

IBM XL C/C++ for Multicore Acceleration for Linux,
V9.0



Language Reference

IBM XL C/C++ for Multicore Acceleration for Linux,
V9.0



Language Reference

Note!

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 425.

First Edition

This edition applies to IBM XL C/C++ for Multicore Acceleration for Linux on System p, V9.0 and IBM XL C/C++ for Multicore Acceleration for Linux on x86 Systems, V9.0. (Programs 5724-T42 & 5724-T43)

© Copyright International Business Machines Corporation 1998, 2007. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this document vii

Who should read this document	vii
How to use this document	vii
Language levels and language extensions	vii
How this document is organized	ix
Conventions used in this document	ix
Related information	xii
IBM XL C/C++ publications.	xii
Standards and specifications documents	xiii
Other IBM publications	xiv
Other publications	xiv
How to send your comments	xiv

Chapter 1. Scope and linkage 1

Scope.	2
Block/local scope.	2
Function scope	3
Function prototype scope	3
File/global scope	3
Examples of scope in C.	4
Class scope (C++ only)	5
Namespaces of identifiers	6
Name hiding (C++ only)	7
Program linkage	7
Internal linkage	8
External linkage	8
No linkage	9
Language linkage (C++ only)	10

Chapter 2. Lexical Elements 13

Tokens	13
Keywords	13
Identifiers	15
Literals	19
Punctuators and operators	30
Source program character set	32
Multibyte characters	32
Escape sequences	33
The Unicode standard	34
Digraph characters	36
Trigraph sequences	37
Comments.	37

Chapter 3. Data objects and declarations 41

Overview of data objects and declarations	41
Overview of data objects	41
Overview of data declarations and definitions.	43
Storage class specifiers.	45
The auto storage class specifier	46
The static storage class specifier	46
The extern storage class specifier	48
The mutable storage class specifier (C++ only).	49
The register storage class specifier	49
The <code>__thread</code> storage class specifier (PPU only)	52

Type specifiers	52
Integral types.	53
Boolean types	54
Floating-point types	54
Character types	55
The void type	56
Compatibility of arithmetic types (C only)	56
Vector types	57
User-defined types	59
Structures and unions	59
Enumerations.	67
Compatibility of structures, unions, and enumerations (C only).	70
typedef definitions	71
Type qualifiers	72
The <code>__align</code> type qualifier	74
The const type qualifier	76
The restrict type qualifier.	77
The volatile type qualifier	78
Type attributes	78
The aligned type attribute	79
The packed type attribute	80
The transparent_union type attribute (C only)	80

Chapter 4. Declarators 83

Overview of declarators	83
Examples of declarators	84
Type names	85
Pointers	86
Pointer arithmetic	87
Type-based aliasing.	88
Compatibility of pointers (C only).	89
Arrays	90
Variable length arrays	91
Compatibility of arrays	92
References (C++ only).	93
Initializers	93
Initialization and storage classes	94
Designated initializers for aggregate types (C only)	95
Initialization of vectors	97
Initialization of structures and unions	98
Initialization of enumerations	100
Initialization of pointers	101
Initialization of arrays	101
Initialization of references (C++ only)	104
Variable attributes	105
The aligned variable attribute	107
The common and nocommon variable attributes	108
The packed variable attribute	108
The <code>init_priority</code> variable attribute (C++ only)	109
The mode variable attribute	109
The section variable attribute	109
The weak variable attribute.	110

Chapter 5. Type conversions 113

Arithmetic conversions and promotions	113
Integral conversions	114
Boolean conversions	114
Floating-point conversions	114
Integral and floating-point promotions	115
Lvalue-to-rvalue conversions	116
Pointer conversions	117
Conversion to void*	118
Reference conversions (C++ only)	118
Qualification conversions (C++ only)	119
Function argument conversions	119

Chapter 6. Expressions and operators 121

Lvalues and rvalues	121
Primary expressions	123
Names	123
Literals	124
Integer constant expressions	124
Identifier expressions (C++ only)	124
Parenthesized expressions ()	125
Scope resolution operator :: (C++ only)	126
Function call expressions	127
Member expressions	128
Dot operator	128
Arrow operator ->	128
Unary expressions	128
Increment operator ++	129
Decrement operator --	130
Unary plus operator +	130
Unary minus operator -	130
Logical negation operator !	131
Bitwise negation operator ~	131
Address operator &	132
Indirection operator *	133
The typeid operator (C++ only) (PPU only)	133
The __alignof__ operator	135
The sizeof operator	135
The typeof operator	137
The __real__ and __imag__ operators	138
The vec_step operator	139
Binary expressions	139
Assignment operators	140
Multiplication operator *	142
Division operator /	142
Remainder operator %	143
Addition operator +	143
Subtraction operator -	143
Bitwise left and right shift operators << >>	144
Relational operators < > <= >=	144
Equality and inequality operators == !=	146
Bitwise AND operator &	147
Bitwise exclusive OR operator ^	147
Bitwise inclusive OR operator 	148
Logical AND operator &&	148
Logical OR operator 	149
Array subscripting operator []	150
Comma operator ,	151
Pointer to member operators .* ->* (C++ only)	153
Conditional expressions	153
Types in conditional C expressions	154

Types in conditional C++ expressions	154
Examples of conditional expressions	155
Cast expressions	156
Cast operator ().	156
The static_cast operator (C++ only)	159
The reinterpret_cast operator (C++ only)	160
The const_cast operator (C++ only)	161
The dynamic_cast operator (C++ only) (PPU only)	162
Compound literal expressions	164
new expressions (C++ only)	165
Placement syntax	166
Initialization of objects created with the new operator	168
Handling new allocation failure	168
delete expressions (C++ only)	169
throw expressions (C++ only) (PPU only)	170
Label value expressions	170
Operator precedence and associativity	170

Chapter 7. Statements 175

Labeled statements	175
Locally declared labels	176
Labels as values	176
Expression statements	177
Resolution of ambiguous statements	177
Block statements	178
Example of blocks	179
Statement expressions	179
Selection statements	180
The if statement	180
The switch statement	181
Iteration statements	185
The while statement	185
The do statement	186
The for statement	187
Jump statements	189
The break statement	189
The continue statement	190
The return statement	191
The goto statement	192
Null statement	194
Inline assembly statements	194
Examples of inline assembly statements	198
Restrictions on inline assembly statements	198

Chapter 8. Functions 201

Function declarations and definitions	201
Function declarations	202
Function definitions	202
Examples of function declarations	203
Examples of function definitions	204
Compatible functions	204
Multiple function declarations	205
Function storage class specifiers	206
The static storage class specifier	206
The extern storage class specifier	206
Function specifiers	208
The inline function specifier	208
Function return type specifiers	213

Function return values	213
Function declarators	214
Parameter declarations	215
Function attributes	218
The alias function attribute	220
The always_inline function attribute	220
The const function attribute	220
The constructor and destructor function attributes	221
The format function attribute	221
The format_arg function attribute	222
The noinline function attribute	222
The noreturn function attribute	223
The pure function attribute	223
The section function attribute	223
The weak function attribute	224
The main() function	224
Function calls	225
Pass by value	226
Pass by reference	227
Allocation and deallocation functions (C++ only)	228
Default arguments in C++ functions	229
Restrictions on default arguments	230
Evaluation of default arguments	231
Pointers to functions	232
Nested functions	232

Chapter 9. Namespaces (C++ only) 235

Defining namespaces	235
Declaring namespaces	235
Creating a namespace alias	235
Creating an alias for a nested namespace	236
Extending namespaces	236
Namespaces and overloading	237
Unnamed namespaces	237
Namespace member definitions	239
Namespaces and friends	239
The using directive	240
The using declaration and namespaces	240
Explicit access	241

Chapter 10. Overloading (C++ only) 243

Overloading functions	243
Restrictions on overloaded functions	244
Overloading operators	245
Overloading unary operators	248
Overloading increment and decrement operators	249
Overloading binary operators	250
Overloading assignments	251
Overloading function calls	252
Overloading subscripting	253
Overloading class member access	254
Overload resolution	255
Implicit conversion sequences	256
Resolving addresses of overloaded functions	257

Chapter 11. Classes (C++ only) 259

Declaring class types	260
Using class objects	260
Classes and structures	262

Scope of class names	263
Incomplete class declarations	264
Nested classes	264
Local classes	266
Local type names	267

Chapter 12. Class members and friends (C++ only) 269

Class member lists	269
Data members	270
Member functions	271
Inline member functions	271
Constant and volatile member functions	272
Virtual member functions	272
Special member functions	273
Member scope	273
Pointers to members	274
The this pointer	276
Static members	279
Using the class access operators with static members	279
Static data members	280
Static member functions	282
Member access	283
Friends	286
Friend scope	288
Friend access	290

Chapter 13. Inheritance (C++ only) 291

Derivation	293
Inherited member access	296
Protected members	296
Access control of base class members	297
The using declaration and class members	298
Overloading member functions from base and derived classes	299
Changing the access of a class member	301
Multiple inheritance	302
Virtual base classes	303
Multiple access	304
Ambiguous base classes	305
Virtual functions	309
Ambiguous virtual function calls	312
Virtual function access	314
Abstract classes	314

Chapter 14. Special member functions (C++ only) 317

Overview of constructors and destructors	317
Constructors	319
Default constructors	319
Explicit initialization with constructors	320
Initialization of base classes and members	322
Construction order of derived class objects	325
Destructors	326
Pseudo-destructors	328
User-defined conversions	329
Conversion constructors	330
The explicit specifier	332
Conversion functions	332

Copy constructors	333
Copy assignment operators.	335

Chapter 15. Templates (C++ only). . . . 337

Template parameters	338
Type template parameters	338
Non-type template parameters	338
Template template parameters.	339
Default arguments for template parameters	339
Template arguments	340
Template type arguments	340
Template non-type arguments	341
Template template arguments	343
Class templates.	344
Class template declarations and definitions	345
Static data members and templates	346
Member functions of class templates	346
Friends and templates	347
Function templates	348
Template argument deduction.	349
Overloading function templates	354
Partial ordering of function templates	355
Template instantiation	356
Implicit instantiation	356
Explicit instantiation	358
Template specialization	359
Explicit specialization	359
Partial specialization	364
Name binding and dependent names	366
The typename keyword	367
The template keyword as qualifier	368

Chapter 16. Exception handling (C++ only) (PPU only) 371

try blocks	371
Nested try blocks	372
catch blocks	373
Function try block handlers	374
Arguments of catch blocks	377
Matching between exceptions thrown and caught.	377
Order of catching	377
throw expressions	379
Rethrowing an exception	379
Stack unwinding	380
Exception specifications	382
Special exception handling functions	384
The unexpected() function	385
The terminate() function.	386
The set_unexpected() and set_terminate() functions	387

Example using the exception handling functions	387
--	-----

Chapter 17. Preprocessor directives 391

Macro definition directives	391
The #define directive	392
The #undef directive	396
The # operator	397
The ## operator	398
Standard predefined macro names	398
File inclusion directives	400
The #include directive	400
The #include_next directive.	401
Conditional compilation directives	402
The #if and #elif directives	403
The #ifdef directive	404
The #ifndef directive	405
The #else directive.	405
The #endif directive	405
Message generation directives	406
The #error directive	407
The #warning directive	407
The #line directive.	407
Assertion directives	409
Predefined assertions.	409
The null directive (#)	410
Pragma directives	410
The _Pragma preprocessing operator	410
Standard pragmas	411

Appendix A. The IBM XL C language extensions 413

C99 features as extensions to C89.	413
Extensions for Unicode support	415
Extensions for GNU C compatibility.	415
Extensions for vector processing support	417

Appendix B. The IBM XL C++ language extensions 419

General IBM extensions	419
Extensions for C99 compatibility	419
Extensions for Unicode support	420
Extensions for GNU C compatibility.	420
Extensions for GNU C++ compatibility.	422
Extensions for vector processing support	423

Notices 425

Trademarks and service marks	427
Industry standards	427

Index 429

About this document

This document describes the syntax, semantics, and IBM® XL C/C++ for Multicore Acceleration for Linux® implementation of the C and C++ programming languages. Although the XL C and XL C++ compilers conform to the specifications maintained by the ISO standards for the C and C++ programming languages, the compilers also incorporate many extensions to the core languages. These extensions have been implemented with the aims of enhancing usability in specific operating environments, supporting compatibility with other compilers, and supporting new hardware capabilities. For example, on UNIX® platforms, many language constructs have been added for compatibility with the GNU Compiler Collection (GCC), to maximize portability between the two development environments.

Who should read this document

This document is a reference for users who already have experience programming applications in C or C++. Users new to C or C++ can still use this document to find information on the language and features unique to XL C/C++; however, this reference does not aim to teach programming concepts nor to promote specific programming practices.

How to use this document

Unless indicated otherwise, all of the text in this reference pertains to both C and C++ languages. Where there are differences between languages, these are indicated through qualifying text and other graphical elements (see below for the conventions used).

XL C/C++ provides separate invocation commands depending on whether you are targeting the PowerPC® Processor Unit (PPU) or the Synergistic Processor Unit (SPU). PPU-targeted compilation uses commands prefixed with **ppu**; SPU-targeted compilation uses commands prefixed with **spu**. For convenience, this document refers to the compiler invocation commands with the **ppu** and **spu** prefixes. Additionally, all of the text in this document pertains to compilation targeting both the PPU and SPU, unless indicated otherwise.

While this document covers both standard and implementation-specific features, it does not include the following topics:

- Standard C and C++ library functions and headers. For information on the standard C and C++ libraries, refer to your operating system documentation.
- Compiler pragmas, predefined macros, and built-in functions. These are described in the *XL C/C++ Compiler Reference*.

Language levels and language extensions

The C and C++ languages described in this reference are based on the standards listed in “Standards and specifications documents” on page xiii.

We refer to the following language specifications as “base language levels” in order to introduce the notion of an extension to a base. In this context the base language levels refer to the following specifications:

- Standard C++

- C++98
- C99
- C89

This document uses the term *K&R C* to refer to the C language plus the generally accepted extensions produced by Brian Kernighan and Dennis Ritchie that were in use prior to the ISO standardization of C.

In addition to the features supported by the base levels, XL C/C++ contains language extensions that enhance usability and facilitate porting programs to different platforms, including:

- Extensions to C++ to support C99 standard features
- Extensions related to GNU C and C++
- Extensions supporting the AltiVec Programming Interface
- Extensions supporting Cell Broadband Engine processors

You can control the language level to be used for compilation through several mechanisms, including:

- various invocation commands
- the `-qlanglvl` option

With a few exceptions, almost all of the language extensions are supported when you compile using the basic invocation commands `ppuxlc` or `spuxlc` (for C) and `ppuxlc++` or `spuxlc++` or `ppuxlc` or `spuxlc` (for C++).

The default language level for the `ppuxlc` or `spuxlc` invocation command is `extc99`, which includes all of the features introduced by the C99 standard, and most of the IBM extensions described in this document. For a complete listing of the C extensions and various methods for enabling them, see Appendix A, “The IBM XL C language extensions,” on page 413.

The language level for the `ppuxlc` or `spuxlc` or `ppuxlc++` or `spuxlc++` invocation command is `extended`, which includes most of the IBM extensions described in this document, as well as many C99 features. For a complete listing of the C++ extensions and methods for enabling them, see Appendix B, “The IBM XL C++ language extensions,” on page 419.

For information on the various methods for controlling the language level for compilation, see “Invoking the compiler” and `-qlanglvl` in the *XL C/C++ Compiler Reference*.

Extensions to C++ to support C99 standard features

The Standard C++ language specification does not include many of the features specified in the C99 language standard. To promote compatibility and portability between C99 and C++, the XL C++ compiler enables many of the C99 features that are supported by the XL C compiler. Since these features extend Standard C++, they are considered extensions to the base language. In this reference, unless the text is marked to indicate that a feature is supported in C or C99 only, C99 features also apply to C++. A complete list of C99 features supported in XL C++ is also provided in “Extensions for C99 compatibility” on page 419.

Extensions related to GNU C and C++

Certain language extensions that correspond to GNU C and C++ features are implemented to facilitate portability. These include extensions to C89, C99, C++98,

and Standard C++. Throughout this document, the text indicates the IBM extensions that have been implemented for compatibility with GNU C and C++; a complete list of these is also provided in “Extensions for GNU C compatibility” on page 415 and “Extensions for GNU C++ compatibility” on page 422.

Extensions supporting the AltiVec Programming Interface

XL C/C++ supports AltiVec vector types for the PPU. These language extensions exploit the SIMD and parallel processing capabilities of the PowerPC processor, and facilitate the associated optimization techniques. The IBM implementation of the AltiVec Programming Interface specification is an extended implementation, which, for the most part, matches the syntax and semantics of the GNU C implementation. In addition to the text provided throughout this document which describes the behavior of the vector extensions, a list of the IBM extensions to the AltiVec Programming Interface is also provided in “Extensions for vector processing support” on page 417.

Extensions supporting Cell Broadband Engine processors

In addition to the vector types mandated by the AltiVec API, XL C/C++ supports special vector types for code targeting the Cell/B.E.[™] Synergistic Processor Unit (SPU). These language extensions exploit the hardware capabilities of the CPU to improve performance. The IBM extensions to AltiVec also apply to the SPU extensions; see “Extensions for vector processing support” on page 417 for a list of these.

How this document is organized

This document is organized to loosely follow the structure of the ISO standard language specifications and topics are grouped into similar headings.

Chapters 1 through 8 discuss language elements that are common to both C and C++, including lexical elements, data types, declarations, declarators, type conversions, expressions, operators, statements, and functions. Throughout these chapters, both standard features and extensions are discussed. Chapters 9 through 16 discuss standard C++ features exclusively, including classes, overloading, inheritance, templates, and exception handling. Chapter 17 discusses directives to the preprocessor.

The appendices provide summary lists of all the extended features supported by each language.

Conventions used in this document

Typographical conventions

The following table explains the typographical conventions used in this document.

Table 1. Typographical conventions

Typeface	Indicates	Example
bold	Lowercase commands, executable names, compiler options and directives.	If you specify -O3 , the compiler assumes -qhot=level=0 . To prevent all HOT optimizations with -O3 , you must specify -qnohot .

Table 1. *Typographical conventions (continued)*

Typeface	Indicates	Example
<i>italics</i>	Parameters or variables whose actual names or values are to be supplied by the user. Italics are also used to introduce new terms.	Make sure that you update the <i>size</i> parameter if you return more than the <i>size</i> requested.
monospace	Programming keywords and library functions, compiler built-in functions, examples of program code, command strings, or user-defined names.	If one or two cases of a <code>switch</code> statement are typically executed much more frequently than other cases, break out those cases by handling them separately before the <code>switch</code> statement.

Qualifying elements (icons and bracket separators)

This document uses marked bracket separators to delineate large blocks of text and icons to delineate small segments of text as follows:

Table 2. *Qualifying elements*

Qualifier/Icon	Meaning
C only 	The text describes a feature that is supported in the C language only; or describes behavior that is specific to the C language.
C++ only 	The text describes a feature that is supported in the C++ language only; or describes behavior that is specific to the C++ language.
IBM extension 	The text describes a feature that is an IBM XL C/C++ compiler extension to the standard language specifications.


Syntax diagrams

Throughout this document, diagrams illustrate XL C/C++ syntax. This section will help you to interpret and use those diagrams.

- Read the syntax diagrams from left to right, from top to bottom, following the path of the line.
 - The  symbol indicates the beginning of a command, directive, or statement.
 - The  symbol indicates that the command, directive, or statement syntax is continued on the next line.
 - The  symbol indicates that a command, directive, or statement is continued from the previous line.
 - The  symbol indicates the end of a command, directive, or statement.
- Fragments, which are diagrams of syntactical units other than complete commands, directives, or statements, start with the  symbol and end with the  symbol.
- Required items are shown on the horizontal line (the main path):



- Optional items are shown below the main path:



- If you can choose from two or more items, they are shown vertically, in a stack. If you *must* choose one of the items, one item of the stack is shown on the main path.



If choosing one of the items is optional, the entire stack is shown below the main path.



- An arrow returning to the left above the main line (a repeat arrow) indicates that you can make more than one choice from the stacked items or repeat an item. The separator character, if it is other than a blank, is also indicated:




- The item that is the default is shown above the main path.



- Keywords are shown in nonitalic letters and should be entered exactly as shown.
- Variables are shown in italicized lowercase letters. They represent user-supplied names or values.
- If punctuation marks, parentheses, arithmetic operators, or other such symbols are shown, you must enter them as part of the syntax.

Sample syntax diagram

The following syntax diagram example shows the syntax for the **#pragma comment** directive.



Notes:

- 1 This is the start of the syntax diagram.
- 2 The symbol # must appear first.
- 3 The keyword pragma must appear following the # symbol.
- 4 The name of the pragma comment must appear following the keyword pragma.

- 5 An opening parenthesis must be present.
 - 6 The comment type must be entered only as one of the types indicated: compiler, date, timestamp, copyright, or user.
 - 7 A comma must appear between the comment type copyright or user, and an optional character string.
 - 8 A character string must follow the comma. The character string must be enclosed in double quotation marks.
 - 9 A closing parenthesis is required.
 - 10 This is the end of the syntax diagram.
- The following examples of the **#pragma comment** directive are syntactically correct according to the diagram shown above:

```
#pragma comment(date)
#pragma comment(user)
#pragma comment(copyright,"This text will appear in the module")
```

Examples

The examples in this document, except where otherwise noted, are coded in a simple style that does not try to conserve storage, check for errors, achieve fast performance, or demonstrate all possible methods to achieve a specific result.

Related information

The following sections provide information on documentation related to XL C/C++:

- “IBM XL C/C++ publications”
- “Other IBM publications” on page xiv
- “Other publications” on page xiv

IBM XL C/C++ publications

XL C/C++ provides product documentation in the following formats:

- Installable man pages

Man pages are provided for the compiler invocations and all command-line utilities provided with the product. Instructions for installing and accessing the man pages are provided in the *XL C/C++ Installation Guide*.

- PDF documents

PDF documents are located by default in the `/opt/ibmcmp/xlc/cbe/9.0/doc/en_US/pdf/` directory.

The following files comprise the full set of XL C/C++ product manuals:

Table 3. XL C/C++ PDF files

Document title	PDF file name	Description
<i>IBM XL C/C++ for Multicore Acceleration for Linux, V9.0 Installation Guide, GC23-8520-00</i>	install.pdf	Contains information for installing XL C/C++ and configuring your environment for basic compilation and program execution.

Table 3. XL C/C++ PDF files (continued)

Document title	PDF file name	Description
<i>Getting Started with IBM XL C/C++ for Multicore Acceleration for Linux, V9.0, GC23-8518-00</i>	getstart.pdf	Contains an introduction to the XL C/C++ product, with information on setting up and configuring your environment, compiling and linking programs, and troubleshooting compilation errors.
<i>IBM XL C/C++ for Multicore Acceleration for Linux, V9.0 Compiler Reference, SC23-8516-00</i>	compiler.pdf	Contains information about the various compiler options, pragmas, macros, environment variables, and built-in functions.
<i>IBM XL C/C++ for Multicore Acceleration for Linux, V9.0 Language Reference, SC23-8519-00</i>	langref.pdf	Contains information about the C and C++ programming languages, as supported by IBM, including language extensions for portability and conformance to non-proprietary standards.
<i>IBM XL C/C++ for Multicore Acceleration for Linux, V9.0 Programming Guide, SC23-8517-00</i>	proguide.pdf	Contains information on advanced programming topics, such as application porting, library development, application optimization, and the XL C/C++ high-performance libraries.

To read a PDF file, use the Adobe® Reader. If you do not have the Adobe Reader, you can download it (subject to license terms) from the Adobe Web site at <http://www.adobe.com>.

More documentation related to XL C/C++ including redbooks, white papers, tutorials, and other articles, is available on the Web at:

<http://www.ibm.com/software/awdtools/xlcpp/library>

Standards and specifications documents

XL C/C++ is designed to support the following standards and specifications. You can refer to these standards for precise definitions of some of the features found in this document.

- *Information Technology – Programming languages – C, ISO/IEC 9899:1990*, also known as C89.
- *Information Technology – Programming languages – C, ISO/IEC 9899:1999*, also known as C99.
- *Information Technology – Programming languages – C++, ISO/IEC 14882:1998*, also known as C++98.
- *Information Technology – Programming languages – C++, ISO/IEC 14882:2003(E)*, also known as *Standard C++*.
- *Information Technology – Programming languages – Extensions for the programming language C to support new character data types, ISO/IEC DTR 19769*. This draft technical report has been accepted by the C standards committee, and is available at <http://www.open-std.org/JTC1/SC22/WG14/www/docs/n1040.pdf>.
- *Draft Technical Report on C++ Library Extensions, ISO/IEC DTR 19768*. This draft technical report has been submitted to the C++ standards committee, and is available at <http://www.open-std.org/JTC1/SC22/WG21/docs/papers/2005/n1836.pdf>.

- *AltiVec Technology Programming Interface Manual*, Motorola Inc. This specification for vector data types, to support vector processing technology, is available at http://www.freescale.com/files/32bit/doc/ref_manual/ALTIVECPIM.pdf.

Other IBM publications

- *IBM C/C++ Language Extensions for Cell Broadband Engine Architecture*, available at <http://www.ibm.com/developerworks/power/cell/documents.html>
- Specifications, white papers, and other technical documents for the Cell Broadband Engine™ architecture are available at http://www.ibm.com/chips/techlib/techlib.nsf/products/Cell_Broadband_Engine.
- The Cell Broadband Engine resource center, at <http://www.ibm.com/developerworks/power/cell>, is the central repository for technical information, including articles, tutorials, programming guides, and educational resources.

Other publications

- *Using the GNU Compiler Collection* available at <http://gcc.gnu.org/onlinedocs>

How to send your comments

Your feedback is important in helping to provide accurate and high-quality information. If you have any comments about this document or any other XL C/C++ documentation, send your comments by e-mail to compinfo@ca.ibm.com.

Be sure to include the name of the document, the part number of the document, the version of XL C/C++, and, if applicable, the specific location of the text you are commenting on (for example, a page number or table number).

Chapter 1. Scope and linkage

Scope is the largest region of program text in which a name can potentially be used without qualification to refer to an entity; that is, the largest region in which the name is potentially valid. Broadly speaking, scope is the general context used to differentiate the meanings of entity names. The rules for scope combined with those for name resolution enable the compiler to determine whether a reference to an identifier is legal at a given point in a file.

The scope of a declaration and the visibility of an identifier are related but distinct concepts. Scope is the mechanism by which it is possible to limit the visibility of declarations in a program. The *visibility* of an identifier is the region of program text from which the object associated with the identifier can be legally accessed. Scope can exceed visibility, but visibility cannot exceed scope. Scope exceeds visibility when a duplicate identifier is used in an inner declarative region, thereby hiding the object declared in the outer declarative region. The original identifier cannot be used to access the first object until the scope of the duplicate identifier (the lifetime of the second object) has ended.

Thus, the scope of an identifier is interrelated with the *storage duration* of the identified object, which is the length of time that an object remains in an identified region of storage. The lifetime of the object is influenced by its storage duration, which in turn is affected by the scope of the object identifier.

Linkage refers to the use or availability of a name across multiple translation units or within a single translation unit. The term *translation unit* refers to a source code file plus all the header and other source files that are included after preprocessing with the `#include` directive, minus any source lines skipped because of conditional preprocessing directives. Linkage allows the correct association of each instance of an identifier with one particular object or function.

Scope and linkage are distinguishable in that scope is for the benefit of the compiler, whereas linkage is for the benefit of the linker. During the translation of a source file to object code, the compiler keeps track of the identifiers that have external linkage and eventually stores them in a table within the object file. The linker is thereby able to determine which names have external linkage, but is unaware of those with internal or no linkage.

The distinctions between the different types of scopes are discussed in “Scope” on page 2. The different types of linkages are discussed in “Program linkage” on page 7.

Related information

- “Storage class specifiers” on page 45
- Chapter 9, “Namespaces (C++ only),” on page 235

Scope

The *scope* of an identifier is the largest region of the program text in which the identifier can potentially be used to refer to its object. In C++, the object being referred to must be unique. However, the name to access the object, the identifier itself, can be reused. The meaning of the identifier depends upon the context in which the identifier is used. Scope is the general context used to distinguish the meanings of names.

The scope of an identifier is possibly noncontiguous. One of the ways that breakage occurs is when the same name is reused to declare a different entity, thereby creating a contained declarative region (inner) and a containing declarative region (outer). Thus, point of declaration is a factor affecting scope. Exploiting the possibility of a noncontiguous scope is the basis for the technique called *information hiding*.

The concept of scope that exists in C was expanded and refined in C++. The following table shows the kinds of scopes and the minor differences in terminology.

Table 4. Kinds of scope

C	C++
block	local
function	function
Function prototype	Function prototype
file	global namespace
	namespace
	class

In all declarations, the identifier is in scope before the initializer. The following example demonstrates this:

```
int x;
void f() {
    int x = x;
}
```

The `x` declared in function `f()` has local scope, not global scope.

Related information

- Chapter 9, “Namespaces (C++ only),” on page 235

Block/local scope

A name has *local scope* or *block scope* if it is declared in a block. A name with local scope can be used in that block and in blocks enclosed within that block, but the name must be declared before it is used. When the block is exited, the names declared in the block are no longer available.

Parameter names for a function have the scope of the outermost block of that function. Also, if the function is declared and not defined, these parameter names have function prototype scope.

When one block is nested inside another, the variables from the outer block are usually visible in the nested block. However, if the declaration of a variable in a nested block has the same name as a variable that is declared in an enclosing block, the declaration in the nested block hides the variable that was declared in the enclosing block. The original declaration is restored when program control returns to the outer block. This is called *block visibility*.

Name resolution in a local scope begins in the immediately enclosing scope in which the name is used and continues outward with each enclosing scope. The order in which scopes are searched during name resolution causes the phenomenon of information hiding. A declaration in an enclosing scope is hidden by a declaration of the same identifier in a nested scope.

Related information

- “Block statements” on page 178

Function scope

The only type of identifier with *function scope* is a label name. A label is implicitly declared by its appearance in the program text and is visible throughout the function that declares it.

A label can be used in a `goto` statement before the actual label is seen.

Related information

- “Labeled statements” on page 175

Function prototype scope

In a function declaration (also called a *function prototype*) or in any function declarator—except the declarator of a function definition—parameter names have *function prototype scope*. Function prototype scope terminates at the end of the nearest enclosing function declarator.

Related information

- “Function declarations” on page 202

File/global scope

C only

A name has *file scope* if the identifier’s declaration appears outside of any block. A name with file scope and internal linkage is visible from the point where it is declared to the end of the translation unit.

End of C only

C++ only

Global scope or *global namespace scope* is the outermost namespace scope of a program, in which objects, functions, types and templates can be defined. A name has *global namespace scope* if the identifier’s declaration appears outside of all blocks, namespaces, and classes.

A name with global namespace scope and internal linkage is visible from the point where it is declared to the end of the translation unit.

A name with global (namespace) scope is also accessible for the initialization of global variables. If that name is declared extern, it is also visible at link time in all object files being linked.

A user-defined namespace can be nested within the global scope using namespace definitions, and each user-defined namespace is a different scope, distinct from the global scope.

End of C++ only

Related information

- Chapter 9, “Namespaces (C++ only),” on page 235
- “Internal linkage” on page 8

Examples of scope in C

The following example declares the variable `x` on line 1, which is different from the `x` it declares on line 2. The declared variable on line 2 has function prototype scope and is visible only up to the closing parenthesis of the prototype declaration. The variable `x` declared on line 1 resumes visibility after the end of the prototype declaration.

```
1  int x = 4;           /* variable x defined with file scope */
2  long myfunc(int x, long y); /* variable x has function    */
3                               /* prototype scope           */
4  int main(void)
5  {
6      /* . . . */
7  }
```

The following program illustrates blocks, nesting, and scope. The example shows two kinds of scope: file and block. The `main` function prints the values 1, 2, 3, 0, 3, 2, 1 on separate lines. Each instance of `i` represents a different variable.

```

#include <stdio.h>
int i = 1;                                /* i defined at file scope */

int main(int argc, char * argv[])
{
1
1     printf("%d\n", i);                    /* Prints 1 */
1
1     {
1 2     int i = 2, j = 3;                    /* i and j defined at block scope */
1 2     printf("%d\n%d\n", i, j);           /* global definition of i is hidden */
1 2     /* Prints 2, 3 */
1 2     {
1 2 3     int i = 0; /* i is redefined in a nested block */
1 2 3     /* previous definitions of i are hidden */
1 2 3     printf("%d\n%d\n", i, j); /* Prints 0, 3 */
1 2     }
1 2     printf("%d\n", i);                    /* Prints 2 */
1 2     }
1
1     printf("%d\n", i);                    /* Prints 1 */
1
1     return 0;
1
}

```

Class scope (C++ only)

A name declared within a member function hides a declaration of the same name whose scope extends to or past the end of the member function's class.

When the scope of a declaration extends to or past the end of a class definition, the regions defined by the member definitions of that class are included in the scope of the class. Members defined lexically outside of the class are also in this scope. In addition, the scope of the declaration includes any portion of the declarator following the identifier in the member definitions.

The name of a class member has *class scope* and can only be used in the following cases:

- In a member function of that class
- In a member function of a class derived from that class
- After the . (dot) operator applied to an instance of that class
- After the . (dot) operator applied to an instance of a class derived from that class, as long as the derived class does not hide the name
- After the -> (arrow) operator applied to a pointer to an instance of that class
- After the -> (arrow) operator applied to a pointer to an instance of a class derived from that class, as long as the derived class does not hide the name
- After the :: (scope resolution) operator applied to the name of a class
- After the :: (scope resolution) operator applied to a class derived from that class

Related information

- Chapter 11, "Classes (C++ only)," on page 259
- "Scope of class names" on page 263

- “Member scope” on page 273
- “Friend scope” on page 288
- “Access control of base class members” on page 297
- “Scope resolution operator :: (C++ only)” on page 126

Namespaces of identifiers

Namespaces are the various syntactic contexts within which an identifier can be used. Within the same context and the same scope, an identifier must uniquely identify an entity. Note that the term *namespace* as used here applies to C as well as C++ and does not refer to the C++ namespace language feature. The compiler sets up *namespaces* to distinguish among identifiers referring to different kinds of entities. Identical identifiers in different namespaces do not interfere with each other, even if they are in the same scope.

The same identifier can declare different objects as long as each identifier is unique within its namespace. The syntactic context of an identifier within a program lets the compiler resolve its namespace without ambiguity.

Within each of the following four namespaces, the identifiers must be unique.

- *Tags* of these types must be unique within a single scope:
 - Enumerations
 - Structures and unions
- *Members* of structures, unions, and classes must be unique within a single structure, union, or class type.
- *Statement labels* have function scope and must be unique within a function.
- All other *ordinary identifiers* must be unique within a single scope:
 - C function names (C++ function names can be overloaded)
 - Variable names
 - Names of function parameters
 - Enumeration constants
 - typedef names.

You can redefine identifiers in the same namespace but within enclosed program blocks.

Structure tags, structure members, variable names, and statement labels are in four different namespaces. No name conflict occurs among the items named `student` in the following example:

```
int get_item()
{
    struct student      /* structure tag */
    {
        char student[20]; /* structure member */
        int section;
        int id;
    } student;          /* structure variable */

    goto student;
    student;           /* null statement label */
    return 0;
}
```

The compiler interprets each occurrence of `student` by its context in the program: when `student` appears after the keyword `struct`, it is a structure tag; when it

appears in the block defining the `student` type, it is a structure member variable; when it appears at the end of the structure definition, it declares a structure variable; and when it appears after the `goto` statement, it is a label.

Name hiding (C++ only)

If a class name or enumeration name is in scope and not hidden, it is *visible*. A class name or enumeration name can be hidden by an explicit declaration of that same name — as an object, function, or enumerator — in a nested declarative region or derived class. The class name or enumeration name is hidden wherever the object, function, or enumerator name is visible. This process is referred to as *name hiding*.

In a member function definition, the declaration of a local name hides the declaration of a member of the class with the same name. The declaration of a member in a derived class hides the declaration of a member of a base class of the same name.

Suppose a name `x` is a member of namespace `A`, and suppose that the members of namespace `A` are visible in a namespace `B` because of a `using` declaration. A declaration of an object named `x` in namespace `B` will hide `A::x`. The following example demonstrates this:

```
#include <iostream>
#include <typeinfo>
using namespace std;

namespace A {
    char x;
};

namespace B {
    using namespace A;
    int x;
};

int main() {
    cout << typeid(B::x).name() << endl;
}
```

The following is the output of the above example:

```
int
```

The declaration of the integer `x` in namespace `B` hides the character `x` introduced by the `using` declaration.

Related information

- Chapter 11, “Classes (C++ only),” on page 259
- “Member functions” on page 271
- “Member scope” on page 273
- Chapter 9, “Namespaces (C++ only),” on page 235

Program linkage

Linkage determines whether identifiers that have identical names refer to the same object, function, or other entity, even if those identifiers appear in different translation units. The linkage of an identifier depends on how it was declared. There are three types of linkages:

- Internal linkage : identifiers can only be seen within a translation unit.
- External linkage : identifiers can be seen (and referred to) in other translation units.
- No linkage: identifiers can only be seen in the scope in which they are defined.

Linkage does not affect scoping, and normal name lookup considerations apply.

C++ only

You can also have linkage between C++ and non-C++ code fragments, which is called *language linkage*. Language linkage enables the close relationship between C++ and C by allowing C++ code to link with that written in C. All identifiers have a language linkage, which by default is C++. Language linkage must be consistent across translation units, and non-C++ language linkage implies that the identifier has external linkage.

End of C++ only

Related information

- “The static storage class specifier” on page 46
- “The extern storage class specifier” on page 48
- “Function storage class specifiers” on page 206
- “Type qualifiers” on page 72
- “Anonymous unions” on page 67

Internal linkage

The following kinds of identifiers have internal linkage:

- Objects, references, or functions explicitly declared `static`
- Objects or references declared in namespace scope (or global scope in C) with the specifier `const` and neither explicitly declared `extern`, nor previously declared to have external linkage
- Data members of an anonymous union
-  Function templates explicitly declared `static`
-  Identifiers declared in the unnamed namespace

A function declared inside a block will usually have external linkage. An object declared inside a block will usually have external linkage if it is specified `extern`. If a variable that has `static` storage is defined outside a function, the variable has internal linkage and is available from the point where it is defined to the end of the current translation unit.

If the declaration of an identifier has the keyword `extern` and if a previous declaration of the identifier is visible at namespace or global scope, the identifier has the same linkage as the first declaration.

External linkage

C only

In global scope, identifiers for the following kinds of entities declared without the `static` storage class specifier have external linkage:

- An object

- A function

If an identifier in C is declared with the `extern` keyword and if a previous declaration of an object or function with the same identifier is visible, the identifier has the same linkage as the first declaration. For example, a variable or function that is first declared with the keyword `static` and later declared with the keyword `extern` has internal linkage. However, a variable or function that has no linkage and was later declared with a linkage specifier will have the linkage that was expressly specified.

End of C only

C++ only

In namespace scope, the identifiers for the following kinds of entities have external linkage:

- A reference or an object that does not have internal linkage
- A function that does not have internal linkage
- A named class or enumeration
- An unnamed class or enumeration defined in a typedef declaration
- An enumerator of an enumeration that has external linkage
- A template, unless it is a function template with internal linkage
- A namespace, unless it is declared in an unnamed namespace

If the identifier for a class has external linkage, then, in the implementation of that class, the identifiers for the following will also have external linkage:

- A member function.
- A static data member.
- A class of class scope.
- An enumeration of class scope.

End of C++ only

No linkage

The following kinds of identifiers have no linkage:

- Names that have neither external or internal linkage
- Names declared in local scopes (with exceptions like certain entities declared with the `extern` keyword)
- Identifiers that do not represent an object or a function, including labels, enumerators, typedef names that refer to entities with no linkage, type names, function parameters, and template names

You cannot use a name with no linkage to declare an entity with linkage. For example, you cannot use the name of a structure or enumeration or a typedef name referring to an entity with no linkage to declare an entity with linkage. The following example demonstrates this:

```
int main() {
    struct A { };
    // extern A a1;
    typedef A myA;
    // extern myA a2;
}
```


The compiler will not allow the declaration of `a1` with external linkage. Structure `A` has no linkage. The compiler will not allow the declaration of `a2` with external linkage. The typedef name `myA` has no linkage because `A` has no linkage.

Language linkage (C++ only)

Linkage between C++ and non-C++ code fragments is called *language linkage*. All function types, function names, and variable names have a language linkage, which by default is C++.

You can link C++ object modules to object modules produced using other source languages such as C by using a *linkage specification*.

Linkage specification syntax



The *string_literal* is used to specify the linkage associated with a particular function. String literals used in linkage specifications should be considered as case-sensitive. All platforms support the following values for *string_literal*:

"C++" Unless otherwise specified, objects and functions have this default linkage specification.

"C" Indicates linkage to a C procedure

Calling shared libraries that were written before C++ needed to be taken into account requires the `#include` directive to be within an `extern "C" { }` declaration.

```
extern "C" {
#include "shared.h"
}
```

The following example shows a C printing function that is called from C++.

```
// in C++ program
extern "C" int displayfoo(const char *);
int main() {
    return displayfoo("hello");
}

/* in C program */
#include <stdio.h>
extern int displayfoo(const char * str) {
    while (*str) {
        putchar(*str);
        putchar(' ');
        ++str;
    }
    putchar('\n');
}
```

Name mangling (C++ only)

Name mangling is the encoding of function and variable names into unique names so that linkers can separate common names in the language. Type names may also be mangled. Name mangling is commonly used to facilitate the overloading feature and visibility within different scopes. The compiler generates function names with an encoding of the types of the function arguments when the module

is compiled. If a variable is in a namespace, the name of the namespace is mangled into the variable name so that the same variable name can exist in more than one namespace. The C++ compiler also mangles C variable names to identify the namespace in which the C variable resides.

The scheme for producing a mangled name differs with the object model used to compile the source code: the mangled name of an object of a class compiled using one object model will be different from that of an object of the same class compiled using a different object model. The object model is controlled by compiler option or by pragma.

Name mangling is not desirable when linking C modules with libraries or object files compiled with a C++ compiler. To prevent the C++ compiler from mangling the name of a function, you can apply the extern "C" linkage specifier to the declaration or declarations, as shown in the following example:

```
extern "C" {  
    int f1(int);  
    int f2(int);  
    int f3(int);  
};
```

This declaration tells the compiler that references to the functions f1, f2, and f3 should not be mangled.

