XDR: External Data Representation Standard

Status of This Memo

Copyright Notice

Abstract

   This document describes the External Data Representation Standard
   (XDR) protocol as it is currently deployed and accepted. This
   document obsoletes RFC 1832.

Table of Contents

1. Introduction

   XDR is a standard for the description and encoding of data. It is
   useful for transferring data between different computer
   architectures, and it has been used to communicate data between such
   diverse machines as the SUN WORKSTATION*, VAX*, IBM-PC*, and Cray*.
   XDR fits into the ISO presentation layer and is roughly analogous in
   purpose to X.409, ISO Abstract Syntax Notation. The major difference
   between these two is that XDR uses implicit typing, while X.409 uses
   explicit typing.

   XDR uses a language to describe data formats. The language can be
   used only to describe data; it is not a programming language. This
   language allows one to describe intricate data formats in a concise
   manner. The alternative of using graphical representations (itself
   an informal language) quickly becomes incomprehensible when faced
   with complexity. The XDR language itself is similar to the C
   language [KERN], just as Courier [COUR] is similar to Mesa.
   Protocols such as ONC RPC (Remote Procedure Call) and the NFS*
   (Network File System) use XDR to describe the format of their data.

   The XDR standard makes the following assumption: that bytes (or
   octets) are portable, where a byte is defined as 8 bits of data. A
   given hardware device should encode the bytes onto the various media
   in such a way that other hardware devices may decode the bytes
   without loss of meaning. For example, the Ethernet* standard
   suggests that bytes be encoded in "little-endian" style [COHE], or
   least significant bit first.

2. Changes from RFC 1832

   This document makes no technical changes to RFC 1832 and is published
   for the purposes of noting IANA considerations, augmenting security
   considerations, and distinguishing normative from informative
   references.

3. Basic Block Size

   The representation of all items requires a multiple of four bytes (or
   32 bits) of data. The bytes are numbered 0 through n-1. The bytes
   are read or written to some byte stream such that byte m always
   precedes byte m+1. If the n bytes needed to contain the data are not
   a multiple of four, then the n bytes are followed by enough (0 to 3)
   residual zero bytes, r, to make the total byte count a multiple of 4.

   We include the familiar graphic box notation for illustration and
   comparison. In most illustrations, each box (delimited by a plus
   sign at the 4 corners and vertical bars and dashes) depicts a byte.

Ellipses (...) between boxes show zero or more additional bytes where
required.

```
   +--------+--------+...+--------+--------+...+--------+
   | byte 0 | byte 1 |...|byte n-1|   0    |...|   0    |   BLOCK
   +--------+--------+...+--------+--------+...+--------+
   |<-----------n bytes---------->|<------r bytes------>|
   |<-----------n+r (where (n+r) mod 4 = 0)----------->|
```
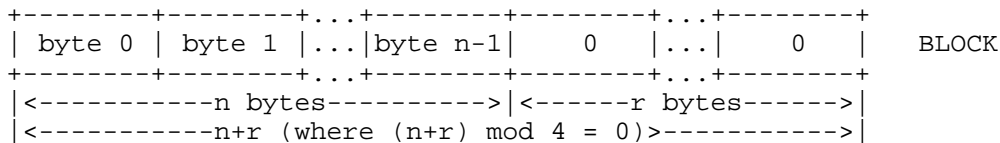
4. XDR Data Types

   Each of the sections that follow describes a data type defined in the
   XDR standard, shows how it is declared in the language, and includes
   a graphic illustration of its encoding.

   For each data type in the language we show a general paradigm
   declaration. Note that angle brackets (< and >) denote variable-
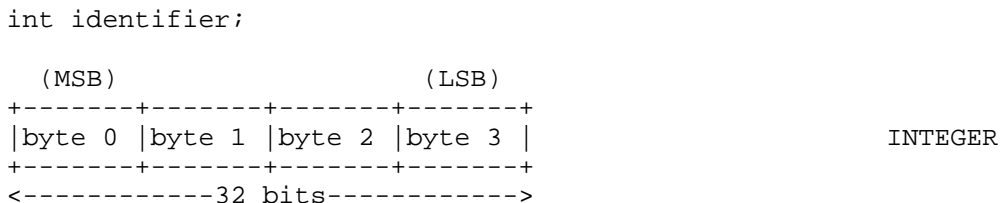   length sequences of data and that square brackets ([and]) denote
   fixed-length sequences of data. "n", "m", and "r" denote integers.
   For the full language specification and more formal definitions of
   terms such as "identifier" and "declaration", refer to Section 6,
   "The XDR Language Specification".

   For some data types, more specific examples are included. A more
   extensive example of a data description is in Section 7, "An Example
   of an XDR Data Description".

4.1. Integer

   An XDR signed integer is a 32-bit datum that encodes an integer in
   the range [-2147483648,2147483647]. The integer is represented in
   two's complement notation. The most and least significant bytes are
   0 and 3, respectively. Integers are declared as follows:

        int identifier;

```
     (MSB)                      (LSB)
   +-------+-------+-------+-------+
   |byte 0 |byte 1 |byte 2 |byte 3 |                INTEGER
   +-------+-------+-------+-------+
   <-----------32 bits------------>
```

4.2. Unsigned Integer

   An XDR unsigned integer is a 32-bit datum that encodes a non-negative
   integer in the range [0,4294967295]. It is represented by an
   unsigned binary number whose most and least significant bytes are 0
   and 3, respectively. An unsigned integer is declared as follows:

```
        unsigned int identifier;

          (MSB)                     (LSB)
          +-------+-------+-------+-------+
          |byte 0 |byte 1 |byte 2 |byte 3 |          UNSIGNED INTEGER
          +-------+-------+-------+-------+
          <-----------32 bits------------>
```
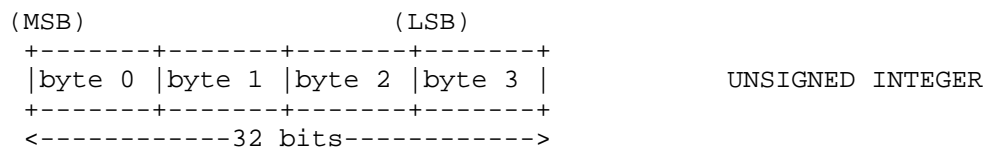
4.3. Enumeration

