

---

---

**Information technology — Programming languages, their environments and system software interfaces — Extensions for the programming language C++ to support decimal floating-point arithmetic**

*Technologies de l'information — Langages de programmation, leurs environnements et interfaces de logiciel système — Extensions pour le langage de programmation C++ pour supporter l'arithmétique flottante décimale*

This is a preview of "ISO/IEC TR 24733:201...". Click here to purchase the full version from the ANSI store.



**COPYRIGHT PROTECTED DOCUMENT**

© ISO/IEC 2011

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office  
Case postale 56 • CH-1211 Geneva 20  
Tel. + 41 22 749 01 11  
Fax + 41 22 749 09 47  
E-mail [copyright@iso.org](mailto:copyright@iso.org)  
Web [www.iso.org](http://www.iso.org)

Published in Switzerland

<b>Contents</b>	<b>Page</b>
Foreword .....	vi
Introduction .....	vii
0.1 General .....	vii
0.2 Arithmetic model .....	vii
0.3 The Formats .....	viii
1 Scope .....	1
2 Normative references .....	1
3 Conventions .....	2
3.1 General .....	2
3.2 Relation to C++ Standard Library Introduction .....	2
3.3 Relation to "Technical Report on C++ Library Extensions" .....	3
3.4 Categories of extensions .....	3
3.5 Namespaces and headers .....	4
4 Decimal floating-point types .....	5
4.1 Characteristics of decimal floating-point types .....	5
4.2 Decimal Types .....	6
4.2.1 Header <decimal> synopsis .....	6
4.2.2 Class decimal32 .....	10
4.2.2.1 Class summary .....	10
4.2.2.2 Construct/copy/destroy .....	11
4.2.2.3 Conversion from floating-point type .....	11
4.2.2.4 Conversion from integral type .....	11
4.2.2.5 Conversion to integral type .....	12
4.2.2.6 Increment and decrement operators .....	12
4.2.2.7 Compound assignment .....	12
4.2.3 Class decimal64 .....	14
4.2.3.1 Class summary .....	14
4.2.3.2 Construct/copy/destroy .....	14
4.2.3.3 Conversion from floating-point type .....	14
4.2.3.4 Conversion from integral type .....	15
4.2.3.5 Conversion to integral type .....	15
4.2.3.6 Increment and decrement operators .....	15
4.2.3.7 Compound assignment .....	16
4.2.4 Class decimal128 .....	17
4.2.4.1 Class summary .....	17
4.2.4.2 Construct/copy/destroy .....	17
4.2.4.3 Conversion from floating-point type .....	17
4.2.4.4 Conversion from integral type .....	18
4.2.4.5 Conversion to integral type .....	18
4.2.4.6 Increment and decrement operators .....	18
4.2.4.7 Compound assignment .....	19
4.2.5 Initialization from coefficient and exponent .....	20

**ISO/IEC TR 24733:2011(E)**

4.2.6 Conversion to generic floating-point type .....	20
4.2.7 Unary arithmetic operators .....	21
4.2.8 Binary arithmetic operators.....	21
4.2.9 Comparison operators .....	22
4.2.10 Formatted input.....	25
4.2.11 Formatted output.....	26
4.3 Additions to header <limits> .....	27
4.4 Headers <cfloat> and <float.h> .....	30
4.4.1 General.....	30
4.4.2 Additions to header <float.h> synopsis.....	30
4.4.3 Additions to header <float.h> synopsis.....	31
4.4.4 Maximum finite value.....	31
4.4.5 Epsilon .....	31
4.4.6 Minimum positive normal value.....	32
4.4.7 Minimum positive subnormal value .....	32
4.4.8 Evaluation format.....	32
4.5 Additions to <cfenv> and <fenv.h> .....	33
4.5.1 General.....	33
4.5.2 Additions to <cfenv> synopsis .....	33
4.5.3 Rounding modes .....	34
4.5.4 The fe_dec_getround function.....	34
4.5.5 The fe_dec_setround function.....	35
4.5.6 Changes to <fenv.h> .....	35
4.6 Additions to <cmath> and <math.h> .....	35
4.6.1 General.....	35
4.6.2 Additions to header <cmath> synopsis .....	36
4.6.3 <cmath> macros .....	41
4.6.4 Evaluation formats.....	41
4.6.5 samequantum functions .....	42
4.6.6 quantexp functions .....	42
4.6.7 quantize functions .....	43
4.6.8 Elementary functions .....	43
4.6.9 abs function overloads .....	45
4.6.10 Changes to <math.h> .....	45
4.6.10.1 General.....	45
4.6.10.2 Additions to header <math.h> synopsis.....	45
4.7 Additions to <cstdio> and <stdio.h> .....	45
4.8 Additions to <cstdlib> and <stdlib.h> .....	45
4.8.1 Additions to header <cstdlib> synopsis.....	45
4.8.2 strtod functions.....	45
4.8.3 Changes to <stdlib.h> .....	45
4.9 Additions to <wchar> and <wchar.h> .....	46
4.9.1 Additions to <wchar> synopsis.....	46
4.9.2 wcstod functions .....	46
4.9.3 Changes to <wchar.h> .....	46
4.10 Facets .....	46
4.10.1 General.....	46