The extern "C" linkage specifier can also be used to prevent mangling of functions that are defined in C++ so that they can be called from C. For example,

```
extern "C" {  
    void p(int){  
        /* not mangled */  
    }  
};
```

In multiple levels of nested extern declarations, the innermost extern specification prevails.

```
extern "C" {  
    extern "C++" {  
        void func();  
    }  
}
```

In this example, func has C++ linkage.

Chapter 2. Lexical Elements

A *lexical element* refers to a character or groupings of characters that may legally appear in a source file. This section contains discussions of the basic lexical elements and conventions of the C and C++ programming languages:

- “Tokens”
- “Source program character set” on page 32
- “Comments” on page 37

Tokens

Source code is treated during preprocessing and compilation as a sequence of *tokens*. A token is the smallest independent unit of meaning in a program, as defined by the compiler. There are four different types of tokens:

- Keywords
- Identifiers
- Literals
- Punctuators and operators

Adjacent identifiers, keywords, and literals must be separated with white space. Other tokens should be separated by white space to make the source code more readable. White space includes blanks, horizontal and vertical tabs, new lines, form feeds, and comments.

Keywords

Keywords are identifiers reserved by the language for special use. Although you can use them for preprocessor macro names, it is considered poor programming style. Only the exact spelling of keywords is reserved. For example, `auto` is reserved but `AUTO` is not.

Table 5. C and C++ keywords

<code>auto</code>	<code>double</code>	<code>int</code>	<code>struct</code>
<code>break</code>	<code>else</code>	<code>long</code>	<code>switch</code>
<code>case</code>	<code>enum</code>	<code>register</code>	<code>typedef</code>
<code>char</code>	<code>extern</code>	<code>return</code>	<code>union</code>
<code>const</code>	<code>float</code>	<code>short</code>	<code>unsigned</code>
<code>continue</code>	<code>for</code>	<code>signed</code>	<code>void</code>
<code>default</code>	<code>goto</code>	<code>sizeof</code>	<code>volatile</code>
<code>do</code>	<code>if</code>	<code>static</code>	<code>while</code>

C only

Standard C at the C99 level also reserves the following keywords:

Table 6. C99 keywords

<code>_Bool</code>	<code>_Imaginary¹</code>
<code>_Complex</code>	<code>inline</code>
	<code>restrict</code>

Notes:

1. The keyword `_Imaginary` is reserved for possible future use. For complex number functionality, use `_Complex`; see “Complex literals” on page 24 for details.

End of C only

C++ only

The C++ language also reserves the following keywords:

Table 7. C++ keywords

<code>asm</code>	<code>export</code>	<code>private</code>	<code>true</code>
<code>bool</code>	<code>false</code>	<code>protected</code>	<code>try</code>
<code>catch</code>	<code>friend</code>	<code>public</code>	<code>typeid</code>
<code>class</code>	<code>inline</code>	<code>reinterpret_cast</code>	<code>typename</code>
<code>const_cast</code>	<code>mutable</code>	<code>static_cast</code>	<code>using</code>
<code>delete</code>	<code>namespace</code>	<code>template</code>	<code>virtual</code>
<code>dynamic_cast</code>	<code>new</code>	<code>this</code>	<code>wchar_t</code>
<code>explicit</code>	<code>operator</code>	<code>throw</code>	

End of C++ only

Keywords for language extensions

IBM extension

In addition to standard language keywords, XL C/C++ reserves the following keywords for use in language extensions:

Table 8. Keywords for C and C++ language extensions

<code>__alignof</code>	<code>__extension__</code>	<code>__restrict</code>
<code>__alignof__</code>	<code>__label__</code>	<code>__restrict__</code>
<code>__asm (C only)</code>	<code>__imag__</code>	<code>__signed__</code>
<code>__asm__ (C only)</code>	<code>__inline__²</code>	<code>__signed</code>
<code>__attribute__</code>	<code>pixel¹</code>	<code>__volatile__</code>
<code>__attribute</code>	<code>__pixel¹</code>	<code>__thread</code>
<code>bool (C only)¹</code>	<code>__real__</code>	<code>typeof³</code>
<code>__complex__</code>		<code>__typeof__</code>
<code>__const__</code>		<code>vector¹</code>
		<code>__vector¹</code>

Notes:

1. These keywords are recognized only in a vector declaration context. `pixel`, `__pixel`, and `bool` are only valid for compilation targeting the PPU.
2.  The `__inline__` keyword uses the GNU C semantics for inline functions. For details, see “Linkage of inline functions” on page 209.
3.  `typeof` is only recognized when `-qkeyword=typeof` is in effect.

C++ only

XL C++ reserves the following keywords as language extensions for compatibility with C99.

Table 9. Keywords for C++ language extensions related to C99

_Complex
 _Imaginary¹
 _Pragma
 restrict

Notes:

1. The keyword `_Imaginary` is reserved for possible future use. For complex number functionality, use `_Complex`; see “Complex literals” on page 24 for details.

End of C++ only

More detailed information regarding the compilation contexts in which extension keywords are valid is provided in the sections of this document that describe each keyword.

Related information

- `-qclanglvl` and `-qkeyword` in the *XL C/C++ Compiler Reference*
- “Vector types” on page 57

End of IBM extension

Identifiers

Identifiers provide names for the following language elements:

- Functions
- Objects
- Labels
- Function parameters
- Macros and macro parameters
- Type definitions
- Enumerated types and enumerators
- Structure and union names
-  Classes and class members
-  Templates
-  Template parameters
-  Namespaces

An identifier consists of an arbitrary number of letters, digits, or the underscore character in the form:



Related information

- “Identifier expressions (C++ only)” on page 124
- “The Unicode standard” on page 34
- “Keywords” on page 13

Characters in identifiers

The first character in an identifier must be a letter or the `_` (underscore) character; however, beginning identifiers with an underscore is considered poor programming style.

The compiler distinguishes between uppercase and lowercase letters in identifiers. For example, `PROFIT` and `profit` represent different identifiers. If you specify a lowercase `a` as part of an identifier name, you cannot substitute an uppercase `A` in its place; you must use the lowercase letter.

The universal character names for letters and digits outside of the basic source character set are allowed in C++ and at the C99 language level.  In C++, you must compile with the `-qlanglvl=ucs` option for universal character name support.

 The dollar sign can appear in identifier names when compiled using the `-qdollar` compiler option or at one of the extended language levels that encompasses this option.

Related information

- `-qdollar` and `-qlanglvl` in the *XL C/C++ Compiler Reference*

Reserved identifiers

Identifiers with two initial underscores or an initial underscore followed by an uppercase letter are reserved globally for use by the compiler.

 Identifiers that begin with a single underscore are reserved as identifiers with file scope in both the ordinary and tag namespaces.

 Identifiers that begin with a single underscore are reserved in the global namespace.

Although the names of system calls and library functions are not reserved words if you do not include the appropriate headers, avoid using them as identifiers. Duplication of a predefined name can lead to confusion for the maintainers of your code and can cause errors at link time or run time. If you include a library in a program, be aware of the function names in that library to avoid name duplications. You should always include the appropriate headers when using standard library functions.

The `__func__` predefined identifier

The C99 predefined identifier `__func__` makes a function name available for use within the function. Immediately following the opening brace of each function definition, `__func__` is implicitly declared by the compiler. The resulting behavior is as if the following declaration had been made:

```
static const char __func__[] = "function-name";
```

where *function-name* is the name of the lexically-enclosing function. The function name is not mangled.

C++ only

The function name is qualified with the enclosing class name or function name. For example, if `foo` is a member function of class `X`, the predefined identifier of `foo` is `X::foo`. If `foo` is defined within the body of `main`, the predefined identifier of `foo` is `main::X::foo`.

The names of template functions or member functions reflect the instantiated type. For example, the predefined identifier for the template function `foo` instantiated with `int`, `template<class T> void foo()` is `foo<int>`.

End of C++ only

For debugging purposes, you can explicitly use the `__func__` identifier to return the name of the function in which it appears. For example:

```
#include <stdio.h>

void myfunc(void) {
    printf("%s\n", __func__);
    printf("size of __func__ = %d\n", sizeof(__func__));
}

int main() {
    myfunc();
}
```

The output of the program is:

```
myfunc
size of __func__=7
```

When the `assert` macro is used inside a function definition, the macro adds the name of the enclosing function on the standard error stream.

Related information

- “Function declarations and definitions” on page 201


Assembly labels

IBM extension

The compiler binds each non-static external variable and function name in the source code to a name that it generates in the object file and any assembly code that is emitted. For compatibility with GCC, the compiler implements an extension to standard C and C++ that allows you to specify the name to be used in the object file and assembly code, by applying an assembly *label* to the declaration of a global variable or function prototype. You can also define names that do not start with an underscore even on systems where an underscore is normally prepended to the name of a function or variable.

▶ C++ You can use assembly labels with member functions, and functions and variables that are declared in namespaces other than the global namespace.

Assembly label syntax



The *string_literal* is a valid assembly name that is to be bound to the given object or function. For a label applied to a function declaration, the name must specify an existing function that is defined in any compilation unit; if no definition is available, a link-time error will occur. For a label applied to a variable declaration, no other definition is required.

The following are examples of assembly label specifications:

```
void func3() __asm__("foo3");
int i __asm("abc");
char c asm("abcs") = 'a';
```

C++ only

To distinguish between overloaded functions, XL C++ mangles function names in the object file. Therefore, if you use an assembly label to map a function name, you must use the mangled name of the target function. Furthermore, you must ensure that an assembly label name that you specify for a variable does not conflict with any mangled name. Alternatively, you can prevent name mangling on a target function by declaring it as having C linkage; for more information, see “Name mangling (C++ only)” on page 10.

End of C++ only

The following are restrictions on the use of assembly labels:

- Assembly labels cannot be specified on local or static variables.
- The same assembly label name cannot be applied to multiple identifiers (in C++, this is the name after mangling) in the same compilation unit.
- The assembly label name cannot be the same as any other global identifier name (in C++, the name after mangling) in the same compilation unit, unless the label name and identifier name are used for the same variable or function declaration.
- The assembly label cannot be specified on typedef declarations.
- An assembly label cannot be the same as a name specified on a different variable or function by a previous **#pragma map** directive. Similarly, the map name specified by a **#pragma map** directive cannot be the same as a name specified by a previous assembly label on a different variable or function.
- You cannot apply an assembly label to an identifier that has been mapped to a different name by a **#pragma map** directive on a previous declaration of that variable or function. Similarly, you cannot specify a **#pragma map** directive on an identifier that has previously been remapped by an assembly label.
- **C** If you apply different labels to multiple declarations of the same variable or function, the first specification is honored, and all subsequent assembly labels are ignored with a warning.
- **C++** You cannot apply an assembly label to any of the following:
 - member variable declarations
 - friend declarations

- template function and member declarations, or any declarations contained within a template
- virtual member functions
- constructors and destructors

Related information

- **#pragma map** in the *XL C/C++ Compiler Reference*
- “The alias function attribute” on page 220
- “Variables in specified registers” on page 50
- “Inline assembly statements” on page 194
- **-qreserved_reg** in the *XL C/C++ Compiler Reference*

End of IBM extension

Literals

The term *literal constant*, or *literal*, refers to a value that occurs in a program and cannot be changed. The C language uses the term *constant* in place of the noun *literal*. The adjective *literal* adds to the concept of a constant the notion that we can speak of it only in terms of its value. A literal constant is nonaddressable, which means that its value is stored somewhere in memory, but we have no means of accessing that address.

Every *literal* has a value and a data type. The value of any literal does not change while the program runs and must be in the range of representable values for its type. The following are the available types of literals:

- Integer literals
- Boolean literals
- Floating-point literals
-  Vector literals
- Character literals
- String literals


Integer literals

Integer literals are numbers that do not have a decimal point or an exponential part. They can be represented as:

- Decimal integer literals
- Hexadecimal integer literals
- Octal integer literals

An integer literal may have a prefix that specifies its base, or a suffix that specifies its type.

Integer literal syntax



The data type of an integer literal is determined by its form, value, and suffix. The following table lists the integer literals and shows the possible data types. The smallest data type that can represent the constant value is used to store the constant.

Integer literal	Possible data types
unsuffixed decimal	int, long int, C long long int ¹
unsuffixed octal or hexadecimal	int, unsigned int, long int, unsigned long int, C long long int ¹ , C unsigned long long int ¹
decimal, octal, or hexadecimal suffixed by u or U	unsigned int, unsigned long int, C unsigned long long int ¹
decimal suffixed by l or L	long int, C long long int ¹
octal or hexadecimal suffixed by l or L	long int, unsigned long int, C long long int ¹ , C unsigned long long int ¹
decimal, octal, or hexadecimal suffixed by both u or U, and l or L	unsigned long int, C unsigned long long int ¹
decimal suffixed by ll or LL	long long int
octal or hexadecimal suffixed by ll or LL	long long int, unsigned long long int
decimal, octal, or hexadecimal suffixed by both u or U, and ll or LL	unsigned long long int
Notes:	
1. Assumes compilation with ppuxlc or spuxlc or ppuc99 or with -qlanglvl=stdc99 or -qlanglvl=extc99 .	

Related information

- “Integral types” on page 53
- “Integral conversions” on page 114
- **-qlanglvl** in the *XL C/C++ Compiler Reference*

Decimal integer literals: A *decimal integer literal* contains any of the digits 0 through 9. The first digit cannot be 0. Integer literals beginning with the digit 0 are interpreted as an octal integer literal rather than as a decimal integer literal.

Decimal integer literal syntax



A plus (+) or minus (-) symbol can precede a decimal integer literal. The operator is treated as a unary operator rather than as part of the literal.

The following are examples of decimal literals:

```
485976
-433132211
+20
5
```

Hexadecimal integer literals: A *hexadecimal integer literal* begins with the 0 digit followed by either an x or X, followed by any combination of the digits 0 through 9 and the letters a through f or A through F. The letters A (or a) through F (or f) represent the values 10 through 15, respectively.

Hexadecimal integer literal syntax



The following are examples of hexadecimal integer literals:

```
0x3b24
0XF96
0x21
0x3AA
0X29b
0X4bd
```

Octal integer literals: An *octal integer literal* begins with the digit 0 and contains any of the digits 0 through 7.

Octal integer literal syntax



The following are examples of octal integer literals:

```
0
0125
034673
03245
```

Boolean literals

C At the C99 level, C defines `true` and `false` as macros in the header file `stdbool.h`.

C++ There are only two Boolean literals: `true` and `false`.

Related information

- “Boolean types” on page 54
- “Boolean conversions” on page 114

Floating-point literals

Floating-point literals are numbers that have a decimal point or an exponential part. They can be represented as:


- Real literals
 - Binary floating-point literals
 - Hexadecimal floating-point literals
- Complex literals

Binary floating-point literals: A real binary floating-point constant consists of the following:

- An integral part
- A decimal point
- A fractional part
- An exponent part
- An optional suffix

Both the integral and fractional parts are made up of decimal digits. You can omit either the integral part or the fractional part, but not both. You can omit either the decimal point or the exponent part, but not both.

Binary floating-point literal syntax



Exponent:



The suffix `f` or `F` indicates a type of float, and the suffix `l` or `L` indicates a type of long double. If a suffix is not specified, the floating-point constant has a type double.

A plus (+) or minus (-) symbol can precede a floating-point literal. However, it is not part of the literal; it is interpreted as a unary operator.

The following are examples of floating-point literals:

Floating-point constant	Value
5.3876e4	53,876
4e-11	0.00000000004
1e+5	100000
7.321E-3	0.007321
3.2E+4	32000
0.5e-6	0.0000005
0.45	0.45
6.e10	60000000000

Related information

- “Floating-point types” on page 54
- “Floating-point conversions” on page 114
- “Unary expressions” on page 128

Hexadecimal floating-point literals: Real hexadecimal floating constants, which are a C99 feature, consist of the following:

- a hexadecimal prefix
- a significant part
- a binary exponent part
- an optional suffix


The significant part represents a rational number and is composed of the following:


- a sequence of hexadecimal digits (whole-number part)
- an optional fraction part

The optional fraction part is a period followed by a sequence of hexadecimal digits.

The exponent part indicates the power of 2 to which the significant part is raised, and is an optionally signed decimal integer. The type suffix is optional. The full syntax is as follows:

Hexadecimal floating-point literal syntax






Exponent:



The suffix *f* or *F* indicates a type of float, and the suffix *l* or *L* indicates a type of long double. If a suffix is not specified, the floating-point constant has a type double. You can omit either the whole-number part or the fraction part, but not both. The binary exponent part is required to avoid the ambiguity of the type suffix *F* being mistaken for a hexadecimal digit.

Complex literals: Complex literals, which are a C99 feature, are constructed in two parts: the real part, and the imaginary part.

Complex literal syntax



Real part:



The *floating-point constant* can be specified as a decimal or hexadecimal floating-point constant (including optional suffixes), in any of the formats described in the previous sections.

`_Complex_I` is a macro defined in the `complex.h` header file, representing the imaginary unit *i*, the square root of -1.

For example, the declaration:


```
varComplex = 2.0f + 2.0f * _Complex_I;
```

initializes the complex variable `varComplex` to a value of $2.0 + 2.0i$.

IBM extension

For ease of porting applications developed with GNU C, XL C/C++ also allows you to indicate the imaginary part of a complex literal with a suffix, in addition to the standard suffixes that indicate the type of the complex number (float, double, or long double).

The simplified syntax for a complex literal using the GNU suffixes is as follows:



real part:



imaginary part:



floating-point constant can be specified as a decimal or hexadecimal floating-point constant (including optional suffixes), in any of the formats described in the previous sections.

imaginary-suffix is one of the suffixes *i*, *I*, *j*, or *J*, representing the imaginary unit.

For example, the declaration

```
varComplex = 3.0f + 4.0fi;
```

initializes the complex variable `varComplex` to a value of `3.0 + 4.0i`.

End of IBM extension

Related information


- “Complex floating-point types ” on page 55

Vector literals

IBM extension

A vector literal is a constant expression for which the value is interpreted as a vector type. The data type of a vector literal is represented by a parenthesized vector type, and its value is represented by a parenthesized set of constant expressions that represent the vector elements. When all vector elements have the same value, the value of the literal can be represented by a single parenthesized constant expression. Vector literals allow the initialization of vector types.

Vector literal syntax



The *vector_type* is a supported vector type; see “Vector types” on page 57 for a list of these.

The *constant_expression* can be either of the following:

- A single expression, which will initialize all elements of the vector to the same value
- A comma-separated list of expressions, the number of which is determined by the type of the vector. The number of constant expressions must be exactly:

- 1 For the `qword` type.¹

- 2 For vector long long and vector double types².
- 4 For vector int, vector long³, and vector float types.
- 8 For vector short and vector pixel⁴ types.
- 16 For vector char types.

Notes:

1. Only valid for compilation targeting the SPU, when the spu_intrinsics.h header is included in program source.
2. Only valid for compilation targeting the SPU.
3. Deprecated; also only valid for compilation targeting the PPU.
4. Only valid for compilation targeting the PPU.

The following table shows the supported vector literals and how the compiler interprets them to determine their values.

Table 10. Vector literals

Syntax	Interpreted by the compiler as
(vector unsigned char)(<i>unsigned int</i>)	A set of 16 unsigned 8-bit quantities that all have the value of the single integer.
(vector unsigned char)(<i>unsigned int, ...</i>)	A set of 16 unsigned 8-bit quantities with the value specified by each of 16 integers.
(vector signed char)(<i>int</i>)	A set of 16 signed 8-bit quantities that all have the value of the single integer.
(vector signed char)(<i>int, ...</i>)	A set of 16 signed 8-bit quantities with the value specified by each of 16 integers.
(vector bool char)(<i>unsigned int</i>)	A set of 16 unsigned 8-bit quantities that all have the value of the single integer.
(vector bool char)(<i>unsigned int, ...</i>)	A set of 16 unsigned 8-bit quantities with a value specified by each of 16 integers.
(vector unsigned short)(<i>unsigned int</i>)	A set of 8 unsigned 16-bit quantities that all have the value of the single integer.
(vector unsigned short)(<i>unsigned int, ...</i>)	A set of 8 unsigned 16-bit quantities with a value specified by each of 8 integers.
(vector signed short)(<i>int</i>)	A set of 8 signed 16-bit quantities that all have the value of the single integer.
(vector signed short)(<i>int, ...</i>)	A set of 8 signed 16-bit quantities with a value specified by each of 8 integers.
(vector bool short)(<i>unsigned int</i>)	A set of 8 unsigned 16-bit quantities that all have the value of the single integer.
(vector bool short)(<i>unsigned int, ...</i>)	A set of 8 unsigned 16-bit quantities with a value specified by each of 8 integers.
(vector unsigned int)(<i>unsigned int</i>)	A set of 4 unsigned 32-bit quantities that all have the value of the single integer.
(vector unsigned int)(<i>unsigned int, ...</i>)	A set of 4 unsigned 32-bit quantities with a value specified by each of 4 integers.
(vector signed int)(<i>int</i>)	A set of 4 signed 32-bit quantities that all have the value of the single integer.
(vector signed int)(<i>int, ...</i>)	A set of 4 signed 32-bit quantities with a value specified by each of 4 integers.

Table 10. Vector literals (continued)

Syntax	Interpreted by the compiler as
(vector bool int)(unsigned int)	A set of 4 unsigned 32-bit quantities that all have the value of the single integer.
(vector bool int)(unsigned int, ...)	A set of 4 unsigned 32-bit quantities with a value specified by each of 4 integers.
(vector unsigned long long)(unsigned long long)	A set of 2 unsigned 64-bit quantities that both have the value of the single long long.
(vector unsigned long long)(unsigned long long, ...)	A set of 2 unsigned 64-bit quantities with a value specified by each of 2 long longs.
(vector signed long long)(signed long long)	A set of 2 signed 64-bit quantities that both have the value of the single long long.
(vector signed long long)(unsigned long long, ...)	A set of 2 signed 64-bit quantities with a value specified by each of 2 long longs.
(vector float)(float)	A set of 4 32-bit single-precision floating-point quantities that all have the value of the single float.
(vector float)(float, ...)	A set of 4 32-bit single-precision floating-point quantities with a value specified by each of 4 floats.
(vector double)(double)	A set of 2 64-bit double-precision floating-point quantities that both have the value of the single double.
(vector double)(double, ...)	A set of 2 64-bit double-precision floating-point quantities with a value specified by each of 2 doubles.
(vector pixel)(unsigned int)	A set of 8 unsigned 16-bit quantities that all have the value of the single integer.
(vector pixel)(unsigned int, ...)	A set of 8 unsigned 16-bit quantities with a value specified by each of 8 integers.

Note: For code targeting the SPU, you can also use curly braces to enclose multiple literal values for vector types supported on the SPU.

For example, for an unsigned integer vector type, the literal could be either of the following:

```
(vector unsigned int)(10) /* initializes all four elements to a value of 10 */
(vector unsigned int)(14, 82, 73, 700) /* initializes the first element
to 14, the second element to 82,
the third element to 73, and the
fourth element to 700 */
(vector unsigned int){14, 82, 73, 700} /* This syntax is also supported for the SPU */
```

A vector literal can be cast to another vector type. A vector literal cast does not change the bit pattern of the operand: the 128 bits representing the value remains the same before and after the cast.

Related information

- “Vector types” on page 57
- “Initialization of vectors” on page 97


Character literals

A *character literal* contains a sequence of characters or escape sequences enclosed in single quotation mark symbols, for example 'c'. A character literal may be prefixed with the letter L, for example L'c'. A character literal without the L prefix is an *ordinary character literal* or a *narrow character literal*. A character literal with the L prefix is a *wide character literal*. An ordinary character literal that contains more than one character or escape sequence (excluding single quotes ('), backslashes (\) or new-line characters) is a *multicharacter literal*.

C The type of a narrow character literal is `int`. The type of a wide character literal is `wchar_t`. The type of a multicharacter literal is `int`.

C++ The type of a character literal that contains only one character is `char`, which is an integral type. The type of a wide character literal is `wchar_t`. The type of a multicharacter literal is `int`.

Character literal syntax



At least one character or escape sequence must appear in the character literal, and the character literal must appear on a single logical source line.

The characters can be from the source program character set. You can represent the double quotation mark symbol by itself, but to represent the single quotation mark symbol, you must use the backslash symbol followed by a single quotation mark symbol (\' escape sequence). (See “Escape sequences” on page 33 for a list of other characters that are represented by escape characters.)

Outside of the basic source character set, the universal character names for letters and digits are allowed in C++ and at the C99 language level. **C++** In C++, you must compile with the `-q[lang]l=ucs` option for universal character name support.

The following are examples of character literals:

```
'a'
'\ '
L'0'
'('
```

Related information

- “Source program character set” on page 32
- “The Unicode standard” on page 34
- “Character types” on page 55

String literals


A *string literal* contains a sequence of characters or escape sequences enclosed in double quotation mark symbols. A string literal with the prefix L is a *wide string literal*. A string literal without the prefix L is an *ordinary* or *narrow string literal*.

► **C** The type of a narrow string literal is array of `char`. The type of a wide string literal is array of `wchar_t`.

► **C++** The type of a narrow string literal is array of `const char`. The type of a wide string literal is array of `const wchar_t`. Both types have static storage duration.

A null (`'\0'`) character is appended to each string. For a wide string literal, the value `'\0'` of type `wchar_t` is appended. By convention, programs recognize the end of a string by finding the null character.

String literal syntax



Multiple spaces contained within a string literal are retained.

Use the escape sequence `\n` to represent a new-line character as part of the string. Use the escape sequence `\\` to represent a backslash character as part of the string. You can represent a single quotation mark symbol either by itself or with the escape sequence `\'`. You must use the escape sequence `\"` to represent a double quotation mark.

Outside of the basic source character set, the universal character names for letters and digits are allowed in C++ and at the C99 language level. ► **C++** In C++, you must compile with the `-qflagl=ucs` option for universal character name support.

The following are examples of string literals:

```
char titles[ ] = "Handel's \"Water Music\"";
char *temp_string = "abc" "def" "ghi"; /* *temp_string = "abcdefghi\0" */
wchar_t *wide_string = L"longstring";
```

This example illustrates escape sequences in string literals:

```
#include <iostream> using namespace std;

int main () {
    char *s = "Hi there! \n";
    cout << s;
    char *p = "The backslash character \\";
    cout << p << endl;
    char *q = "The double quotation mark \".\n";
    cout << q ;
}
```

This program produces the following output:

```
Hi there! The backslash character \. The double quotation mark ".
```

To continue a string on the next line, use the line continuation character (`\` symbol) followed by optional whitespace and a new-line character (required). For example:

```
char *mail_addr = "Last Name   First Name   MI   Street Address \
                   893   City   Province   Postal code ";
```

In the following example, the string literal second causes a compile-time error.

```

char *first = "This string continues onto the next\
             line, where it ends.";           /* compiles successfully. */

char *second = "The comment makes the \
               */ invisible to the compiler."; /* continuation symbol
                                               /* compilation error.      */

```

Note: When a string literal appears more than once in the program source, how that string is stored depends on whether strings are read-only or writeable. By default, the compiler considers strings to be read-only. XL C/C++ may allocate only one location for a read-only string; all occurrences will refer to that one location. However, that area of storage is potentially write-protected. If strings are writeable, then each occurrence of the string will have a separate, distinct storage location that is always modifiable. You can use the **#pragma strings** directive or the **-qro** compiler option to change the default storage for string literals.

Related information

- “Character types” on page 55
- “Source program character set” on page 32
- “The Unicode standard” on page 34
- **-qro** in the *XL C/C++ Compiler Reference*
- **#pragma strings** in the *XL C/C++ Compiler Reference*
- **-qsource** in the *XL C/C++ Compiler Reference*
- **-qmacpstr** in the *XL C/C++ Compiler Reference*

String concatenation: Another way to continue a string is to have two or more consecutive strings. Adjacent string literals will be concatenated to produce a single string. For example:

```

"hello " "there"    /* is equivalent to "hello there" */
"hello" "there"    /* is equivalent to "hellothere"      */

```

Characters in concatenated strings remain distinct. For example, the strings `"\xab"` and `"3"` are concatenated to form `"\xab3 "`. However, the characters `\xab` and `3` remain distinct and are not merged to form the hexadecimal character `\xab3`.

If a wide string literal and a narrow string literal are adjacent, as in the following:

```
"hello " L"there"
```

the result is a wide string literal.

Following any concatenation, `'\0'` of type `char` is appended at the end of each string. For a wide string literal, `'\0'` of type `wchar_t` is appended. C++ programs find the end of a string by scanning for this value. For example:

```

char *first = "Hello ";           /* stored as "Hello \0" */
char *second = "there";          /* stored as "there\0" */
char *third = "Hello " "there";  /* stored as "Hello there\0" */

```

Punctuators and operators

A *punctuator* is a token that has syntactic and semantic meaning to the compiler, but the exact significance depends on the context. A punctuator can also be a token that is used in the syntax of the preprocessor.

C99 and C++ define the following tokens as punctuators, operators, or preprocessing tokens:

Table 11. C and C++ punctuators

[]	()	{ }	,	:	;
*	=	...	#		
.	->	++	--	##	
&	+	-	~	!	
/	%	<<	>>	!=	
<	>	<=	>=	==	
^		&&		?	
*=	/=	%=	+=	-=	
<<=	>>=	&=	^=	=	

C++ only

In addition to the C99 preprocessing tokens, operators, and punctuators, C++ allows the following tokens as punctuators:

Table 12. C++ punctuators

::	.*	->*	new	delete	
and	and_eq	bitand	bitor	comp	
not	not_eq	or	or_eq	xor	xor_eq

End of C++ only

Related information

- Chapter 6, “Expressions and operators,” on page 121

Alternative tokens

Both C and C++ provide the following alternative representations for some operators and punctuators. The alternative representations are also known as *digraphs*.

Operator or punctuator	Alternative representation
{	<%
}	%>
[<:
]	:>
#	%:
##	%.%:

In addition to the operators and punctuators listed above, C++ and C at the C99 language level provide the following alternative representations. In C, they are defined as macros in the header file `iso646.h`.

Operator or punctuator	Alternative representation
&&	and
	bitor
	or
^	xor
~	compl
&	bitand

Operator or punctuation	Alternative representation
&=	and_eq
=	or_eq
^=	xor_eq
!	not
!=	not_eq

Related information

- “Digraph characters” on page 36

Source program character set

The following lists the basic source character sets that are available at both compile time and run time:

- The uppercase and lowercase letters of the English alphabet:
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
- The decimal digits:
0 1 2 3 4 5 6 7 8 9
- The following graphic characters:
! " # % & ' () * + , - . / : ; < = > ? [\] _ { } ~
– The caret (^) character in ASCII (bitwise exclusive OR symbol).
– The split vertical bar (|) character in ASCII.
- The space character
- The control characters representing new-line, horizontal tab, vertical tab, form feed, end of string (NULL character), alert, backspace, and carriage return.

 Depending on the compiler option, other specialized identifiers, such as the dollar sign (\$) or characters in national character sets, may be allowed to appear in an identifier.

Related information

- “Characters in identifiers” on page 16

Multibyte characters

The compiler recognizes and supports the additional characters (the extended character set) which you can meaningfully use in string literals and character constants. The support for extended characters includes *multibyte character* sets. A *multibyte character* is a character whose bit representation fits into more than one byte. To instruct the compiler to recognize multibyte character sets as source input, be sure to compile with the **-qmbcs** option.

Multibyte characters can appear in any of the following contexts:

- String literals and character constants. To declare a multibyte literal, use a wide-character representation, prefixed by L. For example:

```
wchar_t *a = L"wide_char_string";
wchar_t b = L'wide_char';
```

Strings containing multibyte characters are treated essentially the same way as strings without multibyte characters. Generally, wide characters are permitted

anywhere multibyte characters are, but they are incompatible with multibyte characters in the same string because their bit patterns differ. Wherever permitted, you can mix single-byte and multibyte characters in the same string.

- Preprocessor directives. The following preprocessor directives permit multibyte-character constants and string literals:

- #define
- #pragma comment
- #include

A file name specified in an #include directive can contain multibyte characters. For example:

```
#include <multibyte_char/mydir/mysource/multibyte_char.h>
#include "multibyte_char.h"
```

- Macro definitions. Because string literals and character constants can be part of #define statements, multibyte characters are also permitted in both object-like and function-like macro definitions.
- The # and ## operators
- Program comments

The following are restrictions on the use of multibyte characters:

- Multibyte characters are not permitted in identifiers.
- Hexadecimal values for multibyte characters must be in the range of the code page being used.
- You cannot mix wide characters and multibyte characters in macro definitions. For example, a macro expansion that concatenates a wide string and a multibyte string is not permitted.
- Assignment between wide characters and multibyte characters is not permitted.
- Concatenating wide character strings and multibyte character strings is not permitted.

Related information

- “Character literals” on page 28
- “The Unicode standard” on page 34
- “Character types” on page 55
- -qmbcs in the *XL C/C++ Compiler Reference*

Escape sequences

You can represent any member of the execution character set by an *escape sequence*. They are primarily used to put nonprintable characters in character and string literals. For example, you can use escape sequences to put such characters as tab, carriage return, and backspace into an output stream.

Escape character syntax



An escape sequence contains a backslash (\) symbol followed by one of the escape sequence characters or an octal or hexadecimal number. A hexadecimal escape sequence contains an x followed by one or more hexadecimal digits (0-9, A-F, a-f).

An octal escape sequence uses up to three octal digits (0-7). The value of the hexadecimal or octal number specifies the value of the desired character or wide character.

Note: The line continuation sequence (`\` followed by a new-line character) is not an escape sequence. It is used in character strings to indicate that the current line of source code continues on the next line.

The escape sequences and the characters they represent are:

Escape sequence	Character represented
<code>\a</code>	Alert (bell, alarm)
<code>\b</code>	Backspace
<code>\f</code>	Form feed (new page)
<code>\n</code>	New-line
<code>\r</code>	Carriage return
<code>\t</code>	Horizontal tab
<code>\v</code>	Vertical tab
<code>\'</code>	Single quotation mark
<code>\"</code>	Double quotation mark
<code>\?</code>	Question mark
<code>\\</code>	Backslash

The value of an escape sequence represents the member of the character set used at run time. Escape sequences are translated during preprocessing. For example, on a system using the ASCII character codes, the value of the escape sequence `\x56` is the letter V. On a system using EBCDIC character codes, the value of the escape sequence `\xE5` is the letter V.

Use escape sequences only in character constants or in string literals. An error message is issued if an escape sequence is not recognized.

In string and character sequences, when you want the backslash to represent itself (rather than the beginning of an escape sequence), you must use a `\\` backslash escape sequence. For example:

```
cout << "The escape sequence \\n." << endl;
```

This statement results in the following output:

```
The escape sequence \n.
```

The Unicode standard

The *Unicode Standard* is the specification of an encoding scheme for written characters and text. It is a universal standard that enables consistent encoding of multilingual text and allows text data to be interchanged internationally without conflict. The ISO standards for C and C++ refer to *Information technology – Programming Languages – Universal Multiple-Octet Coded Character Set (UCS), ISO/IEC 10646:2003*. (The term *octet* is used by ISO to refer to a byte.) The ISO/IEC 10646 standard is more restrictive than the Unicode Standard in the number of encoding forms: a character set that conforms to ISO/IEC 10646 is also conformant to the Unicode Standard.

The Unicode Standard specifies a unique numeric value and name for each character and defines three encoding forms for the bit representation of the numeric value. The name/value pair creates an identity for a character. The

hexadecimal value representing a character is called a *code point*. The specification also describes overall character properties, such as case, directionality, alphabetic properties, and other semantic information for each character. Modeled on ASCII, the Unicode Standard treats alphabetic characters, ideographic characters, and symbols, and allows implementation-defined character codes in reserved code point ranges. According to the Unicode Standard, the encoding scheme of the standard is therefore sufficiently flexible to handle all known character encoding requirements, including coverage of all the world's historical scripts.

C99 and C++ allow the universal character name construct defined in ISO/IEC 10646 to represent characters outside the basic source character set. Both languages permit universal character names in identifiers, character constants, and string literals.  In C++, you must compile with the `-qclanglvl=ucs` option for universal character name support.

The following table shows the generic universal character name construct and how it corresponds to the ISO/IEC 10646 short name.

Universal character name	ISO/IEC 10646 short name
<i>where N is a hexadecimal digit</i>	
<code>\UNNNNNNNN</code>	NNNNNNNN
<code>\uNNNN</code>	000NNNN

C99 and C++ disallow the hexadecimal values representing characters in the basic character set (base source code set) and the code points reserved by ISO/IEC 10646 for control characters.

The following characters are also disallowed:

- Any character whose short identifier is less than 00A0. The exceptions are 0024 (\$), 0040 (@), or 0060 (').
- Any character whose short identifier is in the code point range D800 through DFFF inclusive.

UTF literals

IBM extension

The C Standards Committee has approved the implementation of *u-literals* and *U-literals* to support Unicode UTF-16 and UTF-32 character literals, respectively. To enable support for UTF literals in your source code, you must compile with the option `-qutf` enabled. The following table shows the syntax for UTF literals.

Table 13. UTF literals

Syntax	Explanation
<code>u'character'</code>	Denotes a UTF-16 character.
<code>u"character-sequence"</code>	Denotes an array of UTF-16 characters.
<code>U'character'</code>	Denotes a UTF-32 character.
<code>U"character-sequence"</code>	Denotes an array of UTF-32 characters.

XL C/C++ implements the macros `uint_least16_t` and `uint_least32_t`, which are defined in the header file `stdint.h`, as data types for UTF-16 and UTF-32

characters. The following example defines an array of `uint_least16_t`, including the characters represented by code points 1234 and 8180:

```
#include <stdint.h>

uint_least16_t msg[] = u"ucs characters \u1234 and \u8180 ";
```

Related information

- `-qutf` in the *XL C/C++ Compiler Reference*

String concatenation of u-literals: The u-literals and U-literals follow the same concatenation rule as wide character literals: the normal character string is widened if they are present. The following shows the allowed combinations. All other combinations are invalid.

Combination	Result
<code>u"a" u"b"</code>	<code>u"ab"</code>
<code>u"a" "b"</code>	<code>u"ab"</code>
<code>"a" u"b"</code>	<code>u"ab"</code>
<code>U"a" U"b"</code>	<code>U"ab"</code>
<code>U"a" "b"</code>	<code>U"ab"</code>
<code>"a" U"b"</code>	<code>U"ab"</code>

Multiple concatenations are allowed, with these rules applied recursively.

Related information

- “String concatenation” on page 30

End of IBM extension

Digraph characters

You can represent unavailable characters in a source program by using a combination of two keystrokes that are called a *digraph character*. The preprocessor reads digraphs as tokens during the preprocessor phase.

The digraph characters are:

<code>%:</code> or <code>%%</code>	<code>#</code>	number sign
<code><</code>	<code>[</code>	left bracket
<code>:</code> <code>></code>	<code>]</code>	right bracket
<code><%</code>	<code>{</code>	left brace
<code>%></code>	<code>}</code>	right brace
<code>%%:</code> or <code>%%%%</code>	<code>##</code>	preprocessor macro concatenation operator

You can create digraphs by using macro concatenation. XL C/C++ does not replace digraphs in string literals or in character literals. For example:

```
char *s = "<%%>"; // stays "<%%>"

switch (c) {
  case '<%': { /* ... */ } // stays '<% '
  case '%>': { /* ... */ } // stays '%> '
}
```

Related information

- `-qdigraph` in the *XL C/C++ Compiler Reference*

Trigraph sequences

Some characters from the C and C++ character set are not available in all environments. You can enter these characters into a C or C++ source program using a sequence of three characters called a *trigraph*. The trigraph sequences are:

Trigraph	Single character	Description
??=	#	pound sign
??([left bracket
??)]	right bracket
??<	{	left brace
??>	}	right brace
??/	\	backslash
??'	^	caret
??!		vertical bar
??-	~	tilde

The preprocessor replaces trigraph sequences with the corresponding single-character representation. For example,

```
some_array??(i??) = n;
```

Represents:

```
some_array[i] = n;
```

Comments

A *comment* is text replaced during preprocessing by a single space character; the compiler therefore ignores all comments.

There are two kinds of comments:

- The `/*` (slash, asterisk) characters, followed by any sequence of characters (including new lines), followed by the `*/` characters. This kind of comment is commonly called a *C-style comment*.
- The `//` (two slashes) characters followed by any sequence of characters. A new line not immediately preceded by a backslash terminates this form of comment. This kind of comment is commonly called a *single-line comment* or a *C++ comment*. A C++ comment can span more than one physical source line if it is joined into one logical source line with line-continuation (`\`) characters. The backslash character can also be represented by a trigraph.

You can put comments anywhere the language allows white space. You cannot nest C-style comments inside other C-style comments. Each comment ends at the first occurrence of `*/`.

You can also include multibyte characters; to instruct the compiler to recognize multibyte characters in the source code, compile with the `-qmbcs` option.

Note: The `/*` or `*/` characters found in a character constant or string literal do not start or end comments.

In the following program, the second `printf()` is a comment:

```

#include <stdio.h>

int main(void)
{
    printf("This program has a comment.\n");
    /* printf("This is a comment line and will not print.\n"); */
    return 0;
}

```

Because the second `printf()` is equivalent to a space, the output of this program is:

This program has a comment.

Because the comment delimiters are inside a string literal, `printf()` in the following program is not a comment.

```

#include <stdio.h>

int main(void)
{
    printf("This program does not have \
/* NOT A COMMENT */ a comment.\n");
    return 0;
}

```

The output of the program is:

This program does not have
/* NOT A COMMENT */ a comment.

In the following example, the comments are highlighted:

/* A program with nested comments. */

```

#include <stdio.h>

int main(void)
{
    test_function();
    return 0;
}

int test_function(void)
{
    int number;
    char letter;
    /*
    number = 55;
    letter = 'A';
    /* number = 44; */
    */
    return 999;
}

```

In `test_function`, the compiler reads the first `/*` through to the first `*/`. The second `*/` causes an error. To avoid commenting over comments already in the source code, you should use conditional compilation preprocessor directives to cause the compiler to bypass sections of a program. For example, instead of commenting out the above statements, change the source code in the following way:

```

/* A program with conditional compilation to avoid nested comments. */

#define TEST_FUNCTION 0
#include <stdio.h>

int main(void)

```

```

{
    test_function();
    return 0;
}

int test_function(void)
{
    int number;
    char letter;
    #if TEST_FUNCTION
        number = 55;
        letter = 'A';
        /*number = 44;*/
    #endif /*TEST_FUNCTION */
}

```

You can nest single line comments within C-style comments. For example, the following program will not output anything:

```
#include <stdio.h>
```

```

int main(void)
{
    /*
    printf("This line will not print.\n");
    // This is a single line comment
    // This is another single line comment
    printf("This line will also not print.\n");
    */
    return 0;
}

```

Note: You can also use the **#pragma comment** directive to place comments into an object module.