   Enumerations have the same representation as signed integers.
   Enumerations are handy for describing subsets of the integers.
   Enumerated data is declared as follows:

        enum { name-identifier = constant, ... } identifier;

   For example, the three colors red, yellow, and blue could be
   described by an enumerated type:

        enum { RED = 2, YELLOW = 3, BLUE = 5 } colors;

   It is an error to encode as an enum any integer other than those that
   have been given assignments in the enum declaration.

4.4. Boolean

   Booleans are important enough and occur frequently enough to warrant
   their own explicit type in the standard. Booleans are declared as
   follows:

        bool identifier;

   This is equivalent to:

        enum { FALSE = 0, TRUE = 1 } identifier;

4.5. Hyper Integer and Unsigned Hyper Integer

   The standard also defines 64-bit (8-byte) numbers called hyper
   integers and unsigned hyper integers. Their representations are the
   obvious extensions of integer and unsigned integer defined above.
   They are represented in two's complement notation. The most and
   least significant bytes are 0 and 7, respectively. Their
   declarations:

   hyper identifier; unsigned hyper identifier;

```
   (MSB)                                                      (LSB)
   +-------+-------+-------+-------+-------+-------+-------+-------+
   |byte 0 |byte 1 |byte 2 |byte 3 |byte 4 |byte 5 |byte 6 |byte 7 |
   +-------+-------+-------+-------+-------+-------+-------+-------+
   <---------------------------64 bits---------------------------->
                                                 HYPER INTEGER
                                                 UNSIGNED HYPER INTEGER
```
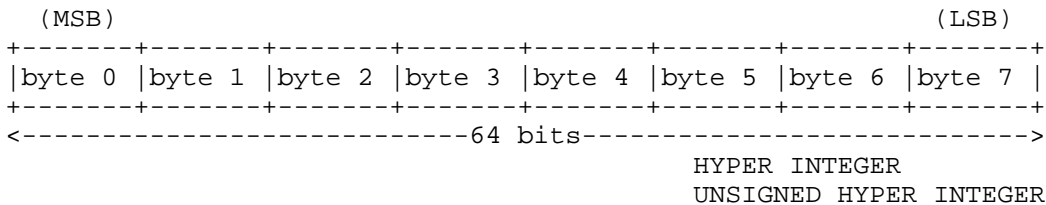
4.6. Floating-Point

   The standard defines the floating-point data type "float" (32 bits or
   4 bytes). The encoding used is the IEEE standard for normalized
   single-precision floating-point numbers [IEEE]. The following three
   fields describe the single-precision floating-point number:

      S: The sign of the number. Values 0 and 1 represent positive and
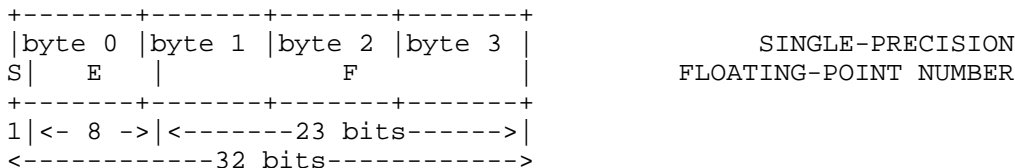   negative, respectively. One bit.

      E: The exponent of the number, base 2. 8 bits are devoted to this
         field. The exponent is biased by 127.

      F: The fractional part of the number's mantissa, base 2. 23 bits
         are devoted to this field.

   Therefore, the floating-point number is described by:

         $(-1)^{S} * 2^{(E-Bias)} * 1.F$

   It is declared as follows:

         float identifier;

```
         +-------+-------+-------+-------+
         |byte 0 |byte 1 |byte 2 |byte 3 |        SINGLE-PRECISION
         S|   E   |            F          |     FLOATING-POINT NUMBER
         +-------+-------+-------+-------+
         1|<- 8 ->|<-------23 bits------>|
         <-----------32 bits------------>
```

   Just as the most and least significant bytes of a number are 0 and 3,
   the most and least significant bits of a single-precision floating-
   point number are 0 and 31. The beginning bit (and most significant
   bit) offsets of S, E, and F are 0, 1, and 9, respectively. Note that
   these numbers refer to the mathematical positions of the bits, and
   NOT to their actual physical locations (which vary from medium to
   medium).

The IEEE specifications should be consulted concerning the encoding
for signed zero, signed infinity (overflow), and denormalized numbers
(underflow) [IEEE]. According to IEEE specifications, the "NaN" (not
a number) is system dependent and should not be interpreted within
XDR as anything other than "NaN".

4.7. Double-Precision Floating-Point

The standard defines the encoding for the double-precision floating-
point data type "double" (64 bits or 8 bytes). The encoding used is
the IEEE standard for normalized double-precision floating-point
numbers [IEEE]. The standard encodes the following three fields,
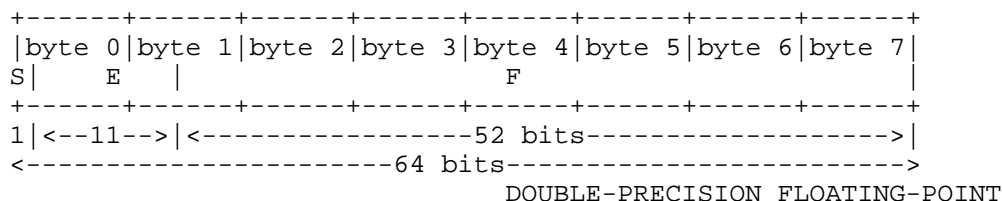which describe the double-precision floating-point number:

    S: The sign of the number. Values 0 and 1 represent positive and
        negative, respectively. One bit.