4.10.2	Additions to header <locale> synopsis .....	47
4.10.3	Class template extended_num_get .....	47
4.10.3.1	extended_num_get members .....	49
4.10.3.2	extended_num_get virtual functions .....	50
4.10.4	Class template extended_num_put .....	51
4.10.4.1	extended_num_put members .....	52
4.10.4.2	extended_num_put virtual functions .....	53
4.11	Type traits .....	53
4.11.1	Addition to header <type_traits> synopsis .....	53
4.11.2	is_decimal_floating_point_type_trait .....	54
4.12	Hash functions .....	54
4.12.1	Additions to header <functional> synopsis .....	54
4.12.2	Hash function specializations .....	54
5	Notes on C compatibility .....	55
5.1	General .....	55
5.2	Literals .....	55
5.3	Conversions .....	55

## ISO/IEC TR 24733:2011(E)

### Foreword

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

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

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

In exceptional circumstances, when the joint technical committee has collected data of a different kind from that which is normally published as an International Standard ("state of the art", for example), it may decide to publish a Technical Report. A Technical Report is entirely informative in nature and shall be subject to review every five years in the same manner as an International Standard.

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

ISO/IEC TR 24733 was prepared jointly by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces*.

## Introduction

### 0.1 General

Most of today's general purpose computing architectures provide binary floating-point arithmetic in hardware. Binary float-point is an efficient representation that minimizes memory use, and is simpler to implement than floating-point arithmetic using other bases. It has therefore become the norm for scientific computations, with almost all implementations following the IEEE-754 standard for binary floating-point arithmetic.

However, human computation and communication of numeric values almost always uses decimal arithmetic, and decimal notations. Laboratory notes, scientific papers, legal documents, business reports and financial statements all record numeric values in decimal form. When numeric data are given to a program or are displayed to a user, binary to-and-from decimal conversion is required. There are inherent rounding errors involved in such conversions; decimal fractions cannot, in general, be represented exactly by floating-point values. These errors often cause usability and efficiency problems, depending on the application.

These problems are minor when the application domain accepts, or requires results to have, associated error estimates (as is the case with scientific applications). However, in business and financial applications, computations are required either to be exact (with no rounding errors) unless explicitly rounded, or to be supported by detailed analyses that are auditable to be correct. Such applications therefore have to take special care in handling any rounding errors introduced by the computations.

The most efficient way to avoid conversion error is to use decimal arithmetic. Recognizing this, the IEEE 754-2008 Standard for Floating-Point Arithmetic specifies decimal floating-point encodings and arithmetic. This technical report specifies extensions to the International Standard for the C++ programming language to permit the use of decimal arithmetic in a manner consistent with the IEEE 754-2008 standard.

### 0.2 Arithmetic model

This Technical Report is based on a model of decimal arithmetic which is a formalization of the decimal system of numeration (algorism) as further defined and constrained by the relevant standards, IEEE 854, ANSI X3-274, and IEEE 754-2008.

There are three components to the model:

- *data*- numbers and NaNs, which can be manipulated by, or be the results of, the core operations defined in the model
- *operations*- (such as addition, multiplication, etc.) which can be carried out on data
- *context* - the status of operations (namely, exceptions flags), and controls to govern the results of operations (for example, rounding modes).

The model defines these components in the abstract. It defines neither the way in which operations are expressed (which might vary depending on the computer language or other interface being used), nor the concrete representation (specific layout in storage, or in a processor's register, for example) of data or context.

## ISO/IEC TR 24733:2011(E)

From the perspective of the C++ language, *data* are represented by data types, *operations* are defined within expressions, and *context* is the floating environment specified in `<cfenv>` and `<fenv.h>`. This Technical Report specifies how the C++ language implements these components.

### **0.3 The formats**

In the C++ International Standard, the representation of a floating-point number is specified in an abstract form where the constituent components of the representation are defined (sign, exponent, significand) but the internals of these components are not. In particular, the exponent range, significand size and the base (or radix), are implementation defined. This allows flexibility for an implementation to take advantage of its underlying hardware architecture. Furthermore, certain behaviors of operations are also implementation defined, for example in the area of handling of special numbers and in exceptions.

This approach was inherited from the C programming language. At the time that C was first standardized, there were already various hardware implementations of floating-point arithmetic in common use. Specifying the exact details of a representation would make most of the existing C implementations at the time not conforming.

This Technical Report specifies decimal floating-point arithmetic according to the IEEE 754-2008 standard, with the constituent components of the representation defined. This is more stringent than the approach taken for the floating-point types in the C++ standard. Since it is expected that all decimal floating-point hardware implementations will conform to the IEEE 754-2008 standard, binding to this standard directly benefits both implementers and programmers.