Related information

- **-qmbcs**, **-qlanglvl**, and **-qcpluscmt** in the *XL C/C++ Compiler Reference*
- “Multibyte characters” on page 32

Chapter 3. Data objects and declarations

This section discusses the various elements that constitute a declaration of a data object, and includes the following topics:

- “Overview of data objects and declarations”
- “Storage class specifiers” on page 45
- “Type specifiers” on page 52
- “User-defined types” on page 59
- “Type qualifiers” on page 72
- “Type attributes” on page 78

Topics are sequenced to loosely follow the order in which elements appear in a declaration. The discussion of the additional elements of data declarations is also continued in Chapter 4, “Declarators,” on page 83.

Overview of data objects and declarations

The following sections introduce some fundamental concepts regarding data objects and data declarations that will be used throughout this reference.

Overview of data objects

A data *object* is a region of storage that contains a value or group of values. Each value can be accessed using its identifier or a more complex expression that refers to the object. In addition, each object has a unique *data type*. The data type of an object determines the storage allocation for that object and the interpretation of the values during subsequent access. It is also used in any type checking operations. Both the identifier and data type of an object are established in the object *declaration*.

► **C++** An instance of a class type is commonly called a *class object*. The individual class members are also called objects. The set of all member objects comprises a class object.

Data types are often grouped into type categories that overlap, such as:

Fundamental types versus derived types

Fundamental data types are also known as “basic”, “fundamental” or “built-in” to the language. These include integers, floating-point numbers, and characters. *Derived* types, also known as “compound” types in Standard C++, are created from the set of basic types, and include arrays, pointers, structures, unions, and enumerations. All C++ classes are considered compound types.

Built-in types versus user-defined types

Built-in data types include all of the fundamental types, plus types that refer to the addresses of basic types, such as arrays and pointers. *User-defined* types are created by the user from the set of basic types, in typedef, structure, union, and enumeration definitions. C++ classes are considered user-defined types.

Scalar types versus aggregate types

Scalar types represent a single data value, while *aggregate* types represent

multiple values, of the same type or of different types. Scalars include the arithmetic types and pointers. Aggregate types include arrays and structures. C++ classes are considered aggregate types.

The following matrix lists the supported data types and their classification into fundamental, derived, scalar, and aggregate types.

Table 14. C/C++ data types

Data object	Basic	Compound	Built-in	User-defined	Scalar	Aggregate
integer types	+		+		+	
floating-point types ¹	+		+		+	
character types			+		+	
Booleans	+		+		+	
void type	+ ²		+		+	
pointers		+	+		+	
arrays		+	+			+
structures		+		+		+
unions		+		+		
enumerations		+		+	see note ³	
 classes		+		+		+
typedef types		+		+		

Notes:

1. Although complex floating-point types are represented internally as an array of two elements, they behave in the same way as real floating-pointing types in terms of alignment and arithmetic operations, and can therefore be considered scalar types.
2. The void type is really an incomplete type, as discussed in “Incomplete types.” Nevertheless, Standard C++ defines it as a fundamental type.
3.  The C standard does not classify enumerations as either scalar or aggregate.  Standard C++ classifies enumerations as scalars.

Related information

- Chapter 11, “Classes (C++ only),” on page 259

Incomplete types

The following are incomplete types:

- The void type
- Arrays of unknown size
- Arrays of elements that are of incomplete type
- Structure, union, or enumerations that have no definition
-  Pointers to class types that are declared but not defined
-  Classes that are declared but not defined

However, if an array size is specified by `[*]`, indicating a variable length array, the size is considered as having been specified, and the array type is then considered a complete type. For more information, see “Variable length arrays ” on page 91.

The following examples illustrate incomplete types:

```
void *incomplete_ptr;
struct dimension linear; /* no previous definition of dimension */
```

Related information

- “The void type” on page 56
- “Incomplete class declarations” on page 264

Compatible and composite types

C only

In C, compatible types are defined as:

- two types that can be used together without modification (as in an assignment expression)
- two types that can be substituted one for the other without modification

When two compatible types are combined, the result is a *composite type*. Determining the resultant composite type for two compatible types is similar to following the usual binary conversions of integral types when they are combined with some arithmetic operators.

Obviously, two types that are identical are compatible; their composite type is the same type. Less obvious are the rules governing type compatibility of non-identical types, user-defined types, type-qualified types, and so on. “Type specifiers” on page 52 discusses compatibility for basic and user-defined types in C.

End of C only

C++ only

A separate notion of type compatibility as distinct from being of the same type does not exist in C++. Generally speaking, type checking in C++ is stricter than in C: identical types are required in situations where C would only require compatible types.

End of C++ only

Related information

- “Compatibility of arrays ” on page 92
- “Compatibility of pointers (C only)” on page 89
- “Compatible functions” on page 204

Overview of data declarations and definitions

A *declaration* establishes the names and characteristics of data objects used in a program. A *definition* allocates storage for data objects, and associates an identifier with that object. When you declare or define a *type*, no storage is allocated.

The following table shows examples of declarations and definitions. The identifiers declared in the first column do not allocate storage; they refer to a corresponding definition. The identifiers declared in the second column allocate storage; they are both declarations and definitions.

Declarations	Declarations and definitions
extern double pi;	double pi = 3.14159265;
struct payroll;	struct payroll { char *name; float salary; } employee;

Note:  The C99 standard no longer requires that all declarations appear at the beginning of a function before the first statement. As in C++, you can mix declarations with other statements in your code.

Declarations determine the following properties of data objects and their identifiers:

- Scope, which describes the region of program text in which an identifier can be used to access its object
- Visibility, which describes the region of program text from which legal access can be made to the identifier's object
- Duration, which defines the period during which the identifiers have real, physical objects allocated in memory
- Linkage, which describes the correct association of an identifier to one particular object
- Type, which determines how much memory is allocated to an object and how the bit patterns found in the storage allocation of that object should be interpreted by the program

The elements of a declaration for a data object are as follows:

- Storage class specifiers, which specify storage duration and linkage
- Type specifiers, which specify data types
- Type qualifiers, which specify the mutability of data values
- Declarators, which introduce and include identifiers
- Initializers, which initialize storage with initial values


IBM extension

In addition, for compatibility with GCC, XL C/C++ allows you to use *attributes* to modify the properties of data objects. *Type attributes*, which can be used to modify the definition of user-defined types, are described in “Type attributes” on page 78. *Variable attributes*, which can be used to modify the declaration of variables, are described in “Variable attributes” on page 105.

End of IBM extension

All declarations have the form:

Data declaration syntax





Related information

- “Function declarations and definitions” on page 201

Tentative definitions

C only

A *tentative definition* is any external data declaration that has no storage class specifier and no initializer. A tentative definition becomes a full definition if the end of the translation unit is reached and no definition has appeared with an initializer for the identifier. In this situation, the compiler reserves uninitialized space for the object defined.

The following statements show normal definitions and tentative definitions.

```
int i1 = 10;          /* definition, external linkage */
static int i2 = 20;  /* definition, internal linkage */
extern int i3 = 30;  /* definition, external linkage */
int i4;              /* tentative definition, external linkage */
static int i5;       /* tentative definition, internal linkage */

int i1;              /* valid tentative definition */
int i2;              /* not legal, linkage disagreement with previous */
int i3;              /* valid tentative definition */
int i4;              /* valid tentative definition */
int i5;              /* not legal, linkage disagreement with previous */
```

End of C only

C++ only

C++ does not support the concept of a tentative definition: an external data declaration without a storage class specifier is always a definition.

End of C++ only

Storage class specifiers

A storage class specifier is used to refine the declaration of a variable, a function, and parameters. Storage classes determine whether:

- The object has internal, external, or no linkage
- The object is to be stored in memory or in a register, if available
- The object receives the default initial value of 0 or an indeterminate default initial value
- The object can be referenced throughout a program or only within the function, block, or source file where the variable is defined
- The storage duration for the object is maintained throughout program run time or only during the execution of the block where the object is defined

For a variable, its default storage duration, scope, and linkage depend on where it is declared: whether inside or outside a block statement or the body of a function. When these defaults are not satisfactory, you can use a storage class specifier to explicitly set its storage class. The storage class specifiers in C and C++ are:

- auto
- static
- extern
-  mutable
- register
-  __thread

Related information

- “Function storage class specifiers” on page 206
- “Initializers” on page 93

The auto storage class specifier

The auto storage class specifier lets you explicitly declare a variable with *automatic storage*. The auto storage class is the default for variables declared inside a block. A variable *x* that has automatic storage is deleted when the block in which *x* was declared exits.

You can only apply the auto storage class specifier to names of variables declared in a block or to names of function parameters. However, these names by default have automatic storage. Therefore the storage class specifier `auto` is usually redundant in a data declaration.

Storage duration of automatic variables

Objects with the auto storage class specifier have automatic storage duration. Each time a block is entered, storage for auto objects defined in that block is made available. When the block is exited, the objects are no longer available for use. An object declared with no linkage specification and without the `static` storage class specifier has automatic storage duration.

If an auto object is defined within a function that is recursively invoked, memory is allocated for the object at each invocation of the block.

Linkage of automatic variables

An auto variable has block scope and no linkage.

Related information

- “Initialization and storage classes” on page 94
- “Block statements” on page 178
- “The goto statement” on page 192

The static storage class specifier

Objects declared with the `static` storage class specifier have *static storage duration*, which means that memory for these objects is allocated when the program begins running and is freed when the program terminates. Static storage duration for a variable is different from file or global scope: a variable can have static duration but local scope.

▶ **C** The keyword `static` is the major mechanism in C to enforce information hiding.

▶ **C++** C++ enforces information hiding through the namespace language feature and the access control of classes. The use of the keyword `static` to limit the scope of external variables is deprecated for declaring objects in namespace scope.

The `static` storage class specifier can be applied to the following declarations:

- Data objects
- ▶ **C++** Class members
- Anonymous unions

You cannot use the `static` storage class specifier with the following:

- Type declarations
- Function parameters

C only

At the C99 language level, the `static` keyword can be used in the declaration of an array parameter to a function. The `static` keyword indicates that the argument passed into the function is a pointer to an array of at least the specified size. In this way, the compiler is informed that the pointer argument is never null. See “Static array indices in function parameter declarations (C only)” on page 217 for more information.

End of C only

Related information

- “The `static` storage class specifier” on page 206
- “Static members” on page 279

Linkage of static variables

A declaration of an object that contains the `static` storage class specifier and has file scope gives the identifier internal linkage. Each instance of the particular identifier therefore represents the same object within one file only. For example, if a static variable `x` has been declared in function `f`, when the program exits the scope of `f`, `x` is not destroyed:

```
#include <stdio.h>

int f(void) {
    static int x = 0;
    x++;
    return x;
}

int main(void) {
    int j;
    for (j = 0; j < 5; j++) {
        printf("Value of f(): %d\n", f());
    }
    return 0;
}
```

The following is the output of the above example:

```
Value of f(): 1
Value of f(): 2
Value of f(): 3
Value of f(): 4
Value of f(): 5
```

Because `x` is a static variable, it is not reinitialized to 0 on successive calls to `f`.

Related information

- “Initialization and storage classes” on page 94
- “Internal linkage” on page 8
- Chapter 9, “Namespaces (C++ only),” on page 235

The extern storage class specifier

The extern storage class specifier lets you declare objects that several source files can use. An extern declaration makes the described variable usable by the succeeding part of the current source file. This declaration does not replace the definition. The declaration is used to describe the variable that is externally defined.

An extern declaration can appear outside a function or at the beginning of a block. If the declaration describes a function or appears outside a function and describes an object with external linkage, the keyword `extern` is optional.

If a declaration for an identifier already exists at file scope, any extern declaration of the same identifier found within a block refers to that same object. If no other declaration for the identifier exists at file scope, the identifier has external linkage.

C++ C++ restricts the use of the extern storage class specifier to the names of objects or functions. Using the extern specifier with type declarations is illegal. An extern declaration cannot appear in class scope.

Storage duration of external variables

All extern objects have static storage duration. Memory is allocated for extern objects before the `main` function begins running, and is freed when the program terminates. The scope of the variable depends on the location of the declaration in the program text. If the declaration appears within a block, the variable has block scope; otherwise, it has file scope.

Linkage of external variables

C Like the scope, the linkage of a variable declared `extern` depends on the placement of the declaration in the program text. If the variable declaration appears outside of any function definition and has been declared `static` earlier in the file, the variable has internal linkage; otherwise, it has external linkage in most cases. All object declarations that occur outside a function and that do not contain a storage class specifier declare identifiers with external linkage.

C++ For objects in the unnamed namespace, the linkage may be external, but the name is unique, and so from the perspective of other translation units, the name effectively has internal linkage.

Related information

- “External linkage” on page 8
- “Initialization and storage classes” on page 94

- “The extern storage class specifier” on page 206
- “Class scope (C++ only)” on page 5
- Chapter 9, “Namespaces (C++ only),” on page 235

The mutable storage class specifier (C++ only)

The mutable storage class specifier is used only on a class data member to make it modifiable even though the member is part of an object declared as `const`. You cannot use the mutable specifier with names declared as `static` or `const`, or reference members.

In the following example:

```
class A
{
    public:
        A() : x(4), y(5) { };
        mutable int x;
        int y;
};

int main()
{
    const A var2;
    var2.x = 345;
    // var2.y = 2345;
}
```

the compiler would not allow the assignment `var2.y = 2345` because `var2` has been declared as `const`. The compiler will allow the assignment `var2.x = 345` because `A::x` has been declared as `mutable`.

Related information

- “Type qualifiers” on page 72
- “References (C++ only)” on page 93

The register storage class specifier

The register storage class specifier indicates to the compiler that the object should be stored in a machine register. The register storage class specifier is typically specified for heavily used variables, such as a loop control variable, in the hopes of enhancing performance by minimizing access time. However, the compiler is not required to honor this request. Because of the limited size and number of registers available on most systems, few variables can actually be put in registers. If the compiler does not allocate a machine register for a register object, the object is treated as having the storage class specifier `auto`.

An object having the register storage class specifier must be defined within a block or declared as a parameter to a function.

The following restrictions apply to the register storage class specifier:

-  You cannot use pointers to reference objects that have the register storage class specifier.
-  You cannot use the register storage class specifier when declaring objects in global scope.
-  A register does not have an address. Therefore, you cannot apply the address operator (`&`) to a register variable.

- **C++** You cannot use the register storage class specifier when declaring objects in namespace scope.

C++ only

Unlike C, C++ lets you take the address of an object with the register storage class. For example:

```
register int i;
int* b = &i;    // valid in C++, but not in C
```

End of C++ only

Storage duration of register variables

Objects with the register storage class specifier have automatic storage duration. Each time a block is entered, storage for register objects defined in that block is made available. When the block is exited, the objects are no longer available for use.

If a register object is defined within a function that is recursively invoked, memory is allocated for the variable at each invocation of the block.

Linkage of register variables

Since a register object is treated as the equivalent to an object of the auto storage class, it has no linkage.

Related information

- “Initialization and storage classes” on page 94
- “Block/local scope” on page 2
- “References (C++ only)” on page 93

Variables in specified registers


IBM extension

You can specify that a particular hardware register is dedicated to a variable by using an *asm register variable* declaration. This language extension is provided for compatibility with GNU C.

Global register variables reserve registers throughout the program; stores into the reserved register are never deleted.

Local register variables do not actually reserve the registers, except when the variables are used as input or output operands in an inline assembly statement. In this case, using the variable as an *asm* operand guarantees that the specified register is used for the operand and is a convenient way to control which register is used.

Register variable declaration syntax



The *register_specifier* is a string representing a hardware register. The register name is CPU-specific. The following are valid register names for the PPU:

r0 to r31
General purpose registers

f0 to f31
Floating-point registers

v0 to v31
Vector registers

Note: You cannot use registers r30 or r14 for register variable declarations, as these are reserved for use by the XL C++ runtime environment. For the SPU, valid register names are 0 to 127, or v0 to v127.

The following are the rules of use for register variables:

- On the PPU, you cannot reserve registers for the following types of variables:
 - long long types
 - aggregate types
 -  class types
 -  reference types
 - void types
 - `_Complex` types
 - 128-bit long double types

On the SPU, you cannot reserve registers for the following types of variables:

- aggregate types
 -  class types
 -  reference types
 - void types
- General purpose registers can only be reserved for variables of integer or pointer type on the PPU.
 - Floating point registers can only be reserved for variables of float, double, or 64-bit long double type on the PPU.
 - Vector registers can only be reserved for variables of vector type on the PPU.
 - A global register variable cannot be initialized.
 - The register dedicated for a global register variable should not be a volatile register, or the value stored into the global variable might not be preserved across a function call.
 - More than one register variable can reserve the same register; however, the two variables become aliases of each other, and this is diagnosed with a warning.
 - The same global register variable cannot reserve more than one register.
 - The register specified in the global register declaration is reserved for the declared variable only in the compilation unit in which the register declaration is specified. The register is not reserved in other compilation units unless you place the global register declaration in a common header file, or use the `-qreserved_reg` compiler option.


Related information

- “Assembly labels” on page 17
- “Inline assembly statements” on page 194
- `-qreserved_reg` in the *XL C/C++ Compiler Reference*

The `__thread` storage class specifier (PPU only)

IBM extension

The `__thread` storage class marks a static variable as having *thread-local* storage duration. This means that, in a multi-threaded application, a unique instance of the variable is created for each thread that uses it, and destroyed when the thread terminates. The `__thread` storage class specifier can provide a convenient way of assuring thread-safety: declaring an object as per-thread allows multiple threads to access the object without the concern of race conditions, while avoiding the need for low-level programming of thread synchronization or significant program restructuring. The specifier can be applied to any of the following:

- global variables
- file-scoped static variables
- function-scoped static variables
-  static data members of a class

It cannot be applied to block-scoped automatic variables or non-static data members.

The specifier can appear on its own, or preceded immediately by either the `static` or `extern` specifier, as in the following examples:

```
__thread int i;
extern __thread struct state s;
static __thread char *p;
```

Variables marked with the `__thread` specifier can be initialized or uninitialized.

 `__thread` variables must be initialized with a constant expression, and must not have a static constructor.

Applying the address-of operator (`&`) to a thread-local variable returns the runtime address of the current thread's instance of the variable. That thread can pass this address to any other thread; however, when the first thread terminates, any pointers to its thread-local variables become invalid.

Related information

- `-qtls` in the *XL C/C++ Compiler Reference*

Type specifiers

Type specifiers indicate the type of the object being declared. The following are the available kinds of type specifiers:

- Fundamental or built-in types:
 - Arithmetic types
 - Integral types
 - Boolean types
 - Floating-point types

- Character types
- The void type
-  Vector types
- User-defined types.

Related information

- “Function return type specifiers” on page 213

Integral types

Integer types fall into the following categories:

- Signed integer types:
 - signed char
 - short int
 - int
 - long int
 - long long int
- Unsigned integer types:
 - unsigned char
 - unsigned short int
 - unsigned int
 - unsigned long int
 - unsigned long long int

The unsigned prefix indicates that the object is a nonnegative integer. Each unsigned type provides the same size storage as its signed equivalent. For example, `int` reserves the same storage as `unsigned int`. Because a signed type reserves a sign bit, an unsigned type can hold a larger positive integer value than the equivalent signed type.

The declarator for a simple integer definition or declaration is an identifier. You can initialize a simple integer definition with an integer constant or with an expression that evaluates to a value that can be assigned to an integer.

C++ only

When the arguments in overloaded functions and overloaded operators are integer types, two integer types that both come from the same group are not treated as distinct types. For example, you cannot overload an `int` argument against a signed `int` argument.

End of C++ only

Related information

- “Integer literals” on page 19
- “Integral conversions” on page 114
- “Arithmetic conversions and promotions” on page 113
- Chapter 10, “Overloading (C++ only),” on page 243

Boolean types

A Boolean variable can be used to hold the integer values 0 or 1, or the literals `true` or `false`, which are implicitly promoted to the integers 0 and 1 whenever an arithmetic value is necessary. The Boolean type is unsigned and has the lowest ranking in its category of standard unsigned integer types; it may not be further qualified by the specifiers `signed`, `unsigned`, `short`, or `long`. In simple assignments, if the left operand is a Boolean type, then the right operand must be either an arithmetic type or a pointer.

C only

Boolean types are a C99 feature. To declare a Boolean variable, use the `_Bool` type specifier.  The token `bool` is recognized as a keyword in C only when used in a vector declaration context.

End of C only

C++ only

To declare a Boolean variable in C++, use the `bool` type specifier. The result of the equality, relational, and logical operators is of type `bool`: either of the Boolean constants `true` or `false`.

End of C++ only

You can use Boolean types make *Boolean logic tests*. A Boolean logic test is used to express the results of a logical operation. For example:

```
_Bool f(int a, int b)
{
    return a==b;
}
```

If `a` and `b` have the same value, `f` returns `true`. If not, `f` returns `false`.

Related information

- “Boolean literals” on page 21
- “Boolean conversions” on page 114
- “Vector types” on page 57

Floating-point types

Floating-point type specifiers fall into the following categories:

- “Real floating-point types”
- Complex floating-point types

Real floating-point types

Generic, or binary, floating-point types consist of the following:

- `float`
- `double`
- `long double`

The magnitude ranges of the real floating-point types are given in the following table.

Table 15. Magnitude ranges of real floating-point types

Type	Range
float	approximately 1.2^{-38} to 3.4^{38}
double, long double	approximately 2.2^{-308} to 1.8^{308}

If a floating-point constant is too large or too small, the result is undefined by the language.

The declarator for a simple floating-point declaration is an identifier. Initialize a simple floating-point variable with a float constant or with a variable or expression that evaluates to an integer or floating-point number.

Related information

- “Floating-point literals” on page 22
- “Floating-point conversions” on page 114
- “Arithmetic conversions and promotions” on page 113

Complex floating-point types

The complex type specifiers are:

- `float _Complex`
- `double _Complex`
- `long double _Complex`

The representation and alignment requirements of a complex type are the same as an array type containing two elements of the corresponding real type. The real part is equal to the first element; the imaginary part is equal to the second element.

The equality and inequality operators have the same behavior as for real types. None of the relational operators may have a complex type as an operand.

 As an extension to C99 and Standard C++, complex numbers may also be operands to the unary operators `++` (increment), `--` (decrement), and `~` (bitwise negation).

Related information

- “Complex literals” on page 24
- “Arithmetic conversions and promotions” on page 113
- “The `__real__` and `__imag__` operators” on page 138

Character types

Character types fall into the following categories:

- Narrow character types:
 - `char`
 - `signed char`
 - `unsigned char`
- Wide character type `wchar_t`

The `char` specifier is an integral type. The `wchar_t` type specifier is an integral type that has enough storage to represent a wide character literal. (A wide character literal is a character literal that is prefixed with the letter `L`, for example `L'x'`)

C A `char` is a distinct type from signed `char` and unsigned `char`, and the three types are not compatible.

C++ For the purposes of distinguishing overloaded functions, a C++ `char` is a distinct type from signed `char` and unsigned `char`.

If it does not matter if a `char` data object is signed or unsigned, you can declare the object as having the data type `char`. Otherwise, explicitly declare signed `char` or unsigned `char` to declare numeric variables that occupy a single byte. When a `char` (signed or unsigned) is widened to an `int`, its value is preserved.

By default, `char` behaves like an unsigned `char`. To change this default, you can use the `-qchars` option or the `#pragma chars` directive. See `-qchars` in the *XL C/C++ Compiler Reference* for more information.

Related information

- “Character literals” on page 28
- “String literals” on page 28
- “Arithmetic conversions and promotions” on page 113

The void type

The `void` data type always represents an empty set of values. The only object that can be declared with the type specifier `void` is a pointer.

You cannot declare a variable of type `void`, but you can explicitly convert any expression to type `void`. The resulting expression can only be used as one of the following:

- An expression statement
- The left operand of a comma expression
- The second or third operand in a conditional expression.

Related information

- “Pointers” on page 86
- “Comma operator ,” on page 151
- “Conditional expressions” on page 153
- “Function declarations and definitions” on page 201

Compatibility of arithmetic types (C only)

Two arithmetic types are compatible only if they are the same type.

The presence of type specifiers in various combinations for arithmetic types may or may not indicate different types. For example, the type signed `int` is the same as `int`, except when used as the types of bit fields; but `char`, signed `char`, and unsigned `char` are different types.

The presence of a type qualifier changes the type. That is, `const int` is not the same type as `int`, and therefore the two types are not compatible.


Vector types

IBM extension

XL C/C++ supports vector processing technologies through language extensions. XL C/C++ implements the AltiVec Programming Interface specification with an extended syntax that allows type qualifiers and storage class specifiers to precede the keyword vector (or its alternate spelling, `__vector`) in a declaration.

Most of the legal forms of the syntax are captured in the following diagram. Some variations have been omitted from the diagram for the sake of clarity: type qualifiers such as `const` and storage class specifiers such as `static` can appear in any order within the declaration, as long as neither immediately follows the keyword vector (or `__vector`).

Vector declaration syntax



Notes:

1. The keyword `vector` is recognized in a declaration context only when used as a type specifier. The keywords `pixel`, `__pixel` and `bool` are recognized as valid type specifiers only when preceded by the keyword `vector` or `__vector`; they are only valid for compilation targeting the PPU.
2. The `long long`, `double`, and `qword` types are only valid for compilation targeting the SPU. The `qword` type is only used for specific SPU built-in functions; it is only recognized when you include the header `spu_intrinsics.h`.
3. The `long` type specifier is deprecated in a vector context, and is treated as an `int`. It is also only valid for compilation targeting the PPU.
4. Duplicate type specifiers are ignored in a vector declaration context. For compilation targeting the PPU, `long long` is treated as `long`.

The following table lists the supported vector data types and the size and possible values for each type.

Table 16. Vector data types

Type	Interpretation of content	Range of values
vector unsigned char	16 unsigned char	0..255
vector signed char	16 signed char	-128..127
vector bool char	16 unsigned char	0, 255
vector unsigned short	8 unsigned short	0..65535
vector unsigned short int		
vector signed short	8 signed short	-32768..32767
vector signed short int		
vector bool short	8 unsigned short	0, 65535
vector bool short int		
vector unsigned int	4 unsigned int	0..2 ³² -1
vector unsigned long		
vector unsigned long int		
vector signed int	4 signed int	-2 ³¹ ..2 ³¹ -1
vector signed long		
vector signed long int		
vector bool int	4 unsigned int	0, 2 ³² -1
vector bool long		
vector bool long int		
vector unsigned long long	2 unsigned long	0..2 ⁶⁴ -1
vector signed long long	2 signed long	-2 ⁶³ ..2 ⁶³ -1
vector float	4 float	For the PPU, IEEE-754 values; for the SPU, 1.17549435 * 10 ⁻³⁸ ..6.80564694 * 10 ⁺³⁸
vector double	2 double	IEEE-754 values
vector pixel	8 unsigned short	1/5/5/5 pixel
qword	1 quadword (16 bytes)	not applicable

For code targeting the SPU, if you include the `spu_intrinsics.h` header, you can also refer to the following data types by single-token typedef names:

Table 17. SPU single-token vector typedef names

Vector data type keyword	Single-token typedef name
vector unsigned char	<code>vec_uchar16</code>
vector signed char	<code>vec_char16</code>
vector unsigned short	<code>vec_ushort8</code>
vector signed short	<code>vec_short8</code>
vector unsigned int	<code>vec_uint4</code>
vector signed int	<code>vec_int4</code>
vector unsigned long long	<code>vec_ullong2</code>
vector signed long long	<code>vec_llong2</code>
vector float	<code>vec_float4</code>

Table 17. SPU single-token vector typedef names (continued)

Vector data type keyword	Single-token typedef name
vector double	vec_double2

All vector types are aligned on a 16-byte boundary. An aggregate that contains one or more vector types is aligned on a 16-byte boundary, and padded, if necessary, so that each member of vector type is also 16-byte aligned.

Related information

- “Vector literals” on page 25
- “Initialization of vectors” on page 97

Vector type casts

Vector types can be cast to other vector types. The cast does not perform a conversion: it preserves the 128-bit pattern, but not necessarily the value. A cast between a vector type and a scalar type is not allowed.

Vector pointers and pointers to non-vector types can be cast back and forth to each other. When a pointer to a non-vector type is cast to a vector pointer, the address should be 16-byte aligned. The referenced object of the pointer to a non-vector type can be aligned on a sixteen-byte boundary by using either the `__align` specifier or `__attribute__((aligned(16)))`.

Related information

- “The `__align` type qualifier” on page 74
- “The aligned variable attribute” on page 107

End of IBM extension

User-defined types

The following are user-defined types:

- Structures and unions
- Enumerations
- typedef definitions
-  Classes
-  Elaborated type specifiers

C++ classes are discussed in Chapter 11, “Classes (C++ only),” on page 259. Elaborated type specifiers are discussed in “Scope of class names” on page 263.

Related information

- “Type attributes” on page 78

Structures and unions

A *structure* contains an ordered group of data objects. Unlike the elements of an array, the data objects within a structure can have varied data types. Each data object in a structure is a *member* or *field*.

A *union* is an object similar to a structure except that all of its members start at the same location in memory. A union variable can represent the value of only one of its members at a time.

▶ **C++** In C++, structures and unions are the same as classes except that their members and inheritance are public by default.

You can declare a structure or union type separately from the definition of variables of that type, as described in “Structure and union type definition” and “Structure and union variable declarations” on page 65; or you can define a structure or union data type and all variables that have that type in one statement, as described in “Structure and union type and variable definitions in a single statement” on page 65.

Structures and unions are subject to alignment considerations. For a complete discussion of alignment, see “Aligning data” in the *XL C/C++ Programming Guide*.


Related information

- “Classes and structures” on page 262

Structure and union type definition

A structure or union *type definition* contains the `struct` or `union` keyword followed by an optional identifier (the structure tag) and a brace-enclosed list of members.

Structure or union type definition syntax



The *tag_identifier* gives a name to the type. If you do not provide a tag name, you must put all variable definitions that refer to the type within the declaration of the type, as described in “Structure and union type and variable definitions in a single statement” on page 65. Similarly, you cannot use a type qualifier with a structure or union definition; type qualifiers placed in front of the `struct` or `union` keyword can only apply to variables that are declared within the type definition.

Related information

- “The aligned type attribute” on page 79
- “The packed type attribute” on page 80

Member declarations

The list of members provides a structure or union data type with a description of the values that can be stored in the structure or union. The definition of a member has the form of a standard variable declaration. The names of member variables must be distinct within a single structure or union, but the same member name may be used in another structure or union type that is defined within the same scope, and may even be the same as a variable, function, or type name.

A structure or union member may be of any type except:

- any variably modified type
- any void type
- ▶ **C** a function

- any incomplete type

Because incomplete types are not allowed as members, a structure or union type may not contain an instance of itself as a member, but is allowed to contain a pointer to an instance of itself. As a special case, the last element of a structure with more than one member may have an incomplete array type, which is called a *flexible array member*, as described in “Flexible array members .”

IBM As an extension to Standard C and C++ for compatibility with GNU C, XL C/C++ also allows zero-extent arrays as members of structures and unions, as described in “Zero-extent array members” on page 62.

C++ A union member cannot be a class object that has a constructor, destructor, or overloaded copy assignment operator, nor can it be of reference type. A union member cannot be declared with the keyword `static`.

A member that does not represent a bit field can be qualified with either of the type qualifiers `volatile` or `const`. The result is an lvalue.

Structure members are assigned to memory addresses in increasing order, with the first component starting at the beginning address of the structure name itself. To allow proper alignment of components, padding bytes may appear between any consecutive members in the structure layout.

The storage allocated for a union is the storage required for the largest member of the union (plus any padding that is required so that the union will end at a natural boundary of its member having the most stringent requirements). All of a union’s components are effectively overlaid in memory: each member of a union is allocated storage starting at the beginning of the union, and only one member can occupy the storage at a time.

Flexible array members : A *flexible array member* is permitted as the last element of a structure even though it has incomplete type, provided that the structure has more than one named member. A flexible array member is a C99 feature and can be used to access a variable-length object. It is declared with an empty index, as follows:

```
array_identifier[ ];
```

For example, `b` is a flexible array member of `Foo`.

```
struct Foo{
    int a;
    int b[];
};
```

Since a flexible array member has incomplete type, you cannot apply the `sizeof` operator to a flexible array.

Any structure containing a flexible array member cannot be a member of another structure or array.

IBM extension

For compatibility with GNU C, XL C/C++ extends Standard C and C++, to ease the restrictions on flexible arrays and allow the following:

- Flexible array members can be declared in any part of a structure, not just as the last member. C++ The type of any member that follows the flexible array member must be compatible with the type of the flexible array member. C The type of any member following the flexible array member is not required to be compatible with the type of the flexible array member; however, a warning is issued in this case.
- Structures containing flexible array members can be members of other structures.
- Flexible array members can be statically initialized.

In the following example:

```
struct Foo{
    int a;
    int b[];
};

struct Foo foo1 = { 55, {6, 8, 10} };
struct Foo foo2 = { 55, {15, 6, 14, 90} };
```

foo1 creates an array b of 3 elements, which are initialized to 6, 8, and 10; while foo2 creates an array of 4 elements, which are initialized to 15, 6, 14, and 90.

Flexible array members can only be initialized if they are contained in the outermost part of nested structures. Members of inner structures cannot be initialized.

End of IBM extension

Related information

- “Variable length arrays ” on page 91

Zero-extent array members:

IBM extension

A zero-extent array is an array with no dimensions. Like a flexible array member, a zero-extent array can be used to access a variable-length object. Unlike a flexible array member, a zero-extent array is not a C99 feature, but is provided for GNU C compatibility.

A zero-extent array must be explicitly declared with zero as its dimension:

```
array_identifier[0]
```

Like a flexible array member, a zero-extent array can be declared in any part of a structure, not just as the last member. C++ The type of any member that follows the zero-extent array must be compatible with the type of the zero-extent array. C The type of any member following the zero-extent array is not required to be compatible with the type of the zero-extent array; however, a warning is issued in this case.

Unlike a flexible array member, a structure containing a zero-extent array can be a member of another array. Also, the `sizeof` operator can be applied to a zero-extent array; the value returned is 0.

A zero-extent array can only be statically initialized with an empty set. For example:

```

struct foo{
    int a;
    char b[0];
}; bar = { 100, { } };

```

Otherwise, it must be initialized as a dynamically-allocated array.

Zero-extent array members can only be initialized if they are contained in the outermost part of nested structures. Members of inner structures cannot be initialized.

|----- End of IBM extension -----|

Bit field members: Both C and C++ allow integer members to be stored into memory spaces smaller than the compiler would ordinarily allow. These space-saving structure members are called *bit fields*, and their width in bits can be explicitly declared. Bit fields are used in programs that must force a data structure to correspond to a fixed hardware representation and are unlikely to be portable.

Bit field member declaration syntax

▶—*type_specifier*—declarator—:*constant_expression*—;—▶

The *constant_expression* is a constant integer expression that indicates the field width in bits. A bit field declaration may not use either of the type qualifiers `const` or `volatile`.

|----- C only -----|

In C99, the allowable data types for a bit field include qualified and unqualified `_Bool`, signed int, and unsigned int. The default integer type for a bit field is signed.

|----- IBM extension -----|

XL C supports extended bit fields, which can be of type `char`, `long`, or `long long`.

|----- End of IBM extension -----|

|----- End of C only -----|

|----- C++ only -----|

A bit field can be any integral type or enumeration type.

|----- End of C++ only -----|

The maximum bit-field length is 64 bits. To increase portability, do not use bit fields greater than 32 bits in size.

The following structure has three bit-field members `kingdom`, `phylum`, and `genus`, occupying 12, 6, and 2 bits respectively:

```

struct taxonomy {
    int kingdom : 12;
    int phylum : 6;
    int genus : 2;
};

```

When you assign a value that is out of range to a bit field, the low-order bit pattern is preserved and the appropriate bits are assigned.

The following restrictions apply to bit fields. You cannot:

- Define an array of bit fields
- Take the address of a bit field
- Have a pointer to a bit field
- Have a reference to a bit field

If a series of bit fields does not add up to the size of an `int`, padding can take place. The amount of padding is determined by the alignment characteristics of the members of the structure. In some instances, bit fields can cross word boundaries.

Bit fields with a length of 0 must be unnamed. Unnamed bit fields cannot be referenced or initialized.

The following example demonstrates padding, and is valid for all implementations. Suppose that an `int` occupies 4 bytes. The example declares the identifier `kitchen` to be of type `struct on_off`:

```

struct on_off {
    unsigned light : 1;
    unsigned toaster : 1;
    int count; /* 4 bytes */
    unsigned ac : 4;
    unsigned : 4;
    unsigned clock : 1;
    unsigned : 0;
    unsigned flag : 1;
} kitchen ;

```

The structure `kitchen` contains eight members totalling 16 bytes. The following table describes the storage that each member occupies:

Member name	Storage occupied
<code>light</code>	1 bit
<code>toaster</code>	1 bit
(padding — 30 bits)	To the next <code>int</code> boundary
<code>count</code>	The size of an <code>int</code> (4 bytes)
<code>ac</code>	4 bits
(unnamed field)	4 bits
<code>clock</code>	1 bit
(padding — 23 bits)	To the next <code>int</code> boundary (unnamed field)
<code>flag</code>	1 bit
(padding — 31 bits)	To the next <code>int</code> boundary


Related information

- "Alignment of bit fields" in the *XL C/C++ Programming Guide*

Structure and union variable declarations

A structure or union *declaration* has the same form as a definition except the declaration does not have a brace-enclosed list of members. You must declare the structure or union data type before you can define a variable having that type.

Structure or union variable declaration syntax



The *tag_identifier* indicates the previously-defined data type of the structure or union.

► C++ The keyword `struct` is optional in structure variable declarations.

You can declare structures or unions having any storage class. The storage class specifier and any type qualifiers for the variable must appear at the beginning of the statement. Structures or unions declared with the `register` storage class specifier are treated as automatic variables.

The following example defines structure type address:

```
struct address {
    int street_no;
    char *street_name;
    char *city;
    char *prov;
    char *postal_code;
};
```

The following examples declare two structure variables of type address:

```
struct address perm_address;
struct address temp_address;
```

Related information

- “The aligned variable attribute” on page 107
- “The `__align` type qualifier” on page 74
- “The packed variable attribute” on page 108
- “Initialization of structures and unions” on page 98
- “Compatibility of structures, unions, and enumerations (C only)” on page 70
- “Dot operator `.`” on page 128
- “Arrow operator `->`” on page 128

Structure and union type and variable definitions in a single statement

You can define a structure (or union) type and a structure (or union) variable in one statement, by putting a declarator and an optional initializer after the variable definition. The following example defines a union data type (not named) and a union variable (named `length`):

```

union {
    float meters;
    double centimeters;
    long inches;
} length;

```

Note that because this example does not name the data type, `length` is the only variable that can have this data type. Putting an identifier after `struct` or `union` keyword provides a name for the data type and lets you declare additional variables of this data type later in the program.

To specify a storage class specifier for the variable or variables, you must put the storage class specifier at the beginning of the statement. For example:

```

static struct {
    int street_no;
    char *street_name;
    char *city;
    char *prov;
    char *postal_code;
} perm_address, temp_address;

```

In this case, both `perm_address` and `temp_address` are assigned static storage.

Type qualifiers can be applied to the variable or variables declared in a type definition. Both of the following examples are valid:

```

volatile struct class1 {
    char descript[20];
    long code;
    short complete;
} file1, file2;

struct class1 {
    char descript[20];
    long code;
    short complete;
} volatile file1, file2;

```

In both cases, the structures `file1` and `file2` are qualified as `volatile`.

Related information

- “Initialization of structures and unions” on page 98
- “Storage class specifiers” on page 45
- “Type qualifiers” on page 72

Access to structure and union members

Once structure or union variables have been declared, members are referenced by specifying the variable name with the dot operator (`.`) or a pointer with the arrow operator (`->`) and the member name. For example, both of the following:

```

perm_address.prov = "Ontario";
p_perm_address -> prov = "Ontario";

```

assign the string "Ontario" to the pointer `prov` that is in the structure `perm_address`.

All references to members of structures and unions, including bit fields, must be fully qualified. In the previous example, the fourth field cannot be referenced by `prov` alone, but only by `perm_address.prov`.

Related information

- “Dot operator .” on page 128
- “Arrow operator ->” on page 128

Anonymous unions

An *anonymous union* is a union without a name. It cannot be followed by a declarator. An anonymous union is not a type; it defines an unnamed object.

The member names of an anonymous union must be distinct from other names within the scope in which the union is declared. You can use member names directly in the union scope without any additional member access syntax.

For example, in the following code fragment, you can access the data members `i` and `cptr` directly because they are in the scope containing the anonymous union. Because `i` and `cptr` are union members and have the same address, you should only use one of them at a time. The assignment to the member `cptr` will change the value of the member `i`.

```
void f()
{
    union { int i; char* cptr ; };
    /* . . . */
    i = 5;
    cptr = "string_in_union"; // overrides the value 5
}
```

 An anonymous union cannot have protected or private members, and it cannot have member functions. A global or namespace anonymous union must be declared with the keyword `static`.

Related information

- “The static storage class specifier” on page 46
- “Member functions” on page 271

Enumerations

An *enumeration* is a data type consisting of a set of named values that represent integral constants, known as *enumeration constants*. An enumeration also referred to as an *enumerated type* because you must list (enumerate) each of the values in creating a name for each of them. In addition to providing a way of defining and grouping sets of integral constants, enumerations are useful for variables that have a small number of possible values.

You can declare an enumeration type separately from the definition of variables of that type, as described in “Enumeration type definition” and “Enumeration variable declarations” on page 69; or you can define an enumeration data type and all variables that have that type in one statement, as described in “Enumeration type and variable definitions in a single statement” on page 69.


Related information

- “Arithmetic conversions and promotions” on page 113

Enumeration type definition

An enumeration type definition contains the `enum` keyword followed by an optional identifier (the enumeration tag) and a brace-enclosed list of enumerators. A comma separates each enumerator in the enumerator list. C99 allows a trailing comma between the last enumerator and the closing brace. XL C++ also supports this feature, for compatibility with C99.


