

N/A
Keiser
Internet-Draft
Ed.
Intended status: Informational
Nomine
Expires: March 14, 2013
2012

T.
A. Deason,
Sine
September 10,

**AFS-3 Rx RPC XDR Primitive Type Definitions
draft-keiser-afs3-xdr-primitive-types-01**

Abstract

AFS-3 embeds a set of XDR primitive type definitions, which are referenced throughout the various AFS-3 protocol specifications. This memo defines the mapping between these AFS-3 primitive types, and the underlying XDR primitives.

Internet Draft Comments

Comments regarding this draft are solicited. Please include the AFS-3 protocol standardization mailing list (afs3-standardization@openafs.org) as a recipient of any comments.

AFS-3 Document State

This document is in state "draft", as per the document state definitions set forth in [[I-D.wilkinson-afs3-standardisation](#)].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 14, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

<u>1.</u>	Introduction	3
<u>1.1.</u>	Purpose	3
<u>1.2.</u>	Abbreviations	3
<u>2.</u>	Conventions	4
<u>3.</u>	Primitive Integer Types	4
<u>3.1.</u>	char	4
<u>3.2.</u>	unsigned char	5
<u>3.3.</u>	short	5
<u>3.4.</u>	unsigned short	5
<u>3.5.</u>	1- and 2-octet integer types	5
<u>4.</u>	afsUUID	5
<u>4.1.</u>	Encoding	6
<u>4.2.</u>	Decoding	7
<u>5.</u>	IANA Considerations	8
<u>6.</u>	AFS Assign Numbers Registrar Considerations	8
<u>7.</u>	Security Considerations	8
<u>8.</u>	References	9
<u>8.1.</u>	Normative References	9
<u>8.2.</u>	Informative References	9

[Appendix A](#). Base Type Definitions
[10](#)
Authors' Addresses
[10](#)

1. Introduction

AFS-3 [[CMU-ITC-88-062](#)] [[CMU-ITC-87-068](#)] is a distributed file system that has its origins in the VICE project [[CMU-ITC-84-020](#)] [[CMU-ITC-85-039](#)] at the Carnegie Mellon University Information Technology Center [[CMU-ITC-83-025](#)], a joint venture between CMU and IBM. VICE later became AFS when CMU moved development to a new commercial venture called Transarc Corporation, which later became IBM Pittsburgh Labs. AFS-3 is a suite of un-standardized network protocols based on a remote procedure call (RPC) suite known as Rx [[AFS3-RX](#)]. While de jure standards for AFS-3 fail to exist, the various AFS-3 implementations have agreed upon certain de facto standards, largely helped by the existence of an open source fork called OpenAFS that has served the role of reference implementation. In addition to using OpenAFS as a reference, IBM wrote and donated developer documentation that contains somewhat outdated specifications for the Rx protocol and all AFS-3 remote procedure calls, as well as a detailed description of the AFS-3 system architecture.

1.1. Purpose

The Rx RPC protocol utilizes XDR [[RFC4506](#)] as its means of encoding RPC call and response payloads. While XDR provides type definitions for each popular bit-length integer, it does so using relatively ambiguous names (e.g., hyper). To improve readability, the AFS-3 RPC-L grammar references custom XDR types that embed the bit length within the type name. This memo standardizes those primitive types, as well as the encoding for the AFS-3 UUID.

1.2. Abbreviations

- AFS - Historically, AFS stood for the Andrew File System; AFS no longer stands for anything
- DCE - The Distributed Computing Environment
- LSB - Least-Significant Bit
- MSB - Most-Significant Bit
- RPC - Remote Procedure Call
- RPC-L - Rx RPC Interface Definition Language (fork of ONC RPC [[RFC5531](#)] .x file format)

- Rx - The Remote Procedure Call mechanism utilized by AFS-3 [[AFS3-RX](#)]
- UUID - Universally Unique Identifier
- XDR - eXternal Data Representation [[RFC4506](#)]

2. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

3. Primitive Integer Types

AFS-3 defines a number of special types which are direct mappings onto existing XDR types. Thus, these types are essentially XDR typedefs:

AFS-3 type name	->	XDR primitive type	[Reference]
-----		-----	-----
char		int	RFC 4506 Section 4.1
unsigned char		unsigned int	RFC 4506 Section 4.2
afs_int8		int	RFC 4506 Section 4.1
afs_uint8		unsigned int	RFC 4506 Section 4.2
short		int	RFC 4506 Section 4.1
unsigned short		unsigned int	RFC 4506 Section 4.2
afs_int16		int	RFC 4506 Section 4.1
afs_uint16		unsigned int	RFC 4506 Section 4.2
afs_int32		int	RFC 4506 Section 4.1
afs_uint32		unsigned int	RFC 4506 Section 4.2
afs_int64		hyper	RFC 4506 Section 4.5
afs_uint64		unsigned hyper	RFC 4506 Section 4.5

AFS-3 common typedefs

Figure 1

3.1. char

This type is considered deprecated; future protocol specifications should reference the "afs_int8" type instead.

3.2. unsigned char

This type is considered deprecated; future protocol specifications should reference the "afs_uint8" type instead.

3.3. short

This type is considered deprecated; future protocol specifications should reference the "afs_int16" type instead.

3.4. unsigned short

This type is considered deprecated; future protocol specifications should reference the "afs_uint16" type instead.

3.5. 1- and 2-octet integer types

Please note that XDR uses a 4-octet alignment, and thus these 1- and 2-octet types will take 4 octets on the wire. Consequently, this is merely a way of defining data structures that have an optimized in-memory footprint, without imbuing any wire-encoding advantage.

4. afsUUID

AFS-3, being closely related to DCE, relies heavily upon a UUID type.

The AFS-3 UUID type is identical to the DCE-variant, version 1 UUID type (see [Appendix A](#) of [[DCE-RPC](#)]). The C data structure definition for such a UUID is as follows:


```
/*
 * Copyright 2000, International Business Machines Corporation
 * and others. All Rights Reserved.
 *
 * This software has been released under the terms of the IBM
 * Public License. For details, see the LICENSE file in the
 * top-level source directory or online at
 * http://www.openafs.org/dl/license10.html
 */

struct afsUUID {
    afs_uint32 time_low;
    afs_uint16 time_mid;
    afs_uint16 time_hi_and_version;
    char clock_seq_hi_and_reserved;
    char clock_seq_low;
    char node[6];
};
```



```
typedef struct afsUUID afsUUID;
```

An afsUUID type is XDR encoded on the wire as follows:



- {1} clock_seq_hi_and_reserved
- {2} clock_seq_low
- {3} sign extended padding: 0, or 0xffffffff depending on value of MSB in field to be sign-extended and padded

afsUUID XDR encoding

4.1. Encoding

XDR encoding an "afsUUID" type SHALL involve the following sequence of operations:

1. Encode "time_low" field as an XDR unsigned integer (afs_uint32)
2. Encode "time_mid" field as an XDR unsigned integer

3. Encode "time_hi_and_version" field as an XDR unsigned integer
4. Encode "clock_seq_hi_and_reserved" field as an XDR signed integer
5. Encode "clock_seq_low" field as an XDR signed integer
6. Encode "node[0]" field as an XDR signed integer
7. Encode "node[1]" field as an XDR signed integer
8. Encode "node[2]" field as an XDR signed integer
9. Encode "node[3]" field as an XDR signed integer
10. Encode "node[4]" field as an XDR signed integer
11. Encode "node[5]" field as an XDR signed integer

Many of the fields which are encoded in this data structure are smaller than four octets. In order to XDR encode these fields, they must first be resized. Since many of these fields are signed, this involves sign extension. This process depends upon the machine architecture. Virtually all machines in existence today utilize 2s-complement integer arithmetic, where this process merely involves padding with zeros if the MSB is zero or ones if the MSB is one.

4.2. Decoding

XDR decoding an "afsUUID" type SHALL involve the following sequence of operations:

1. Decode an XDR unsigned integer into the "time_low" field
2. Decode an XDR unsigned integer. If the integer is greater than 65535, the decoding SHALL fail. Copy the least-significant 16 bits into the "time_mid" field.
3. Decode an XDR unsigned integer. If the integer is greater than 65535, the decoding SHALL fail. Copy the least-significant 16 bits into the "time_hi_and_version" field.
4. Decode an XDR signed integer. If the integer is greater than 32767, or less than -32768, the decoding SHALL fail. Copy the least-significant 16 bits into the "clock_seq_hi_and_reserved" field.