    E: The exponent of the number, base 2. 11 bits are devoted to
        this field. The exponent is biased by 1023.

    F: The fractional part of the number's mantissa, base 2. 52 bits
        are devoted to this field.

Therefore, the floating-point number is described by:

    $(-1)^{**}S * 2^{**}(E-Bias) * 1.F$

It is declared as follows:

    double identifier;

```
    +------+------+------+------+------+------+------+------+
    |byte 0|byte 1|byte 2|byte 3|byte 4|byte 5|byte 6|byte 7|
    S|    E    |                    F                        |
    +------+------+------+------+------+------+------+------+
    1|<--11-->|<----------------52 bits------------------->|
    <----------------------64 bits------------------------>
                                DOUBLE-PRECISION FLOATING-POINT
```

Just as the most and least significant bytes of a number are 0 and 3,
the most and least significant bits of a double-precision floating-
point number are 0 and 63. The beginning bit (and most significant
bit) offsets of S, E, and F are 0, 1, and 12, respectively. Note
that these numbers refer to the mathematical positions of the bits,
and NOT to their actual physical locations (which vary from medium to
medium).

The IEEE specifications should be consulted concerning the encoding
for signed zero, signed infinity (overflow), and denormalized numbers
(underflow) [IEEE]. According to IEEE specifications, the "NaN" (not
a number) is system dependent and should not be interpreted within
XDR as anything other than "NaN".

4.8. Quadruple-Precision Floating-Point

   The standard defines the encoding for the quadruple-precision
   floating-point data type "quadruple" (128 bits or 16 bytes). The
   encoding used is designed to be a simple analog of the encoding used
   for single- and double-precision floating-point numbers using one
   form of IEEE double extended precision. The standard encodes the
   following three fields, which describe the quadruple-precision
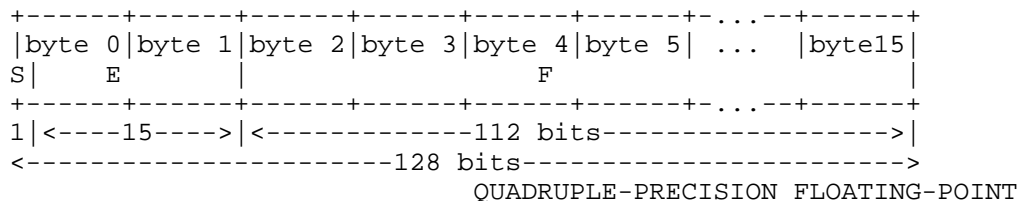   floating-point number:

      S: The sign of the number. Values 0 and 1 represent positive and
         negative, respectively. One bit.

      E: The exponent of the number, base 2. 15 bits are devoted to
         this field. The exponent is biased by 16383.

      F: The fractional part of the number's mantissa, base 2. 112 bits
         are devoted to this field.

   Therefore, the floating-point number is described by:

         $(-1)^{**}S * 2^{**}(E-Bias) * 1.F$

   It is declared as follows:

         quadruple identifier;

         +------+------+------+------+------+------+-...--+------+
         |byte 0|byte 1|byte 2|byte 3|byte 4|byte 5| ... |byte15|
         S| E | F |
         +------+------+------+------+------+------+-...--+------+
         1|<----15---->|<------------112 bits------------------>|
         <----------------------128 bits----------------------->
                                          QUADRUPLE-PRECISION FLOATING-POINT

   Just as the most and least significant bytes of a number are 0 and 3,
   the most and least significant bits of a quadruple-precision
   floating-point number are 0 and 127. The beginning bit (and most
   significant bit) offsets of S, E , and F are 0, 1, and 16,
   respectively. Note that these numbers refer to the mathematical
   positions of the bits, and NOT to their actual physical locations
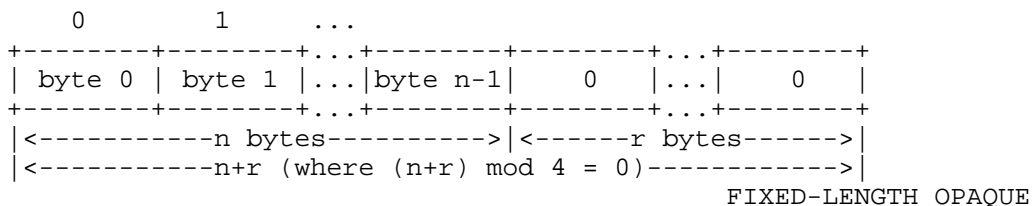   (which vary from medium to medium).

The encoding for signed zero, signed infinity (overflow), and
denormalized numbers are analogs of the corresponding encodings for
single and double-precision floating-point numbers [SPAR], [HPRE].
The "NaN" encoding as it applies to quadruple-precision floating-
point numbers is system dependent and should not be interpreted
within XDR as anything other than "NaN".