Enumeration definition syntax



The *tag_identifier* gives a name to the enumeration type. If you do not provide a tag name, you must put all variable definitions that refer to the enumeration type within the declaration of the type, as described in “Enumeration type and variable definitions in a single statement” on page 69. Similarly, you cannot use a type qualifier with an enumeration definition; type qualifiers placed in front of the `enum` keyword can only apply to variables that are declared within the type definition.

Enumeration members: The list of enumeration members, or *enumerators*, provides the data type with a set of values.

Enumeration member declaration syntax



C In C, an *enumeration constant* is of type `int`. If a constant expression is used as an initializer, the value of the expression cannot exceed the range of `int` (that is, `INT_MIN` to `INT_MAX` as defined in the header `limits.h`).

C++ In C++, each enumeration constant has a value that can be promoted to a signed or unsigned integer value and a distinct type that does not have to be integral. You can use an enumeration constant anywhere an integer constant is allowed, or anywhere a value of the enumeration type is allowed.

The value of an enumeration constant is determined in the following way:

1. An equal sign (=) and a constant expression after the enumeration constant gives an explicit value to the enumeration constant. The enumeration constant represents the value of the constant expression.
2. If no explicit value is assigned, the leftmost enumeration constant in the list receives the value zero (0).
3. Enumeration constants with no explicitly assigned values receive the integer value that is one greater than the value represented by the previous enumeration constant.

The following data type declarations list oats, wheat, barley, corn, and rice as enumeration constants. The number under each constant shows the integer value.

```
enum grain { oats, wheat, barley, corn, rice };
/*      0      1      2      3      4      */

enum grain { oats=1, wheat, barley, corn, rice };
/*      1      2      3      4      5      */

enum grain { oats, wheat=10, barley, corn=20, rice };
/*      0      10     11     20     21     */
```

It is possible to associate the same integer with two different enumeration constants. For example, the following definition is valid. The identifiers `suspend` and `hold` have the same integer value.

```
enum status { run, clear=5, suspend, resume, hold=6 };
/*      0      5      6      7      6      */
```

Each enumeration constant must be unique within the scope in which the enumeration is defined. In the following example, the second declarations of `average` and `poor` cause compiler errors:

```
func()
{
    enum score { poor, average, good };
    enum rating { below, average, above };
    int poor;
}
```


Related information

- “Integral types” on page 53

Enumeration variable declarations

You must declare the enumeration data type before you can define a variable having that type.

Enumeration variable declaration syntax



The `tag_identifier` indicates the previously-defined data type of the enumeration.

C++ The keyword `enum` is optional in enumeration variable declarations.

Related information

- “Initialization of enumerations” on page 100
- “Compatibility of structures, unions, and enumerations (C only)” on page 70

Enumeration type and variable definitions in a single statement

You can define a type and a variable in one statement by using a declarator and an optional initializer after the variable definition. To specify a storage class specifier for the variable, you must put the storage class specifier at the beginning of the declaration. For example:

```
register enum score { poor=1, average, good } rating = good;
```

C++ only

C++ also lets you put the storage class immediately before the declarator list. For example:

```
enum score { poor=1, average, good } register rating = good;
```

End of C++ only

Either of these examples is equivalent to the following two declarations:

```
enum score { poor=1, average, good };
register enum score rating = good;
```

Both examples define the enumeration data type `score` and the variable `rating`. `rating` has the storage class specifier `register`, the data type `enum score`, and the initial value `good`.

Combining a data type definition with the definitions of all variables having that data type lets you leave the data type unnamed. For example:

```
enum { Sunday, Monday, Tuesday, Wednesday, Thursday, Friday,
      Saturday } weekday;
```

defines the variable `weekday`, which can be assigned any of the specified enumeration constants. However, you can not declare any additional enumeration variables using this set of enumeration constants.

Compatibility of structures, unions, and enumerations (C only)

Within a single source file, each structure or union definition creates a new type that is neither the same as nor compatible with any other structure or union type. However, a type specifier that is a reference to a previously defined structure or union type is the same type. The tag associates the reference with the definition, and effectively acts as the type name. To illustrate this, only the types of structures `j` and `k` are compatible in this example:

```
struct { int a; int b; } h;
struct { int a; int b; } i;
struct S { int a; int b; } j;
struct S k;
```

Compatible structures may be assigned to each other.

Structures or unions with identical members but different tags are not compatible and cannot be assigned to each other. Structures and unions with identical members but using different alignments are not also compatible and cannot be assigned to each other.

Since the compiler treats enumeration variables and constants as integer types, you can freely mix the values of different enumerated types, regardless of type compatibility. Compatibility between an enumerated type and the integer type that represents it is controlled by compiler options and related pragmas. For a full discussion of the `-qenum` compiler option and related pragmas, see `-qenum` and `#pragma enum` in the *XL C/C++ Compiler Reference*.

Related information

- “Arithmetic conversions and promotions” on page 113
- Chapter 11, “Classes (C++ only),” on page 259
- “Structure and union type definition” on page 60
- “Incomplete types” on page 42

Compatibility across separate source files

When the definitions for two structures, unions, or enumerations are defined in separate source files, each file can theoretically contain a different definition for an object of that type with the same name. The two declarations must be compatible, or the run time behavior of the program is undefined. Therefore, the compatibility rules are more restrictive and specific than those for compatibility within the same

source file. For structure, union, and enumeration types defined in separately compiled files, the composite type is the type in the current source file.

The requirements for compatibility between two structure, union, or enumerated types declared in separate source files are as follows:

- If one is declared with a tag, the other must also be declared with the same tag.
- If both are completed types, their members must correspond exactly in number, be declared with compatible types, and have matching names.

For enumerations, corresponding members must also have the same values.

For structures and unions, the following additional requirements must be met for type compatibility:

- Corresponding members must be declared in the same order (applies to structures only).
- Corresponding bit fields must have the same widths.

typedef definitions

A typedef declaration lets you define your own identifiers that can be used in place of type specifiers such as `int`, `float`, and `double`. A typedef declaration does not reserve storage. The names you define using typedef are not new data types, but synonyms for the data types or combinations of data types they represent.

The namespace for a typedef name is the same as other identifiers. The exception to this rule is if the typedef name specifies a variably modified type. In this case, it has block scope.

When an object is defined using a typedef identifier, the properties of the defined object are exactly the same as if the object were defined by explicitly listing the data type associated with the identifier.

IBM extension

typedef definitions are extended to handle vector types. A vector type can be used in a typedef definition, and the new type name can be used in the usual ways, except for declaring other vectors. In a vector declaration context, a typedef name is disallowed as a type specifier. The following example illustrates a typical usage of typedef with vector types:

```
typedef vector unsigned short vint16;  
vint16 v1;
```

End of IBM extension

Related information

- “Type names” on page 85
- “Type specifiers” on page 52
- “Structures and unions” on page 59
- Chapter 11, “Classes (C++ only),” on page 259

Examples of typedef definitions

The following statements define `LENGTH` as a synonym for `int` and then use this typedef to declare `length`, `width`, and `height` as integer variables:

```
typedef int LENGTH;
LENGTH length, width, height;
```

The following declarations are equivalent to the above declaration:

```
int length, width, height;
```

Similarly, typedef can be used to define a structure, union, or C++ class. For example:

```
typedef struct {
    int scruples;
    int drams;
    int grains;
} WEIGHT;
```

The structure WEIGHT can then be used in the following declarations:

```
WEIGHT chicken, cow, horse, whale;
```

In the following example, the type of yds is "pointer to function with no parameter specified, returning int".

```
typedef int SCROLL();
extern SCROLL *yds;
```

In the following typedefs, the token struct is part of the type name: the type of ex1 is struct a; the type of ex2 is struct b.

```
typedef struct a { char x; } ex1, *ptr1;
typedef struct b { char x; } ex2, *ptr2;
```

Type ex1 is compatible with the type struct a and the type of the object pointed to by ptr1. Type ex1 is not compatible with char, ex2, or struct b.

C++ only

In C++, a typedef name must be different from any class type name declared within the same scope. If the typedef name is the same as a class type name, it can only be so if that typedef is a synonym of the class name. This condition is not the same as in C. The following can be found in standard C headers:

```
typedef class C { /* data and behavior */ } C;
```

A C++ class defined in a typedef without being named is given a dummy name and the typedef name for linkage. Such a class cannot have constructors or destructors. For example:

```
typedef class {
    Trees();
} Trees;
```

Here the function Trees() is an ordinary member function of a class whose type name is unspecified. In the above example, Trees is an alias for the unnamed class, not the class type name itself, so Trees() cannot be a constructor for that class.

End of C++ only

Type qualifiers

A type qualifier is used to refine the declaration of a variable, a function, and parameters, by specifying whether:

- The value of an object can be changed

- The value of an object must always be read from memory rather than from a register
- More than one pointer can access a modifiable memory address

XL C/C++ recognizes the following type qualifiers:

-  `__align`
- `const`
- `restrict`
- `volatile`

Standard C++ refers to the type qualifiers `const` and `volatile` as *cv-qualifiers*. In both languages, the cv-qualifiers are only meaningful in expressions that are lvalues.

When the `const` and `volatile` keywords are used with pointers, the placement of the qualifier is critical in determining whether it is the pointer itself that is to be qualified, or the object to which the pointer points. For a pointer that you want to qualify as `volatile` or `const`, you must put the keyword between the `*` and the identifier. For example:

```
int * volatile x;      /* x is a volatile pointer to an int */
int * const y = &z;   /* y is a const pointer to the int variable z */
```

For a pointer to a `volatile` or `const` data object, the type specifier and qualifier can be in any order, provided that the qualifier does not follow the `*` operator. For example:

```
volatile int *x;      /* x is a pointer to a volatile int
or
int volatile *x;     /* x is a pointer to a volatile int */

const int *y;        /* y is a pointer to a const int
or
int const *y;       /* y is a pointer to a const int */
```

The following examples contrast the semantics of these declarations:

Declaration	Description
<code>const int * ptr1;</code>	Defines a pointer to a constant integer: the value pointed to cannot be changed.
<code>int * const ptr2;</code>	Defines a constant pointer to an integer: the integer can be changed, but <code>ptr2</code> cannot point to anything else.
<code>const int * const ptr3;</code>	Defines a constant pointer to a constant integer: neither the value pointed to nor the pointer itself can be changed.

You can put more than one qualifier on a declaration: the compiler ignores duplicate type qualifiers.

A type qualifier cannot apply to user-defined types, but only to objects created from a user-defined type. Therefore, the following declaration is illegal:

```
volatile struct omega {
    int limit;
    char code;
}
```

However, if a variable or variables are declared within the same definition of the type, a type qualifier can be applied to the variable or variables by placing at the beginning of the statement or before the variable declarator or declarators.

Therefore:

```
volatile struct omega {
    int limit;
    char code;
} group;
```

provides the same storage as:

```
struct omega {
    int limit;
    char code;
} volatile group;
```

In both examples, the `volatile` qualifier only applies to the structure variable `group`.

When type qualifiers are applied to a structure, class, or union, or class variable, they also apply to the members of the structure, class or union.

Related information

- “Pointers” on page 86
- “Constant and volatile member functions” on page 272

The `__align` type qualifier

IBM extension

The `__align` qualifier is a language extension that allows you to specify an explicit alignment for an aggregate or a static (or global) variable. The specified byte boundary affects the alignment of an aggregate as a whole, not that of its members. The `__align` qualifier can be applied to an aggregate definition nested within another aggregate definition, but not to individual elements of an aggregate. The alignment specification is ignored for parameters and automatic variables.

A declaration takes one of the following forms:

`__align` qualifier syntax for simple variables

```
▶▶ type specifier __align (int_constant) declarator ▶▶
```

`__align` qualifier syntax for structures or unions

```
▶▶ __align (int_constant) struct tag_identifier ▶▶
union
▶▶ { member_declaration_list } ; ▶▶
```

where *int_constant* is a positive integer value indicating the byte-alignment boundary. The legal values are 1, 2, 4, 8, or 16.

The following restrictions and limitations apply:

- The `__align` qualifier cannot be used where the size of the variable alignment is smaller than the size of the type alignment.
- Not all alignments may be representable in an object file.
- The `__align` qualifier cannot be applied to the following:
 - Individual elements within an aggregate definition.
 - Individual elements of an array.
 - Variables of incomplete type.
 - Aggregates declared but not defined.
 - Other types of declarations or definitions, such as a typedef, a function, or an enumeration.

Examples using the `__align` qualifier

Applying `__align` to static or global variables:

```
int __align(1024) varA; /* varA is aligned on a 1024-byte boundary
main()                and padded with 1020 bytes                */
{...}

static int __align(512) varB; /* varB is aligned on a 512-byte boundary
                               and padded with 508 bytes                */

int __align(128) functionB( ); /* An error                          */

typedef int __align(128) T; /* An error                          */

__align enum C {a, b, c}; /* An error                          */
```

Applying `__align` to align and pad aggregate tags without affecting aggregate members:

```
__align(1024) struct structA {int i; int j;}; /* struct structA is aligned
                                               on a 1024-byte boundary
                                               with size including padding
                                               of 1024 bytes                */

__align(1024) union unionA {int i; int j;}; /* union unionA is aligned
                                               on a 1024-byte boundary
                                               with size including padding
                                               of 1024 bytes                */
```

Applying `__align` to a structure or union, where the size and alignment of the aggregate using the structure or union is affected:

```
__align(128) struct S {int i;}; /* sizeof(struct S) == 128                */

struct S sarray[10]; /* sarray is aligned on 128-byte boundary
                    with sizeof(sarray) == 1280                */

struct S __align(64) svar; /* error - alignment of variable is
                           smaller than alignment of type      */

struct S2 {struct S s1; int a;} s2; /* s2 is aligned on 128-byte boundary
                                   with sizeof(s2) == 256                */
```

Applying `__align` to an array:

```
AnyType __align(64) arrayA[10]; /* Only arrayA is aligned on a 64-byte
                                boundary, and elements within that array
                                are aligned according to the alignment
```

```
of AnyType. Padding is applied after the
back of the array and does not affect
the size of the array member itself.  */
```

Applying `__align` where the size of the variable alignment differs from the size of the type alignment:

```
__align(64) struct S {int i;};

struct S __align(32) s1;      /* error, alignment of variable is smaller
                             than alignment of type          */

struct S __align(128) s2;   /* s2 is aligned on 128-byte boundary */

struct S __align(16) s3[10]; /* error                               */
int __align(1) s4;         /* error                               */
__align(1) struct S {int i;}; /* error                               */
```

Related information

- "The aligned variable attribute" on page 107
- "The `__alignof__` operator" on page 135
- "Aligning data" in the *XL C/C++ Programming Guide*

End of IBM extension

The const type qualifier

The `const` qualifier explicitly declares a data object as something that cannot be changed. Its value is set at initialization. You cannot use `const` data objects in expressions requiring a modifiable lvalue. For example, a `const` data object cannot appear on the lefthand side of an assignment statement.

C only

A `const` object cannot be used in constant expressions. A global `const` object without an explicit storage class is considered `extern` by default.

End of C only

C++ only

In C++, all `const` declarations must have initializers, except those referencing externally defined constants. A `const` object can appear in a constant expression if it is an integer and it is initialized to a constant. The following example demonstrates this:

```
const int k = 10;
int ary[k];    /* allowed in C++, not legal in C */
```

In C++ a global `const` object without an explicit storage class is considered `static` by default, with internal linkage.

```
const int k = 12; /* Different meanings in C and C++ */

static const int k2 = 120; /* Same meaning in C and C++ */
extern const int k3 = 121; /* Same meaning in C and C++ */
```

Because its linkage is assumed to be internal, a `const` object can be more easily defined in header files in C++ than in C.

End of C++ only

An item can be both `const` and `volatile`. In this case the item cannot be legitimately modified by its own program but can be modified by some asynchronous process.

Related information

- “The `#define` directive” on page 392
- “The `this` pointer” on page 276

The `restrict` type qualifier

A pointer is the address of a location in memory. More than one pointer can access the same chunk of memory and modify it during the course of a program. The `restrict` (or `__restrict` or `__restrict__`)¹ type qualifier is an indication to the compiler that, if the memory addressed by the `restrict`-qualified pointer is modified, no other pointer will access that same memory. The compiler may choose to optimize code involving `restrict`-qualified pointers in a way that might otherwise result in incorrect behavior. It is the responsibility of the programmer to ensure that `restrict`-qualified pointers are used as they were intended to be used. Otherwise, undefined behavior may result.

If a particular chunk of memory is not modified, it can be aliased through more than one restricted pointer. The following example shows restricted pointers as parameters of `foo()`, and how an unmodified object can be aliased through two restricted pointers.

```
void foo(int n, int * restrict a, int * restrict b, int * restrict c)
{
    int i;
    for (i = 0; i < n; i++)
        a[i] = b[i] + c[i];
}
```

Assignments between restricted pointers are limited, and no distinction is made between a function call and an equivalent nested block.

```
{
    int * restrict x;
    int * restrict y;
    x = y; // undefined
    {
        int * restrict x1 = x; // okay
        int * restrict y1 = y; // okay
        x = y1; // undefined
    }
}
```

In nested blocks containing restricted pointers, only assignments of restricted pointers from outer to inner blocks are allowed. The exception is when the block in which the restricted pointer is declared finishes execution. At that point in the program, the value of the restricted pointer can be carried out of the block in which it was declared.

Notes:

1. The `restrict` qualifier is represented by the following keywords (all have the same semantics):

-  The `restrict` keyword is recognized under compilation with `ppuxlc` or `spuxlc` or `ppuc99` or with the `-qlanglvl=stdc99` or `-qlanglvl=extc99` options or `-qkeyword=restrict`. The `__restrict` and `__restrict__` keywords are recognized at all language levels.
-   The `restrict`, `__restrict` and `__restrict__` keywords are recognized by default.

Related information

- `-qlanglvl` and `-qkeyword` in the *XL C/C++ Compiler Reference*

The volatile type qualifier

The `volatile` qualifier declares a data object that can have its value changed in ways outside the control or detection of the compiler (such as a variable updated by the system clock or by another program). This prevents the compiler from optimizing code referring to the object by storing the object's value in a register and re-reading it from there, rather than from memory, where it may have changed.

Accessing any lvalue expression that is `volatile`-qualified produces a side effect. A side effect means that the state of the execution environment changes.

References to an object of type "pointer to `volatile`" may be optimized, but no optimization can occur to references to the object to which it points. An explicit cast must be used to assign a value of type "pointer to `volatile T`" to an object of type "pointer to `T`". The following shows valid uses of `volatile` objects.

```
volatile int * pvol;
int *ptr;
pvol = ptr;           /* Legal */
ptr = (int *)pvol;   /* Explicit cast required */
```

 A signal-handling function may store a value in a variable of type `sig_atomic_t`, provided that the variable is declared `volatile`. This is an exception to the rule that a signal-handling function may not access variables with static storage duration.

An item can be both `const` and `volatile`. In this case the item cannot be legitimately modified by its own program but can be modified by some asynchronous process.

Type attributes


IBM extension

Type attributes are language extensions provided to facilitate compilation of programs developed with the GNU C/C++ compilers. These language features allow you to use named attributes to specify special properties of data objects. *Type* attributes apply to the definitions of user-defined types, such as structures, unions, enumerations, classes, and typedef definitions. Any variables that are declared as having that type will have the attribute applied to them.

A type attribute is specified with the keyword `__attribute__` followed by the attribute name and any additional arguments the attribute name requires. Although there are variations, the syntax of a type attribute is of the general form:

of the union type. Because of this, all members of the union must have the same machine representation. Transparent unions are useful in library functions that use multiple interfaces to resolve issues of compatibility.

transparent_union type attribute syntax



The union must be a complete union type. The `transparent_union` type attribute can be applied to anonymous unions with tag names.

When the `transparent_union` type attribute is applied to the outer union of a nested union, the size of the inner union (that is, its largest member) is used to determine if it has the same machine representation as the other members of the outer union. For example,

```
union __attribute__((transparent_union)) u_t {
    union u2_t {
        char a;
        short b;
        char c;
        char d;
    };
    int a;
};
```

the attribute is ignored because the first member of union `u_t`, which is itself a union, has a machine representation of 2 bytes, whereas the other member of union `u_t` is of type `int`, which has a machine representation of 4 bytes.

The same rationale applies to members of a union that are structures. When a member of a union to which type attribute `transparent_union` has been applied is a struct, the machine representation of the entire struct is considered, rather than members.

All members of the union must have the same machine representation as the first member of the union. This means that all members must be representable by the same amount memory as the first member of the union. The machine representation of the first member represents the maximum memory size for any remaining union members. For instance, if the first member of a union to which type attribute `transparent_union` has been applied is of type `int`, then all following members must be representable by at most 4 bytes. Members that are representable by 1, 2, or 4 bytes are considered valid for this transparent union.

Floating-point types (`float`, `double`, `float _Complex`, or `double _Complex`) types or vector types can be members of a transparent union, but they cannot be the first member. The restriction that all members of the transparent union have the same machine representation as the first member still applies.

End of IBM extension

Chapter 4. Declarators

This section continues the discussion of data declarations and includes the following topics:


- “Overview of declarators”
- “Type names” on page 85
- “Pointers” on page 86
- “Arrays” on page 90
- “References (C++ only)” on page 93
- “Initializers” on page 93
- “Variable attributes” on page 105

Overview of declarators


A *declarator* designates a data object or function. A declarator can also include an initialization. Declarators appear in most data definitions and declarations and in some type definitions.

For data declarations, a declarator has the form:

Declarator syntax




Direct declarator:



C only


Pointer operator:



Declarator name:



End of C only



The *type_qualifiers* represent one or a combination of `const` and `volatile`.

C++ A *nested_name_specifier* is a qualified identifier expression. An *identifier_expression* can be a qualified or unqualified identifier.

Initializers are discussed in “Initializers” on page 93.

The following are known as *derived declarator* types, and are therefore discussed in this section:

- Pointers
- Arrays
- References (C++ only)

IBM In addition, for compatibility with GNU C and C++, XL C/C++ allows you to use *variable attributes* to modify the properties of data objects. As they are normally specified as part of the declarator in a declaration, they are described in this section, in “Variable attributes” on page 105.

Related information

- “Type qualifiers” on page 72

Examples of declarators

The following table indicates the declarators within the declarations:

Declaration	Declarator	Description
<code>int owner;</code>	<code>owner</code>	<code>owner</code> is an integer data object.
<code>int *node;</code>	<code>*node</code>	<code>node</code> is a pointer to an integer data object.
<code>int names[126];</code>	<code>names[126]</code>	<code>names</code> is an array of 126 integer elements.
<code>volatile int min;</code>	<code>min</code>	<code>min</code> is a volatile integer.

Declaration	Declarator	Description
<code>int * volatile volume;</code>	<code>* volatile volume</code>	volume is a volatile pointer to an integer.
<code>volatile int * next;</code>	<code>*next</code>	next is a pointer to a volatile integer.
<code>volatile int * sequence[5];</code>	<code>*sequence[5]</code>	sequence is an array of five pointers to volatile integer data objects.
<code>extern const volatile int clock;</code>	<code>clock</code>	clock is a constant and volatile integer with static storage duration and external linkage.

Related information

- “Type qualifiers” on page 72
- “Array subscripting operator []” on page 150
- “Scope resolution operator :: (C++ only)” on page 126
- “Function declarators” on page 214

Type names

A *type name*, is required in several contexts as something that you must specify without declaring an object; for example, when writing an explicit cast expression or when applying the `sizeof` operator to a type. Syntactically, the name of a data type is the same as a declaration of a function or object of that type, but without the identifier.

To read or write a type name correctly, put an “imaginary” identifier within the syntax, splitting the type name into simpler components. For example, `int` is a type specifier, and it always appears to the left of the identifier in a declaration. An imaginary identifier is unnecessary in this simple case. However, `int *[5]` (an array of 5 pointers to `int`) is also the name of a type. The type specifier `int *` always appears to the left of the identifier, and the array subscripting operator always appears to the right. In this case, an imaginary identifier is helpful in distinguishing the type specifier.

As a general rule, the identifier in a declaration always appears to the left of the subscripting and function call operators, and to the right of a type specifier, type qualifier, or indirection operator. Only the subscripting, function call, and indirection operators may appear in a type name declaration. They bind according to normal operator precedence, which is that the indirection operator is of lower precedence than either the subscripting or function call operators, which have equal ranking in the order of precedence. Parentheses may be used to control the binding of the indirection operator.

It is possible to have a type name within a type name. For example, in a function type, the parameter type syntax nests within the function type name. The same rules of thumb still apply, recursively.

The following constructions illustrate applications of the type naming rules.

Table 18. Type names

Syntax	Description
<code>int *[5]</code>	array of 5 pointers to <code>int</code>
<code>int (*)[5]</code>	pointer to an array of 5 integers
<code>int (*)[*]</code>	pointer to an variable length array of an unspecified number of integers
<code>int *()</code>	function with no parameter specification returning a pointer to <code>int</code>
<code>int *(void)</code>	function with no parameters returning an <code>int</code>
<code>int (*const [])(unsigned int, ...)</code>	array of an unspecified number of constant pointers to functions returning an <code>int</code> . Each function takes one parameter of type <code>unsigned int</code> and an unspecified number of other parameters.

The compiler turns any function designator into a pointer to the function. This behavior simplifies the syntax of function calls.

```
int foo(float); /* foo is a function designator */
int (*p)(float); /* p is a pointer to a function */
p=&foo; /* legal, but redundant */
p=foo; /* legal because the compiler turns foo into a function pointer */
```

 In C++, the keywords `typename` and `class`, which are interchangeable, indicate the name of the type.

Related information

- “Operator precedence and associativity” on page 170
- “Examples of expressions and precedence” on page 173
- “The `typename` keyword” on page 367
- “Parenthesized expressions ()” on page 125

Pointers

A *pointer* type variable holds the address of a data object or a function. A pointer can refer to an object of any one data type; it cannot refer to a bit field or a reference.

Some common uses for pointers are:

- To access dynamic data structures such as linked lists, trees, and queues.
- To access elements of an array or members of a structure or C++ class.
- To access an array of characters as a string.
- To pass the address of a variable to a function. (In C++, you can also use a reference to do this.) By referencing a variable through its address, a function can change the contents of that variable.

Note that the placement of the type qualifiers `volatile` and `const` affects the semantics of a pointer declaration. If either of the qualifiers appears before the `*`, the declarator describes a pointer to a type-qualified object. If either of the qualifiers appears between the `*` and the identifier, the declarator describes a type-qualified pointer.

The following table provides examples of pointer declarations.

Table 19. Pointer declarations

Declaration	Description
<code>long *pcoat;</code>	<code>pcoat</code> is a pointer to an object having type <code>long</code>
<code>extern short * const pvolt;</code>	<code>pvolt</code> is a constant pointer to an object having type <code>short</code>
<code>extern int volatile *pnut;</code>	<code>pnut</code> is a pointer to an <code>int</code> object having the <code>volatile</code> qualifier
<code>float * volatile psoup;</code>	<code>psoup</code> is a <code>volatile</code> pointer to an object having type <code>float</code>
<code>enum bird *pfowl;</code>	<code>pfowl</code> is a pointer to an enumeration object of type <code>bird</code>
<code>char (*pvish)(void);</code>	<code>pvish</code> is a pointer to a function that takes no parameters and returns a <code>char</code>

Related information

- “Type qualifiers” on page 72
- “Initialization of pointers” on page 101
- “Compatibility of pointers (C only)” on page 89
- “Pointer conversions” on page 117
- “Address operator &” on page 132
- “Indirection operator *” on page 133
- “Pointers to functions” on page 232

Pointer arithmetic

You can perform a limited number of arithmetic operations on pointers. These operations are:

- Increment and decrement
- Addition and subtraction
- Comparison
- Assignment

The increment (`++`) operator increases the value of a pointer by the size of the data object the pointer refers to. For example, if the pointer refers to the second element in an array, the `++` makes the pointer refer to the third element in the array.

The decrement (`--`) operator decreases the value of a pointer by the size of the data object the pointer refers to. For example, if the pointer refers to the second element in an array, the `--` makes the pointer refer to the first element in the array.

You can add an integer to a pointer but you cannot add a pointer to a pointer.

If the pointer `p` points to the first element in an array, the following expression causes the pointer to point to the third element in the same array:

```
p = p + 2;
```

If you have two pointers that point to the same array, you can subtract one pointer from the other. This operation yields the number of elements in the array that separate the two addresses that the pointers refer to.

You can compare two pointers with the following operators: `==`, `!=`, `<`, `>`, `<=`, and `>=`.

Pointer comparisons are defined only when the pointers point to elements of the same array. Pointer comparisons using the `==` and `!=` operators can be performed even when the pointers point to elements of different arrays.

You can assign to a pointer the address of a data object, the value of another compatible pointer or the NULL pointer.

IBM extension

Pointer arithmetic is defined for pointer to vector types. Given:

```
vector unsigned int *v;
```

the expression `v + 1` represents a pointer to the vector following `v`.

End of IBM extension

Related information

- “Increment operator `++`” on page 129
- “Arrays” on page 90
- “Decrement operator `--`” on page 130
- Chapter 6, “Expressions and operators,” on page 121

Type-based aliasing

The compiler follows the type-based aliasing rule in the C and C++ standards when the `-qalias=ansi` option is in effect (which it is by default). This rule, also known as the ANSI aliasing rule, states that a pointer can only be dereferenced to an object of the same type or a compatible type.¹The common coding practice of

-
1. The C Standard states that an object shall have its stored value accessed only by an lvalue that has one of the following types:
 - the declared type of the object,
 - a qualified version of the declared type of the object,
 - a type that is the signed or unsigned type corresponding to the declared type of the object,
 - a type that is the signed or unsigned type corresponding to a qualified version of the declared type of the object,
 - an aggregate or union type that includes one of the aforementioned types among its members (including, recursively, a member of a subaggregate or contained union), or
 - a character type

The C++ standard states that if a program attempts to access the stored value of an object through an lvalue of other than one of the following types, the behavior is undefined:

- the dynamic type of the object,
- a cv-qualified version of the dynamic type of the object,
- a type that is the signed or unsigned type corresponding to the dynamic type of the object,
- a type that is the signed or unsigned type corresponding to a cv-qualified version of the dynamic type of the object,
- an aggregate or union type that includes one of the aforementioned types among its members (including, recursively, a member of a subaggregate or contained union),
- a type that is a (possible cv-qualified) base class type of the dynamic type of the object,
- a char or unsigned char type.

casting a pointer to an incompatible type and then dereferencing it violates this rule. (Note that char pointers are an exception to this rule.)

The compiler uses the type-based aliasing information to perform optimizations to the generated code. Contravening the type-based aliasing rule can lead to unexpected behavior, as demonstrated in the following example:

```
int *p;
double d = 0.0;

int *faa(double *g);          /* cast operator inside the function */

void foo(double f) {
    p = faa(&f);              /* turning &f into a int ptr */
    f += 1.0;                 /* compiler may discard this statement */
    printf("f=%x\n", *p);
}

int *faa(double *g) { return (int*) g; } /* questionable cast; */
                                        /* the function can be in */
                                        /* another translation unit */

int main() {
    foo(d);
}
```

In the above `printf` statement, `*p` cannot be dereferenced to a double under the ANSI aliasing rule. The compiler determines that the result of `f += 1.0;` is never used subsequently. Thus, the optimizer may discard the statement from the generated code. If you compile the above example with optimization enabled, the `printf` statement may output 0 (zero).

Related information

- “The `reinterpret_cast` operator (C++ only)” on page 160
- `-qalias=ansi` in the *XL C/C++ Compiler Reference*

Compatibility of pointers (C only)

Two pointer types with the same type qualifiers are compatible if they point to objects of compatible types. The composite type for two compatible pointer types is the similarly qualified pointer to the composite type.

The following example shows compatible declarations for the assignment operation:

```
float subtotal;
float * sub_ptr;
/* ... */
sub_ptr = &subtotal;
printf("The subtotal is %f\n", *sub_ptr);
```

The next example shows incompatible declarations for the assignment operation:

```
double league;
int * minor;
/* ... */
minor = &league;    /* error */
```

Arrays

An *array* is a collection of objects of the same data type, allocated contiguously in memory. Individual objects in an array, called *elements*, are accessed by their position in the array. The subscripting operator (`[]`) provides the mechanics for creating an index to array elements. This form of access is called *indexing* or *subscripting*. An array facilitates the coding of repetitive tasks by allowing the statements executed on each element to be put into a loop that iterates through each element in the array.

The C and C++ languages provide limited built-in support for an array type: reading and writing individual elements. Assignment of one array to another, the comparison of two arrays for equality, returning self-knowledge of size are not supported by either language.

The type of an array is derived from the type of its elements, in what is called *array type derivation*. If array objects are of incomplete type, the array type is also considered incomplete. Array elements may not be of type `void` or of function type. However, arrays of pointers to functions are allowed.  Array elements may not be of reference type or of an abstract class type.

The array declarator contains an identifier followed by an optional *subscript declarator*. An identifier preceded by an asterisk (`*`) is an array of pointers.

Array subscript declarator syntax



The *constant_expression* is a constant integer expression, indicating the size of the array, which must be positive.

If the declaration appears in block or function scope, a nonconstant expression can be specified for the array subscript declarator, and the array is considered a *variable-length array*, as described in “Variable length arrays” on page 91.

The subscript declarator describes the number of dimensions in the array and the number of elements in each dimension. Each bracketed expression, or subscript, describes a different dimension and must be a constant expression.

The following example defines a one-dimensional array that contains four elements having type `char`:

```
char  
list[4];
```

The first subscript of each dimension is 0. The array `list` contains the elements:

```
list[0]  
list[1]  
list[2]  
list[3]
```


The following example defines a two-dimensional array that contains six elements of type `int`:

```
int  
roster[3][2];
```

Multidimensional arrays are stored in row-major order. When elements are referred to in order of increasing storage location, the last subscript varies the fastest. For example, the elements of array `roster` are stored in the order:

```
roster[0][0]
roster[0][1]
roster[1][0]
roster[1][1]
roster[2][0]
roster[2][1]
```

In storage, the elements of `roster` would be stored as:



You can leave the first (and only the first) set of subscript brackets empty in:

- Array definitions that contain initializations
- extern declarations
- Parameter declarations

In array definitions that leave the first set of subscript brackets empty, the initializer determines the number of elements in the first dimension. In a one-dimensional array, the number of initialized elements becomes the total number of elements. In a multidimensional array, the initializer is compared to the subscript declarator to determine the number of elements in the first dimension.


Related information

- “Array subscripting operator []” on page 150
- “Initialization of arrays” on page 101

Variable length arrays

A variable length array, which is a C99 feature, is an array of automatic storage duration whose length is determined at run time.

Variable length array declarator syntax



If the size of the array is indicated by `*` instead of an expression, the variable length array is considered to be of unspecified size. Such arrays are considered complete types, but can only be used in declarations of function prototype scope.

A variable length array and a pointer to a variable length array are considered *variably modified types*. Declarations of variably modified types must be at either block scope or function prototype scope. Array objects declared with the extern storage class specifier cannot be of variable length array type. Array objects

declared with the `static` storage class specifier can be a pointer to a variable length array, but not an actual variable length array. A variable length array cannot be initialized.

Note:  In C++ applications, storage allocated for use by variable length arrays is not released until the function they reside in completes execution.

A variable length array can be the operand of a `sizeof` expression. In this case, the operand is evaluated at run time, and the size is neither an integer constant nor a constant expression, even though the size of each instance of a variable array does not change during its lifetime.

A variable length array can be used in a `typedef` statement. The `typedef` name will have only block scope. The length of the array is fixed when the `typedef` name is defined, not each time it is used.

A function parameter can be a variable length array. The necessary size expressions must be provided in the function definition. The compiler evaluates the size expression of a variably modified parameter on entry to the function. For a function declared with a variable length array as a parameter, as in the following, `void f(int x, int a[][x]);`

the size of the variable length array argument must match that of the function definition.

  The C++ extension does not include support for references to a variable length array type; neither may a function parameter be a reference to a variable length array type.

Related information

- “Flexible array members ” on page 61

Compatibility of arrays

Two array types that are similarly qualified are compatible if the types of their elements are compatible. For example,

```
char ex1[25];  
const char ex2[25];
```

are not compatible.

The composite type of two compatible array types is an array with the composite element type. The sizes of both original types must be equivalent if they are known. If the size of only one of the original array types is known, then the composite type has that size. For example:

```
char ex3[];  
char ex4[42];
```

The composite type of `ex3` and `ex4` is `char[42]`. If one of the original types is a variable length array, the composite type is that type.

Related information

- “External linkage” on page 8

References (C++ only)

A *reference* is an alias or an alternative name for an object. All operations applied to a reference act on the object to which the reference refers. The address of a reference is the address of the aliased object.

A reference type is defined by placing the reference modifier `&` after the type specifier. You must initialize all references except function parameters when they are defined. Once defined, a reference cannot be reassigned because it is an alias to its target. What happens when you try to reassign a reference turns out to be the assignment of a new value to the target.

Because arguments of a function are passed by value, a function call does not modify the actual values of the arguments. If a function needs to modify the actual value of an argument or needs to return more than one value, the argument must be *passed by reference* (as opposed to being *passed by value*). Passing arguments by reference can be done using either references or pointers. Unlike C, C++ does not force you to use pointers if you want to pass arguments by reference. The syntax of using a reference is somewhat simpler than that of using a pointer. Passing an object by reference enables the function to change the object being referred to without creating a copy of the object within the scope of the function. Only the address of the actual original object is put on the stack, not the entire object.

For example:

```
int f(int&);
int main()
{
    extern int i;
    f(i);
}
```

You cannot tell from the function call `f(i)` that the argument is being passed by reference.

References to NULL are not allowed.

Related information

- “Initialization of references (C++ only)” on page 104
- “Pointers” on page 86
- “Reference conversions (C++ only)” on page 118
- “Address operator `&`” on page 132
- “Pass by reference” on page 227

Initializers

An *initializer* is an optional part of a data declaration that specifies an initial value of a data object. The initializers that are legal for a particular declaration depend on the type and storage class of the object to be initialized.

The initializer consists of the `=` symbol followed by an initial *expression* or a brace-enclosed list of initial expressions separated by commas. Individual expressions must be separated by commas, and groups of expressions can be enclosed in braces and separated by commas. Braces (`{ }`) are optional if the initializer for a character string is a string literal. The number of initializers must

not be greater than the number of elements to be initialized. The initial expression evaluates to the first value of the data object.

To assign a value to an arithmetic or pointer type, use the simple initializer: `= expression`. For example, the following data definition uses the initializer `= 3` to set the initial value of `group` to 3:

```
int group = 3;
```

You initialize a variable of character type with a character literal (consisting of one character) or with an expression that evaluates to an integer.

 You can initialize variables at namespace scope with nonconstant expressions.  You cannot initialize variables at global scope with nonconstant expressions.

“Initialization and storage classes” discusses the rules for initialization according to the storage class of variables.

“Designated initializers for aggregate types (C only)” on page 95 describes designated initializers, which are a C99 feature that can be used to initialize arrays, structures, and unions.

The following sections discuss initializations for derived types:

- “Initialization of vectors” on page 97
- “Initialization of structures and unions” on page 98
- “Initialization of pointers” on page 101
- “Initialization of arrays” on page 101
- “Initialization of references (C++ only)” on page 104

Related information

- “Using class objects” on page 260

Initialization and storage classes

Initialization of automatic variables

You can initialize any `auto` variable except function parameters. If you do not explicitly initialize an automatic object, its value is indeterminate. If you provide an initial value, the expression representing the initial value can be any valid C or C++ expression. The object is then set to that initial value each time the program block that contains the object’s definition is entered.

Note that if you use the `goto` statement to jump into the middle of a block, automatic variables within that block are not initialized.

Related information

- “The `auto` storage class specifier” on page 46

Initialization of static variables

You initialize a `static` object with a constant expression, or an expression that reduces to the address of a previously declared `extern` or `static` object, possibly modified by a constant expression. If you do not explicitly initialize a `static` (or `external`) variable, it will have a value of zero of the appropriate type, unless it is a pointer, in which case it will be initialized to `NULL`.

A static variable in a block is initialized only one time, prior to program execution, whereas an auto variable that has an initializer is initialized every time it comes into existence.

▶ **C++** A static object of class type will use the default constructor if you do not initialize it. Automatic and register variables that are not initialized will have undefined values.

Related information

- “The static storage class specifier” on page 46

Initialization of external variables

You can initialize any object with the extern storage class specifier at global scope in C or at namespace scope in C++. The initializer for an extern object must either:

- Appear as part of the definition and the initial value must be described by a constant expression; or
- Reduce to the address of a previously declared object with static storage duration. You may modify this object with pointer arithmetic. (In other words, you may modify the object by adding or subtracting an integral constant expression.)

If you do not explicitly initialize an extern variable, its initial value is zero of the appropriate type. Initialization of an extern object is completed by the time the program starts running.

Related information

- “The extern storage class specifier” on page 48

Initialization of register variables

You can initialize any register object except function parameters. If you do not initialize an automatic object, its value is indeterminate. If you provide an initial value, the expression representing the initial value can be any valid C or C++ expression. The object is then set to that initial value each time the program block that contains the object’s definition is entered.

Related information

- “The register storage class specifier” on page 49

Designated initializers for aggregate types (C only)

Designated initializers, a C99 feature, are supported for aggregate types, including arrays, structures, and unions. A designated initializer, or *designator*, points out a particular element to be initialized. A *designator list* is a comma-separated list of one or more designators. A designator list followed by an equal sign constitutes a *designation*.

Designated initializers allow for the following flexibility:

- Elements within an aggregate can be initialized in any order.
- The initializer list can omit elements that are declared anywhere in the aggregate, rather than only at the end. Elements that are omitted are initialized as if they are static objects: arithmetic types are initialized to 0; pointers are initialized to NULL.

- Where inconsistent or incomplete bracketing of initializers for multi-dimensional arrays or nested aggregates may be difficult to understand, designators can more clearly identify the element or member to be initialized.

