Not Identify the Host

This section describes how to generate a version 1 UUID if an IEEE 802 address is not available, or its use is not desired.

One approach is to contact the IEEE and get a separate block of addresses. At the time of writing, the application could be found at <http://standards.ieee.org/regauth/oui/pilot-ind.html>, and the cost was US\$550.

A better solution is to obtain a 47-bit cryptographic quality random number and use it as the low 47 bits of the node ID, with the least significant bit of the first octet of the node ID set to one. This bit is the unicast/multicast bit, which will never be set in IEEE 802 addresses obtained from network cards. Hence, there can never be a conflict between UUIDs generated by machines with and without network cards. (Recall that the IEEE 802 spec talks about transmission order, which is the opposite of the in-memory representation that is discussed in this document.)

For compatibility with earlier specifications, note that this document uses the unicast/multicast bit, instead of the arguably more correct local/global bit.

Advice on generating cryptographic-quality random numbers can be found in [RFC1750](#) [5].

In addition, items such as the computer's name and the name of the operating system, while not strictly speaking random, will help differentiate the results from those obtained by other systems.

The exact algorithm to generate a node ID using these data is system specific, because both the data available and the functions to obtain them are often very system specific. A generic approach, however, is to accumulate as many sources as possible into a buffer, use a message digest such as [MD5](#) [4] or [SHA-1](#) [8], take an arbitrary 6 bytes from the hash value, and set the multicast bit as described above.

5. Community Considerations

The use of UUIDs is extremely pervasive in computing. They comprise the core identifier infrastructure for many operating systems (Microsoft Windows) and applications (the Mozilla browser) and in many cases, become exposed to the Web in many non-standard ways. This specification attempts to standardize that practice as openly as possible and in a way that attempts to benefit the entire Internet.

6. Security Considerations

Do not assume that UUIDs are hard to guess; they should not be used as security capabilities (identifiers whose mere possession grants access), for example. A predictable random number source will exacerbate the situation.

Do not assume that it is easy to determine if a UUID has been slightly transposed in order to redirect a reference to another object. Humans do not have the ability to easily check the integrity of a UUID by simply glancing at it.

Distributed applications generating UUIDs at a variety of hosts must be willing to rely on the random number source at all hosts. If this is not feasible, the namespace variant should be used.

7. Acknowledgments

This document draws heavily on the OSF DCE specification for UUIDs. Ted Ts'o provided helpful comments, especially on the byte ordering section which we mostly plagiarized from a proposed wording he supplied (all errors in that section are our responsibility, however).

We are also grateful to the careful reading and bit-twiddling of Ralf S. Engelschall, John Larmouth, and Paul Thorpe. Professor Larmouth was also invaluable in achieving coordination with ISO/IEC.

8 Normative References

- [1] Zahn, L., Dineen, T., and P. Leach, "Network Computing Architecture", ISBN 0-13-611674-4, January 1990.
- [2] "DCE: Remote Procedure Call", Open Group CAE Specification C309, ISBN 1-85912-041-5, August 1994.
- [3] ISO/IEC 9834-8:2004 Information Technology, "Procedures for the operation of OSI Registration Authorities: Generation and registration of Universally Unique Identifiers (UUIDs) and their use as ASN.1 Object Identifier components", ITU-T Rec. X.687, 2004.
- [4] Rivest, R., "[The MD5 Message-Digest Algorithm](#)", RFC 1321, April 1992.
- [5] Eastlake, D., Schiller, J., and S. Crocker, "[Randomness Requirements for Security](#)", BCP 106, RFC 4086, June 2005.
- [6] Moats, R., "[URN Syntax](#)", RFC 2141, May 1997.
- [7] Crocker, D., Ed. and P. Overell, "[Augmented BNF for Syntax Specifications: ABNF](#)", RFC 2234, November 1997.
- [8] National Institute of Standards and Technology, "[Secure Hash Standard](#)", FIPS PUB 180-1, April 1995, <<http://www.itl.nist.gov/fipspubs/fip180-1.htm>>.

A. Appendix A - Sample Implementation

This implementation consists of 5 files: `uuid.h`, `uuid.c`, `sysdep.h`, `sysdep.c` and `utest.c`. The `uuid.*` files are the system independent implementation of the UUID generation algorithms described above, with all the optimizations described above except efficient state sharing across processes included. The code has been tested on Linux (Red Hat 4.0) with GCC (2.7.2), and Windows NT 4.0 with VC++ 5.0. The code assumes 64-bit integer support, which makes it much clearer.