4.9. Fixed-Length Opaque Data

   At times, fixed-length uninterpreted data needs to be passed among
   machines. This data is called "opaque" and is declared as follows:

        opaque identifier[n];

   where the constant n is the (static) number of bytes necessary to
   contain the opaque data. If n is not a multiple of four, then the n
   bytes are followed by enough (0 to 3) residual zero bytes, r, to make
   the total byte count of the opaque object a multiple of four.

```
         0          1     ...
      +--------+--------+...+--------+--------+...+--------+
      | byte 0 | byte 1 |...|byte n-1|    0   |...|    0   |
      +--------+--------+...+--------+--------+...+--------+
      |<-----------n bytes---------->|<------r bytes------>|
      |<-----------n+r (where (n+r) mod 4 = 0)------------>|
                                             FIXED-LENGTH OPAQUE
```
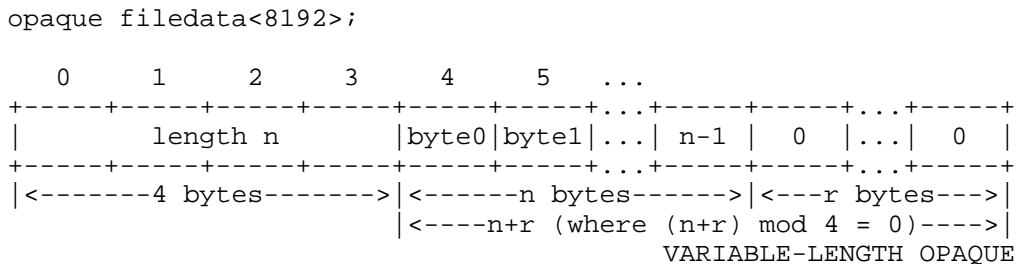
4.10. Variable-Length Opaque Data

   The standard also provides for variable-length (counted) opaque data,
   defined as a sequence of n (numbered 0 through n-1) arbitrary bytes
   to be the number n encoded as an unsigned integer (as described
   below), and followed by the n bytes of the sequence.

   Byte m of the sequence always precedes byte m+1 of the sequence, and
   byte 0 of the sequence always follows the sequence's length (count).
   If n is not a multiple of four, then the n bytes are followed by
   enough (0 to 3) residual zero bytes, r, to make the total byte count
   a multiple of four. Variable-length opaque data is declared in the
   following way:

        opaque identifier<m>;
    or
        opaque identifier<>;

   The constant m denotes an upper bound of the number of bytes that the
   sequence may contain. If m is not specified, as in the second
   declaration, it is assumed to be (2**32) - 1, the maximum length.

The constant m would normally be found in a protocol specification.
For example, a filing protocol may state that the maximum data
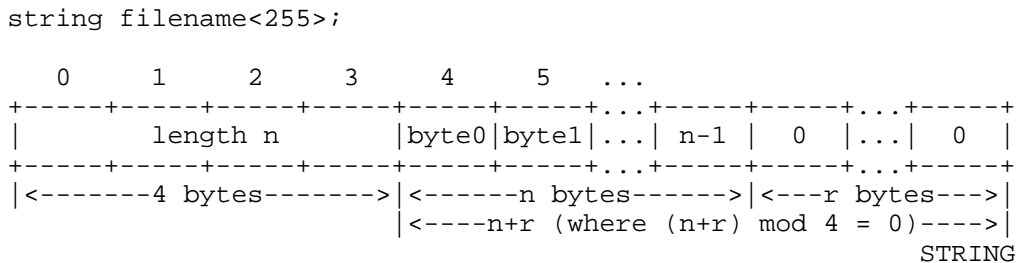transfer size is 8192 bytes, as follows:

```
    opaque filedata<8192>;

      0     1     2     3     4     5   ...
   +-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
   |        length n       |byte0|byte1|...| n-1 |  0  |...|  0  |
   +-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
   |<-------4 bytes------->|<------n bytes------>|<---r bytes--->|
                           |<----n+r (where (n+r) mod 4 = 0)---->|
                                                 VARIABLE-LENGTH OPAQUE
```

It is an error to encode a length greater than the maximum described
in the specification.

4.11. String

The standard defines a string of n (numbered 0 through n-1) ASCII
bytes to be the number n encoded as an unsigned integer (as described
above), and followed by the n bytes of the string. Byte m of the
string always precedes byte m+1 of the string, and byte 0 of the
string always follows the string's length. If n is not a multiple of
four, then the n bytes are followed by enough (0 to 3) residual zero
bytes, r, to make the total byte count a multiple of four. Counted
byte strings are declared as follows:

```
       string object<m>;
   or
       string object<>;
```

The constant m denotes an upper bound of the number of bytes that a
string may contain. If m is not specified, as in the second
declaration, it is assumed to be (2**32) - 1, the maximum length.
The constant m would normally be found in a protocol specification.
For example, a filing protocol may state that a file name can be no
longer than 255 bytes, as follows:

```
    string filename<255>;

      0     1     2     3     4     5   ...
   +-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
   |        length n       |byte0|byte1|...| n-1 |  0  |...|  0  |
   +-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
   |<-------4 bytes------->|<------n bytes------>|<---r bytes--->|
                           |<----n+r (where (n+r) mod 4 = 0)---->|
                                                 STRING
```
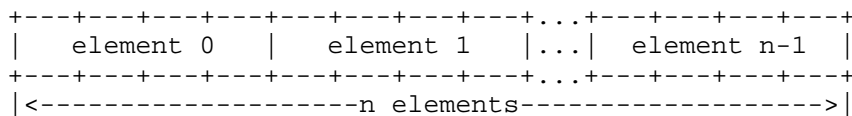
It is an error to encode a length greater than the maximum described
in the specification.

4.12. Fixed-Length Array

Declarations for fixed-length arrays of homogeneous elements are in
the following form:

        type-name identifier[n];

Fixed-length arrays of elements numbered 0 through n-1 are encoded by
individually encoding the elements of the array in their natural
order, 0 through n-1. Each element's size is a multiple of four
bytes. Though all elements are of the same type, the elements may
have different sizes. For example, in a fixed-length array of
strings, all elements are of type "string", yet each element will
vary in its length.

```
        +---+---+---+---+---+---+---+---+...+---+---+---+---+
        |   element 0   |   element 1   |...|  element n-1  |
        +---+---+---+---+---+---+---+---+...+---+---+---+---+
        |<------------------n elements------------------->|
```