Designator list syntax for structures and unions



Designator list syntax for arrays



In the following example, the designator is `.any_member` and the designated initializer is `.any_member = 13`:

```
union { /* ... */ } caw = { .any_member = 13 };
```

The following example shows how the second and third members `b` and `c` of structure variable `k1m` are initialized with designated initializers:

```
struct xyz {
    int a;
    int b;
    int c;
} k1m = { .a = 99, .c = 100 };
```

In the following example, the third and second elements of the one-dimensional array `aa` are initialized to 3 and 6, respectively:

```
int aa[4] = { [2] = 3, [1] = 6 };
```

The following example initializes the first four and last four elements, while omitting the middle four:

```
static short grid[3][4] = { [0][0]=8, [0][1]=6,
                           [0][2]=4, [0][3]=1,
                           [2][0]=9, [2][1]=3,
                           [2][2]=1, [2][3]=1 };
```

The omitted four elements of `grid` are initialized to zero:

Element	Value	Element	Value
grid[0][0]	8	grid[1][2]	0
grid[0][1]	6	grid[1][3]	0
grid[0][2]	4	grid[2][0]	9
grid[0][3]	1	grid[2][1]	3
grid[1][0]	0	grid[2][2]	1
grid[1][1]	0	grid[2][3]	1

Designated initializers can be combined with regular initializers, as in the following example:

```
int a[10] = {2, 4, [8]=9, 10}
```

In this example, `a[0]` is initialized to 2, `a[1]` is initialized to 4, `a[2]` to `a[7]` are initialized to 0, and `a[9]` is initialized to 10.

In the following example, a single designator is used to "allocate" space from both ends of an array:

```
int a[MAX] = {
    1, 3, 5, 7, 9, [MAX-5] = 8, 6, 4, 2, 0
};
```

The designated initializer, `[MAX-5] = 8`, means that the array element at subscript `MAX-5` should be initialized to the value 8. If `MAX` is 15, `a[5]` through `a[9]` will be initialized to zero. If `MAX` is 7, `a[2]` through `a[4]` will first have the values 5, 7, and 9, respectively, which are overridden by the values 8, 6, and 4. In other words, if `MAX` is 7, the initialization would be the same as if the declaration had been written:

```
int a[MAX] = {
    1, 3, 8, 6, 4, 2, 0
};
```

You can also use designators to represent members of nested structures. For example:

```
struct a {
    struct b {
        int c;
        int d;
    } e;
    float f;
} g = {.e.c = 3};
```

initializes member `c` of structure variable `e`, which is a member of structure variable `g`, to the value of 3.

Related information

- "Initialization of structures and unions" on page 98
- "Initialization of arrays" on page 101

Initialization of vectors

IBM extension

A vector type is initialized by a vector literal or any expression having the same vector type. For example:

```
vector unsigned int v1;
vector unsigned int v2 = (vector unsigned int)(10);
v1 = v2;
```

XL C/C++ extends the `Altivec` specification to allow a vector type to be initialized by an initializer list. This feature is an extension for compatibility with GNU C.

Vector initializer list syntax

► `vector_type` `identifier` = { `initializer` } ;

The number of values in a braced initializer list must be less than or equal to the number of elements of the vector type. Any uninitialized element will be initialized to zero.

The following are examples of vector initialization using initializer lists:

```
vector<unsigned int> v1 = {1}; // initialize the first 4 bytes of v1 with 1
                          // and the remaining 12 bytes with zeros
```

```
vector<unsigned int> v2 = {1,2}; // initialize the first 8 bytes of v2 with 1 and 2
                               // and the remaining 8 bytes with zeros
```

```
vector<unsigned int> v3 = {1,2,3,4}; // equivalent to the vector literal
                                   // (vector<unsigned int>)(1,2,3,4)
```

Unlike vector literals, the values in the initializer list do not have to be constant expressions unless the initialized vector variable has static duration. Thus, the following is legal:

```
int i=1;
int foo() { return 2; }
int main()
{
    vector<unsigned int> v1 = {i, foo()};
    return 0;
}
```

End of IBM extension

Initialization of structures and unions

An initializer for a structure is a brace-enclosed comma-separated list of values, and for a union, a brace-enclosed single value. The initializer is preceded by an equal sign (=).

C99 and C++ allow the initializer for an automatic member variable of a union or structure type to be a constant or non-constant expression.

► **C++** The initializer for a static member variable of a union or structure type must be a constant expression or string literal. See “Static data members” on page 280 for more information.

There are two ways to specify initializers for structures and unions:

- With C89-style initializers, structure members must be initialized in the order declared, and only the first member of a union can be initialized.
- ► **C** Using *designated* initializers, a C99 feature which allows you to *name* members to be initialized, structure members can be initialized in any order, and any (single) member of a union can be initialized. Designated initializers are described in detail in “Designated initializers for aggregate types (C only)” on page 95.

Using C89-style initialization, the following example shows how you would initialize the first union member `birthday` of the union variable `people`:

```

union {
    char birthday[9];
    int age;
    float weight;
} people = {"23/07/57"};

```

C only

Using a designated initializer in the same example, the following initializes the second union member age :

```

union {
    char birthday[9];
    int age;
    float weight;
} people = { .age = 14 };

```

End of C only

The following definition shows a completely initialized structure:

```

struct address {
    int street_no;
    char *street_name;
    char *city;
    char *prov;
    char *postal_code;
};
static struct address perm_address =
    { 3, "Savona Dr.", "Dundas", "Ontario", "L4B 2A1"};

```

The values of perm_address are:

Member	Value
perm_address.street_no	3
perm_address.street_name	address of string "Savona Dr."
perm_address.city	address of string "Dundas"
perm_address.prov	address of string "Ontario"
perm_address.postal_code	address of string "L4B 2A1"

Unnamed structure or union members do not participate in initialization and have indeterminate value after initialization. Therefore, in the following example, the bit field is not initialized, and the initializer 3 is applied to member b:

```

struct {
    int a;
    int :10;
    int b;
} w = { 2, 3 };

```

You do not have to initialize all members of a structure or union; the initial value of uninitialized structure members depends on the storage class associated with the structure or union variable. In a structure declared as static, any members that are not initialized are implicitly initialized to zero of the appropriate type; the members of a structure with automatic storage have no default initialization. The default initializer for a union with static storage is the default for the first component; a union with automatic storage has no default initialization.

The following definition shows a partially initialized structure:

```

struct address {
    int street_no;
    char *street_name;
    char *city;
    char *prov;
    char *postal_code;
};
struct address temp_address =
    { 44, "Knyvet Ave.", "Hamilton", "Ontario" };

```

The values of `temp_address` are:

Member	Value
<code>temp_address.street_no</code>	44
<code>temp_address.street_name</code>	address of string "Knyvet Ave."
<code>temp_address.city</code>	address of string "Hamilton"
<code>temp_address.prov</code>	address of string "Ontario"
<code>temp_address.postal_code</code>	Depends on the storage class of the <code>temp_address</code> variable; if it is static, the value would be NULL.

C only

To initialize only the third and fourth members of the `temp_address` variable, you could use a designated initializer list, as follows:

```

struct address {
    int street_no;
    char *street_name;
    char *city;
    char *prov;
    char *postal_code;
};
struct address temp_address =
    { .city = "Hamilton", .prov = "Ontario" };

```

End of C only

Related information

- “Structure and union variable declarations” on page 65
- “Explicit initialization with constructors” on page 320
- “Assignment operators” on page 140

Initialization of enumerations

The initializer for an enumeration variable contains the `=` symbol followed by an expression *enumeration_constant*.

 In C++, the initializer must have the same type as the associated enumeration type.

The first line of the following example declares the enumeration `grain`. The second line defines the variable `g_food` and gives `g_food` the initial value of `barley` (2).

```

enum grain { oats, wheat, barley, corn, rice };
enum grain g_food = barley;

```

Related information

- “Enumeration variable declarations” on page 69

Initialization of pointers

The initializer is an = (equal sign) followed by the expression that represents the address that the pointer is to contain. The following example defines the variables `time` and `speed` as having type `double` and `amount` as having type `pointer to a double`. The pointer `amount` is initialized to point to `total`:

```
double total, speed, *amount = &total;
```

The compiler converts an unsubscripted array name to a pointer to the first element in the array. You can assign the address of the first element of an array to a pointer by specifying the name of the array. The following two sets of definitions are equivalent. Both define the pointer `student` and initialize `student` to the address of the first element in `section`:

```
int section[80];
int *student = section;
```

is equivalent to:

```
int section[80];
int *student = &section[0];
```

You can assign the address of the first character in a string constant to a pointer by specifying the string constant in the initializer. The following example defines the pointer variable `string` and the string constant `"abcd"`. The pointer `string` is initialized to point to the character `a` in the string `"abcd"`.

```
char *string = "abcd";
```

The following example defines `weekdays` as an array of pointers to string constants. Each element points to a different string. The pointer `weekdays[2]`, for example, points to the string `"Tuesday"`.

```
static char *weekdays[ ] =
{
    "Sunday", "Monday", "Tuesday", "Wednesday",
    "Thursday", "Friday", "Saturday"
};
```

A pointer can also be initialized to null using any integer constant expression that evaluates to 0, for example `char * a=0;`. Such a pointer is a *null pointer*. It does not point to any object.

Related information

- “Pointers” on page 86

Initialization of arrays

The initializer for an array is a comma-separated list of constant expressions enclosed in braces (`{ }`). The initializer is preceded by an equal sign (`=`). You do not need to initialize all elements in an array. If an array is partially initialized, elements that are not initialized receive the value 0 of the appropriate type. The same applies to elements of arrays with static storage duration. (All file-scope variables and function-scope variables declared with the `static` keyword have static storage duration.)

There are two ways to specify initializers for arrays:

- With C89-style initializers, array elements must be initialized in subscript order.
-  Using *designated* initializers, which allow you to specify the values of the subscript elements to be initialized, array elements can be initialized in any order. Designated initializers are described in detail in “Designated initializers for aggregate types (C only)” on page 95.

Using C89-style initializers, the following definition shows a completely initialized one-dimensional array:

```
static int number[3] = { 5, 7, 2 };
```

The array `number` contains the following values: `number[0]` is 5, `number[1]` is 7; `number[2]` is 2. When you have an expression in the subscript declarator defining the number of elements (in this case 3), you cannot have more initializers than the number of elements in the array.

The following definition shows a partially initialized one-dimensional array:

```
static int number1[3] = { 5, 7 };
```

The values of `number1[0]` and `number1[1]` are the same as in the previous definition, but `number1[2]` is 0.

C only

The following definition shows how you can use designated initializers to skip over elements of the array that you don't want to initialize explicitly:

```
static int number[3] = { [0] = 5, [2] = 7 };
```

The array `number` contains the following values: `number[0]` is 5; `number[1]` is implicitly initialized to 0; `number[2]` is 7.

End of C only

Instead of an expression in the subscript declarator defining the number of elements, the following one-dimensional array definition defines one element for each initializer specified:

```
static int item[ ] = { 1, 2, 3, 4, 5 };
```

The compiler gives `item` the five initialized elements, because no size was specified and there are five initializers.

Initialization of character arrays

You can initialize a one-dimensional character array by specifying:

- A brace-enclosed comma-separated list of constants, each of which can be contained in a character
- A string constant (braces surrounding the constant are optional)

Initializing a string constant places the null character (`\0`) at the end of the string if there is room or if the array dimensions are not specified.

The following definitions show character array initializations:

```
static char name1[ ] = { 'J', 'a', 'n' };
static char name2[ ] = { "Jan" };
static char name3[4] = "Jan";
```

These definitions create the following elements:

Element	Value	Element	Value	Element	Value
name1[0]	J	name2[0]	J	name3[0]	J
name1[1]	a	name2[1]	a	name3[1]	a
name1[2]	n	name2[2]	n	name3[2]	n
		name2[3]	\0	name3[3]	\0

Note that the following definition would result in the null character being lost:

```
static char name3[3]="Jan";
```

C++ When you initialize an array of characters with a string, the number of characters in the string — including the terminating '\0' — must not exceed the number of elements in the array.

Initialization of multidimensional arrays

You can initialize a multidimensional array using any of the following techniques:

- Listing the values of all elements you want to initialize, in the order that the compiler assigns the values. The compiler assigns values by increasing the subscript of the last dimension fastest. This form of a multidimensional array initialization looks like a one-dimensional array initialization. The following definition completely initializes the array `month_days`:

```
static month_days[2][12] =
{
    31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31,
    31, 29, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31
};
```

- Using braces to group the values of the elements you want initialized. You can put braces around each element, or around any nesting level of elements. The following definition contains two elements in the first dimension (you can consider these elements as rows). The initialization contains braces around each of these two elements:

```
static int month_days[2][12] =
{
    { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 },
    { 31, 29, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 }
};
```

- Using nested braces to initialize dimensions and elements in a dimension selectively. In the following example, only the first eight elements of the array `grid` are explicitly initialized. The remaining four elements that are not explicitly initialized are automatically initialized to zero.

```
static short grid[3][4] = {8, 6, 4, 1, 9, 3, 1, 1};
```

The initial values of `grid` are:

Element	Value	Element	Value
grid[0][0]	8	grid[1][2]	1
grid[0][1]	6	grid[1][3]	1
grid[0][2]	4	grid[2][0]	0

Element	Value	Element	Value
grid[0] [3]	1	grid[2] [1]	0
grid[1] [0]	9	grid[2] [2]	0
grid[1] [1]	3	grid[2] [3]	0

C only

- Using *designated* initializers. The following example uses designated initializers to explicitly initialize only the last four elements of the array. The first eight elements that are not explicitly initialized are automatically initialized to zero.

```
static short grid[3] [4] = { [2][0] = 8, [2][1] = 6,
                           [2][2] = 4, [2][3] = 1 };
```

The initial values of grid are:

Element	Value	Element	Value
grid[0] [0]	0	grid[1] [2]	0
grid[0] [1]	0	grid[1] [3]	0
grid[0] [2]	0	grid[2] [0]	8
grid[0] [3]	0	grid[2] [1]	6
grid[1] [0]	0	grid[2] [2]	4
grid[1] [1]	0	grid[2] [3]	1

End of C only

Related information

- “Arrays” on page 90
- “Designated initializers for aggregate types (C only)” on page 95

Initialization of references (C++ only)

The object that you use to initialize a reference must be of the same type as the reference, or it must be of a type that is convertible to the reference type. If you initialize a reference to a constant using an object that requires conversion, a temporary object is created. In the following example, a temporary object of type float is created:

```
int i;
const float& f = i; // reference to a constant float
```

When you initialize a reference with an object, you *bind* that reference to that object.

Attempting to initialize a nonconstant reference with an object that requires a conversion is an error.

Once a reference has been initialized, it cannot be modified to refer to another object. For example:

```

int num1 = 10;
int num2 = 20;

int &RefOne = num1;           // valid
int &RefOne = num2;           // error, two definitions of RefOne
RefOne = num2;                // assign num2 to num1
int &RefTwo;                  // error, uninitialized reference
int &RefTwo = num2;           // valid

```

Note that the initialization of a reference is not the same as an assignment to a reference. Initialization operates on the actual reference by initializing the reference with the object it is an alias for. Assignment operates through the reference on the object referred to.

A reference can be declared without an initializer:

- When it is used in an parameter declaration
- In the declaration of a return type for a function call
- In the declaration of class member within its class declaration
- When the extern specifier is explicitly used

You cannot have references to any of the following:

- Other references
- Bit fields
- Arrays of references
- Pointers to references

Direct binding

Suppose a reference *r* of type *T* is initialized by an expression *e* of type *U*.

The reference *r* is *bound directly* to *e* if the following statements are true:

- Expression *e* is an lvalue
- *T* is the same type as *U*, or *T* is a base class of *U*
- *T* has the same, or more, `const` or `volatile` qualifiers than *U*

The reference *r* is also bound directly to *e* if *e* can be implicitly converted to a type such that the previous list of statements is true.

Related information

- “References (C++ only)” on page 93
- “Pass by reference” on page 227

Variable attributes


IBM extension

Variable attributes are language extensions provided to facilitate the compilation of programs developed with the GNU C/C++ compilers. These language features allow you to use named attributes to specify special properties of data objects. *Variable* attributes apply to the declarations of simple variables, aggregates, and member variables of aggregates.


A variable attribute is specified with the keyword `__attribute__` followed by the attribute name and any additional arguments the attribute name requires. A

variable `__attribute__` specification is included in the declaration of a variable, and can be placed before or after the declarator. Although there are variations, the syntax generally takes either of the following forms:

Variable attribute syntax: post-declarator



Variable attribute syntax: pre-declarator



The *attribute name* can be specified with or without leading and trailing double underscore characters; however, using the double underscore reduces the likelihood of a name conflict with a macro of the same name. For unsupported attribute names, the XL C/C++ compiler issues diagnostics and ignores the attribute specification. Multiple attribute names can be specified in the same attribute specification.

In a comma-separated list of declarators on a single declaration line, if a variable attribute appears before all the declarators, it applies to all declarators in the declaration. If the attribute appears after a declarator, it only applies to the immediately preceding declarator. For example:

```
struct A {
    int b __attribute__((aligned));           /* typical placement of variable */
                                           /* attribute */

    int __attribute__((aligned)) __c = 10; /* variable attribute can also be */
                                           /* placed here */

    int d, e, f __attribute__((aligned));    /* attribute applies to f only */

    int g __attribute__((aligned)), h, i;    /* attribute applies to g only */

    int __attribute__((aligned)) j, k, l;    /* attribute applies to j, k, and l */
};
```

The following variable attributes are supported:

- The aligned variable attribute
- The packed variable attribute
- “The common and nocommon variable attributes” on page 108
- The `init_priority` variable attribute (C++ only)
- The mode variable attribute

- The section variable attribute
- The weak variable attribute

Related information

- “Type attributes” on page 78
- “Function attributes” on page 218

The aligned variable attribute

The aligned variable attribute allows you to override the default alignment mode to specify a minimum alignment value, expressed as a number of bytes, for any of the following:

- a non-aggregate variable
- an aggregate variable (such as a structure, class, or union)
- selected member variables

The attribute is typically used to increase the alignment of the given variable.

aligned variable attribute syntax

►► `__attribute__((aligned(__aligned_))` ◀◀

└──────────┘ └──────────┘

`(-alignment_factor-)`

The *alignment_factor* is the number of bytes, specified as a constant expression that evaluates to a positive power of 2. You can specify a value up to a maximum of 1048576 bytes. If you omit the alignment factor (and its enclosing parentheses the compiler automatically uses 16 bytes. If you specify an alignment factor greater than the maximum, the attribute specification is ignored, and the compiler simply uses the default alignment in effect.

When you apply the aligned attribute to a bit field structure member variable, the attribute specification is applied to the bit field *container*. If the default alignment of the container is greater than the alignment factor, the default alignment is used.

In the following example, the structures `first_address` and `second_address` are set to an alignment of 16 bytes:

```
struct address {
    int street_no;
    char *street_name;
    char *city;
    char *prov;
    char *postal_code;
} first_address __attribute__((aligned(16))) ;

struct address second_address __attribute__((aligned(16))) ;
```

In the following example, only the members `first_address.prov` and `first_address.postal_code` are set to an alignment of 16 bytes:

```
struct address {
    int street_no;
    char *street_name;
    char *city;
    char *prov __attribute__((aligned(16))) ;
    char *postal_code __attribute__((aligned(16))) ;
} first_address ;
```

Related information


- “The `__align` type qualifier” on page 74

- "Aligning data" in the *XL C/C++ Programming Guide*
- "The `__alignof__` operator" on page 135
- "The aligned type attribute" on page 79

The common and nocommon variable attributes

The variable attribute `common` allows you to specify that an uninitialized global variable or a global variable explicitly initialized to 0 should be allocated in the common section of the object file. The variable attribute `nocommon` specifies that an uninitialized global variable should be allocated in the data section of the object file. The variable is automatically initialized to zero.

nocommon and common variable attribute syntax



For example:

```

int i __attribute__((nocommon));    /* allocate i at .data */
int k __attribute__((common));     /* allocate k at .comm */

```

You can only apply the variable attributes to global scalar or aggregate variables. If you try to assign either attribute to a static or automatic variable or structure or union member, the attribute is ignored and a warning is issued.

Note that using `nocommon` to allocate uninitialized global variables in the data section can dramatically increase the size of the generated object. Also, specifying `nocommon` on a global variable that is simultaneously defined in different object files will cause an error at link time; such variables should be defined in one file and referred to in other files with an `extern` declaration.

The attributes take precedence over the `-qcommon|nocommon` compiler option.

If multiple specifications of the attribute appear in the same attribute statement, the last one specified will take effect. For example:

```

int i __attribute__((common, nocommon));    /* allocate i at .data */
int k __attribute__((common, nocommon, common)); /* allocate k at .comm */

```

If both the `common` or `nocommon` attribute and the `section` attribute are applied to the same variable, the `section` attribute takes precedence.

Related References

- `-qcommon` in the *XL C/C++ Compiler Reference*

The packed variable attribute

The variable attribute `packed` allows you to override the default alignment mode, to reduce the alignment for all members of an aggregate, or selected members of an aggregate to the smallest possible alignment: one byte for a member and one bit for a bit field member.

section variable attribute syntax

```
▶ declarator →  
▶ _attribute_ (( section ( section_name ) ) )
```

The *section_name* specifies a named section as a string literal, maximum length of 16 characters, not counting spaces. Spaces in the string are ignored.

The section variable attribute can be applied to a declaration or definition of the following types of variables:

- initialized or static global or namespace variables
- static local variables
-  uninitialized global or namespace variables
-  static structure or class member variables

A section attribute applied to a local variable with automatic storage duration is ignored with a warning because such variables are stored on the stack.

 A section attribute applied to a structure member is ignored with a warning. A section attribute applied to an uninitialized global variable is ignored without a warning; the symbols for uninitialized global variables are always placed in the common section.

When multiple section attributes are applied to a variable declaration, the last specification prevails. The section indicated in the prevailing variable declaration should match that of the variable definition because a variable definition cannot be overwritten. Each defined variable can reside in only one section.

The section attribute overrides the `-qcommon|nocommon` compiler option and the `common|nocommon` attribute. That is, if both attributes are specified for the same variable, the section attribute takes priority.

A named section can be used for multiple variables, but not for both variables and functions in the same compilation unit.

Related information

- “The section function attribute” on page 223
- “The common and nocommon variable attributes” on page 108
- `-qcommon` in the *XL C/C++ Compiler Reference*

The weak variable attribute

The weak variable attribute causes the symbol resulting from the variable declaration to appear in the object file as a weak symbol, rather than a global one. The language feature provides the programmer writing library functions with a way to allow variable definitions in user code to override the library declaration without causing duplicate name errors.

weak variable attribute syntax

►► `__attribute__((weak
 __weak__))` ◀◀

Related information

- `#pragma weak` in the *XL C/C++ Compiler Reference*
- “The weak function attribute” on page 224

End of IBM extension

Chapter 5. Type conversions

An expression of a given type is *implicitly converted* in the following situations:

- The expression is used as an operand of an arithmetic or logical operation.
- The expression is used as a condition in an if statement or an iteration statement (such as a for loop). The expression will be converted to a Boolean (or an integer in C89).
- The expression is used in a switch statement. The expression will be converted to an integral type.
- The expression is used as an initialization. This includes the following:
 - An assignment is made to an lvalue that has a different type than the assigned value.
 - A function is provided an argument value that has a different type than the parameter.
 - The value specified in the return statement of a function has a different type from the defined return type for the function.

You can perform *explicit* type conversions using a *cast* expression, as described in “Cast expressions” on page 156. The following sections discuss the conversions that are allowed by either implicit or explicit conversion, and the rules governing type promotions:

- “Arithmetic conversions and promotions”
- “Lvalue-to-rvalue conversions” on page 116
- “Pointer conversions” on page 117
- “Reference conversions (C++ only)” on page 118
- “Qualification conversions (C++ only)” on page 119
- “Function argument conversions” on page 119

Related information

- “User-defined conversions” on page 329
- “Conversion constructors” on page 330
- “Conversion functions” on page 332
- “The switch statement” on page 181
- “The if statement” on page 180
- “The return statement” on page 191

Arithmetic conversions and promotions

The following sections discuss the rules for the standard conversions for arithmetic types:

- “Integral conversions” on page 114
- “Floating-point conversions” on page 114
- “Boolean conversions” on page 114

If two operands in an expression have different types, they are subject to the rules of the *usual arithmetic conversions*, as described in “Integral and floating-point promotions” on page 115.

Integral conversions

Unsigned integer to unsigned integer or signed integer to signed integer

If the types are identical, there is no change. If the types are of a different size, and the value can be represented by the new type, the value is not changed; if the value cannot be represented by the new type, truncation or sign shifting will occur.

Signed integer to unsigned integer

The resulting value is the smallest unsigned integer type congruent to the source integer. If the value cannot be represented by the new type, truncation or sign shifting will occur.

Unsigned integer to signed integer

If the signed type is large enough to hold the original value, there is no change. If the value can be represented by the new type, the value is not changed; if the value cannot be represented by the new type, truncation or sign shifting will occur.

Signed and unsigned character types to integer

If the original value can be represented by `int`, it is represented as `int`. If the value cannot be represented by `int`, it is promoted to unsigned `int`.

Wide character type `wchar_t` to integer

If the original value can be represented by `int`, it is represented as `int`. If the value cannot be represented by `int`, it is promoted to the smallest type that can hold it: unsigned `int`, `long`, or unsigned `long`.

Signed and unsigned integer bit field to integer

If the original value can be represented by `int`, it is represented as `int`. If the value cannot be represented by `int`, it is promoted to unsigned `int`.

Enumeration type to integer

If the original value can be represented by `int`, it is represented as `int`. If the value cannot be represented by `int`, it is promoted to the smallest type that can hold it: unsigned `int`, `long`, or unsigned `long`. Note that an enumerated type can be converted to an integral type, but an integral type cannot be converted to an enumeration.

Boolean conversions

Boolean to integer

C If the Boolean value is 0, the result is an `int` with a value of 0. If the Boolean value is 1, the result is an `int` with a value of 1. **C++** If the Boolean value is `false`, the result is an `int` with a value of 0. If the Boolean value is `true`, the result is an `int` with a value of 1.

Scalar to Boolean

C If the scalar value is equal to 0, the Boolean value is 0; otherwise the Boolean value is 1. **C++** A zero, null pointer, or null member pointer value is converted to `false`. All other values are converted to `true`.

Floating-point conversions

The standard rule for converting between real floating-point types is as follows:

If the value being converted can be represented exactly in the new type, it is unchanged. If the value being converted is in the range of values that can be represented but cannot be represented exactly, the result is rounded, according to

the current compile-time or runtime rounding mode in effect. If the value being converted is outside the range of values that can be represented, the result is dependant on the rounding mode.

Integer to floating-point

If the value being converted can be represented exactly in the new type, it is unchanged. If the value being converted is in the range of values that can be represented but cannot be represented exactly, the result is correctly rounded. If the value being converted is outside the range of values that can be represented, the result is quiet NaN.

Floating-point to integer

The fractional part is discarded (i.e., the value is truncated toward zero). If the value of the integral part cannot be represented by the integer type, the result is one of the following:

- If the integer type is unsigned, the result is the largest representable number if the floating-point number is positive, or 0 otherwise.
- If the integer type is signed, the result is the most negative or positive representable number according to the sign of the floating-point number.

Complex conversions

Complex to complex

If the types are identical, there is no change. If the types are of a different size, and the value can be represented by the new type, the value is not changed; if the value cannot be represented by the new type, both real and imaginary parts are converted according to the standard conversion rule given above.

Complex to real

The imaginary part of the complex value is discarded. If necessary, the value of the real part is converted according to the standard conversion rule given above.

Real to complex

The source value is used as the real part of the complex value, and converted, if necessary, according to the standard conversion rule given above. The value of the imaginary part is zero.

Integral and floating-point promotions

When different arithmetic types are used as operands in certain types of expressions, standard conversions known as *usual arithmetic conversions* are applied. These conversions are applied according to the rank of the arithmetic type: the operand with a type of lower rank is converted to the type of the operand with a higher rank. This is known as integral or floating point *promotion*.

For example, when the values of two different integral types are added together, both values are first converted to the same type: when a short int value and an int value are added together, the short int value is converted to the int type. Chapter 6, “Expressions and operators,” on page 121 provides a list of the operators and expressions that participate in the usual arithmetic conversions.

The ranking of arithmetic types, listed from highest to lowest, is as follows:

Table 20. Conversion rankings for floating-point types

Operand type
long double or long double _Complex

Table 20. Conversion rankings for floating-point types (continued)

Operand type
double or double _Complex
float or float _Complex

Table 21. Conversion rankings for integer types

Operand type
unsigned long long or unsigned long long int
long long or long long int ¹
unsigned long int ¹
long int ²
unsigned int ²
int and enumerated types
short int
char, signed char and unsigned char
Boolean

Notes:

1. If one operand has unsigned long int type and the other operand has long long int (or long long) type but the value of the unsigned long int cannot be represented in a long long int (or long long), both operands are converted to unsigned long long int (or unsigned long long).
2. If one operand has unsigned int type and the other operand has long int type but the value of the unsigned int cannot be represented in a long int, both operands are converted to unsigned long int.

Related information

- “Integral types” on page 53
- “Boolean types” on page 54
- “Floating-point types” on page 54
- “Character types” on page 55
- “Enumerations” on page 67
- “Binary expressions” on page 139

Lvalue-to-rvalue conversions

If an lvalue appears in a situation in which the compiler expects an rvalue, the compiler converts the lvalue to an rvalue. The following table lists exceptions to this:

Situation before conversion	Resulting behavior
The lvalue is a function type.	
The lvalue is an array.	
The type of the lvalue is an incomplete type.	compile-time error
The lvalue refers to an uninitialized object.	undefined behavior

Situation before conversion	Resulting behavior
The lvalue refers to an object not of the type of the rvalue, nor of a type derived from the type of the rvalue.	undefined behavior
 The lvalue is a nonclass type, qualified by either <code>const</code> or <code>volatile</code> .	The type after conversion is not qualified by either <code>const</code> or <code>volatile</code> .

Related information

- “Lvalues and rvalues” on page 121

Pointer conversions

Pointer conversions are performed when pointers are used, including pointer assignment, initialization, and comparison.

C only

Conversions that involve pointers must use an explicit type cast. The exceptions to this rule are the allowable assignment conversions for C pointers. In the following table, a `const`-qualified lvalue cannot be used as a left operand of the assignment.

Table 22. Legal assignment conversions for C pointers

Left operand type	Permitted right operand types
pointer to (object) T	<ul style="list-style-type: none"> • the constant <code>0</code> • a pointer to a type compatible with T • a pointer to <code>void</code> (<code>void*</code>)
pointer to (function) F	<ul style="list-style-type: none"> • the constant <code>0</code> • a pointer to a function compatible with F

The referenced type of the left operand must have the same qualifiers as the right operand. An object pointer may be an incomplete type if the other pointer has type `void*`.

End of C only

Zero constant to null pointer

A constant expression that evaluates to zero is a *null pointer constant*. This expression can be converted to a pointer. This pointer will be a null pointer (pointer with a zero value), and is guaranteed not to point to any object.

 A constant expression that evaluates to zero can also be converted to the null pointer to a member.

Array to pointer

An lvalue or rvalue with type “array of N,” where N is the type of a single element of the array, to `N*`. The result is a pointer to the initial element of the array. A conversion cannot be performed if the expression is used as the operand of the `&` (address) operator or the `sizeof` operator.

Function to pointer

An lvalue that is a function can be converted to an rvalue that is a pointer to a function of the same type, except when the expression is used as the operand of the `&` (address) operator, the `()` (function call) operator, or the `sizeof` operator.

Related information

- “Pointers” on page 86
- “Integer constant expressions” on page 124
- “Arrays” on page 90
- “Pointers to functions” on page 232
- “Pointers to members” on page 274
- “Pointer conversions” on page 308

Conversion to void*

C pointers are not necessarily the same size as type `int`. Pointer arguments given to functions should be explicitly cast to ensure that the correct type expected by the function is being passed. The generic object pointer in C is `void*`, but there is no generic function pointer.

Any pointer to an object, optionally type-qualified, can be converted to `void*`, keeping the same `const` or `volatile` qualifications.

C only

The allowable assignment conversions involving `void*` as the left operand are shown in the following table.

Table 23. Legal assignment conversions in C for `void*`

Left operand type	Permitted right operand types
<code>(void*)</code>	<ul style="list-style-type: none">• The constant <code>0</code>.• A pointer to an object. The object may be of incomplete type.• <code>(void*)</code>

End of C only

C++ only

Pointers to functions cannot be converted to the type `void*` with a standard conversion: this can be accomplished explicitly, provided that a `void*` has sufficient bits to hold it.

End of C++ only

Related information

- “The void type” on page 56

Reference conversions (C++ only)

A reference conversion can be performed wherever a reference initialization occurs, including reference initialization done in argument passing and function return values. A reference to a class can be converted to a reference to an accessible base class of that class as long as the conversion is not ambiguous. The result of the conversion is a reference to the base class subobject of the derived class object.

Reference conversion is allowed if the corresponding pointer conversion is allowed.

Related information

- “References (C++ only)” on page 93
- “Initialization of references (C++ only)” on page 104
- “Function calls” on page 225
- “Function return values” on page 213

Qualification conversions (C++ only)

An type-qualified rvalue of any type, containing zero or more `const` or `volatile` qualifications, can be converted to an rvalue of type-qualified type where the second rvalue contains more `const` or `volatile` qualifications than the first rvalue.

An rvalue of type pointer to member of a class can be converted to an rvalue of type pointer to member of a class if the second rvalue contains more `const` or `volatile` qualifications than the first rvalue.

Related information

- “Type qualifiers” on page 72

Function argument conversions

When a function is called, if a function declaration is present and includes declared argument types, the compiler performs type checking. The compiler compares the data types provided by the calling function with the data types that the called function expects and performs necessary type conversions. For example, when function `funct` is called, argument `f` is converted to a `double`, and argument `c` is converted to an `int`:

```
char * funct (double d, int i);
    /* ... */
int main(void)
{
    float f;
    char c;
    funct(f, c) /* f is converted to a double, c is converted to an int */
    return 0;
}
```

If no function declaration is visible when a function is called, or when an expression appears as an argument in the variable part of a prototype argument list, the compiler performs default argument promotions or converts the value of the expression before passing any arguments to the function. The automatic conversions consist of the following:

- Integral and floating-point values are promoted.
- Arrays or functions are converted to pointers.
-  Non-static class member functions are converted to pointers to members.

Related information

- “Integral and floating-point promotions” on page 115
- “The `transparent_union` type attribute (C only)” on page 80
- “Function call expressions” on page 127
- “Function calls” on page 225

Chapter 6. Expressions and operators

Expressions are sequences of operators, operands, and punctuators that specify a computation. The evaluation of expressions is based on the operators that the expressions contain and the context in which they are used. An expression can result in a value and can produce *side effects*. A side effect is a change in the state of the execution environment.

The following sections describe these types of expressions:

- “Lvalues and rvalues”
- “Primary expressions” on page 123
- “Function call expressions” on page 127
- “Member expressions” on page 128
- “Unary expressions” on page 128
- “Binary expressions” on page 139
- “Conditional expressions” on page 153
- “Compound literal expressions” on page 164
- “Cast expressions” on page 156
- “new expressions (C++ only)” on page 165
- “delete expressions (C++ only)” on page 169
- “throw expressions (C++ only) (PPU only)” on page 170
-  “Label value expressions” on page 170

“Operator precedence and associativity” on page 170 provides tables listing the precedence of all the operators described in the various sections listed above.

 C++ operators can be defined to behave differently when applied to operands of class type. This is called operator *overloading*, and is described in “Overloading operators” on page 245.

Lvalues and rvalues

An *object* is a region of storage that can be examined and stored into. An *lvalue* is an expression that refers to such an object. An lvalue does not necessarily permit modification of the object it designates. For example, a *const* object is an lvalue that cannot be modified. The term *modifiable lvalue* is used to emphasize that the lvalue allows the designated object to be changed as well as examined. The following object types are lvalues, but not modifiable lvalues:

- An array type
- An incomplete type
- A *const*-qualified type
- A structure or union type with one of its members qualified as a *const* type

Because these lvalues are not modifiable, they cannot appear on the left side of an assignment statement.

The term *rvalue* refers to a data value that is stored at some address in memory. An rvalue is an expression that cannot have a value assigned to it. Both a literal

constant and a variable can serve as an rvalue. When an lvalue appears in a context that requires an rvalue, the lvalue is implicitly converted to an rvalue. The reverse, however, is not true: an rvalue cannot be converted to an lvalue. Rvalues always have complete types or the void type.

C C defines a *function designator* as an expression that has function type. A function designator is distinct from an object type or an lvalue. It can be the name of a function or the result of dereferencing a function pointer. The C language also differentiates between its treatment of a function pointer and an object pointer.

C++ On the other hand, in C++, a function call that returns a reference is an lvalue. Otherwise, a function call is an rvalue expression. In C++, every expression produces an lvalue, an rvalue, or no value.

In both C and C++, certain operators require lvalues for some of their operands. The table below lists these operators and additional constraints on their usage.

Operator	Requirement
& (unary)	Operand must be an lvalue.
++ --	Operand must be an lvalue. This applies to both prefix and postfix forms.
= += -= *= %= <<= >>= &= ^= =	Left operand must be an lvalue.

For example, all assignment operators evaluate their right operand and assign that value to their left operand. The left operand must be a modifiable lvalue or a reference to a modifiable object.

The address operator (&) requires an lvalue as an operand while the increment (++) and the decrement (--) operators require a modifiable lvalue as an operand. The following example shows expressions and their corresponding lvalues.

Expression	Lvalue
<code>x = 42</code>	<code>x</code>
<code>*ptr = newvalue</code>	<code>*ptr</code>
<code>a++</code>	<code>a</code>
C++ <code>int& f()</code>	The function call to <code>f()</code>

IBM extension

When compiled with the GNU C language extensions enabled, compound expressions, conditional expressions, and casts are allowed as lvalues, provided that their operands are lvalues. The use of this language extension is deprecated for C++ code.

A compound expression can be assigned if the last expression in the sequence is an lvalue. The following expressions are equivalent:

```
(x + 1, y) *= 42;
x + 1, (y *=42);
```

The address operator can be applied to a compound expression, provided the last expression in the sequence is an lvalue. The following expressions are equivalent:

```
&(x + 1, y);
x + 1, &y;
```

A conditional expression can be a valid lvalue if its type is not void and both of its branches for true and false are valid lvalues. Casts are valid lvalues if the operand is an lvalue. The primary restriction is that you cannot take the address of an lvalue cast.

End of IBM extension

Related information

- “Arrays” on page 90
- “Lvalue-to-rvalue conversions” on page 116

Primary expressions

Primary expressions fall into the following general categories:

- Names (identifiers)
- Literals (constants)
- Integer constant expressions
- Identifier expressions (C++ only)
- Parenthesized expressions ()
-  The `this` pointer (described in “The `this` pointer” on page 276)
-  Names qualified by the scope resolution operator (`::`)

Names

The value of a name depends on its type, which is determined by how that name is declared. The following table shows whether a name is an lvalue expression.

Table 24. *Primary expressions: Names*

Name declared as	Evaluates to	Is an lvalue?
Variable of arithmetic, pointer, enumeration, structure, or union type	An object of that type	yes
Enumeration constant	The associated integer value	no
Array	That array. In contexts subject to conversions, a pointer to the first object in the array, except where the name is used as the argument to the <code>sizeof</code> operator.	 no  yes
Function	That function. In contexts subject to conversions, a pointer to that function, except where the name is used as the argument to the <code>sizeof</code> operator, or as the function in a function call expression.	 no  yes

As an expression, a name may not refer to a label, typedef name, structure member, union member, structure tag, union tag, or enumeration tag. Names used

for these purposes reside in a namespace that is separate from that of names used in expressions. However, some of these names may be referred to within expressions by means of special constructs: for example, the dot or arrow operators may be used to refer to structure and union members; typedef names may be used in casts or as an argument to the sizeof operator.

Literals

A literal is a numeric constant or string literal. When a literal is evaluated as an expression, its value is a constant. A lexical constant is never an lvalue. However, a string literal is an lvalue.

Related information

- “Literals” on page 19
- “The this pointer” on page 276

Integer constant expressions

An *integer compile-time constant* is a value that is determined during compilation and cannot be changed at run time. An *integer compile-time constant expression* is an expression that is composed of constants and evaluated to a constant.

An integer constant expression is an expression that is composed of only the following:

- literals
- enumerators
- const variables
- static data members of integral or enumeration types
- casts to integral types
- sizeof expressions, where the operand is not a variable length array

The sizeof operator applied to a variable length array type is evaluated at run time, and therefore is not a constant expression.

You must use an integer constant expression in the following situations:

- In the subscript declarator as the description of an array bound.
- After the keyword case in a switch statement.
- In an enumerator, as the numeric value of an enumeration constant.
- In a bit-field width specifier.
- In the preprocessor #if statement. (Enumeration constants, address constants, and sizeof cannot be specified in a preprocessor #if statement.)

Related information

- “The sizeof operator” on page 135

Identifier expressions (C++ only)


An identifier expression, or *id-expression*, is a restricted form of primary expression. Syntactically, an *id-expression* requires a higher level of complexity than a simple identifier to provide a name for all of the language elements of C++.

An *id-expression* can be either a qualified or unqualified identifier. It can also appear after the dot and arrow operators.


Identifier expression syntax



unqualified_id:



qualified_id:




Related information


- “Identifiers” on page 15
- Chapter 4, “Declarators,” on page 83

Parenthesized expressions ()

Use parentheses to explicitly force the order of expression evaluation. The following expression does not use parentheses to group operands and operators. The parentheses surrounding `weight`, `zipcode` are used to form a function call. Note how the compiler groups the operands and operators in the expression according to the rules for operator precedence and associativity:



The following expression is similar to the previous expression, but it contains parentheses that change how the operands and operators are grouped:



In an expression that contains both associative and commutative operators, you can use parentheses to specify the grouping of operands with operators. The parentheses in the following expression guarantee the order of grouping operands with the operators:

```
x = f + (g + h);
```

Related information

- “Operator precedence and associativity” on page 170

Scope resolution operator :: (C++ only)

The :: (scope resolution) operator is used to qualify hidden names so that you can still use them. You can use the unary scope operator if a namespace scope or global scope name is hidden by an explicit declaration of the same name in a block or class. For example:

```
int count = 0;

int main(void) {
    int count = 0;
    ::count = 1; // set global count to 1
    count = 2;  // set local count to 2
    return 0;
}
```

The declaration of count declared in the main function hides the integer named count declared in global namespace scope. The statement ::count = 1 accesses the variable named count declared in global namespace scope.

You can also use the class scope operator to qualify class names or class member names. If a class member name is hidden, you can use it by qualifying it with its class name and the class scope operator.