All the following source files should have the following copyright notice included:

```
copyrt.h
```

```
/*
** Copyright (c) 1990- 1993, 1996 Open Software Foundation, Inc.
** Copyright (c) 1989 by Hewlett-Packard Company, Palo Alto, Ca. &
** Digital Equipment Corporation, Maynard, Mass.
** Copyright (c) 1998 Microsoft.
** To anyone who acknowledges that this file is provided "AS IS"
** without any express or implied warranty: permission to use, copy,
** modify, and distribute this file for any purpose is hereby
** granted without fee, provided that the above copyright notices and
** this notice appears in all source code copies, and that none of
** the names of Open Software Foundation, Inc., Hewlett-Packard
** Company, or Digital Equipment Corporation be used in advertising
** or publicity pertaining to distribution of the software without
** specific, written prior permission. Neither Open Software
** Foundation, Inc., Hewlett-Packard Company, Microsoft, nor Digital
** Equipment Corporation makes any representations about the suitability
** of this software for any purpose.
*/
```

```
uuid.h
```

```
#include "copyrt.h"
#undef uuid_t
typedef struct {
    unsigned32  time_low;
    unsigned16  time_mid;
    unsigned16  time_hi_and_version;
    unsigned8   clock_seq_hi_and_reserved;
    unsigned8   clock_seq_low;
    byte        node[6];
} uuid_t;

/* uuid_create -- generate a UUID */
int uuid_create(uuid_t * uuid);

/* uuid_create_md5_from_name -- create a version 3 (MD5) UUID using a
   "name" from a "name space" */
void uuid_create_md5_from_name(
    uuid_t *uuid,          /* resulting UUID */
    uuid_t nsid,          /* UUID of the namespace */
    void *name,           /* the name from which to generate a UUID */
    int namelen           /* the length of the name */
);

/* uuid_create_shal_from_name -- create a version 5 (SHA-1) UUID
   using a "name" from a "name space" */
void uuid_create_shal_from_name(
    uuid_t *uuid,          /* resulting UUID */
    uuid_t nsid,          /* UUID of the namespace */
    void *name,           /* the name from which to generate a UUID */
    int namelen           /* the length of the name */
);

/* uuid_compare -- Compare two UUID's "lexically" and return
   -1  u1 is lexically before u2
   0   u1 is equal to u2
   1   u1 is lexically after u2
```

B. Appendix B - Sample Output of utest

```
uuid_create(): 7d444840-9dc0-11d1-b245-5ffdce74fad2
uuid_compare(u,u): 0
uuid_compare(u, NameSpace_DNS): 1
uuid_compare(NameSpace_DNS, u): -1
uuid_create_md5_from_name(): e902893a-9d22-3c7e-a7b8-d6e313b71d9f
```

C. Appendix C - Some Name Space IDs

This appendix lists the name space IDs for some potentially interesting name spaces, as initialized C structures and in the string representation defined above.

```
/* Name string is a fully-qualified domain name */
uuid_t NameSpace_DNS = { /* 6ba7b810-9dad-11d1-80b4-00c04fd430c8 */
    0x6ba7b810,
    0x9dad,
    0x11d1,
    0x80, 0xb4, 0x00, 0xc0, 0x4f, 0xd4, 0x30, 0xc8
};

/* Name string is a URL */
uuid_t NameSpace_URL = { /* 6ba7b811-9dad-11d1-80b4-00c04fd430c8 */
    0x6ba7b811,
    0x9dad,
    0x11d1,
    0x80, 0xb4, 0x00, 0xc0, 0x4f, 0xd4, 0x30, 0xc8
};

/* Name string is an ISO OID */
uuid_t NameSpace_OID = { /* 6ba7b812-9dad-11d1-80b4-00c04fd430c8 */
    0x6ba7b812,
    0x9dad,
    0x11d1,
    0x80, 0xb4, 0x00, 0xc0, 0x4f, 0xd4, 0x30, 0xc8
};

/* Name string is an X.500 DN (in DER or a text output format) */
uuid_t NameSpace_X500 = { /* 6ba7b814-9dad-11d1-80b4-00c04fd430c8 */
    0x6ba7b814,
    0x9dad,
    0x11d1,
    0x80, 0xb4, 0x00, 0xc0, 0x4f, 0xd4, 0x30, 0xc8
};
```

Authors' Addresses

Paul J. Leach

Microsoft
1 Microsoft Way
Redmond, WA 98052
US

Phone: [+1 425-882-8080](tel:+14258828080)

E-Mail: paulle@microsoft.com

Michael Mealling

Refactored Networks, LLC
1635 Old Hwy 41
Suite 112, Box 138
Kennesaw, GA 30152
US

Phone: [+1-678-581-9656](tel:+16785819656)

E-Mail: michael@refactored-networks.com

URI: <http://www.refactored-networks.com>

Rich Salz

DataPower Technology, Inc.
1 Alewife Center
Cambridge, MA 02142
US

Phone: [+1 617-864-0455](tel:+16178640455)

E-Mail: rsalz@datapower.com

URI: <http://www.datapower.com>

Full Copyright Statement

Copyright © The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an “AS IS” basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <<http://www.ietf.org/ipr>>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org¹.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

¹ <mailto:ietf-ipr@ietf.org>