                                            FIXED-LENGTH ARRAY
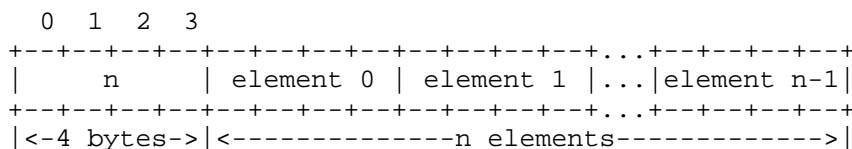
4.13. Variable-Length Array

Counted arrays provide the ability to encode variable-length arrays
of homogeneous elements. The array is encoded as the element count n
(an unsigned integer) followed by the encoding of each of the array's
elements, starting with element 0 and progressing through element
n-1. The declaration for variable-length arrays follows this form:

        type-name identifier<m>;
    or
        type-name identifier<>;

The constant m specifies the maximum acceptable element count of an
array; if m is not specified, as in the second declaration, it is
assumed to be (2**32) - 1.

```
          0  1  2  3
        +--+--+--+--+--+--+--+--+--+--+--+--+...+--+--+--+--+
        |     n     | element 0 | element 1 |...|element n-1|
        +--+--+--+--+--+--+--+--+--+--+--+--+...+--+--+--+--+
        |<-4 bytes->|<-------------n elements------------->|
```
                                                  COUNTED ARRAY

It is an error to encode a value of n that is greater than the
maximum described in the specification.

4.14. Structure

Structures are declared as follows:

```
    struct {
       component-declaration-A;
       component-declaration-B;
       ...
    } identifier;
```

The components of the structure are encoded in the order of their
declaration in the structure. Each component's size is a multiple of
four bytes, though the components may be different sizes.

```
        +-------------+-------------+...
        | component A | component B |...                  STRUCTURE
        +-------------+-------------+...
```

4.15. Discriminated Union

A discriminated union is a type composed of a discriminant followed
by a type selected from a set of prearranged types according to the
value of the discriminant. The type of discriminant is either "int",
"unsigned int", or an enumerated type, such as "bool". The component
types are called "arms" of the union and are preceded by the value of
the discriminant that implies their encoding. Discriminated unions
are declared as follows:

```
    union switch (discriminant-declaration) {
    case discriminant-value-A:
       arm-declaration-A;
    case discriminant-value-B:
       arm-declaration-B;
    ...
    default: default-declaration;
    } identifier;
```

Each "case" keyword is followed by a legal value of the discriminant.
The default arm is optional. If it is not specified, then a valid
encoding of the union cannot take on unspecified discriminant values.
The size of the implied arm is always a multiple of four bytes.

The discriminated union is encoded as its discriminant followed by
the encoding of the implied arm.

```
   0   1   2   3
 +---+---+---+---+---+---+---+---+
 | discriminant  |  implied arm  |        DISCRIMINATED UNION
 +---+---+---+---+---+---+---+---+
 |<---4 bytes--->|
```
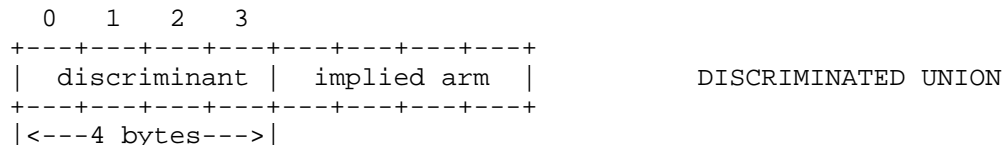
4.16. Void

   An XDR void is a 0-byte quantity. Voids are useful for describing
   operations that take no data as input or no data as output. They are
   also useful in unions, where some arms may contain data and others do
   not. The declaration is simply as follows:

         void;

   Voids are illustrated as follows:

```
    ++
    ||                                              VOID
    ++
  --><-- 0 bytes
```

4.17. Constant

   The data declaration for a constant follows this form:

         const name-identifier = n;

   "const" is used to define a symbolic name for a constant; it does not
   declare any data. The symbolic constant may be used anywhere a
   regular constant may be used. For example, the following defines a
   symbolic constant DOZEN, equal to 12.

         const DOZEN = 12;

4.18. Typedef

   "typedef" does not declare any data either, but serves to define new
   identifiers for declaring data. The syntax is:

         typedef declaration;

   The new type name is actually the variable name in the declaration
   part of the typedef. For example, the following defines a new type
   called "eggbox" using an existing type called "egg":

         typedef egg eggbox[DOZEN];

Variables declared using the new type name have the same type as the
new type name would have in the typedef, if it were considered a
variable. For example, the following two declarations are equivalent
in declaring the variable "fresheggs":

```
eggbox  fresheggs; egg     fresheggs[DOZEN];
```

When a typedef involves a struct, enum, or union definition, there is
another (preferred) syntax that may be used to define the same type.
In general, a typedef of the following form:

```
typedef <<struct, union, or enum definition>> identifier;
```

may be converted to the alternative form by removing the "typedef"
part and placing the identifier after the "struct", "union", or
"enum" keyword, instead of at the end. For example, here are the two
ways to define the type "bool":

```
typedef enum {     /* using typedef */
   FALSE = 0,
   TRUE = 1
} bool;

enum bool {        /* preferred alternative */
   FALSE = 0,
   TRUE = 1
};
```

This syntax is preferred because one does not have to wait until the
end of a declaration to figure out the name of the new type.