In the following example, the declaration of the variable `X` hides the class type `X`, but you can still use the static class member `count` by qualifying it with the class type `X` and the scope resolution operator.

```
#include <iostream>
using namespace std;

class X
{
public:
    static int count;
};
int X::count = 10;           // define static data member

int main ()
{
    int X = 0;              // hides class type X
    cout << X::count << endl; // use static member of class X
}
```

Related information

- “Scope of class names” on page 263
- Chapter 9, “Namespaces (C++ only),” on page 235

Function call expressions

A *function call* is an expression containing the function name followed by the function call operator, `()`. If the function has been defined to receive parameters, the values that are to be sent into the function are listed inside the parentheses of the function call operator. The argument list can contain any number of expressions separated by commas. It can also be empty.

The type of a function call expression is the return type of the function. This type can either be a complete type, a reference type, or the type `void`.  A function call expression is always an rvalue.  A function call is an lvalue if and only if the type of the function is a reference.

Here are some examples of the function call operator:

```
stub()
overdue(account, date, amount)
notify(name, date + 5)
report(error, time, date, ++num)
```

The order of evaluation for function call arguments is not specified. In the following example:

```
method(sample1, batch.process--, batch.process);
```

the argument `batch.process--` might be evaluated last, causing the last two arguments to be passed with the same value.

Related information

- “Function argument conversions” on page 119
- “Function calls” on page 225

Member expressions

Member expressions indicate members of classes, structures, or unions. The member operators are:

- Dot operator `.`
- Arrow operator `->`

Dot operator `.`

The `.` (dot) operator is used to access class, structure, or union members. The member is specified by a postfix expression, followed by a `.` (dot) operator, followed by a possibly qualified identifier or a pseudo-destructor name. (A *pseudo-destructor* is a destructor of a nonclass type.) The postfix expression must be an object of type `class`, `struct` or `union`. The name must be a member of that object.

The value of the expression is the value of the selected member. If the postfix expression and the name are lvalues, the expression value is also an lvalue. If the postfix expression is type-qualified, the same type qualifiers will apply to the designated member in the resulting expression.

Related information

- “Access to structure and union members” on page 66
- “Pseudo-destructors” on page 328

Arrow operator `->`

The `->` (arrow) operator is used to access class, structure or union members using a pointer. A postfix expression, followed by an `->` (arrow) operator, followed by a possibly qualified identifier or a pseudo-destructor name, designates a member of the object to which the pointer points. (A *pseudo-destructor* is a destructor of a nonclass type.) The postfix expression must be a pointer to an object of type `class`, `struct` or `union`. The name must be a member of that object.

The value of the expression is the value of the selected member. If the name is an lvalue, the expression value is also an lvalue. If the expression is a pointer to a qualified type, the same type-qualifiers will apply to the designated member in the resulting expression.

Related information

- “Pointers” on page 86
- “Access to structure and union members” on page 66
- Chapter 12, “Class members and friends (C++ only),” on page 269
- “Pseudo-destructors” on page 328

Unary expressions

A *unary expression* contains one operand and a unary operator.

The supported unary operators are:

- Increment operator `++`
- Decrement operator `--`
- Unary plus operator `+`
- Unary minus operator `-`

- Logical negation operator !
- Bitwise negation operator ~
- Address operator &
- Indirection operator *
-  typeid
-  alignof
- sizeof
-  typeof
-  __real__ and __imag__
-  vec_step

All unary operators have the same precedence and have right-to-left associativity, as shown in Table 28 on page 172.

As indicated in the descriptions of the operators, the usual arithmetic conversions are performed on the operands of most unary expressions.

Related information

- “Pointer arithmetic” on page 87
- “Lvalues and rvalues” on page 121
- “Arithmetic conversions and promotions” on page 113

Increment operator ++

The ++ (increment) operator adds 1 to the value of a scalar operand, or if the operand is a pointer, increments the operand by the size of the object to which it points. The operand receives the result of the increment operation. The operand must be a modifiable lvalue of arithmetic or pointer type.

You can put the ++ before or after the operand. If it appears before the operand, the operand is incremented. The incremented value is then used in the expression. If you put the ++ after the operand, the value of the operand is used in the expression *before* the operand is incremented. For example:

```
play = ++play1 + play2++;
```

is similar to the following expressions; play2 is altered before play:

```
int temp, temp1, temp2;

temp1 = play1 + 1;
temp2 = play2;
play1 = temp1;
temp = temp1 + temp2;
play2 = play2 + 1;
play = temp;
```

The result has the same type as the operand after integral promotion.

The usual arithmetic conversions on the operand are performed.

IBM extension

The increment operator has been extended to handle complex types. The operator works in the same manner as it does on a real type, except that only the real part

of the operand is incremented, and the imaginary part is unchanged.

End of IBM extension

Decrement operator --

The -- (decrement) operator subtracts 1 from the value of a scalar operand, or if the operand is a pointer, decreases the operand by the size of the object to which it points. The operand receives the result of the decrement operation. The operand must be a modifiable lvalue.

You can put the -- before or after the operand. If it appears before the operand, the operand is decremented, and the decremented value is used in the expression. If the -- appears after the operand, the current value of the operand is used in the expression and the operand is decremented.

For example:

```
play = --play1 + play2--;
```

is similar to the following expressions; play2 is altered before play:

```
int temp, temp1, temp2;

temp1 = play1 - 1;
temp2 = play2;
play1 = temp1;
temp = temp1 + temp2;
play2 = play2 - 1;
play = temp;
```

The result has the same type as the operand after integral promotion, but is not an lvalue.

The usual arithmetic conversions are performed on the operand.

IBM extension

The decrement operator has been extended to handle complex types, for compatibility with GNU C. The operator works in the same manner as it does on a real type, except that only the real part of the operand is decremented, and the imaginary part is unchanged.

End of IBM extension

Unary plus operator +

The + (unary plus) operator maintains the value of the operand. The operand can have any arithmetic type or pointer type. The result is not an lvalue.

The result has the same type as the operand after integral promotion.

Note: Any plus sign in front of a constant is not part of the constant.

Unary minus operator -

The - (unary minus) operator negates the value of the operand. The operand can have any arithmetic type. The result is not an lvalue.

For example, if quality has the value 100, -quality has the value -100.

The result has the same type as the operand after integral promotion.

Note: Any minus sign in front of a constant is not part of the constant.

Logical negation operator !

The ! (logical negation) operator determines whether the operand evaluates to 0 (false) or nonzero (true).

C The expression yields the value 1 (true) if the operand evaluates to 0, and yields the value 0 (false) if the operand evaluates to a nonzero value.

C++ The expression yields the value true if the operand evaluates to false (0), and yields the value false if the operand evaluates to true (nonzero). The operand is implicitly converted to bool, and the type of the result is bool.

The following two expressions are equivalent:

```
!right;  
right == 0;
```

Related information

- “Boolean types” on page 54

Bitwise negation operator ~

The ~ (bitwise negation) operator yields the bitwise complement of the operand. In the binary representation of the result, every bit has the opposite value of the same bit in the binary representation of the operand. The operand must have an integral type. The result has the same type as the operand but is not an lvalue.

Suppose x represents the decimal value 5. The 16-bit binary representation of x is:
0000000000000101

The expression ~x yields the following result (represented here as a 16-bit binary number):

```
1111111111111010
```

Note that the ~ character can be represented by the trigraph ??-.

The 16-bit binary representation of ~0 is:

```
1111111111111111
```

IBM extension

The bitwise negation operator has been extended to handle complex types. With a complex type, the operator computes the complex conjugate of the operand by reversing the sign of the imaginary part.

End of IBM extension

Related information

- “Trigraph sequences” on page 37

Address operator &

The & (address) operator yields a pointer to its operand. The operand must be an lvalue, a function designator, or a qualified name. It cannot be a bit field, nor can it have the storage class register.

If the operand is an lvalue or function, the resulting type is a pointer to the expression type. For example, if the expression has type `int`, the result is a pointer to an object having type `int`.

If the operand is a qualified name and the member is not static, the result is a pointer to a member of class and has the same type as the member. The result is not an lvalue.

If `p_to_y` is defined as a pointer to an `int` and `y` as an `int`, the following expression assigns the address of the variable `y` to the pointer `p_to_y` :

```
p_to_y = &y;
```

IBM extension

The address operator has been extended to handle vector types. The result of the address operator applied to a vector type can be stored in a pointer to a compatible vector type. The address of a vector type can be used to initialize a pointer to vector type if both sides of the initialization have compatible types. A pointer to `void` can also be initialized with the address of a vector type.

End of IBM extension

C++ only

The ampersand symbol & is used in C++ as a reference declarator in addition to being the address operator. The meanings are related but not identical.

```
int target;
int &rTarg = target; // rTarg is a reference to an integer.
                    // The reference is initialized to refer to target.
void f(int*& p);    // p is a reference to a pointer
```

If you take the address of a reference, it returns the address of its target. Using the previous declarations, `&rTarg` is the same memory address as `&target`.

You may take the address of a register variable.

You can use the & operator with overloaded functions only in an initialization or assignment where the left side uniquely determines which version of the overloaded function is used.

End of C++ only

 The address of a label can be taken using the GNU C address operator `&&`. The label can thus be used as a value.

Related information

- “Indirection operator `*`” on page 133
- “Pointers” on page 86
- “References (C++ only)” on page 93

If *expr* is a reference or a dereferenced pointer to a polymorphic class, `typeid` will return a `type_info` object that represents the object that the reference or pointer denotes at run time. If it is not a polymorphic class, `typeid` will return a `type_info` object that represents the type of the reference or dereferenced pointer. The following example demonstrates this:

```
#include <iostream>
#include <typeinfo>
using namespace std;

struct A { virtual ~A() { } };
struct B : A { };

struct C { };
struct D : C { };

int main() {
    B bobj;
    A* ap = &bobj;
    A& ar = bobj;
    cout << "ap: " << typeid(*ap).name() << endl;
    cout << "ar: " << typeid(ar).name() << endl;

    D dobj;
    C* cp = &dobj;
    C& cr = dobj;
    cout << "cp: " << typeid(*cp).name() << endl;
    cout << "cr: " << typeid(cr).name() << endl;
}
```

The following is the output of the above example:

```
ap: B
ar: B
cp: C
cr: C
```

Classes A and B are polymorphic; classes C and D are not. Although `cp` and `cr` refer to an object of type D, `typeid(*cp)` and `typeid(cr)` return objects that represent class C.

Lvalue-to-rvalue, array-to-pointer, and function-to-pointer conversions will not be applied to *expr*. For example, the output of the following example will be `int [10]`, not `int *`:

```
#include <iostream>
#include <typeinfo>
using namespace std;

int main() {
    int myArray[10];
    cout << typeid(myArray).name() << endl;
}
```

If *expr* is a class type, that class must be completely defined.

The `typeid` operator ignores top-level `const` or `volatile` qualifiers.

Related information

- “Type names” on page 85
- “The `typeof` operator” on page 137

The `__alignof__` operator

IBM extension

The `__alignof__` operator is a language extension to C99 and Standard C++ that returns the number of bytes used in the alignment of its operand. The operand can be an expression or a parenthesized type identifier. If the operand is an expression representing an lvalue, the number returned by `__alignof__` represents the alignment that the lvalue is known to have. The type of the expression is determined at compile time, but the expression itself is not evaluated. If the operand is a type, the number represents the alignment usually required for the type on the target platform.

The `__alignof__` operator may not be applied to the following:

- An lvalue representing a bit field
- A function type
- An undefined structure or class
- An incomplete type (such as `void`)

`__alignof__` operator syntax

► `__alignof__` *unary_expression*
 (*type-id*) ◄

If *type-id* is a reference or a referenced type, the result is the alignment of the referenced type. If *type-id* is an array, the result is the alignment of the array element type. If *type-id* is a fundamental type, the result is implementation-defined.

For example, on Linux, `__alignof__(long)` returns 4 in 32-bit mode, and 8 in 64-bit mode.

The operand of `__alignof__` can be a vector type. For example,

```
vector unsigned int v1 = (vector unsigned int)(10);
vector unsigned int *pv1 = &v1;
__alignof__(v1); // vector type alignment: 16.
__alignof__(&v1); // address of vector alignment: 4.
__alignof__(*pv1); // dereferenced pointer to vector alignment: 16.
__alignof__(pv1); // pointer to vector alignment: 4.
__alignof__(vector signed char); // vector type alignment: 16.
```

When `__attribute__((aligned))` is used to increase the alignment of a variable of vector type, the value returned by the `__alignof__` operator is the alignment factor specified by `__attribute__((aligned))`.

Related information

- “The aligned variable attribute” on page 107

End of IBM extension

The `sizeof` operator

The `sizeof` operator yields the size in bytes of the operand, which can be an expression or the parenthesized name of a type.

sizeof operator syntax

► sizeof $\left[\begin{array}{l} \text{expr} \\ \text{(-type-name-)} \end{array} \right]$ ◄

The result for either kind of operand is not an lvalue, but a constant integer value. The type of the result is the unsigned integral type `size_t` defined in the header file `stddef.h`.

Except in preprocessor directives, you can use a `sizeof` expression wherever an integral constant is required. One of the most common uses for the `sizeof` operator is to determine the size of objects that are referred to during storage allocation, input, and output functions.

Another use of `sizeof` is in porting code across platforms. You can use the `sizeof` operator to determine the size that a data type represents. For example:

```
sizeof(int);
```

The `sizeof` operator applied to a type name yields the amount of memory that would be used by an object of that type, including any internal or trailing padding.

IBM extension

The operand of the `sizeof` operator can be a vector type or the result of dereferencing a pointer to vector type. In these cases, the return value of `sizeof` is always 16.

```
vector<bool> v1;
vector<bool> *pv1 = &v1;
sizeof(v1); // vector type: 16.
sizeof(&v1); // address of vector: 4.
sizeof(*pv1); // dereferenced pointer to vector: 16.
sizeof(pv1); // pointer to vector: 4.
sizeof(vector<bool>); // vector type: 16.
```

End of IBM extension

For compound types, results are as follows:

Operand	Result
An array	The result is the total number of bytes in the array. For example, in an array with 10 elements, the size is equal to 10 times the size of a single element. The compiler does not convert the array to a pointer before evaluating the expression.
 A class	The result is always nonzero, and is equal to the number of bytes in an object of that class including any padding required for placing class objects in an array.
 A reference	The result is the size of the referenced object.

The `sizeof` operator may not be applied to:

- A bit field
- A function type
- An undefined structure or class

- An incomplete type (such as `void`)

The `sizeof` operator applied to an expression yields the same result as if it had been applied to only the name of the type of the expression. At compile time, the compiler analyzes the expression to determine its type. None of the usual type conversions that occur in the type analysis of the expression are directly attributable to the `sizeof` operator. However, if the operand contains operators that perform conversions, the compiler does take these conversions into consideration in determining the type. For example, the second line of the following sample causes the usual arithmetic conversions to be performed. Assuming that a `short` uses 2 bytes of storage and an `int` uses 4 bytes,

```
short x; ... sizeof (x)          /* the value of sizeof operator is 2 */
short x; ... sizeof (x + 1)     /* value is 4, result of addition is type int */
```

The result of the expression `x + 1` has type `int` and is equivalent to `sizeof(int)`. The value is also 4 if `x` has type `char`, `short`, or `int` or any enumeration type.

Related information

- “Type names” on page 85
- “Integer constant expressions” on page 124
- “Arrays” on page 90
- “References (C++ only)” on page 93

The `typeof` operator

IBM extension

The `typeof` operator returns the type of its argument, which can be an expression or a type. The language feature provides a way to derive the type from an expression. Given an expression `e`, `__typeof__(e)` can be used anywhere a type name is needed, for example in a declaration or in a cast. The alternate spelling of the keyword, `__typeof__`, is recommended.

The `typeof` operator is extended to accept a vector type as its operand.

`typeof` operator syntax

```
▶▶ [__typeof__] ( [expr] ) ▶▶
   [typeof]      [type-name]
```

A `typeof` construct itself is not an expression, but the name of a type. A `typeof` construct behaves like a type name defined using `typedef`, although the syntax resembles that of `sizeof`.

The following examples illustrate its basic syntax. For an expression `e`:

```
int e;
__typeof__(e + 1) j; /* the same as declaring int j; */
e = (__typeof__(e)) f; /* the same as casting e = (int) f; */
```

Using a `typeof` construct is equivalent to declaring a `typedef` name. Given

```
int T[2];
int i[2];
```

you can write

```

__typeof__(i) a;          /* all three constructs have the same meaning */
__typeof__(int[2]) a;
__typeof__(T) a;

```

The behavior of the code is as if you had declared `int a[2];`.

For a bit field, `typeof` represents the underlying type of the bit field. For example, `int m:2;`, the `typeof(m)` is `int`. Since the bit field property is not reserved, `n` in `typeof(m) n;` is the same as `int n`, but not `int n:2`.

The `typeof` operator can be nested inside `sizeof` and itself. The following declarations of `arr` as an array of pointers to `int` are equivalent:

```

int *arr[10];          /* traditional C declaration          */
__typeof__(__typeof__(int *)[10]) a; /* equivalent declaration */

```

The `typeof` operator can be useful in macro definitions where expression `e` is a parameter. For example,

```

#define SWAP(a,b) { __typeof__(a) temp; temp = a; a = b; b = temp; }

```

Notes:

1. The `typeof` and `__typeof__` keywords are supported as follows:
 - **C** The `__typeof__` keyword is recognized under compilation with the `ppuxlc` or `spuxlc` invocation command or the `-qlanglvl=extc89`, `-qlanglvl=extc99`, or `-qlanglvl=extended` options. The `typeof` keyword is only recognized under compilation with `-qkeyword=typeof`.
 - **C++** The `typeof` and `__typeof__` keywords are recognized by default.

Related information

- “Type names” on page 85
- “typedef definitions” on page 71
- “Vector types” on page 57

End of IBM extension

The `__real__` and `__imag__` operators

IBM extension

XL C/C++ extends the C99 and C++ standards to support the unary operators `__real__` and `__imag__`. These operators provide the ability to extract the real and imaginary parts of a complex type. These extensions have been implemented to ease the porting applications developed with GNU C.

`__real__` and `__imag__` operator syntax

```

➤ [__real__] (—var_identifier—) ➤
  [__imag__]

```

The `var_identifier` is the name of a previously declared complex variable. The `__real__` operator returns the real part of the complex variable, while the `__imag__` operator returns the imaginary part of the variable. If the operand of these operators is an lvalue, the resulting expression can be used in any context where lvalues are allowed. They are especially useful in initializations of complex variables, and as arguments to calls to library functions such as `printf` and `scanf` that have no format specifiers for complex types. For example:

```
float _Complex myvar;
__imag__(myvar) = 2.0f;
__real__(myvar) = 3.0f;
```

initializes the imaginary part of the complex variable `myvar` to $2.0i$ and the real part to 3.0, and

```
printf("myvar = %f + %f * i\n", __real__(myvar), __imag__(myvar));
```

prints:

```
myvar = 2.000000 + 3.000000 * i
```

Related information

- “Complex literals” on page 24
- “Complex floating-point types ” on page 55

End of IBM extension

The `vec_step` operator

IBM extension

The `vec_step` operator takes a vector type argument and returns an integer value representing the amount by which a pointer to a vector element should be incremented to move by 16 bytes. For complete documentation of this operator, see the *AltiVec Technology Programming Interface Manual*, available at http://www.freescale.com/files/32bit/doc/ref_manual/ALTIVECPIM.pdf

End of IBM extension

Binary expressions

A *binary expression* contains two operands separated by one operator. The supported binary operators are:

- “Assignment operators” on page 140
- Multiplication operator `*`
- Division operator `/`
- Remainder operator `%`
- Addition operator `+`
- Subtraction operator `-`
- Bitwise left and right shift operators `<<` `>>`
- Relational operators `<` `>` `<=` `>=`
- Equality and inequality operators `==` `!=`
- Bitwise AND operator `&`
- Bitwise exclusive OR operator `^`
- Bitwise inclusive OR operator `|`
- Logical AND operator `&&`
- Logical OR operator `||`
- Array subscripting operator `[]`
- “Comma operator `,`” on page 151
- Pointer to member operators `.*` `->*` (C++ only)

All binary operators have left-to-right associativity, but not all binary operators have the same precedence. The ranking and precedence rules for binary operators is summarized in Table 29 on page 172.

The order in which the operands of most binary operators are evaluated is not specified. To ensure correct results, avoid creating binary expressions that depend on the order in which the compiler evaluates the operands.

As indicated in the descriptions of the operators, the usual arithmetic conversions are performed on the operands of most binary expressions.

Related information

- “Lvalues and rvalues” on page 121
- “Arithmetic conversions and promotions” on page 113

Assignment operators

An *assignment expression* stores a value in the object designated by the left operand. There are two types of assignment operators:

- Simple assignment operator =
- Compound assignment operators

The left operand in all assignment expressions must be a modifiable lvalue. The type of the expression is the type of the left operand. The value of the expression is the value of the left operand after the assignment has completed.

 The result of an assignment expression is not an lvalue.  The result of an assignment expression is an lvalue.

All assignment operators have the same precedence and have right-to-left associativity.

Related information

- “Lvalues and rvalues” on page 121
- “Pointers” on page 86
- “Type qualifiers” on page 72

Simple assignment operator =

The simple assignment operator has the following form:

lvalue = *expr*

The operator stores the value of the right operand *expr* in the object designated by the left operand *lvalue*.

The left operand must be a modifiable lvalue. The type of an assignment operation is the type of the left operand.

If the left operand is not a class type or a vector type, the right operand is implicitly converted to the type of the left operand. This converted type will not be qualified by `const` or `volatile`.

If the left operand is a class type, that type must be complete. The copy assignment operator of the left operand will be called.

If the left operand is an object of reference type, the compiler will assign the value of the right operand to the object denoted by the reference.

IBM The assignment operator has been extended to permit operands of vector type. Both sides of an assignment expression must be of the same vector type.

Compound assignment operators

The compound assignment operators consist of a binary operator and the simple assignment operator. They perform the operation of the binary operator on both operands and store the result of that operation into the left operand, which must be a modifiable lvalue.

The following table shows the operand types of compound assignment expressions:

Operator	Left operand	Right operand
<code>+=</code> or <code>--</code>	Arithmetic	Arithmetic
<code>+=</code> or <code>--</code>	Pointer	Integral type
<code>*=</code> , <code>/=</code> , and <code>%=</code>	Arithmetic	Arithmetic
<code><<=</code> , <code>>>=</code> , <code>&=</code> , <code>^=</code> , and <code> =</code>	Integral type	Integral type

Note that the expression

```
a *= b + c
```

is equivalent to

```
a = a * (b + c)
```

and *not*

```
a = a * b + c
```

The following table lists the compound assignment operators and shows an expression using each operator:

Operator	Example	Equivalent expression
<code>+=</code>	<code>index += 2</code>	<code>index = index + 2</code>
<code>--</code>	<code>*(pointer++) -= 1</code>	<code>*pointer = *(pointer++) - 1</code>
<code>*=</code>	<code>bonus *= increase</code>	<code>bonus = bonus * increase</code>
<code>/=</code>	<code>time /= hours</code>	<code>time = time / hours</code>
<code>%=</code>	<code>allowance %= 1000</code>	<code>allowance = allowance % 1000</code>
<code><<=</code>	<code>result <<= num</code>	<code>result = result << num</code>
<code>>>=</code>	<code>form >>= 1</code>	<code>form = form >> 1</code>
<code>&=</code>	<code>mask &= 2</code>	<code>mask = mask & 2</code>
<code>^=</code>	<code>test ^= pre_test</code>	<code>test = test ^ pre_test</code>
<code> =</code>	<code>flag = ON</code>	<code>flag = flag ON</code>

Although the equivalent expression column shows the left operands (from the example column) twice, it is in effect evaluated only once.

C++ In addition to the table of operand types, an expression is implicitly converted to the cv-unqualified type of the left operand if it is not of class type. However, if the left operand is of class type, the class becomes complete, and assignment to objects of the class behaves as a copy assignment operation. Compound expressions and conditional expressions are lvalues in C++, which allows them to be a left operand in a compound assignment expression.

IBM extension

C When GNU C language features have been enabled, compound expressions and conditional expressions are allowed as lvalues, provided that their operands are lvalues. The following compound assignment of the compound expression (a, b) is legal under GNU C, provided that expression b, or more generally, the last expression in the sequence, is an lvalue:

```
(a,b) += 5 /* Under GNU C, this is equivalent to
a, (b += 5) */
```

End of IBM extension

Multiplication operator *

The * (multiplication) operator yields the product of its operands. The operands must have an arithmetic or enumeration type. The result is not an lvalue. The usual arithmetic conversions on the operands are performed.

Because the multiplication operator has both associative and commutative properties, the compiler can rearrange the operands in an expression that contains more than one multiplication operator. For example, the expression:

```
sites * number * cost
```

can be interpreted in any of the following ways:

```
(sites * number) * cost
sites * (number * cost)
(cost * sites) * number
```

Division operator /

The / (division) operator yields the algebraic quotient of its operands. If both operands are integers, any fractional part (remainder) is discarded. Throwing away the fractional part is often called *truncation toward zero*. The operands must have an arithmetic or enumeration type. The right operand may not be zero: the result is undefined if the right operand evaluates to 0. For example, expression `7 / 4` yields the value 1 (rather than 1.75 or 2). The result is not an lvalue.

The usual arithmetic conversions on the operands are performed.

C++ If both operands are negative, the sign of the remainder is also negative. Otherwise, the sign of the remainder is the same as the sign of the quotient.

Remainder operator %

The % (remainder) operator yields the remainder from the division of the left operand by the right operand. For example, the expression `5 % 3` yields 2. The result is not an lvalue.

Both operands must have an integral or enumeration type. If the right operand evaluates to 0, the result is undefined. If either operand has a negative value, the result is such that the following expression always yields the value of `a` if `b` is not 0 and `a/b` is representable:

```
( a / b ) * b + a %b;
```

The usual arithmetic conversions on the operands are performed.

 If both operands are negative, the sign of the remainder is also negative. Otherwise, the sign of the remainder is the same as the sign of the quotient.

Addition operator +

The + (addition) operator yields the sum of its operands. Both operands must have an arithmetic type, or one operand must be a pointer to an object type and the other operand must have an integral or enumeration type.

When both operands have an arithmetic type, the usual arithmetic conversions on the operands are performed. The result has the type produced by the conversions on the operands and is not an lvalue.

A pointer to an object in an array can be added to a value having integral type. The result is a pointer of the same type as the pointer operand. The result refers to another element in the array, offset from the original element by the amount of the integral value treated as a subscript. If the resulting pointer points to storage outside the array, other than the first location outside the array, the result is undefined. A pointer to one element past the end of an array cannot be used to access the memory content at that address. The compiler does not provide boundary checking on the pointers. For example, after the addition, `ptr` points to the third element of the array:

```
int array[5];
int *ptr;
ptr = array + 2;
```

Related information

- “Pointer arithmetic” on page 87
- “Pointer conversions” on page 117

Subtraction operator -

The - (subtraction) operator yields the difference of its operands. Both operands must have an arithmetic or enumeration type, or the left operand must have a pointer type and the right operand must have the same pointer type or an integral or enumeration type. You cannot subtract a pointer from an integral value.

When both operands have an arithmetic type, the usual arithmetic conversions on the operands are performed. The result has the type produced by the conversions on the operands and is not an lvalue.

When the left operand is a pointer and the right operand has an integral type, the compiler converts the value of the right to an address offset. The result is a pointer of the same type as the pointer operand.

If both operands are pointers to elements in the same array, the result is the number of objects separating the two addresses. The number is of type `ptrdiff_t`, which is defined in the header file `stddef.h`. Behavior is undefined if the pointers do not refer to objects in the same array.

Related information

- “Pointer arithmetic” on page 87
- “Pointer conversions” on page 117

Bitwise left and right shift operators << >>

The bitwise shift operators move the bit values of a binary object. The left operand specifies the value to be shifted. The right operand specifies the number of positions that the bits in the value are to be shifted. The result is not an lvalue. Both operands have the same precedence and are left-to-right associative.

Operator	Usage
<<	Indicates the bits are to be shifted to the left.
>>	Indicates the bits are to be shifted to the right.

Each operand must have an integral or enumeration type. The compiler performs integral promotions on the operands, and then the right operand is converted to type `int`. The result has the same type as the left operand (after the arithmetic conversions).

The right operand should not have a negative value or a value that is greater than or equal to the width in bits of the expression being shifted. The result of bitwise shifts on such values is unpredictable.

If the right operand has the value 0, the result is the value of the left operand (after the usual arithmetic conversions).

The << operator fills vacated bits with zeros. For example, if `left_op` has the value 4019, the bit pattern (in 16-bit format) of `left_op` is:

```
0000111110110011
```

The expression `left_op << 3` yields:

```
0111110110011000
```

The expression `left_op >> 3` yields:

```
000000011110110
```

Relational operators < > <= >=

The relational operators compare two operands and determine the validity of a relationship. The following table describes the four relational operators:

Operator	Usage
<	Indicates whether the value of the left operand is less than the value of the right operand.
>	Indicates whether the value of the left operand is greater than the value of the right operand.
<=	Indicates whether the value of the left operand is less than or equal to the value of the right operand.
>=	Indicates whether the value of the left operand is greater than or equal to the value of the right operand.

Both operands must have arithmetic or enumeration types or be pointers to the same type.

► **C** The type of the result is `int` and has the values 1 if the specified relationship is true, and 0 if false. ► **C++** The type of the result is `bool` and has the values `true` or `false`.

The result is not an lvalue.

If the operands have arithmetic types, the usual arithmetic conversions on the operands are performed.

When the operands are pointers, the result is determined by the locations of the objects to which the pointers refer. If the pointers do not refer to objects in the same array, the result is not defined.

A pointer can be compared to a constant expression that evaluates to 0. You can also compare a pointer to a pointer of type `void*`. The pointer is converted to a pointer of type `void*`.

If two pointers refer to the same object, they are considered equal. If two pointers refer to nonstatic members of the same object, the pointer to the object declared later is greater, provided that they are not separated by an access specifier; otherwise the comparison is undefined. If two pointers refer to data members of the same union, they have the same address value.

If two pointers refer to elements of the same array, or to the first element beyond the last element of an array, the pointer to the element with the higher subscript value is greater.

You can only compare members of the same object with relational operators.

Relational operators have left-to-right associativity. For example, the expression:

```
a < b <= c
```

is interpreted as:

```
(a < b) <= c
```

If the value of `a` is less than the value of `b`, the first relationship yields 1 in C, or `true` in C++. The compiler then compares the value `true` (or 1) with the value of `c` (integral promotions are carried out if needed).

Equality and inequality operators == !=

The equality operators, like the relational operators, compare two operands for the validity of a relationship. The equality operators, however, have a lower precedence than the relational operators. The following table describes the two equality operators:

Operator	Usage
==	Indicates whether the value of the left operand is equal to the value of the right operand.
!=	Indicates whether the value of the left operand is not equal to the value of the right operand.

Both operands must have arithmetic or enumeration types or be pointers to the same type, or one operand must have a pointer type and the other operand must be a pointer to void or a null pointer.

C The type of the result is `int` and has the values 1 if the specified relationship is true, and 0 if false. **C++** The type of the result is `bool` and has the values `true` or `false`.

If the operands have arithmetic types, the usual arithmetic conversions on the operands are performed.

If the operands are pointers, the result is determined by the locations of the objects to which the pointers refer.

If one operand is a pointer and the other operand is an integer having the value 0, the `==` expression is true only if the pointer operand evaluates to `NULL`. The `!=` operator evaluates to true if the pointer operand does not evaluate to `NULL`.

You can also use the equality operators to compare pointers to members that are of the same type but do not belong to the same object. The following expressions contain examples of equality and relational operators:

```
time < max_time == status < complete
letter != EOF
```

Note: The equality operator (`==`) should not be confused with the assignment (`=`) operator.

For example,

```
if (x == 3)
```

evaluates to true (or 1) if `x` is equal to three. Equality tests like this should be coded with spaces between the operator and the operands to prevent unintentional assignments.

```
while
```

```
if (x = 3)
```

is taken to be true because `(x = 3)` evaluates to a nonzero value (3). The expression also assigns the value 3 to `x`.

Related information

- “Simple assignment operator `=`” on page 140

Bitwise AND operator &

The & (bitwise AND) operator compares each bit of its first operand to the corresponding bit of the second operand. If both bits are 1's, the corresponding bit of the result is set to 1. Otherwise, it sets the corresponding result bit to 0.

Both operands must have an integral or enumeration type. The usual arithmetic conversions on each operand are performed. The result has the same type as the converted operands.

Because the bitwise AND operator has both associative and commutative properties, the compiler can rearrange the operands in an expression that contains more than one bitwise AND operator.

The following example shows the values of a, b, and the result of a & b represented as 16-bit binary numbers:

bit pattern of a	0000000001011100
bit pattern of b	0000000000101110
bit pattern of a & b	0000000000001100

Note: The bitwise AND (&) should not be confused with the logical AND. (&&) operator. For example,

```
1 & 4 evaluates to 0
while
1 && 4 evaluates to true
```

Bitwise exclusive OR operator ^

The bitwise exclusive OR operator (in EBCDIC, the ^ symbol is represented by the \neg symbol) compares each bit of its first operand to the corresponding bit of the second operand. If both bits are 1's or both bits are 0's, the corresponding bit of the result is set to 0. Otherwise, it sets the corresponding result bit to 1.

Both operands must have an integral or enumeration type. The usual arithmetic conversions on each operand are performed. The result has the same type as the converted operands and is not an lvalue.

Because the bitwise exclusive OR operator has both associative and commutative properties, the compiler can rearrange the operands in an expression that contains more than one bitwise exclusive OR operator. Note that the ^ character can be represented by the trigraph '??'.

The following example shows the values of a, b, and the result of a ^ b represented as 16-bit binary numbers:

bit pattern of a	0000000001011100
bit pattern of b	0000000000101110
bit pattern of a ^ b	0000000001110010

Related information

- "Trigraph sequences" on page 37

Bitwise inclusive OR operator |

The `|` (bitwise inclusive OR) operator compares the values (in binary format) of each operand and yields a value whose bit pattern shows which bits in either of the operands has the value 1. If both of the bits are 0, the result of that bit is 0; otherwise, the result is 1.

Both operands must have an integral or enumeration type. The usual arithmetic conversions on each operand are performed. The result has the same type as the converted operands and is not an lvalue.

Because the bitwise inclusive OR operator has both associative and commutative properties, the compiler can rearrange the operands in an expression that contains more than one bitwise inclusive OR operator. Note that the `|` character can be represented by the trigraph `??|`.

The following example shows the values of `a`, `b`, and the result of `a | b` represented as 16-bit binary numbers:

bit pattern of <code>a</code>	0000000001011100
bit pattern of <code>b</code>	000000000101110
bit pattern of <code>a b</code>	000000000111110

Note: The bitwise OR (`|`) should not be confused with the logical OR (`||`) operator. For example,

```
1 | 4 evaluates to 5
while
1 || 4 evaluates to true
```

Related information

- “Trigraph sequences” on page 37

Logical AND operator &&

The `&&` (logical AND) operator indicates whether both operands are true.

C If both operands have nonzero values, the result has the value 1. Otherwise, the result has the value 0. The type of the result is `int`. Both operands must have a arithmetic or pointer type. The usual arithmetic conversions on each operand are performed.

C++ If both operands have values of `true`, the result has the value `true`. Otherwise, the result has the value `false`. Both operands are implicitly converted to `bool` and the result type is `bool`.

Unlike the `&` (bitwise AND) operator, the `&&` operator guarantees left-to-right evaluation of the operands. If the left operand evaluates to 0 (or `false`), the right operand is not evaluated.

The following examples show how the expressions that contain the logical AND operator are evaluated:

Expression	Result
1 && 0	false or 0
1 && 4	true or 1
0 && 0	false or 0

The following example uses the logical AND operator to avoid division by zero:
`(y != 0) && (x / y)`

The expression `x / y` is not evaluated when `y != 0` evaluates to `0` (or false).

Note: The logical AND (`&&`) should not be confused with the bitwise AND (`&`) operator. For example:

```

1 && 4 evaluates to 1 (or true)
while
1 & 4 evaluates to 0

```

Logical OR operator ||

The `||` (logical OR) operator indicates whether either operand is true.

C If either of the operands has a nonzero value, the result has the value 1. Otherwise, the result has the value 0. The type of the result is `int`. Both operands must have a arithmetic or pointer type. The usual arithmetic conversions on each operand are performed.

C++ If either operand has a value of `true`, the result has the value `true`. Otherwise, the result has the value `false`. Both operands are implicitly converted to `bool` and the result type is `bool`.

Unlike the `|` (bitwise inclusive OR) operator, the `||` operator guarantees left-to-right evaluation of the operands. If the left operand has a nonzero (or true) value, the right operand is not evaluated.

The following examples show how expressions that contain the logical OR operator are evaluated:

Expression	Result
1 0	true or 1
1 4	true or 1
0 0	false or 0

The following example uses the logical OR operator to conditionally increment `y`:
`++x || ++y;`

The expression `++y` is not evaluated when the expression `++x` evaluates to a nonzero (or true) quantity.

Note: The logical OR (`||`) should not be confused with the bitwise OR (`|`) operator.
For example:

```
1 || 4 evaluates to 1 (or true)
while
1 | 4 evaluates to 5
```

Array subscripting operator []

A postfix expression followed by an expression in [] (brackets) specifies an element of an array. The expression within the brackets is referred to as a *subscript*. The first element of an array has the subscript zero.

By definition, the expression `a[b]` is equivalent to the expression `*((a) + (b))`, and, because addition is associative, it is also equivalent to `b[a]`. Between expressions `a` and `b`, one must be a pointer to a type `T`, and the other must have integral or enumeration type. The result of an array subscript is an lvalue. The following example demonstrates this:

```
#include <stdio.h>

int main(void) {
    int a[3] = { 10, 20, 30 };
    printf("a[0] = %d\n", a[0]);
    printf("a[1] = %d\n", 1[a]);
    printf("a[2] = %d\n", *(2 + a));
    return 0;
}
```

The following is the output of the above example:

```
a[0] = 10
a[1] = 20
a[2] = 30
```

 The above restrictions on the types of expressions required by the subscript operator, as well as the relationship between the subscript operator and pointer arithmetic, do not apply if you overload operator `[]` of a class.

The first element of each array has the subscript 0. The expression `contract[35]` refers to the 36th element in the array `contract`.

In a multidimensional array, you can reference each element (in the order of increasing storage locations) by incrementing the right-most subscript most frequently.

For example, the following statement gives the value 100 to each element in the array `code[4][3][6]`:

```
for (first = 0; first < 4; ++first)
{
    for (second = 0; second < 3; ++second)
    {
        for (third = 0; third < 6; ++third)
        {
            code[first][second][third] =
            100;
        }
    }
}
```

C only

C99 allows array subscripting on arrays that are not lvalues. However, using the address of a non-lvalue as an array subscript is still not allowed. The following example is valid in C99:

```
struct trio{int a[3];};
struct trio f();
foo (int index)
{
    return f().a[index];
}
```

End of C only

Related information

- “Pointers” on page 86
- “Integral types” on page 53
- “Lvalues and rvalues” on page 121
- “Arrays” on page 90
- “Overloading subscripting” on page 253
- “Pointer arithmetic” on page 87

Comma operator ,

A *comma expression* contains two operands of any type separated by a comma and has left-to-right associativity. The left operand is fully evaluated, possibly producing side effects, and its value, if there is one, is discarded. The right operand is then evaluated. The type and value of the result of a comma expression are those of its right operand, after the usual unary conversions.

C only

The result of a comma expression is not an lvalue.

End of C only

C++ only

In C++, the result is an lvalue if the right operand is an lvalue. The following statements are equivalent:

```
r = (a,b,...,c);
a; b; r = c;
```

The difference is that the comma operator may be suitable for expression contexts, such as loop control expressions.

Similarly, the address of a compound expression can be taken if the right operand is an lvalue.

```
&(a, b)
a, &b
```

End of C++ only

Any number of expressions separated by commas can form a single expression because the comma operator is associative. The use of the comma operator guarantees that the subexpressions will be evaluated in left-to-right order, and the value of the last becomes the value of the entire expression. In the following example, if `omega` has the value 11, the expression increments `delta` and assigns the value 3 to `alpha`:

```
alpha = (delta++, omega % 4);
```

A sequence point occurs after the evaluation of the first operand. The value of `delta` is discarded. Similarly, in the following example, the value of the expression: `intensity++, shade * increment, rotate(direction);`

is the value of the expression:

```
rotate(direction)
```

In some contexts where the comma character is used, parentheses are required to avoid ambiguity. For example, the function

```
f(a, (t = 3, t + 2), c);
```

has only three arguments: the value of `a`, the value 5, and the value of `c`. Other contexts in which parentheses are required are in field-length expressions in structure and union declarator lists, enumeration value expressions in enumeration declarator lists, and initialization expressions in declarations and initializers.

In the previous example, the comma is used to separate the argument expressions in a function invocation. In this context, its use does not guarantee the order of evaluation (left to right) of the function arguments.

The primary use of the comma operator is to produce side effects in the following situations:

- Calling a function
- Entering or repeating an iteration loop
- Testing a condition
- Other situations where a side effect is required but the result of the expression is not immediately needed

The following table gives some examples of the uses of the comma operator.

Statement	Effects
<code>for (i=0; i<2; ++i, f());</code>	A for statement in which <code>i</code> is incremented and <code>f()</code> is called at each iteration.
<code>if (f(), ++i, i>1) { /* ... */ }</code>	An if statement in which function <code>f()</code> is called, variable <code>i</code> is incremented, and variable <code>i</code> is tested against a value. The first two expressions within this comma expression are evaluated before the expression <code>i>1</code> . Regardless of the results of the first two expressions, the third is evaluated and its result determines whether the if statement is processed.
<code>func((++a, f(a)));</code>	A function call to <code>func()</code> in which <code>a</code> is incremented, the resulting value is passed to a function <code>f()</code> , and the return value of <code>f()</code> is passed to <code>func()</code> . The function <code>func()</code> is passed only a single argument, because the comma expression is enclosed in parentheses within the function argument list.

Pointer to member operators `.*` `->*` (C++ only)

There are two pointer to member operators: `.*` and `->*`.

The `.*` operator is used to dereference pointers to class members. The first operand must be of class type. If the type of the first operand is class type `T`, or is a class that has been derived from class type `T`, the second operand must be a pointer to a member of a class type `T`.

The `->*` operator is also used to dereference pointers to class members. The first operand must be a pointer to a class type. If the type of the first operand is a pointer to class type `T`, or is a pointer to a class derived from class type `T`, the second operand must be a pointer to a member of class type `T`.