5. Decode an XDR signed integer. If the integer is greater than 32767, or less than -32768, the decoding SHALL fail. Copy the least-significant 16 bits into the "clock_seq_low" field.
6. Decode an XDR signed integer. If the integer is greater than 127, or less than -128, the decoding SHALL fail. Copy the least-significant 8 bits into the "node[0]" field.
7. Decode an XDR signed integer. If the integer is greater than 127, or less than -128, the decoding SHALL fail. Copy the least-significant 8 bits into the "node[1]" field.
8. Decode an XDR signed integer. If the integer is greater than 127, or less than -128, the decoding SHALL fail. Copy the least-significant 8 bits into the "node[2]" field.
9. Decode an XDR signed integer. If the integer is greater than 127, or less than -128, the decoding SHALL fail. Copy the least-significant 8 bits into the "node[3]" field.
10. Decode an XDR signed integer. If the integer is greater than 127, or less than -128, the decoding SHALL fail. Copy the least-significant 8 bits into the "node[4]" field.
11. Decode an XDR signed integer. If the integer is greater than 127, or less than -128, the decoding SHALL fail. Copy the least-significant 8 bits into the "node[5]" field.

5. IANA Considerations

This memo includes no request to IANA.

6. AFS Assign Numbers Registrar Considerations

This memo includes no request to the AFS Assigned Numbers Registrar.

7. Security Considerations

This document merely describes various AFS-3 XDR types. Any protocol specification which uses these primitives types must ensure the security of the resulting XDR data streams, e.g., via prescription of a suitable Rx RPC security class.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4506] Eisler, M., "XDR: External Data Representation Standard", STD 67, [RFC 4506](#), May 2006.

8.2. Informative References

- [AFS3-RX] Zayas, E., "AFS-3 Programmer's Reference: Specification for the Rx Remote Procedure Call Facility", Transarc Corp.
Tech. Rep. FS-00-D164, August 1991.
- [CMU-ITC-83-025] Morris, J., Van Houweling, D., and K. Slack, "The Information Technology Center", CMU ITC Tech. Rep. CMU-ITC-83-025, 1983.
- [CMU-ITC-84-020] West, M., "VICE File System Services", CMU ITC Tech. Rep. CMU-ITC-84-020, August 1984.
- [CMU-ITC-85-039] Satyanarayanan, M., Howard, J., Nichols, D., Sidebotham, R., Spector, A., and M. West, "The ITC Distributed File System: Principles and Design", Proc. 10th ACM Symp. Operating Sys. Princ. Vol. 19, No. 5, December 1985.
- [CMU-ITC-87-068] Howard, J., Kazar, M., Menees, S., Nichols, D., Satyanarayanan, M., Sidebotham, R., and M. West, "Scale and Performance in a Distributed File System", ACM Trans. Comp. Sys. Vol. 6, No. 1, pp. 51-81, February 1988.
- [CMU-ITC-88-062] Howard, J., "An Overview of the Andrew File System", Proc. 1988 USENIX Winter Tech. Conf. pp. 23-26, February 1988.
- [DCE-RPC] The Open Group, "CAE Specification, DCE 1.1: Remote Procedure Call", CAE C706, August 1997.
- [I-D.wilkinson-afs3-standardisation] Wilkinson, S., "Options for AFS Standardisation", [draft-wilkinson-afs3-standardisation-00](#) (work in progress), June 2010.

[RFC5531] Thurlow, R., "RPC: Remote Procedure Call Protocol Specification Version 2", [RFC 5531](#), May 2009.

Appendix A. Base Type Definitions

```
typedef afs_int8  int;  
typedef afs_uint8 unsigned int;  
typedef afs_int16 int;  
typedef afs_uint16 unsigned int;  
typedef afs_int32 int;  
typedef afs_uint32 unsigned int;  
typedef afs_int64 hyper;  
typedef afs_uint64 unsigned hyper;
```

Authors' Addresses

Thomas Keiser
Sine Nomine Associates
43596 Blacksmith Square
Ashburn, VA 20147
USA

Email: tkeiser@gmail.com

Andrew Deason (editor)
Sine Nomine Associates
43596 Blacksmith Square
Ashburn, Virginia 20147-4606
USA

Phone: +1 703 723 6673
Email: adeason@sinenomine.net