4.19. Optional-Data

Optional-data is one kind of union that occurs so frequently that we
give it a special syntax of its own for declaring it. It is declared
as follows:

```
type-name *identifier;
```

This is equivalent to the following union:

```
union switch (bool opted) {
case TRUE:
   type-name element;
case FALSE:
  void;
} identifier;
```

It is also equivalent to the following variable-length array
declaration, since the boolean "opted" can be interpreted as the
length of the array:

```
     type-name identifier<1>;
```

Optional-data is not so interesting in itself, but it is very useful
for describing recursive data-structures such as linked-lists and
trees. For example, the following defines a type "stringlist" that
encodes lists of zero or more arbitrary length strings:

```
     struct stringentry {
        string item<>;
        stringentry *next;
     };

     typedef stringentry *stringlist;
```

It could have been equivalently declared as the following union:

```
     union stringlist switch (bool opted) {
     case TRUE:
        struct {
           string item<>;
           stringlist next;
        } element;
     case FALSE:
        void;
     };
```

or as a variable-length array:

```
     struct stringentry {
        string item<>;
        stringentry next<1>;
     };

     typedef stringentry stringlist<1>;
```

Both of these declarations obscure the intention of the stringlist
type, so the optional-data declaration is preferred over both of
them. The optional-data type also has a close correlation to how
recursive data structures are represented in high-level languages
such as Pascal or C by use of pointers. In fact, the syntax is the
same as that of the C language for pointers.

4.20. Areas for Future Enhancement

   The XDR standard lacks representations for bit fields and bitmaps,
   since the standard is based on bytes. Also missing are packed (or
   binary-coded) decimals.

   The intent of the XDR standard was not to describe every kind of data
   that people have ever sent or will ever want to send from machine to
   machine. Rather, it only describes the most commonly used data-types
   of high-level languages such as Pascal or C so that applications
   written in these languages will be able to communicate easily over
   some medium.

   One could imagine extensions to XDR that would let it describe almost
   any existing protocol, such as TCP. The minimum necessary for this
   is support for different block sizes and byte-orders. The XDR
   discussed here could then be considered the 4-byte big-endian member
   of a larger XDR family.

5. Discussion

   (1) Why use a language for describing data? What's wrong with
       diagrams?

   There are many advantages in using a data-description language such
   as XDR versus using diagrams. Languages are more formal than
   diagrams and lead to less ambiguous descriptions of data. Languages
   are also easier to understand and allow one to think of other issues
   instead of the low-level details of bit encoding. Also, there is a
   close analogy between the types of XDR and a high-level language such
   as C or Pascal. This makes the implementation of XDR encoding and
   decoding modules an easier task. Finally, the language specification
   itself is an ASCII string that can be passed from machine to machine
   to perform on-the-fly data interpretation.

   (2) Why is there only one byte-order for an XDR unit?

   Supporting two byte-orderings requires a higher-level protocol for
   determining in which byte-order the data is encoded. Since XDR is
   not a protocol, this can't be done. The advantage of this, though,
   is that data in XDR format can be written to a magnetic tape, for
   example, and any machine will be able to interpret it, since no
   higher-level protocol is necessary for determining the byte-order.

   (3) Why is the XDR byte-order big-endian instead of little-endian?
       Isn't this unfair to little-endian machines such as the VAX(r),
       which has to convert from one form to the other?

   Yes, it is unfair, but having only one byte-order means you have to
   be unfair to somebody. Many architectures, such as the Motorola
   68000* and IBM 370*, support the big-endian byte-order.

   (4) Why is the XDR unit four bytes wide?

   There is a tradeoff in choosing the XDR unit size. Choosing a small
   size, such as two, makes the encoded data small, but causes alignment
   problems for machines that aren't aligned on these boundaries. A
   large size, such as eight, means the data will be aligned on
   virtually every machine, but causes the encoded data to grow too big.
   We chose four as a compromise. Four is big enough to support most
   architectures efficiently, except for rare machines such as the
   eight-byte-aligned Cray*. Four is also small enough to keep the
   encoded data restricted to a reasonable size.

   (5) Why must variable-length data be padded with zeros?

   It is desirable that the same data encode into the same thing on all
   machines, so that encoded data can be meaningfully compared or
   checksummed. Forcing the padded bytes to be zero ensures this.

   (6) Why is there no explicit data-typing?

   Data-typing has a relatively high cost for what small advantages it
   may have. One cost is the expansion of data due to the inserted type
   fields. Another is the added cost of interpreting these type fields
   and acting accordingly. And most protocols already know what type
   they expect, so data-typing supplies only redundant information.
   However, one can still get the benefits of data-typing using XDR.
   One way is to encode two things: first, a string that is the XDR data
   description of the encoded data, and then the encoded data itself.
   Another way is to assign a value to all the types in XDR, and then
   define a universal type that takes this value as its discriminant and
   for each value, describes the corresponding data type.

6. The XDR Language Specification

6.1. Notational Conventions

   This specification uses an extended Back-Naur Form notation for
   describing the XDR language. Here is a brief description of the
   notation:

   (1) The characters '|', '(', ')', '[', ']', '"', and '*' are special.
   (2) Terminal symbols are strings of any characters surrounded by
   double quotes. (3) Non-terminal symbols are strings of non-special
   characters. (4) Alternative items are separated by a vertical bar

("|"). (5) Optional items are enclosed in brackets. (6) Items are
grouped together by enclosing them in parentheses. (7) A '*'
following an item means 0 or more occurrences of that item.

For example, consider the following pattern:

        "a " "very" (", " "very")* [" cold " "and "] " rainy "
        ("day" | "night")

An infinite number of strings match this pattern. A few of them are:

        "a very rainy day"
        "a very, very rainy day"
        "a very cold and rainy day"
        "a very, very, very cold and rainy night"

6.2. Lexical Notes

(1) Comments begin with '/*' and terminate with '*/'. (2) White
space serves to separate items and is otherwise ignored. (3) An
identifier is a letter followed by an optional sequence of letters,
digits, or underbar ('_'). The case of identifiers is not ignored.
(4) A decimal constant expresses a number in base 10 and is a
sequence of one or more decimal digits, where the first digit is not
a zero, and is optionally preceded by a minus-sign ('-'). (5) A
hexadecimal constant expresses a number in base 16, and must be
preceded by '0x', followed by one or hexadecimal digits ('A', 'B',
'C', 'D', E', 'F', 'a', 'b', 'c', 'd', 'e', 'f', '0', '1', '2', '3',
'4', '5', '6', '7', '8', '9'). (6) An octal constant expresses a
number in base 8, always leads with digit 0, and is a sequence of one
or more octal digits ('0', '1', '2', '3', '4', '5', '6', '7').