The `.*` and `->*` operators bind the second operand to the first, resulting in an object or function of the type specified by the second operand.

If the result of `.*` or `->*` is a function, you can only use the result as the operand for the `()` (function call) operator. If the second operand is an lvalue, the result of `.*` or `->*` is an lvalue.

Related information

- “Class member lists” on page 269
- “Pointers to members” on page 274

Conditional expressions

A *conditional expression* is a compound expression that contains a condition that is implicitly converted to type `bool` in C++(`operand1`), an expression to be evaluated if the condition evaluates to true (`operand2`), and an expression to be evaluated if the condition has the value false (`operand3`).

The conditional expression contains one two-part operator. The `?` symbol follows the condition, and the `:` symbol appears between the two action expressions. All expressions that occur between the `?` and `:` are treated as one expression.

The first operand must have a scalar type. The type of the second and third operands must be one of the following:

- An arithmetic type
- A compatible pointer, structure, or union type
- void

The second and third operands can also be a pointer or a null pointer constant.

Two objects are compatible when they have the same type but not necessarily the same type qualifiers (`volatile` or `const`). Pointer objects are compatible if they have the same type or are pointers to void.

The first operand is evaluated, and its value determines whether the second or third operand is evaluated:

- If the value is true, the second operand is evaluated.
- If the value is false, the third operand is evaluated.

The result is the value of the second or third operand.

If the second and third expressions evaluate to arithmetic types, the usual arithmetic conversions are performed on the values. The types of the second and third operands determine the type of the result as shown in the following tables.

Conditional expressions have right-to-left associativity with respect to their first and third operands. The leftmost operand is evaluated first, and then only one of the remaining two operands is evaluated. The following expressions are equivalent:

```
a ? b : c ? d : e ? f : g
a ? b : (c ? d : (e ? f : g))
```

Types in conditional C expressions

C only

In C, a conditional expression is not an lvalue, nor is its result.

Table 25. Types of operands and results in conditional C expressions

Type of one operand	Type of other operand	Type of result
Arithmetic	Arithmetic	Arithmetic type after usual arithmetic conversions
Structure or union type	Compatible structure or union type	Structure or union type with all the qualifiers on both operands
void	void	void
Pointer to compatible type	Pointer to compatible type	Pointer to type with all the qualifiers specified for the type
Pointer to type	NULL pointer (the constant 0)	Pointer to type
Pointer to object or incomplete type	Pointer to void	Pointer to void with all the qualifiers specified for the type

IBM extension

In GNU C, a conditional expression is a valid lvalue, provided that its type is not void and both of its branches are valid lvalues. The following conditional expression (a ? b : c) is legal under GNU C:

```
(a ? b : c) = 5
/* Under GNU C, equivalent to (a ? b = 5 : (c = 5)) */
```

This extension is available when compiling in one of the extended language levels.

End of IBM extension

End of C only

Types in conditional C++ expressions

C++ only

In C++, a conditional expression is a valid lvalue if its type is not void, and its result is an lvalue.

Table 26. Types of operands and results in C++ conditional expressions

Type of one operand	Type of other operand	Type of result
Reference to type	Reference to type	Reference after usual reference conversions
Class T	Class T	Class T
Class T	Class X	Class type for which a conversion exists. If more than one possible conversion exists, the result is ambiguous.
throw expression	Other (type, pointer, reference)	Type of the expression that is not a throw expression

End of C++ only

Examples of conditional expressions

The following expression determines which variable has the greater value, *y* or *z*, and assigns the greater value to the variable *x*:

```
x = (y > z) ? y : z;
```

The following is an equivalent statement:

```
if (y > z)
    x = y;
else
    x = z;
```

The following expression calls the function `printf`, which receives the value of the variable *c*, if *c* evaluates to a digit. Otherwise, `printf` receives the character constant 'x'.

```
printf(" c = %c\n", isdigit(c) ? c : 'x');
```

If the last operand of a conditional expression contains an assignment operator, use parentheses to ensure the expression evaluates properly. For example, the = operator has higher precedence than the ?: operator in the following expression:

```
int i,j,k;
(i == 7) ? j ++ : k = j;
```

The compiler will interpret this expression as if it were parenthesized this way:

```
int i,j,k;
((i == 7) ? j ++ : k) = j;
```

That is, *k* is treated as the third operand, not the entire assignment expression *k* = *j*.

To assign the value of *j* to *k* when *i* == 7 is false, enclose the last operand in parentheses:

```
int i,j,k;
(i == 7) ? j ++ : (k = j);
```

Cast expressions

A cast operator is used for *explicit type conversions*. It converts the value of an expression to a specified type.

The following cast operators are supported:

- “Cast operator ()”
- “The `static_cast` operator (C++ only)” on page 159
- “The `reinterpret_cast` operator (C++ only)” on page 160
- “The `const_cast` operator (C++ only)” on page 161
- “The `dynamic_cast` operator (C++ only) (PPU only)” on page 162

Cast operator ()

Cast expression syntax

►►—(—*type*—)—*expression*—◄◄

C The result of this operation is not an lvalue. **C++** The result of this operation is an lvalue if *type* is a reference; in all other cases, the result is an rvalue.

The following demonstrates the use of the cast operator to dynamically create an integer array of size 10:

```
#include <stdlib.h>

int main(void) {
    int* myArray = (int*) malloc(10 * sizeof(int));
    free(myArray);
    return 0;
}
```

The `malloc` library function returns a `void` pointer that points to memory that will hold an object of the size of its argument. The statement `int* myArray = (int*) malloc(10 * sizeof(int))` does the following:

- Creates a `void` pointer that points to memory that can hold ten integers.
- Converts that `void` pointer into an integer pointer with the use of the cast operator.
- Assigns that integer pointer to `myArray`. Because a name of an array is the same as a pointer to the initial element of the array, `myArray` is an array of ten integers stored in the memory created by the call to `malloc()`.

C++ only

In C++ you can also use the following in cast expressions:

- Function-style casts
- C++ conversion operators, such as `static_cast`.

Function-style notation converts the value of *expression* to the type *type*:

expression(*type*)

The following example shows the same value cast with a C-style cast, the C++ function-style cast, and a C++ cast operator:

```
#include <iostream>
using namespace std;

int main() {
    float num = 98.76;
    int x1 = (int) num;
    int x2 = int(num);
    int x3 = static_cast<int>(num);

    cout << "x1 = " << x1 << endl;
    cout << "x2 = " << x2 << endl;
    cout << "x3 = " << x3 << endl;
}
```

The following is the output of the above example:

```
x1 = 98
x2 = 98
x3 = 98
```

The integer `x1` is assigned a value in which `num` has been explicitly converted to an `int` with the C-style cast. The integer `x2` is assigned a value that has been converted with the function-style cast. The integer `x3` is assigned a value that has been converted with the `static_cast` operator.

A cast is a valid lvalue if its operand is an lvalue. In the following simple assignment expression, the right-hand side is first converted to the specified type, then to the type of the inner left-hand side expression, and the result is stored. The value is converted back to the specified type, and becomes the value of the assignment. In the following example, `i` is of type `char *`.

```
(int)i = 8 // This is equivalent to the following expression
(int)(i = (char*) (int)(8))
```

For compound assignment operation applied to a cast, the arithmetic operator of the compound assignment is performed using the type resulting from the cast, and then proceeds as in the case of simple assignment. The following expressions are equivalent. Again, `i` is of type `char *`.

```
(int)i += 8 // This is equivalent to the following expression
(int)(i = (char*) (int)((int)i = 8))
```

For C++, the operand of a cast expression can have class type. If the operand has class type, it can be cast to any type for which the class has a user-defined conversion function. Casts can invoke a constructor, if the target type is a class, or they can invoke a conversion function, if the source type is a class. They can be ambiguous if both conditions hold.

End of C++ only

Related information

- “Type names” on page 85
- “Conversion functions” on page 332
- “Conversion constructors” on page 330
- “Lvalues and rvalues” on page 121

Cast to union type (C only)

IBM extension

Casting to a union type is the ability to cast a union member to the same type as the union to which it belongs. Such a cast does not produce an lvalue, unlike other casts. The feature is supported as an extension to C99, implemented to facilitate porting programs developed with GNU C.

Only a type that explicitly exists as a member of a union type can be cast to that union type. The cast can use either the tag of the union type or a union type name declared in a typedef expression. The type specified must be a complete union type. An anonymous union type can be used in a cast to a union type, provided that it has a tag or type name. A bit field can be cast to a union type, provided that the union contains a bit field member of the same type, but not necessarily of the same length. The following shows an example of a simple cast to union:

```
#include <stdio.h>

union foo {
    char t;
    short u;
    int v;
    long w;
    long long x;
    float y;
    double z;
};

int main() {
    union foo u;
    char a = 1;
    u = (union foo)a;
    printf("u = %i\n", u.t);
}
```

The output of this example is:

```
u = 1
```

Casting to a nested union is also allowed. In the following example, the double type `dd` can be cast to the nested union `u2_t`.

```
int main() {
    union u_t {
        char a;
        short b;
        int c;
        union u2_t {
            double d;
        }u2;
    };
    union u_t U;
    double dd = 1.234;
    U.u2 = (union u2_t) dd;    // Valid.
    printf("U.u2 is %f\n", U.u2);
}
```

The output of this example is:

```
U.u2 is 1.234
```

A union cast is also valid as a function argument, part of a constant expression for initialization of a static or non-static data object, and in a compound literal statement. The following example shows a cast to union used as part of an expression for initializing a static object:

```
struct S
{
    int a;
};

union U {
    struct S *s;
};

struct T {
    union U u;
} t[] = {
    {(union U)&s}
};
```

Related information

- “Structures and unions” on page 59
- “The transparent_union type attribute (C only)” on page 80

End of IBM extension

The `static_cast` operator (C++ only)

The *static_cast operator* converts a given expression to a specified type.

`static_cast` operator syntax

►—`static_cast`—◀—*Type*—>—(—*expression*—)—————▶

The following is an example of the `static_cast` operator.

```
#include <iostream>
using namespace std;

int main() {
    int j = 41;
    int v = 4;
    float m = j/v;
    float d = static_cast<float>(j)/v;
    cout << "m = " << m << endl;
    cout << "d = " << d << endl;
}
```

The following is the output of the above example:

```
m = 10
d = 10.25
```

In this example, `m = j/v;` produces an answer of type `int` because both `j` and `v` are integers. Conversely, `d = static_cast<float>(j)/v;` produces an answer of type `float`. The `static_cast` operator converts variable `j` to a type `float`. This allows the compiler to generate a division with an answer of type `float`. All `static_cast` operators resolve at compile time and do not remove any `const` or `volatile` modifiers.

Applying the `static_cast` operator to a null pointer will convert it to a null pointer value of the target type.

You can explicitly convert a pointer of a type A to a pointer of a type B if A is a base class of B. If A is not a base class of B, a compiler error will result.

You may cast an lvalue of a type A to a type B& if the following are true:

- A is a base class of B
- You are able to convert a pointer of type A to a pointer of type B
- The type B has the same or greater `const` or `volatile` qualifiers than type A
- A is not a virtual base class of B

The result is an lvalue of type B.

A pointer to member type can be explicitly converted into a different pointer to member type if both types are pointers to members of the same class. This form of explicit conversion may also take place if the pointer to member types are from separate classes, however one of the class types must be derived from the other.

Related information

- “User-defined conversions” on page 329

The `reinterpret_cast` operator (C++ only)

A *reinterpret_cast* operator handles conversions between unrelated types.

`reinterpret_cast` operator syntax

►► `reinterpret_cast` <—Type—> (—expression—) ◀◀

The `reinterpret_cast` operator produces a value of a new type that has the same bit pattern as its argument. You cannot cast away a `const` or `volatile` qualification. You can explicitly perform the following conversions:

- A pointer to any integral type large enough to hold it
- A value of integral or enumeration type to a pointer
- A pointer to a function to a pointer to a function of a different type
- A pointer to an object to a pointer to an object of a different type
- A pointer to a member to a pointer to a member of a different class or type, if the types of the members are both function types or object types

A null pointer value is converted to the null pointer value of the destination type.

Given an lvalue expression of type T and an object x, the following two conversions are synonymous:

- `reinterpret_cast<T&>(x)`
- `*reinterpret_cast<T*>(&x)`

C++ also supports C-style casts. The two styles of explicit casts have different syntax but the same semantics, and either way of reinterpreting one type of pointer as an incompatible type of pointer is usually invalid. The `reinterpret_cast` operator, as well as the other named cast operators, is more easily spotted than C-style casts, and highlights the paradox of a strongly typed language that allows explicit casts.

The C++ compiler detects and quietly fixes most but not all violations. It is important to remember that even though a program compiles, its source code may not be completely correct. On some platforms, performance optimizations are predicated on strict adherence to standard aliasing rules. Although the C++ compiler tries to help with type-based aliasing violations, it cannot detect all possible cases.

The following example violates the aliasing rule, but will execute as expected when compiled unoptimized in C++ or in K&R C. It will also successfully compile optimized in C++, but will not necessarily execute as expected. The offending line 7 causes an old or uninitialized value for `x` to be printed.

```

1  extern int y = 7.;
2
3  int main() {
4      float x;
5      int i;
6      x = y;
7      i = *(int *) &x;
8      printf("i=%d. x=%f.\n", i, x);
9  }
```

The next code example contains an incorrect cast that the compiler cannot even detect because the cast is across two different files.

```

1  /* separately compiled file 1 */
2      extern float f;
3      extern int * int_pointer_to_f = (int *) &f; /* suspicious cast */
4
5  /* separately compiled file 2 */
6      extern float f;
7      extern int * int_pointer_to_f;
8      f = 1.0;
9      int i = *int_pointer_to_f;          /* no suspicious cast but wrong */
```

In line 8, there is no way for the compiler to know that `f = 1.0` is storing into the same object that `int i = *int_pointer_to_f` is loading from.

Related information

- “User-defined conversions” on page 329

The `const_cast` operator (C++ only)

A *const_cast operator* is used to add or remove a `const` or `volatile` modifier to or from a type.

`const_cast` operator syntax

►—`const_cast`—◀—*Type*—>—(—*expression*—)—————▶

Type and the type of *expression* may only differ with respect to their `const` and `volatile` qualifiers. Their cast is resolved at compile time. A single `const_cast` expression may add or remove any number of `const` or `volatile` modifiers.

The result of a `const_cast` expression is an rvalue unless *Type* is a reference type. In this case, the result is an lvalue.

Types can not be defined within `const_cast`.

The following demonstrates the use of the `const_cast` operator:

```

#include <iostream>
using namespace std;

void f(int* p) {
    cout << *p << endl;
}

int main(void) {
    const int a = 10;
    const int* b = &a;

    // Function f() expects int*, not const int*
    // f(b);
    int* c = const_cast<int*>(b);
    f(c);

    // Lvalue is const
    // *b = 20;

    // Undefined behavior
    // *c = 30;

    int a1 = 40;
    const int* b1 = &a1;
    int* c1 = const_cast<int*>(b1);

    // Integer a1, the object referred to by c1, has
    // not been declared const
    *c1 = 50;

    return 0;
}

```

The compiler will not allow the function call `f(b)`. Function `f()` expects a pointer to an `int`, not a `const int`. The statement `int* c = const_cast<int*>(b)` returns a pointer `c` that refers to `a` without the `const` qualification of `a`. This process of using `const_cast` to remove the `const` qualification of an object is called *casting away constness*. Consequently the compiler will allow the function call `f(c)`.

The compiler would not allow the assignment `*b = 20` because `b` points to an object of type `const int`. The compiler will allow the `*c = 30`, but the behavior of this statement is undefined. If you cast away the constness of an object that has been explicitly declared as `const`, and attempt to modify it, the results are undefined.

However, if you cast away the constness of an object that has not been explicitly declared as `const`, you can modify it safely. In the above example, the object referred to by `b1` has not been declared `const`, but you cannot modify this object through `b1`. You may cast away the constness of `b1` and modify the value to which it refers.

Related information

- “Type qualifiers” on page 72

The `dynamic_cast` operator (C++ only) (PPU only)

The *dynamic_cast* operator performs type conversions at run time. The `dynamic_cast` operator guarantees the conversion of a pointer to a base class to a pointer to a derived class, or the conversion of an lvalue referring to a base class to

a reference to a derived class. A program can thereby use a class hierarchy safely. This operator and the `typeid` operator provide runtime type information (RTTI) support in C++.

The expression `dynamic_cast<T>(v)` converts the expression `v` to type `T`. Type `T` must be a pointer or reference to a complete class type or a pointer to void. If `T` is a pointer and the `dynamic_cast` operator fails, the operator returns a null pointer of type `T`. If `T` is a reference and the `dynamic_cast` operator fails, the operator throws the exception `std::bad_cast`. You can find this class in the standard library header `<typeinfo>`.

The `dynamic_cast` operator requires runtime type information (RTTI) to be generated, which must be explicitly specified at compile time through a compiler option.

If `T` is a void pointer, then `dynamic_cast` will return the starting address of the object pointed to by `v`. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct A {
    virtual ~A() { };
};

struct B : A { };

int main() {
    B bobj;
    A* ap = &bobj;
    void * vp = dynamic_cast<void *>(ap);
    cout << "Address of vp : " << vp << endl;
    cout << "Address of bobj: " << &bobj << endl;
}
```

The output of this example will be similar to the following. Both `vp` and `&bobj` will refer to the same address:

```
Address of vp : 12FF6C
Address of bobj: 12FF6C
```

The primary purpose for the `dynamic_cast` operator is to perform type-safe *downcasts*. A downcast is the conversion of a pointer or reference to a class `A` to pointer or reference to a class `B`, where class `A` is a base class of `B`. The problem with downcasts is that a pointer of type `A*` can and must point to any object of a class that has been derived from `A`. The `dynamic_cast` operator ensures that if you convert a pointer of class `A` to a pointer of a class `B`, the object that `A` points to belongs to class `B` or a class derived from `B`.

The following example demonstrates the use of the `dynamic_cast` operator:

```
#include <iostream>
using namespace std;

struct A {
    virtual void f() { cout << "Class A" << endl; }
};

struct B : A {
    virtual void f() { cout << "Class B" << endl; }
};

struct C : A {
```

```

    virtual void f() { cout << "Class C" << endl; }
};

void f(A* arg) {
    B* bp = dynamic_cast<B*>(arg);
    C* cp = dynamic_cast<C*>(arg);

    if (bp)
        bp->f();
    else if (cp)
        cp->f();
    else
        arg->f();
};

int main() {
    A aobj;
    C cobj;
    A* ap = &cobj;
    A* ap2 = &aobj;
    f(ap);
    f(ap2);
}

```

The following is the output of the above example:

```

Class C
Class A

```

The function `f()` determines whether the pointer `arg` points to an object of type `A`, `B`, or `C`. The function does this by trying to convert `arg` to a pointer of type `B`, then to a pointer of type `C`, with the `dynamic_cast` operator. If the `dynamic_cast` operator succeeds, it returns a pointer that points to the object denoted by `arg`. If `dynamic_cast` fails, it returns `0`.

You may perform downcasts with the `dynamic_cast` operator only on polymorphic classes. In the above example, all the classes are polymorphic because class `A` has a virtual function. The `dynamic_cast` operator uses the runtime type information generated from polymorphic classes.

Related information


- “Derivation” on page 293
- “User-defined conversions” on page 329

Compound literal expressions

A *compound literal* is a postfix expression that provides an unnamed object whose value is given by an initializer list. The C99 language feature allows you to pass parameters to functions without the need for temporary variables. It is useful for specifying constants of an aggregate type (arrays, structures, and unions) when only one instance of such types is needed.

The syntax for a compound literal resembles that of a cast expression. However, a compound literal is an lvalue, while the result of a cast expression is not. Furthermore, a cast can only convert to scalar types or `void`, whereas a compound literal results in an object of the specified type.

Compound literal syntax



The `type_name` can be any data type, including vectors, and user-defined types. It can be an array of unknown size, but not a variable length array. If the type is an array of unknown size, the size is determined by the initializer list.

The following example passes a constant structure variable of type `point` containing two integer members to the function `drawline`:

```
drawline((struct point){6,7});
```

If the compound literal occurs outside the body of a function, the initializer list must consist of constant expressions, and the unnamed object has static storage duration. If the compound literal occurs within the body of a function, the initializer list need not consist of constant expressions, and the unnamed object has automatic storage duration.

IBM For compatibility with GNU C, a static variable can be initialized with a compound literal of the same type, provided that all the initializers in the initializer list are constant expressions.


Related information

- “String literals” on page 28

new expressions (C++ only)

The `new` operator provides dynamic storage allocation.

new operator syntax



If you prefix `new` with the scope resolution operator (`::`), the global operator `new()` is used. If you specify an `argument_list`, the overloaded `new` operator that corresponds to that `argument_list` is used. The `type` is an existing built-in or user-defined type. A `new_type` is a type that has not already been defined and can include type specifiers and declarators.

An allocation expression containing the `new` operator is used to find storage in free store for the object being created. The `new expression` returns a pointer to the object created and can be used to initialize the object. If the object is an array, a pointer to the initial element is returned.

You cannot use the `new` operator to allocate function types, `void`, or incomplete class types because these are not object types. However, you can allocate pointers to functions with the `new` operator. You cannot create a reference with the `new` operator.

When the object being created is an array, only the first dimension can be a general expression. All subsequent dimensions must be constant integral expressions. The first dimension can be a general expression even when an existing *type* is used. You can create an array with zero bounds with the new operator. For example:

```
char * c = new char[0];
```

In this case, a pointer to a unique object is returned.

An object created with operator `new()` or operator `new[]()` exists until the operator `delete()` or operator `delete[]()` is called to deallocate the object's memory. A delete operator or a destructor will not be implicitly called for an object created with a new that has not been explicitly deallocated before the end of the program.

If parentheses are used within a new type, parentheses should also surround the new type to prevent syntax errors.

In the following example, storage is allocated for an array of pointers to functions:

```
void f();
void g();

int main(void)
{
    void (**p)(), (**q)();
    // declare p and q as pointers to pointers to void functions
    p = new (void (*[3])());
    // p now points to an array of pointers to functions
    q = new void(*[3])(); // error
    // error - bound as 'q = (new void) (*[3])();'
    p[0] = f; // p[0] to point to function f
    q[2] = g; // q[2] to point to function g
    p[0](); // call f()
    q[2](); // call g()
    return (0);
}
```

However, the second use of new causes an erroneous binding of `q = (new void) (*[3])()`.

The type of the object being created cannot contain class declarations, enumeration declarations, or `const` or `volatile` types. It can contain pointers to `const` or `volatile` objects.

For example, `const char*` is allowed, but `char* const` is not.

Related information

- “Allocation and deallocation functions (C++ only)” on page 228

Placement syntax

Additional arguments can be supplied to new by using the *argument_list*, also called the *placement syntax*. If placement arguments are used, a declaration of operator `new()` or operator `new[]()` with these arguments must exist. For example:

```
#include <new>
using namespace std;

class X
{
public:
```

```

        void* operator new(size_t,int, int){ /* ... */ }
};

// ...

int main ()
{
    X* ptr = new(1,2) X;
}

```

The placement syntax is commonly used to invoke the global placement `new` function. The global placement `new` function initializes an object or objects at the location specified by the placement argument in the placement `new` expression. This location must address storage that has previously been allocated by some other means, because the global placement `new` function does not itself allocate memory. In the following example, no new memory is allocated by the calls `new(whole) X(8);`, `new(seg2) X(9);`, or `new(seg3) X(10);`; Instead, the constructors `X(8)`, `X(9)`, and `X(10)` are called to reinitialize the memory allocated to the buffer `whole`.

Because placement `new` does not allocate memory, you should not use `delete` to deallocate objects created with the placement syntax. You can only delete the entire memory pool (`delete whole`). In the example, you can keep the memory buffer but destroy the object stored in it by explicitly calling a destructor.

```

#include <new>
class X
{
public:
    X(int n): id(n){ }
    ~X(){ }
private:
    int id;
    // ...
};

int main()
{
    char* whole = new char[ 3 * sizeof(X) ]; // a 3-part buffer
    X * p1 = new(whole) X(8);                // fill the front
    char* seg2 = &whole[ sizeof(X) ];        // mark second segment
    X * p2 = new(seg2) X(9);                // fill second segment
    char* seg3 = &whole[ 2 * sizeof(X) ];    // mark third segment
    X * p3 = new(seg3) X(10);               // fill third segment

    p2->~X(); // clear only middle segment, but keep the buffer
    // ...
    return 0;
}

```

The placement `new` syntax can also be used for passing parameters to an allocation routine rather than to a constructor.

Related information

- “delete expressions (C++ only)” on page 169
- “Scope resolution operator `::` (C++ only)” on page 126
- “Overview of constructors and destructors” on page 317

Initialization of objects created with the new operator

You can initialize objects created with the new operator in several ways. For nonclass objects, or for class objects without constructors, a *new initializer* expression can be provided in a new expression by specifying (*expression*) or (). For example:

```
double* pi = new double(3.1415926);
int* score = new int(89);
float* unknown = new float();
```

If a class does not have a default constructor, the new initializer must be provided when any object of that class is allocated. The arguments of the new initializer must match the arguments of a constructor.

You cannot specify an initializer for arrays. You can initialize an array of class objects only if the class has a default constructor. The constructor is called to initialize each array element (class object).

Initialization using the new initializer is performed only if new successfully allocates storage.

Related information

- “Overview of constructors and destructors” on page 317

Handling new allocation failure

When the new operator creates a new object, it calls the operator new() or operator new[]() function to obtain the needed storage.

When new cannot allocate storage to create a new object, it calls a *new handler* function if one has been installed by a call to set_new_handler(). The std::set_new_handler() function is declared in the header <new>. Use it to call a new handler you have defined or the default new handler.

Your new handler must perform one of the following:

- obtain more storage for memory allocation, then return
- throw an exception of type std::bad_alloc or a class derived from std::bad_alloc
- call either abort() or exit()

The set_new_handler() function has the prototype:

```
typedef void(*PNH)();
PNH set_new_handler(PNH);
```

set_new_handler() takes as an argument a pointer to a function (the new handler), which has no arguments and returns void. It returns a pointer to the previous new handler function.

If you do not specify your own set_new_handler() function, new throws an exception of type std::bad_alloc.

The following program fragment shows how you could use set_new_handler() to return a message if the new operator cannot allocate storage:

```
#include <iostream>
#include <new>
#include <cstdlib>
```

```

using namespace std;

void no_storage()
{
    std::cerr << "Operator new failed: no storage is
available.\n";
    std::exit(1);
}

int main(void)
{
    std::set_new_handler(&no_storage);
    // Rest of program ...
}

```

If the program fails because new cannot allocate storage, the program exits with the message:

```

Operator new failed:
no storage is available.

```

Related information

- `-qnewcheck` option in the *XL C/C++ Compiler Reference*

delete expressions (C++ only)

The delete operator destroys the object created with new by deallocating the memory associated with the object.

The delete operator has a void return type.

delete operator syntax

```

▶▶ [::] delete object_pointer ▶▶

```

The operand of delete must be a pointer returned by new, and cannot be a pointer to constant. Deleting a null pointer has no effect.

The delete[] operator frees storage allocated for array objects created with new[]. The delete operator frees storage allocated for individual objects created with new.

delete[] operator syntax

```

▶▶ [::] delete [ ] array ▶▶

```

The result of deleting an array object with delete is undefined, as is deleting an individual object with delete[]. The array dimensions do not need to be specified with delete[].

The result of any attempt to access a deleted object or array is undefined.

If a destructor has been defined for a class, delete invokes that destructor. Whether a destructor exists or not, delete frees the storage pointed to by calling the function operator delete() of the class if one exists.

The global ::operator delete() is used if:

- The class has no operator delete().
- The object is of a nonclass type.
- The object is deleted with the ::delete expression.

The global ::operator delete[] () is used if:

- The class has no operator delete[] ()
- The object is of a nonclass type
- The object is deleted with the ::delete[] expression.

The default global operator delete() only frees storage allocated by the default global operator new(). The default global operator delete[] () only frees storage allocated for arrays by the default global operator new[] ().

Related information

- “Overview of constructors and destructors” on page 317
- “The void type” on page 56

throw expressions (C++ only) (PPU only)

A *throw* expression is used to throw exceptions to C++ exception handlers. A throw expression is of type void.

Related information

- Chapter 16, “Exception handling (C++ only) (PPU only),” on page 371
- “The void type” on page 56

Label value expressions

IBM extension

The label value operator && returns the address of its operand, which must be a label defined in the current function or a containing function. The value is a constant of type void* and should be used only in a computed goto statement. The language feature is an extension to C and C++, implemented to facilitate porting programs developed with GNU C.

Related information

- “Labels as values” on page 176
- “Computed goto statement” on page 193

End of IBM extension

Operator precedence and associativity

Two operator characteristics determine how operands group with operators: *precedence* and *associativity*. Precedence is the priority for grouping different types of operators with their operands. Associativity is the left-to-right or right-to-left order for grouping operands to operators that have the same precedence. An operator’s precedence is meaningful only if other operators with higher or lower precedence are present. Expressions with higher-precedence operators are evaluated first. The grouping of operands can be forced by using parentheses.

For example, in the following statements, the value of 5 is assigned to both a and b because of the right-to-left associativity of the = operator. The value of c is assigned to b first, and then the value of b is assigned to a.

```
b = 9;
c = 5;
a = b = c;
```

Because the order of subexpression evaluation is not specified, you can explicitly force the grouping of operands with operators by using parentheses.

In the expression

```
a + b * c / d
```

the * and / operations are performed before + because of precedence. b is multiplied by c before it is divided by d because of associativity.

The following tables list the C and C++ language operators in order of precedence and show the direction of associativity for each operator. Operators that have the same rank have the same precedence.

Table 27. Precedence and associativity of postfix operators

Rank	Right associative?	Operator function	Usage
1	yes	▶ C++ global scope resolution	:: <i>name_or_qualified name</i>
1		▶ C++ class or namespace scope resolution	<i>class_or_namespace :: member</i>
2		member selection	<i>object . member</i>
2		member selection	<i>pointer -> member</i>
2		subscripting	<i>pointer [expr]</i>
2		function call	<i>expr (expr_list)</i>
2		value construction	<i>type (expr_list)</i>
2		postfix increment	<i>lvalue ++</i>
2		postfix decrement	<i>lvalue --</i>
2	yes	▶ C++ type identification	<i>typeid (type)</i>
2	yes	▶ C++ type identification at run time	<i>typeid (expr)</i>
2	yes	▶ C++ conversion checked at compile time	<i>static_cast < type > (expr)</i>
2	yes	▶ C++ conversion checked at run time	<i>dynamic_cast < type > (expr)</i>
2	yes	▶ C++ unchecked conversion	<i>reinterpret_cast < type > (expr)</i>
2	yes	▶ C++ const conversion	<i>const_cast < type > (expr)</i>

Table 28. Precedence and associativity of unary operators

Rank	Right associative?	Operator function	Usage
3	yes	size of object in bytes	<code>sizeof expr</code>
3	yes	size of type in bytes	<code>sizeof (type)</code>
3	yes	prefix increment	<code>++ lvalue</code>
3	yes	prefix decrement	<code>-- lvalue</code>
3	yes	bitwise negation	<code>~ expr</code>
3	yes	not	<code>! expr</code>
3	yes	unary minus	<code>- expr</code>
3	yes	unary plus	<code>+ expr</code>
3	yes	address of	<code>& lvalue</code>
3	yes	indirection or dereference	<code>* expr</code>
3	yes	 create (allocate memory)	<code>new type</code>
3	yes	 create (allocate and initialize memory)	<code>new type (expr_list) type</code>
3	yes	 create (placement)	<code>new type (expr_list) type (expr_list)</code>
3	yes	 destroy (deallocate memory)	<code>delete pointer</code>
3	yes	 destroy array	<code>delete [] pointer</code>
3	yes	type conversion (cast)	<code>(type) expr</code>

Table 29. Precedence and associativity of binary operators

Rank	Right associative?	Operator function	Usage
4		 member selection	<code>object .* ptr_to_member</code>
4		 member selection	<code>object ->* ptr_to_member</code>
5		multiplication	<code>expr * expr</code>
5		division	<code>expr / expr</code>
5		modulo (remainder)	<code>expr % expr</code>
6		binary addition	<code>expr + expr</code>
6		binary subtraction	<code>expr - expr</code>
7		bitwise shift left	<code>expr << expr</code>
7		bitwise shift right	<code>expr >> expr</code>
8		less than	<code>expr < expr</code>
8		less than or equal to	<code>expr <= expr</code>
8		greater than	<code>expr > expr</code>
8		greater than or equal to	<code>expr >= expr</code>
9		equal	<code>expr == expr</code>

Table 29. Precedence and associativity of binary operators (continued)

Rank	Right associative?	Operator function	Usage
9		not equal	<i>expr != expr</i>
10		bitwise AND	<i>expr & expr</i>
11		bitwise exclusive OR	<i>expr ^ expr</i>
12		bitwise inclusive OR	<i>expr expr</i>
13		logical AND	<i>expr && expr</i>
14		logical inclusive OR	<i>expr expr</i>
15		conditional expression	<i>expr ? expr : expr</i>
16	yes	simple assignment	<i>lvalue = expr</i>
16	yes	multiply and assign	<i>lvalue *= expr</i>
16	yes	divide and assign	<i>lvalue /= expr</i>
16	yes	modulo and assign	<i>lvalue %= expr</i>
16	yes	add and assign	<i>lvalue += expr</i>
16	yes	subtract and assign	<i>lvalue -= expr</i>
16	yes	shift left and assign	<i>lvalue <<= expr</i>
16	yes	shift right and assign	<i>lvalue >>= expr</i>
16	yes	bitwise AND and assign	<i>lvalue &= expr</i>
16	yes	bitwise exclusive OR and assign	<i>lvalue ^= expr</i>
16	yes	bitwise inclusive OR and assign	<i>lvalue = expr</i>
17	yes	 throw expression	throw <i>expr</i>
18		comma (sequencing)	<i>expr , expr</i>

Examples of expressions and precedence

The parentheses in the following expressions explicitly show how the compiler groups operands and operators.

```
total = (4 + (5 * 3));
total = (((8 * 5) / 10) / 3);
total = (10 + (5/3));
```

If parentheses did not appear in these expressions, the operands and operators would be grouped in the same manner as indicated by the parentheses. For example, the following expressions produce the same output.

```
total = (4+(5*3));
total = 4+5*3;
```

Because the order of grouping operands with operators that are both associative and commutative is not specified, the compiler can group the operands and operators in the expression:

```
total = price + prov_tax +
city_tax;
```

in the following ways (as indicated by parentheses):

```
total = (price + (prov_tax + city_tax));  
total = ((price + prov_tax) + city_tax);  
total = ((price + city_tax) + prov_tax);
```

The grouping of operands and operators does not affect the result unless one ordering causes an overflow and another does not. For example, if `price = 32767`, `prov_tax = -42`, and `city_tax = 32767`, and all three of these variables have been declared as integers, the third statement `total = ((price + city_tax) + prov_tax)` will cause an integer overflow and the rest will not.

Because intermediate values are rounded, different groupings of floating-point operators may give different results.

In certain expressions, the grouping of operands and operators can affect the result. For example, in the following expression, each function call might be modifying the same global variables.

```
a = b() + c() + d();
```

This expression can give different results depending on the order in which the functions are called.

If the expression contains operators that are both associative and commutative and the order of grouping operands with operators can affect the result of the expression, separate the expression into several expressions. For example, the following expressions could replace the previous expression if the called functions do not produce any side effects that affect the variable `a`.

```
a = b();  
a += c();  
a += d();
```

The order of evaluation for function call arguments or for the operands of binary operators is not specified. Therefore, the following expressions are ambiguous:

```
z = (x * ++y) / func1(y);  
func2(++i, x[i]);
```

If `y` has the value of 1 before the first statement, it is not known whether or not the value of 1 or 2 is passed to `func1()`. In the second statement, if `i` has the value of 1 before the expression is evaluated, it is not known whether `x[1]` or `x[2]` is passed as the second argument to `func2()`.

Chapter 7. Statements

A statement, the smallest independent computational unit, specifies an action to be performed. In most cases, statements are executed in sequence. The following is a summary of the statements available in C and C++:

- Labeled statements
- Expression statements
- Block statements
- Selection statements
- Iteration statements
- Jump statements
- Declaration statements
-  try blocks
- Null statement
-  Inline assembly statements

Related information

- Chapter 3, “Data objects and declarations,” on page 41
- “Function declarations” on page 202
- “try blocks” on page 371

Labeled statements

There are three kinds of labels: identifier, case, and default.

Labeled statement syntax

►—*identifier*—:—*statement*—◄◄

The label consists of the *identifier* and the colon (:) character.

 A label name must be unique within the function in which it appears.

 In C++, an identifier label may only be used as the target of a `goto` statement. A `goto` statement can use a label before its definition. Identifier labels have their own namespace; you do not have to worry about identifier labels conflicting with other identifiers. However, you may not redeclare a label within a function.

Case and default label statements only appear in `switch` statements. These labels are accessible only within the closest enclosing `switch` statement.

case statement syntax

►—*case*—*constant_expression*—:—*statement*—◄◄

default statement syntax

► default : *statement* ◀

The following are examples of labels:

```
comment_complete : ;          /* null statement label */
test_for_null : if (NULL == pointer)
```

Related information

- “The goto statement” on page 192
- “The switch statement” on page 181

Locally declared labels

IBM extension

A locally declared label, or *local label*, is an identifier label that is declared at the beginning of a statement expression and for which the scope is the statement expression in which it is declared and defined. This language feature is an extension of C and C++ to facilitate handling programs developed with GNU C.

A local label can be used as the target of a `goto` statement, jumping to it from within the same block in which it was declared. This language extension is particularly useful for writing macros that contain nested loops, capitalizing on the difference between its statement scope and the function scope of an ordinary label.

Locally declared label syntax

► *__label__* *identifier* ; ◀

In a statement expression, the declaration of a local label must appear immediately after the left parenthesis and left brace, and must precede any ordinary declarations and statements. The label is defined in the usual way, with a name and a colon, within the statements of the statement expression.

Related information

- “Statement expressions” on page 179

End of IBM extension

Labels as values

IBM extension

The address of a label defined in the current function or a containing function can be obtained and used as a value wherever a constant of type `void*` is valid. The address is the return value when the label is the operand of the unary operator `&&`. The ability to use the address of label as a value is an extension to C99 and C++, implemented to facilitate porting programs developed with GNU C.

In the following example, the computed goto statements use the values of `label1` and `label2` to jump to those spots in the function.

```
int main()
{
    void * ptr1, *ptr2;
    ...
    label1: ...
    ...
    label2: ...
    ...
    ptr1 = &label1;
    ptr2 = &label2;
    if (...) {
        goto *ptr1;
    } else {
        goto *ptr2;
    }
    ...
}
```

Related information

- “Label value expressions” on page 170
- “Computed goto statement” on page 193

End of IBM extension

Expression statements

An *expression statement* contains an expression. The expression can be null.

Expression statement syntax

→ expression ; →

An expression statement evaluates *expression*, then discards the value of the expression. An expression statement without an expression is a null statement.

The following are examples of statements:

```
printf("Account Number: \n");          /* call to the printf */
marks = dollars * exch_rate;           /* assignment to marks */
(difference < 0) ? ++losses : ++gain;   /* conditional increment */
```

Related information

- Chapter 6, “Expressions and operators,” on page 121

Resolution of ambiguous statements

C++ only

The C++ syntax does not disambiguate between expression statements and declaration statements. The ambiguity arises when an expression statement has a function-style cast as its left-most subexpression. (Note that, because C does not support function-style casts, this ambiguity does not occur in C programs.) If the statement can be interpreted both as a declaration and as an expression, the statement is interpreted as a declaration statement.

Note: The ambiguity is resolved only on a syntactic level. The disambiguation does not use the meaning of the names, except to assess whether or not they are type names.

The following expressions disambiguate into expression statements because the ambiguous subexpression is followed by an assignment or an operator. `type_spec` in the expressions can be any type specifier:

```
type_spec(i)++;           // expression statement
type_spec(i,3)<<d;       // expression statement
type_spec(i)->l=24;      // expression statement
```

In the following examples, the ambiguity cannot be resolved syntactically, and the statements are interpreted as declarations. `type_spec` is any type specifier:

```
type_spec(*i)(int);      // declaration
type_spec(j)[5];        // declaration
type_spec(m) = { 1, 2 }; // declaration
type_spec(*k) (float(3)); // declaration
```

The last statement above causes a compile-time error because you cannot initialize a pointer with a float value.

Any ambiguous statement that is not resolved by the above rules is by default a declaration statement. All of the following are declaration statements:

```
type_spec(a);           // declaration
type_spec(*b)();        // declaration
type_spec(c)=23;        // declaration
type_spec(d),e,f,g=0;   // declaration
type_spec(h)(e,3);     // declaration
```

Related information


- Chapter 3, “Data objects and declarations,” on page 41
- Chapter 6, “Expressions and operators,” on page 121
- “Function call expressions” on page 127

End of C++ only

Block statements

A *block statement*, or *compound statement*, lets you group any number of data definitions, declarations, and statements into one statement. All definitions, declarations, and statements enclosed within a single set of braces are treated as a single statement. You can use a block wherever a single statement is allowed.

Block statement syntax



A block defines a local scope. If a data object is usable within a block and its identifier is not redefined, all nested blocks can use that data object.

Example of blocks

The following program shows how the values of data objects change in nested blocks:

```
/**
** This example shows how data objects change in nested blocks.
**/
#include <stdio.h>

int main(void)
{
    int x = 1;                /* Initialize x to 1 */
    int y = 3;

    if (y > 0)
    {
        int x = 2;          /* Initialize x to 2 */
        printf("second x = %4d\n", x);
    }
    printf("first x = %4d\n", x);

    return(0);
}
```

The program produces the following output:

```
second x =    2
first x =    1
```

Two variables named *x* are defined in *main*. The first definition of *x* retains storage while *main* is running. However, because the second definition of *x* occurs within a nested block, `printf("second x = %4d\n", x);` recognizes *x* as the variable defined on the previous line. Because `printf("first x = %4d\n", x);` is not part of the nested block, *x* is recognized as the first definition of *x*.

Statement expressions

IBM extension

A compound statement is a sequence of statements enclosed by braces. In GNU C, a compound statement inside parentheses may appear as an expression in what is called a *statement expression*.

Statement expression syntax

```
▶▶ ( { statement ; } ) ▶▶
```

The value of a statement expression is the value of the last simple expression to appear in the entire construct. If the last statement is not an expression, then the construct is of type `void` and has no value.

The statement expression can be combined with the `typeof` operator to create complex function-like macros in which each operand is evaluated only once. For example:

```
#define SWAP(a,b) ( {__typeof__(a) temp; temp=a; a=b; b=temp;} )
```

End of IBM extension

Selection statements

Selection statements consist of the following types of statements:

- The if statement
- The switch statement

The if statement

An if statement is a selection statement that allows more than one possible flow of control.

► **C++** An *if statement* lets you conditionally process a statement when the specified test expression, implicitly converted to `bool`, evaluates to `true`. If the implicit conversion to `bool` fails the program is ill-formed.

► **C** In *C*, an if statement lets you conditionally process a statement when the specified test expression evaluates to a nonzero value. The test expression must be of arithmetic or pointer type.

You can optionally specify an `else` clause on the if statement. If the test expression evaluates to `false` (or in *C*, a zero value) and an `else` clause exists, the statement associated with the `else` clause runs. If the test expression evaluates to `true`, the statement following the expression runs and the `else` clause is ignored.

if statement syntax

► `if (—expression—) —statement—` else —statement— ◀

When if statements are nested and `else` clauses are present, a given `else` is associated with the closest preceding if statement within the same block.

A single statement following any selection statements (`if`, `switch`) is treated as a compound statement containing the original statement. As a result any variables declared on that statement will be out of scope after the if statement. For example:

```
if (x)
int i;
```

is equivalent to:

```
if (x)
{ int i; }
```

Variable `i` is visible only within the if statement. The same rule applies to the `else` part of the if statement.

Examples of if statements

The following example causes grade to receive the value `A` if the value of `score` is greater than or equal to `90`.