6.3. Syntax Information

```
    declaration:
         type-specifier identifier
       | type-specifier identifier "[" value "]"
       | type-specifier identifier "<" [ value ] ">"
       | "opaque" identifier "[" value "]"
       | "opaque" identifier "<" [ value ] ">"
       | "string" identifier "<" [ value ] ">"
       | type-specifier "*" identifier
       | "void"

    value:
         constant
       | identifier
```

```
constant:
   decimal-constant | hexadecimal-constant | octal-constant

type-specifier:
    [ "unsigned" ] "int"
  | [ "unsigned" ] "hyper"
  | "float"
  | "double"
  | "quadruple"
  | "bool"
  | enum-type-spec
  | struct-type-spec
  | union-type-spec
  | identifier

enum-type-spec:
   "enum" enum-body

enum-body:
   "{"
      ( identifier "=" value )
      ( "," identifier "=" value )*
   "}"

struct-type-spec:
   "struct" struct-body

struct-body:
   "{"
      ( declaration ";" )
      ( declaration ";" )*
   "}"

union-type-spec:
   "union" union-body

union-body:
   "switch" "(" declaration ")" "{"
      case-spec
      case-spec *
      [ "default" ":" declaration ";" ]
   "}"

case-spec:
  ( "case" value ":")
  ( "case" value ":") *
  declaration ";"
```

```
      constant-def:
         "const" identifier "=" constant ";"

      type-def:
            "typedef" declaration ";"
          | "enum" identifier enum-body ";"
          | "struct" identifier struct-body ";"
          | "union" identifier union-body ";"

      definition:
            type-def
          | constant-def

      specification:
            definition *
```

6.4. Syntax Notes

   (1) The following are keywords and cannot be used as identifiers:
   "bool", "case", "const", "default", "double", "quadruple", "enum",
   "float", "hyper", "int", "opaque", "string", "struct", "switch",
   "typedef", "union", "unsigned", and "void".

   (2) Only unsigned constants may be used as size specifications for
   arrays. If an identifier is used, it must have been declared
   previously as an unsigned constant in a "const" definition.

   (3) Constant and type identifiers within the scope of a specification
   are in the same name space and must be declared uniquely within this
   scope.

   (4) Similarly, variable names must be unique within the scope of
   struct and union declarations. Nested struct and union declarations
   create new scopes.

   (5) The discriminant of a union must be of a type that evaluates to
   an integer. That is, "int", "unsigned int", "bool", an enumerated
   type, or any typedefed type that evaluates to one of these is legal.
   Also, the case values must be one of the legal values of the
   discriminant. Finally, a case value may not be specified more than
   once within the scope of a union declaration.

7. An Example of an XDR Data Description

   Here is a short XDR data description of a thing called a "file",
   which might be used to transfer files from one machine to another.

```
        const MAXUSERNAME = 32;     /* max length of a user name */
        const MAXFILELEN = 65535;   /* max length of a file      */
        const MAXNAMELEN = 255;     /* max length of a file name */

        /*
         * Types of files:
         */
        enum filekind {
           TEXT = 0,       /* ascii data */
           DATA = 1,       /* raw data   */
           EXEC = 2        /* executable */
        };

        /*
         * File information, per kind of file:
         */
        union filetype switch (filekind kind) {
        case TEXT:
           void;                        /* no extra information */
        case DATA:
           string creator<MAXNAMELEN>;    /* data creator         */
        case EXEC:
           string interpretor<MAXNAMELEN>; /* program interpretor  */
        };

        /*
         * A complete file:
         */
        struct file {
           string filename<MAXNAMELEN>; /* name of file    */
           filetype type;               /* info about file */
           string owner<MAXUSERNAME>;   /* owner of file   */
           opaque data<MAXFILELEN>;     /* file data       */
        };
```

   Suppose now that there is a user named "john" who wants to store his
   lisp program "sillyprog" that contains just the data "(quit)". His
   file would be encoded as follows:

```
     OFFSET  HEX BYTES        ASCII    COMMENTS
     ------  ---------        -----    --------
       0     00 00 00 09      ....     -- length of filename = 9
       4     73 69 6c 6c      sill     -- filename characters
       8     79 70 72 6f      ypro     -- ... and more characters ...
      12     67 00 00 00      g...     -- ... and 3 zero-bytes of fill
      16     00 00 00 02      ....     -- filekind is EXEC = 2
      20     00 00 00 04      ....     -- length of interpretor = 4
      24     6c 69 73 70      lisp     -- interpretor characters
      28     00 00 00 04      ....     -- length of owner = 4
      32     6a 6f 68 6e      john     -- owner characters
      36     00 00 00 06      ....     -- length of file data = 6
      40     28 71 75 69      (qui     -- file data bytes ...
      44     74 29 00 00      t)..     -- ... and 2 zero-bytes of fill
```

8. Security Considerations

   XDR is a data description language, not a protocol, and hence it does
   not inherently give rise to any particular security considerations.
   Protocols that carry XDR-formatted data, such as NFSv4, are
   responsible for providing any necessary security services to secure
   the data they transport.

   Care must be take to properly encode and decode data to avoid
   attacks. Known and avoidable risks include:

   *   Buffer overflow attacks. Where feasible, protocols should be
        defined with explicit limits (via the "<" [value] ">" notation
        instead of "<" ">") on elements with variable-length data types.
        Regardless of the feasibility of an explicit limit on the
        variable length of an element of a given protocol, decoders need
        to ensure the incoming size does not exceed the length of any
        provisioned receiver buffers.

   *   Nul octets embedded in an encoded value of type string. If the
        decoder's native string format uses nul-terminated strings, then
        the apparent size of the decoded object will be less than the
        amount of memory allocated for the string. Some memory
        deallocation interfaces take a size argument. The caller of the
        deallocation interface would likely determine the size of the
        string by counting to the location of the nul octet and adding
        one. This discrepancy can cause memory leakage (because less
        memory is actually returned to the free pool than allocated),
        leading to system failure and a denial of service attack.

   *   Decoding of characters in strings that are legal ASCII
        characters but nonetheless are illegal for the intended
        application. For example, some operating systems treat the '/'

character as a component separator in path names. For a
protocol that encodes a string in the argument to a file
creation operation, the decoder needs to ensure that '/' is not
inside the component name. Otherwise, a file with an illegal
'/' in its name will be created, making it difficult to remove,
and is therefore a denial of service attack.

*   Denial of service caused by recursive decoder or encoder
       subroutines. A recursive decoder or encoder might process data
       that has a structured type with a member of type optional data
       that directly or indirectly refers to the structured type (i.e.,
       a linked list). For example,

```
struct m {
  int x;
  struct m *next;
};
```

An encoder or decoder subroutine might be written to recursively
call itself each time another element of type "struct m" is
found. An attacker could construct a long linked list of
"struct m" elements in the request or response, which then
causes a stack overflow on the decoder or encoder. Decoders and
encoders should be written non-recursively or impose a limit on
list length.

9. IANA Considerations

It is possible, if not likely, that new data types will be added to
XDR in the future. The process for adding new types is via a
standards track RFC and not registration of new types with IANA.
Standards track RFCs that update or replace this document should be
documented as such in the RFC Editor's database of RFCs.

10. Trademarks and Owners

SUN WORKSTATION   Sun Microsystems, Inc.
VAX   Hewlett-Packard Company
IBM-PC   International Business Machines Corporation
Cray   Cray Inc.
NFS   Sun Microsystems, Inc.
Ethernet   Xerox Corporation.
Motorola 68000   Motorola, Inc.
IBM 370   International Business Machines Corporation

11. ANSI/IEEE Standard 754-1985

   The definition of NaNs, signed zero and infinity, and denormalized
   numbers from [IEEE] is reproduced here for convenience. The
   definitions for quadruple-precision floating point numbers are
   analogs of those for single and double-precision floating point
   numbers and are defined in [IEEE].

   In the following, 'S' stands for the sign bit, 'E' for the exponent,
   and 'F' for the fractional part. The symbol 'u' stands for an
   undefined bit (0 or 1).

   For single-precision floating point numbers:

Type	S (1 bit)	E (8 bits)	F (23 bits)
signalling NaN	u	255 (max)	.0uuuuu---u (with at least one 1 bit)
quiet NaN	u	255 (max)	.1uuuuu---u
negative infinity	1	255 (max)	.000000---0
positive infinity	0	255 (max)	.000000---0
negative zero	1	0	.000000---0
positive zero	0	0	.000000---0

   For double-precision floating point numbers:

Type	S (1 bit)	E (11 bits)	F (52 bits)
signalling NaN	u	2047 (max)	.0uuuuu---u (with at least one 1 bit)
quiet NaN	u	2047 (max)	.1uuuuu---u
negative infinity	1	2047 (max)	.000000---0
positive infinity	0	2047 (max)	.000000---0
negative zero	1	0	.000000---0
positive zero	0	0	.000000---0

For quadruple-precision floating point numbers:

Type	S (1 bit)	E (15 bits)	F (112 bits)
signalling NaN	u	32767 (max)	.0uuuuu---u (with at least one 1 bit)
quiet NaN	u	32767 (max)	.1uuuuu---u
negative infinity	1	32767 (max)	.000000---0
positive infinity	0	32767 (max)	.000000---0
negative zero	1	0	.000000---0
positive zero	0	0	.000000---0

Subnormal numbers are represented as follows:

Precision	Exponent	Value
Single	0	$(-1)^{**}S * 2^{**}(-126) * 0.F$
Double	0	$(-1)^{**}S * 2^{**}(-1022) * 0.F$
Quadruple	0	$(-1)^{**}S * 2^{**}(-16382) * 0.F$

12. Normative References

[IEEE] "IEEE Standard for Binary Floating-Point Arithmetic",
        ANSI/IEEE Standard 754-1985, Institute of Electrical and
        Electronics Engineers, August 1985.

13. Informative References

[KERN] Brian W. Kernighan & Dennis M. Ritchie, "The C Programming
        Language", Bell Laboratories, Murray Hill, New Jersey, 1978.

[COHE] Danny Cohen, "On Holy Wars and a Plea for Peace", IEEE
        Computer, October 1981.

[COUR] "Courier: The Remote Procedure Call Protocol", XEROX
        Corporation, XSIS 038112, December 1981.

[SPAR] "The SPARC Architecture Manual: Version 8", Prentice Hall,
        ISBN 0-13-825001-4.

[HPRE] "HP Precision Architecture Handbook", June 1987, 5954-9906.

14. Acknowledgements

   Bob Lyon was Sun's visible force behind ONC RPC in the 1980s. Sun
   Microsystems, Inc., is listed as the author of RFC 1014. Raj
   Srinivasan and the rest of the old ONC RPC working group edited RFC
   1014 into RFC 1832, from which this document is derived. Mike Eisler
   and Bill Janssen submitted the implementation reports for this
   standard. Kevin Coffman, Benny Halevy, and Jon Peterson reviewed
   this document and gave feedback. Peter Astrand and Bryan Olson
   pointed out several errors in RFC 1832 which are corrected in this
   document.

Editor's Address

   Mike Eisler
   5765 Chase Point Circle
   Colorado Springs, CO 80919
   USA

   Phone: 719-599-9026
   EMail: email2mre-rfc4506@yahoo.com

   Please address comments to: nfsv4@ietf.org

Full Copyright Statement

Intellectual Property

Acknowledgement