```
if (score >= 90)
    grade = 'A';
```

The following example displays `Number is positive` if the value of `number` is greater than or equal to `0`. If the value of `number` is less than `0`, it displays `Number is negative`.

```

if (number >= 0)
    printf("Number is positive\n");
else
    printf("Number is negative\n");

```

The following example shows a nested if statement:

```

if (paygrade == 7)
    if (level >= 0 && level <= 8)
        salary *= 1.05;
    else
        salary *= 1.04;
else
    salary *= 1.06;
cout << "salary is " << salary << endl;

```

The following example shows a nested if statement that does not have an else clause. Because an else clause always associates with the closest if statement, braces might be needed to force a particular else clause to associate with the correct if statement. In this example, omitting the braces would cause the else clause to associate with the nested if statement.

```

if (kegs > 0) {
    if (furlongs > kegs)
        fxph = furlongs/kegs;
}
else
    fxph = 0;

```

The following example shows an if statement nested within an else clause. This example tests multiple conditions. The tests are made in order of their appearance. If one test evaluates to a nonzero value, a statement runs and the entire if statement ends.

```

if (value > 0)
    ++increase;
else if (value == 0)
    ++break_even;
else
    ++decrease;

```

Related information

- “Boolean types” on page 54

The switch statement

A *switch statement* is a selection statement that lets you transfer control to different statements within the switch body depending on the value of the switch expression. The switch expression must evaluate to an integral or enumeration value. The body of the switch statement contains *case clauses* that consist of


- A case label
- An optional default label
- A case expression
- A list of statements.

If the value of the switch expression equals the value of one of the case expressions, the statements following that case expression are processed. If not, the default label statements, if any, are processed.

switch statement syntax

►► switch (*expression*) *switch_body* ◀◀

The *switch body* is enclosed in braces and can contain definitions, declarations, *case clauses*, and a *default clause*. Each case clause and default clause can contain statements.



Note: An initializer within a *type_definition*, *file_scope_data_declaration* or *block_scope_data_declaration* is ignored.

A *case clause* contains a *case label* followed by any number of statements. A case clause has the form:

Case clause syntax

►► *case_label* *statement* ◀◀

A *case label* contains the word *case* followed by an integral constant expression and a colon. The value of each integral constant expression must represent a different value; you cannot have duplicate case labels. Anywhere you can put one case label, you can put multiple case labels. A case label has the form:

case label syntax

►► case *integral_constant_expression* : ◀◀

A *default clause* contains a default label followed by one or more statements. You can put a case label on either side of the default label. A switch statement can have only one default label. A *default_clause* has the form:

Default clause statement

►► *case_label* default : *case_label* *statement* ◀◀

The `switch` statement passes control to the statement following one of the labels or to the statement following the `switch` body. The value of the expression that precedes the `switch` body determines which statement receives control. This expression is called the *switch expression*.

The value of the `switch` expression is compared with the value of the expression in each case label. If a matching value is found, control is passed to the statement following the case label that contains the matching value. If there is no matching value but there is a default label in the `switch` body, control passes to the default labelled statement. If no matching value is found, and there is no default label anywhere in the `switch` body, no part of the `switch` body is processed.

When control passes to a statement in the `switch` body, control only leaves the `switch` body when a `break` statement is encountered or the last statement in the `switch` body is processed.

If necessary, an integral promotion is performed on the controlling expression, and all expressions in the case statements are converted to the same type as the controlling expression. The `switch` expression can also be of class type if there is a single conversion to integral or enumeration type.

Compiling with option `-qinfo=gen` finds case labels that fall through when they should not.

Restrictions on switch statements

You can put data definitions at the beginning of the `switch` body, but the compiler does not initialize `auto` and `register` variables at the beginning of a `switch` body. You can have declarations in the body of the `switch` statement.

You cannot use a `switch` statement to jump over initializations.

When the scope of an identifier with a variably modified type includes a case or default label of a `switch` statement, the entire `switch` statement is considered to be within the scope of that identifier. That is, the declaration of the identifier must precede the `switch` statement.

 In C++, you cannot transfer control over a declaration containing an explicit or implicit initializer unless the declaration is located in an inner block that is completely bypassed by the transfer of control. All declarations within the body of a `switch` statement that contain initializers must be contained in an inner block.

Examples of switch statements

The following `switch` statement contains several case clauses and one default clause. Each clause contains a function call and a `break` statement. The `break` statements prevent control from passing down through each statement in the `switch` body.

If the `switch` expression evaluated to `'/'`, the `switch` statement would call the function `divide`. Control would then pass to the statement following the `switch` body.

```
char key;

printf("Enter an arithmetic operator\n");
scanf("%c",&key);

switch (key)
{
```

```

    case '+':
        add();
        break;

    case '-':
        subtract();
        break;

    case '*':
        multiply();
        break;

    case '/':
        divide();
        break;

    default:
        printf("invalid key\n");
        break;
}

```

If the switch expression matches a case expression, the statements following the case expression are processed until a break statement is encountered or the end of the switch body is reached. In the following example, break statements are not present. If the value of text[i] is equal to 'A', all three counters are incremented. If the value of text[i] is equal to 'a', lettera and total are increased. Only total is increased if text[i] is not equal to 'A' or 'a'.

```

char text[100];
int capa, lettera, total;

// ...

for (i=0; i<sizeof(text); i++) {

    switch (text[i])
    {
        case 'A':
            capa++;
        case 'a':
            lettera++;
        default:
            total++;
    }
}

```

The following switch statement performs the same statements for more than one case label:

```

/**
** This example contains a switch statement that performs
** the same statement for more than one case label.
**/

#include <stdio.h>

int main(void)
{
    int month;

    /* Read in a month value */
    printf("Enter month: ");
    scanf("%d", &month);

    /* Tell what season it falls into */
    switch (month)

```

```

{
  case 12:
  case 1:
  case 2:
    printf("month %d is a winter month\n", month);
    break;

  case 3:
  case 4:
  case 5:
    printf("month %d is a spring month\n", month);
    break;

  case 6:
  case 7:
  case 8:
    printf("month %d is a summer month\n", month);
    break;

  case 9:
  case 10:
  case 11:
    printf("month %d is a fall month\n", month);
    break;

  case 66:
  case 99:
  default:
    printf("month %d is not a valid month\n", month);
}

return(0);
}

```

If the expression `month` has the value 3, control passes to the statement:
`printf("month %d is a spring month\n", month);`

The `break` statement passes control to the statement following the `switch` body.

Related information

- `-qinfo=gen` in the *XL C/C++ Compiler Reference*
- “Case and Default Labels” on page 175
- “The break statement” on page 189

Iteration statements

Iteration statements consist of the following types of statements:

- The `while` statement
- The `do` statement
- The `for` statement

Related information

- “Boolean types” on page 54

The while statement

A *while statement* repeatedly runs the body of a loop until the controlling expression evaluates to `false` (or `0` in C).

while statement syntax

► while (*expression*) *statement* ◄

► **C** The *expression* must be of arithmetic or pointer type. ► **C++** The expression must be convertible to bool.

The expression is evaluated to determine whether or not to process the body of the loop. If the expression evaluates to false, the body of the loop never runs. If the expression does not evaluate to false, the loop body is processed. After the body has run, control passes back to the expression. Further processing depends on the value of the condition.

A break, return, or goto statement can cause a while statement to end, even when the condition does not evaluate to false.

► **C++** A throw expression also can cause a while statement to end prior to the condition being evaluated.

In the following example, `item[index]` triples and is printed out, as long as the value of the expression `++index` is less than `MAX_INDEX`. When `++index` evaluates to `MAX_INDEX`, the while statement ends.

```
/**
 ** This example illustrates the while statement.
 **/

#define MAX_INDEX (sizeof(item) / sizeof(item[0]))
#include <stdio.h>

int main(void)
{
    static int item[ ] = { 12, 55, 62, 85, 102 };
    int index = 0;

    while (index < MAX_INDEX)
    {
        item[index] *= 3;
        printf("item[%d] = %d\n", index, item[index]);
        ++index;
    }

    return(0);
}
```

The do statement

A *do statement* repeatedly runs a statement until the test expression evaluates to false (or 0 in C). Because of the order of processing, the statement is run at least once.

do statement syntax

► do *statement* while (*expression*) ; ◄

► **C** The *expression* must be of arithmetic or pointer type. ► **C++** The controlling *expression* must be convertible to type bool.

The body of the loop is run before the controlling while clause is evaluated. Further processing of the do statement depends on the value of the while clause. If the while clause does not evaluate to false, the statement runs again. When the while clause evaluates to false, the statement ends.

A break, return, or goto statement can cause the processing of a do statement to end, even when the while clause does not evaluate to false.

▶ **C++** A throw expression also can cause a while statement to end prior to the condition being evaluated.

The following example keeps incrementing *i* while *i* is less than 5:

```
#include <stdio.h>

int main(void) {
    int i = 0;
    do {
        i++;
        printf("Value of i: %d\n", i);
    }
    while (i < 5);
    return 0;
}
```

The following is the output of the above example:

```
Value of i: 1
Value of i: 2
Value of i: 3
Value of i: 4
Value of i: 5
```

The for statement

A *for statement* lets you do the following:

- Evaluate an expression before the first iteration of the statement (*initialization*)
- Specify an expression to determine whether or not the statement should be processed (the *condition*)
- Evaluate an expression after each iteration of the statement (often used to increment for each iteration)
- Repeatedly process the statement if the controlling part does not evaluate to false (or 0 in C).

for statement syntax

```
▶▶ for ( [expression1] ; [expression2] ; [expression3] ) ▶▶
▶▶ statement ▶▶
```

expression1 is the *initialization expression*. It is evaluated only before the *statement* is processed for the first time. You can use this expression to initialize a variable. You can also use this expression to declare a variable, provided that the variable is not declared as static (it must be automatic and may also be declared as register). If you declare a variable in this expression, or anywhere else in *statement*, that variable goes out of scope at the end of the for loop. If you do not want to evaluate an expression prior to the first iteration of the statement, you can omit this expression.

expression2 is the *conditional expression*. It is evaluated before each iteration of the *statement*.  *expression2* must be of arithmetic or pointer type.  *expression3* must be convertible to type `bool`.

If it evaluates to `false` (or `0` in C), the *statement* is not processed and control moves to the next *statement* following the *for* *statement*. If *expression2* does not evaluate to `false`, the *statement* is processed. If you omit *expression2*, it is as if the *expression* had been replaced by `true`, and the *for* *statement* is not terminated by failure of this condition.

expression3 is evaluated after each iteration of the *statement*. This *expression* is often used for incrementing, decrementing, or assigning to a variable. This *expression* is optional.

A `break`, `return`, or `goto` *statement* can cause a *for* *statement* to end, even when the second *expression* does not evaluate to `false`. If you omit *expression2*, you must use a `break`, `return`, or `goto` *statement* to end the *for* *statement*.

 You can set a compiler option that allows a variable declared in the scope of a *for* *statement* to have a scope that is not local to the *for* *statement*.

Related information

- `-qlanglvl=noansifor` in the *XL C/C++ Compiler Reference*

Examples of for statements

The following *for* *statement* prints the value of `count` 20 times. The *for* *statement* initially sets the value of `count` to 1. After each iteration of the *statement*, `count` is incremented.

```
int count;
for (count = 1; count <= 20; count++)
    printf("count = %d\n", count);
```

The following sequence of *statements* accomplishes the same task. Note the use of the *while* *statement* instead of the *for* *statement*.

```
int count = 1;
while (count <= 20)
{
    printf("count = %d\n", count);
    count++;
}
```

The following *for* *statement* does not contain an initialization *expression*:

```
for (; index > 10; --index)
{
    list[index] = var1 + var2;
    printf("list[%d] = %d\n", index,
        list[index]);
}
```

The following *for* *statement* will continue running until `scanf` receives the letter `e`:

```
for (;;)
{
    scanf("%c", &letter);
    if (letter == '\n')
        continue;
```

```

    if (letter == 'e')
        break;
    printf("You entered the letter %c\n", letter);
}

```

The following for statement contains multiple initializations and increments. The comma operator makes this construction possible. The first comma in the for expression is a punctuator for a declaration. It declares and initializes two integers, `i` and `j`. The second comma, a comma operator, allows both `i` and `j` to be incremented at each step through the loop.

```

for (int i = 0,
     j = 50; i < 10; ++i, j += 50)
{
    cout << "i = " << i << "and j = " << j
        << endl;
}

```

The following example shows a nested for statement. It prints the values of an array having the dimensions [5][3].

```

for (row = 0; row < 5; row++)
    for (column = 0; column < 3; column++)
        printf("%d\n",
            table[row][column]);

```

The outer statement is processed as long as the value of `row` is less than 5. Each time the outer for statement is executed, the inner for statement sets the initial value of `column` to zero and the statement of the inner for statement is executed 3 times. The inner statement is executed as long as the value of `column` is less than 3.

Jump statements

Jump statements consist of the following types of statements:

- The break statement
- The continue statement
- The return statement
- The goto statement
-  Computed goto statement

The break statement

A *break statement* lets you end an *iterative* (do, for, or while) statement or a switch statement and exit from it at any point other than the logical end. A break may only appear on one of these statements.

break statement syntax

▶▶—break—;—————▶▶

In an iterative statement, the break statement ends the loop and moves control to the next statement outside the loop. Within nested statements, the break statement ends only the smallest enclosing do, for, switch, or while statement.

In a switch statement, the break passes control out of the switch body to the next statement outside the switch statement.

The continue statement

A *continue statement* ends the current iteration of a loop. Program control is passed from the continue statement to the end of the loop body.

A continue statement has the form:

```
▶▶—continue—;—————▶▶
```

A continue statement can only appear within the body of an iterative statement, such as do, for, or while.

The continue statement ends the processing of the action part of an iterative statement and moves control to the loop continuation portion of the statement. For example, if the iterative statement is a for statement, control moves to the third expression in the condition part of the statement, then to the second expression (the test) in the condition part of the statement.

Within nested statements, the continue statement ends only the current iteration of the do, for, or while statement immediately enclosing it.

Examples of continue statements

The following example shows a continue statement in a for statement. The continue statement causes processing to skip over those elements of the array rates that have values less than or equal to 1.

```
/**
 ** This example shows a continue statement in a for statement.
 **/

#include <stdio.h>
#define SIZE 5

int main(void)
{
    int i;
    static float rates[SIZE] = { 1.45, 0.05, 1.88, 2.00, 0.75 };

    printf("Rates over 1.00\n");
    for (i = 0; i < SIZE; i++)
    {
        if (rates[i] <= 1.00) /* skip rates <= 1.00 */
            continue;
        printf("rate = %.2f\n", rates[i]);
    }

    return(0);
}
```

The program produces the following output:

```
Rates over 1.00
rate = 1.45
rate = 1.88
rate = 2.00
```

The following example shows a continue statement in a nested loop. When the inner loop encounters a number in the array strings, that iteration of the loop ends. Processing continues with the third expression of the inner loop. The inner loop ends when the '\0' escape sequence is encountered.

```

/**
 ** This program counts the characters in strings that are part
 ** of an array of pointers to characters. The count excludes
 ** the digits 0 through 9.
 **/

#include <stdio.h>
#define SIZE 3

int main(void)
{
    static char *strings[SIZE] = { "ab", "c5d", "e5" };
    int i;
    int letter_count = 0;
    char *pointer;
    for (i = 0; i < SIZE; i++)          /* for each string */
        for (pointer = strings[i]; *pointer != '\0'; /* for each character */
            ++pointer)
        {
            /* if a number */
            if (*pointer >= '0' && *pointer <= '9')
                continue;
            letter_count++;
        }
    printf("letter count = %d\n", letter_count);

    return(0);
}

```


The program produces the following output:

```
letter count = 5
```

The return statement

A *return statement* ends the processing of the current function and returns control to the caller of the function.

return statement syntax



A value-returning function should include a return statement, containing an *expression*.  If an expression is not given on a return statement in a function declared with a non-void return type, the compiler issues a warning message.  If an expression is not given on a return statement in a function declared with a non-void return type, the compiler issues an error message.

If the data type of the expression is different from the function return type, conversion of the return value takes place as if the value of the expression were assigned to an object with the same function return type.

For a function of return type `void`, a return statement is not strictly necessary. If the end of such a function is reached without encountering a return statement, control is passed to the caller as if a return statement without an expression were encountered. In other words, an implicit return takes place upon completion of the final statement, and control automatically returns to the calling function.  If a return statement is used, it must not contain an expression.

Examples of return statements

The following are examples of return statements:

```
return;           /* Returns no value */
return result;   /* Returns the value of result */
return 1;        /* Returns the value 1 */
return (x * x);  /* Returns the value of x * x */
```

The following function searches through an array of integers to determine if a match exists for the variable number. If a match exists, the function match returns the value of i. If a match does not exist, the function match returns the value -1 (negative one).

```
int match(int number, int array[ ], int n)
{
    int i;

    for (i = 0; i < n; i++)
        if (number == array[i])
            return (i);
    return(-1);
}
```

A function can contain multiple return statements. For example:

```
void copy( int *a, int *b, int c)
{
    /* Copy array a into b, assuming both arrays are the same size */

    if (!a || !b)      /* if either pointer is 0, return */
        return;

    if (a == b)        /* if both parameters refer */
        return;        /* to same array, return */

    if (c == 0)        /* nothing to copy */
        return;

    for (int i = 0; i < c; ++i) /* do the copying */
        b[i] = a[i];        /* implicit return */
}
```

In this example, the return statement is used to cause a premature termination of the function, similar to a break statement.

An expression appearing in a return statement is converted to the return type of the function in which the statement appears. If no implicit conversion is possible, the return statement is invalid.

Related information

- “Function return type specifiers” on page 213
- “Function return values” on page 213

The goto statement

A *goto statement* causes your program to unconditionally transfer control to the statement associated with the label specified on the goto statement.

goto statement syntax

```
▶▶ goto label_identifier; ◀◀
```

Because the `goto` statement can interfere with the normal sequence of processing, it makes a program more difficult to read and maintain. Often, a `break` statement, a `continue` statement, or a function call can eliminate the need for a `goto` statement.

If an active block is exited using a `goto` statement, any local variables are destroyed when control is transferred from that block.

You cannot use a `goto` statement to jump over initializations.

A `goto` statement is allowed to jump within the scope of a variable length array, but not past any declarations of objects with variably modified types.

The following example shows a `goto` statement that is used to jump out of a nested loop. This function could be written without using a `goto` statement.

```
/**
 ** This example shows a goto statement that is used to
 ** jump out of a nested loop.
 **/

#include <stdio.h>
void display(int matrix[3][3]);

int main(void)
{
    int matrix[3][3]= {1,2,3,4,5,2,8,9,10};
    display(matrix);
    return(0);
}

void display(int matrix[3][3])
{
    int i, j;

    for (i = 0; i < 3; i++)
        for (j = 0; j < 3; j++)
        {
            if ( (matrix[i][j] < 1) || (matrix[i][j] > 6) )
                goto out_of_bounds;
            printf("matrix[%d][%d] = %d\n", i, j, matrix[i][j]);
        }
    return;
    out_of_bounds: printf("number must be 1 through 6\n");
}
```

Related information

- “Labeled statements” on page 175

Computed goto statement

IBM extension

A computed `goto` is a `goto` statement for which the target is a label from the same function. The address of the label is a constant of type `void*`, and is obtained by applying the unary label value operator `&&` to the label. The target of a computed `goto` is known at run time, and all computed `goto` statements from the same function will have the same targets. The language feature is an extension to C99 and C++, implemented to facilitate porting programs developed with GNU C.

Computed goto statement syntax

►► `goto *expression;` ◄◄

The **expression* is an expression of type `void*`.

Related information

- “Labels as values” on page 176
- “Label value expressions” on page 170

End of IBM extension

Null statement

The *null statement* performs no operation. It has the form:

►► `;` ◄◄

A null statement can hold the label of a labeled statement or complete the syntax of an iterative statement.

The following example initializes the elements of the array `price`. Because the initializations occur within the `for` expressions, a statement is only needed to finish the `for` syntax; no operations are required.

```
for (i = 0; i < 3; price[i++] = 0)
    ;
```

A null statement can be used when a label is needed before the end of a block statement. For example:

```
void func(void) {
    if (error_detected)
        goto depart;
    /* further processing */
depart: ; /* null statement required */
}
```

Inline assembly statements


IBM extension

Under extended language levels, the compiler provides full support for embedded assembly code fragments among C and C++ source statements. This extension has been implemented for use in general system programming code, and in the operating system kernel and device drivers, which were originally developed with GNU C.


The keyword `asm` stands for assembly code. When strict language levels are used in compilation, the C compiler recognizes and ignores the keyword `asm` in a declaration. The C++ compiler always recognizes the keyword.

The syntax is as follows:


asm statement syntax — statement in local scope




input:



output:



asm statement syntax — statement in global scope



The qualifier *volatile* instructs the compiler that the assembler instructions may update memory not listed in *output*, *input*, or *clobbers*.

The *code_format_string* is the source text of the asm instructions and is a string literal similar to a `printf` format specifier. Depending on the operands taken by the assembly instruction, the string contains zero or more modifiers, each of which corresponds to an input or output operand, separated by commas. The `%` modifier can take either of the following forms:

- `%integer`, where *integer* refers to the sequential number of the input or output operand.
- `%[symbolic_name]`, where the *symbolic_name* is referenced in the operand list. The symbolic operand names have no relation to any C identifiers. Any name can be used, even those of existing C symbols. However, no two operands in the same assembly statement can use the same symbolic name.

The *input* consists of zero, one or more input operands, separated by commas. Each operand consists of a *constraint* (*C_expression*) pair.

The *output* consists of zero, one or more output operands, separated by commas. Each operand consists of a *constraint* (*C_expression*) pair. The output operand must be constrained by the `=` or `+` modifier (described below), and, optionally, by an additional `%` or `&` modifier.

The *modifier* is one of the following:

- = Indicates that the operand is write-only for this instruction. The previous value is discarded and replaced by output data.
- + Indicates that the operand is both read and written by the instruction.
- & Indicates that the operand may be modified before the instruction is finished using the input operands; a register that is used as input should not be reused here.
- % Declares the instruction to be commutative for this operand and the following operand. This means that the order of this operand and the next may be swapped when generating the instruction. This modifier can be used on an input or output operand, but cannot be specified on the last operand.

The *constraint* is a string literal that describes the kind of operand that is permitted, one character per constraint. The following constraints are supported:

b (PPU only)

Use a general register other than zero.

c (PPU only)

Use the CTR register.

f (PPU only)

Use a floating-point register.

g | X (SPU only)

Use a general register, memory, or immediate operand.

h (PPU only)

Use the CTR or LINK register.

i Use an immediate integer or string literal operand.

l (PPU only)

Use the CTR register.

m Use a memory operand supported by the machine.

n Use an immediate integer.

o Use a memory operand that is offsetable.

r Use a general register.

s Use a string literal operand.

v For the PPU, use a vector register. For the SPU, use a 128-bit register.

0, 1, 2, ...

A matching constraint. Allocate the same register in output as in the corresponding input.

I, J, K, L, M, N, O, P

Constant values. Fold the expression in the operand and substitute the value into the % specifier. For the PPU and SPU, these constraints specify a maximum value for the operand, as follows:

For the SPU:

- I — 7 bits, signed
- J — 7 bits, unsigned
- K — 10 bits, signed
- L — 14 bits, signed

- M — 16 bits, signed
- N — 16 bits, unsigned
- O — 18 bits, signed
- P — 18 bits, unsigned

For the PPU

- I — signed 16-bit
- J — unsigned 16-bit shifted left 16 bits
- K — unsigned 16-bit constant
- L — signed 16-bit shifted left 16 bits
- M — unsigned constant greater than 31
- N — unsigned constant that is an exact power of 2
- O — zero
- P — signed whose negation is a signed 16-bit constant

The *C_expression* is a C or C++ expression whose value is used as the operand for the `asm` instruction. Output operands must be modifiable lvalues. The *C_expression* must be consistent with the constraint specified on it. For example, if `i` is specified, the operand must be an integer constant number.

Note: If pointer expressions are used in *input* or *output*, the assembly instructions should honor the ANSI aliasing rule (see “Type-based aliasing” on page 88 for more information). This means that indirect addressing using values in pointer expression operands should be consistent with the pointer types; otherwise, you must disable the `-qalias=ansi` option during compilation.

clobbers is a comma-separated list of register names enclosed in double quotes. These are registers that can be updated by the `asm` instruction. The following are valid register names for the PPU:

r0 to r31

General purpose registers

f0 to f31

Floating-point registers

lr Link register

ctr Loop count, decrement and branching register

fpscr Floating-point status and control register

xer Fixed-point exception register

cr0 to cr7

Condition registers

v0 to v31

Vector registers

The following are valid register names for the SPU:

[v]0 to [v]127

General purpose registers

fpscr Floating-point status and control register

Related information

- `-qalias=ansi` in the *XL C/C++ Compiler Reference*

Examples of inline assembly statements

In the following example (for the PPU):

```
int a ;
int b = 100 ;
int c = 200 ;
asm("add %0, %1, %2" : "=r"(a) : "r"(b) , "r"(c));
```

`add` is the op code of the instruction, understood by the assembler. `%0`, `%1` and `%2` are the operands, which are to be substituted by the C expressions in the output/input operand fields. The output operand uses the `=` constraint to indicate that a modifiable operand is required; and the `r` constraint to indicate that a general purpose register is required. Likewise, the `r` in the input operands indicates that general purpose registers are required. Within these restrictions, the compiler is free to choose any registers to substitute for `%0`, `%1`, and `%2`.

In the following example, the `%` constraint modifier tells the compiler that operands `a` and `b` can be switched if the compiler thinks it can generate better code in doing so.

```
asm("add %0, %1, %2" : "=r"(c) : "%r"(a), "r"(b));
```

The following example illustrates the use of the symbolic names for input and output operands (for the PPU):

```
int a ;
int b = 1, c = 2, d = 3 ;
__asm("addc %[result],[first],[second]" : [result]"=r"(a) : [first]"r"(b), [second]"r"(d));
```

The following shows an example of an inline assembly statement that uses the SPU assembly instruction `a`:

```
int res, x=22, y=33;
__asm("a %0, %1, %2\n" : "=r"(res) 115 : "r"(x), "r"(y+5));
```

`res` is the output operand, and `x` and `y+5` are the input operands. `r5` is the clobber register.

Restrictions on inline assembly statements

The following are restrictions on the use of inline assembly statements:

- The assembler instructions must be self-contained within an `asm` statement. The `asm` statement can only be used to generate instructions. All connections to the rest of the program must be established through the output and input operand list.
- Branching to a label in another `asm` statement is not supported for the PPU or SPU.
- Referencing an external symbol directly, without going through the operand list, is not supported for the SPU.
- Assembler instructions requiring a pair of registers are not specifiable by any constraints, and are therefore not supported. For example, you can not use the `%f` constraint for a long double operand.

Related information

- “Assembly labels” on page 17
- “Variables in specified registers” on page 50
- `-qasm` in the *XL C/C++ Compiler Reference*

End of IBM extension

Chapter 8. Functions

In the context of programming languages, the term *function* means an assemblage of statements used for computing an output value. The word is used less strictly than in mathematics, where it means a set relating input variables *uniquely* to output variables. Functions in C or C++ programs may not produce consistent outputs for all inputs, may not produce output at all, or may have side effects. Functions can be understood as user-defined operations, in which the parameters of the parameter list, if any, are the operands.

This section discusses the following topics:

- “Function declarations and definitions”
- “Function storage class specifiers” on page 206
- “Function specifiers” on page 208
- “Function return type specifiers” on page 213
- “Function declarators” on page 214
- “Function attributes” on page 218
- “The main() function” on page 224
- “Function calls” on page 225
- “Default arguments in C++ functions” on page 229
- “Pointers to functions” on page 232
- “Nested functions” on page 232

Function declarations and definitions

The distinction between a function *declaration* and function *definition* is similar to that of a data declaration and definition. The declaration establishes the names and characteristics of a function but does not allocate storage for it, while the *definition* specifies the body for a function, associates an identifier with the function, and allocates storage for it. Thus, the identifiers declared in this example:

```
float square(float x);
```

do not allocate storage.

The *function definition* contains a function declaration and the body of a function. The body is a block of statements that perform the work of the function. The identifiers declared in this example allocate storage; they are both declarations and definitions.

```
float square(float x)
{ return x*x; }
```

A function can be declared several times in a program, but all declarations for a given function must be compatible; that is, the return type is the same and the parameters have the same type. However, a function can only have one definition. Declarations are typically placed in header files, while definitions appear in source files.

Function declarations

A function identifier preceded by its return type and followed by its parameter list is called a *function declaration* or *function prototype*. The prototype informs the compiler of the format and existence of a function prior to its use. The compiler checks for mismatches between the parameters of a function call and those in the function declaration. The compiler also uses the declaration for argument type checking and argument conversions.

▶ **C++** Implicit declaration of functions is not allowed: you must explicitly declare every function before you can call it.


▶ **C** If a function declaration is not visible at the point at which a call to the function is made, the compiler assumes an implicit declaration of `extern int func();` However, for conformance to C99, you should explicitly prototype every function before making a call to it.

The elements of a declaration for a function are as follows:

- Function storage class specifiers, which specify linkage
- Function return type specifiers, which specify the data type of a value to be returned
- Function specifiers, which specify additional properties for functions
- Function declarators, which include function identifiers as well as lists of parameters

All function declarations have the form:

Function declaration syntax



▶ **IBM** In addition, for compatibility with GNU C and C++, XL C/C++ allows you to use *attributes* to modify the properties of functions. They are described in “Function attributes” on page 218.

Function definitions


The elements of a function definition are as follows:

- Function storage class specifiers, which specify linkage
- Function return type specifiers, which specify the data type of a value to be returned
- Function specifiers, which specify additional properties for functions
- Function declarators, which include function identifiers as well as lists of parameters
- The *function body*, which is a braces-enclosed series of statements representing the actions that the function performs
- ▶ **C++** Constructor-initializers, which are used only in constructor functions declared in classes; they are described in “Constructors” on page 319.

- **C++** Try blocks, which are used in class functions; they are described in “try blocks” on page 371.

IBM In addition, for compatibility with GNU C and C++, XL C/C++ allows you to use *attributes* to modify the properties of functions. They are described in “Function attributes” on page 218.

Function definitions take the following form:



Examples of function declarations

The following code fragments show several function declarations (or *prototypes*). The first declares a function `f` that takes two integer arguments and has a return type of `void`:

```
void f(int, int);
```

This fragment declares a pointer `p1` to a function that takes a pointer to a constant character and returns an integer:

```
int (*p1) (const char*);
```

The following code fragment declares a function `f1` that takes an integer argument, and returns a pointer to a function that takes an integer argument and returns an integer:

```
int (*f1(int)) (int);
```

Alternatively, a typedef can be used for the complicated return type of function `f1`:

```
typedef int f1_return_type(int);
f1_return_type* f1(int);
```

The following declaration is of an external function `f2` that takes a constant integer as its first argument, can have a variable number and variable types of other arguments, and returns type `int`.

```
int extern f2(const int, ...); /* C version */
int extern f2(const int ...); // C++ version
```

Function `f6` is a `const` class member function of class `X`, takes no arguments, and has a return type of `int`:

```
class X
{
public:
    int f6() const;
};
```

Function `f4` takes no arguments, has return type `void`, and can throw class objects of types `X` and `Y`.

```
class X;
class Y;

// ...

void f4() throw(X,Y);
```

Examples of function definitions

The following example is a definition of the function `sum`:

```
int sum(int x,int y)
{
    return(x + y);
}
```

The function `sum` has external linkage, returns an object that has type `int`, and has two parameters of type `int` declared as `x` and `y`. The function body contains a single statement that returns the sum of `x` and `y`.

The following function `set_date` declares a pointer to a structure of type `date` as a parameter. `date_ptr` has the storage class specifier `register`.

```
void set_date(register struct date *date_ptr)
{
    date_ptr->mon = 12;
    date_ptr->day = 25;
    date_ptr->year = 87;
}
```

Compatible functions

C only

For two function types to be compatible, they must meet the following requirements:

- They must agree in the number of parameters (and use of ellipsis).
- They must have compatible return types.
- The corresponding parameters must be compatible with the type that results from the application of the default argument promotions.

The composite type of two function types is determined as follows:

- If one of the function types has a parameter type list, the composite type is a function prototype with the same parameter type list.
- If both function types have parameter type lists, the composite type of each parameter is determined as follows:
 - The composite of parameters of different rank is the type that results from the application of the default argument promotions.
 - The composite of parameters with array or function type is the adjusted type.
 - The composite of parameters with qualified type is the unqualified version of the declared type.

For example, for the following two function declarations:

```
int f(int (*)(), double (*)[3]);
int f(int (*)(char *), double (*)[]);
```

The resulting composite type would be:

```
int f(int (*)(char *), double (*)[3]);
```

If the function declarator is not part of the function declaration, the parameters may have incomplete type. The parameters may also specify variable length array types by using the [*] notation in their sequences of declarator specifiers. The following are examples of compatible function prototype declarators:

```
double maximum(int n, int m, double a[n][m]);
double maximum(int n, int m, double a[*][*]);
double maximum(int n, int m, double a[ ][*]);
double maximum(int n, int m, double a[ ][m]);
```

Related information

- “Compatible and composite types” on page 43

End of C only

Multiple function declarations

C++ only

All function declarations for a particular function must have the same number and type of parameters, and must have the same return type.

These return and parameter types are part of the function type, although the default arguments and exception specifications are not.

If a previous declaration of an object or function is visible in an enclosing scope, the identifier has the same linkage as the first declaration. However, a variable or function that has no linkage and later declared with a linkage specifier will have the linkage you have specified.

For the purposes of argument matching, ellipsis and linkage keywords are considered a part of the function type. They must be used consistently in all declarations of a function. If the only difference between the parameter types in two declarations is in the use of typedef names or unspecified argument array bounds, the declarations are the same. A `const` or `volatile` type qualifier is also part of the function type, but can only be part of a declaration or definition of a nonstatic member function.

If two function declarations match in both return type and parameter lists, then the second declaration is treated as redeclaration of the first. The following example declares the same function:

```
int foo(const string &bar);
int foo(const string &);
```

Declaring two functions differing only in return type is not valid function overloading, and is flagged as a compile-time error. For example:

```
void f();
int f();           // error, two definitions differ only in
                  // return type
int g()
{
    return f();
}
```

Related information

- “Overloading functions” on page 243

End of C++ only

Function storage class specifiers

For a function, the storage class specifier determines the linkage of the function. By default, function definitions have external linkage, and can be called by functions defined in other files.  An exception is inline functions, which are treated by default as having internal linkage; see “Linkage of inline functions” on page 209 for more information.

A storage class specifier may be used in both function declarations and definitions. The only storage class options for functions are:

- `static`
- `extern`

The static storage class specifier

A function declared with the `static` storage class specifier has internal linkage, which means that it may be called only within the translation unit in which it is defined.

The `static` storage class specifier can be used in a function declaration only if it is at file scope. You cannot declare functions within a block as `static`.

 This use of `static` is deprecated in C++. Instead, place the function in the unnamed namespace.

Related information

- “Internal linkage” on page 8
- Chapter 9, “Namespaces (C++ only),” on page 235

The extern storage class specifier

A function that is declared with the `extern` storage class specifier has external linkage, which means that it can be called from other translation units. The

keyword `extern` is optional; if you do not specify a storage class specifier, the function is assumed to have external linkage.

C++ In XL C++, an `extern` declaration cannot appear in class scope.

C++ only

In XL C++, you can use the `extern` keyword with arguments that specify the type of linkage.

extern function storage class specifier syntax

► `extern` — "*linkage_specification*" — ◀

where *linkage_specification* can be any of the following:

- C
- C++

The following fragments illustrate the use of `extern "C"`:

```
extern "C" int cf();           //declare function cf to have C linkage

extern "C" int (*c_fp)();     //declare a pointer to a function,
                             // called c_fp, which has C linkage

extern "C" {
    typedef void(*cfp_T)();  //create a type pointer to function with C
                             // linkage
    void cfn();              //create a function with C linkage
    void (*cfp)();          //create a pointer to a function, with C
                             // linkage
}
```

Linkage compatibility affects all C library functions that accept a user function pointer as a parameter, such as `qsort`. Use the `extern "C"` linkage specification to ensure that the declared linkages are the same. The following example fragment uses `extern "C"` with `qsort`.

```
#include <stdlib.h>

// function to compare table elements
extern "C" int TableCmp(const void *, const void *); // C linkage
extern void * GenTable();                          // C++ linkage

int main() {
    void *table;

    table = GenTable();           // generate table
    qsort(table, 100, 15, TableCmp); // sort table, using TableCmp
                                    // and C library routine qsort();
}
```

While the C++ language supports overloading, other languages do not. The implications of this are:

- You can overload a function as long as it has C++ (default) linkage. Therefore, XL C++ allows the following series of statements:

```
int func(int);           // function with C++ linkage
int func(char);         // overloaded function with C++ linkage
```

By contrast, you cannot overload a function that has non-C++ linkage:

```
extern "FORTRAN">{int func(int);}
extern "FORTRAN">{int func(int,int);} // not allowed
//compiler will issue an error message
```

- Only one non-C++-linkage function can have the same name as overloaded functions. For example:

```
int func(char);
int func(int);
extern "FORTRAN">{int func(int,int);}
```

However, the non-C++-linkage function cannot have the same parameters as any of the C++ functions with the same name:

```
int func(char); // first function with C++ linkage
int func(int, int); // second function with C++ linkage
extern "FORTRAN">{int func(int,int);} // not allowed since the parameter
// list is the same as the one for
// the second function with C++ linkage
// compiler will issue an error message
```

End of C++ only

Related information

- “External linkage” on page 8
- “Language linkage (C++ only)” on page 10
- “Class scope (C++ only)” on page 5
- Chapter 9, “Namespaces (C++ only),” on page 235

Function specifiers

The available function specifiers for function definitions are:

- `inline`, which instructs the compiler to expand a function definition at the point of a function call.
-  `explicit`, which can only be used for member functions of classes, and is described in “The explicit specifier” on page 332
-  `virtual`, which can only be used for member functions of classes, and is described in “Virtual functions” on page 309

The inline function specifier

An inline function is one for which the compiler copies the code from the function definition directly into the code of the calling function rather than creating a separate set of instructions in memory. Instead of transferring control to and from the function code segment, a modified copy of the function body may be substituted directly for the function call. In this way, the performance overhead of a function call is avoided. Using the `inline` specifier is only a suggestion to the compiler that an inline expansion can be performed; the compiler is free to ignore the suggestion.

 Any function, with the exception of `main`, can be declared or defined as inline with the `inline` function specifier. Static local variables are not allowed to be defined within the body of an inline function.

 C++ functions implemented inside of a class declaration are automatically defined inline. Regular C++ functions and member functions declared outside of a class declaration, with the exception of `main`, can be declared

or defined as inline with the `inline` function specifier. Static locals and string literals defined within the body of an inline function are treated as the same object across translation units; see “Linkage of inline functions” for details.

The following code fragment shows an inline function definition:

```
inline int add(int i, int j) { return i + j; }
```

The use of the `inline` specifier does not change the meaning of the function. However, the inline expansion of a function may not preserve the order of evaluation of the actual arguments.

The most efficient way to code an inline function is to place the inline function definition in a header file, and then include the header in any file containing a call to the function which you would like to inline.

Note: The `inline` specifier is represented by the following keywords:

-  **C** The `inline` keyword is recognized under compilation with `ppuxlc` or `spuxlc` or `ppuc99`, or with the `-qlanglvl=stdc99` or `-qlanglvl=extc99` options or `-qkeyword=inline`. The `__inline__` keyword is recognized at all language levels; however, see “Linkage of inline functions” below for the semantics of this keyword.
-  **C++** The `inline` and `__inline__` keywords are recognized at all language levels.

Related information

- “The `always_inline` function attribute” on page 220
- “The `noinline` function attribute” on page 222
- `-qlanglvl` and `-qkeyword` in the *XL C/C++ Compiler Reference*

Linkage of inline functions

C only

In C, inline functions are treated by default as having *static* linkage; that is, they are only visible within a single translation unit. Therefore, in the following example, even though function `foo` is defined in exactly the same way, `foo` in file `a.c` and `foo` in file `b.c` are treated as separate functions: two function bodies are generated, and assigned two different addresses in memory:

```
// a.c

#include <stdio.h>

inline int foo(){
return 3;
}

void g() {
printf("foo called from g: return value = %d, address = %p\n", foo(), &foo);
}

// b.c

#include <stdio.h>

inline int foo(){
return 3;
}
```

```

}

void g();

int main() {
printf("foo called from main: return value = %d, address = %p\n", foo(), &foo);
g();
}

```

The output from the compiled program is:

```

foo called from main: return value = 3, address = 0x10000580
foo called from g: return value = 3, address = 0x10000500

```

Since inline functions are treated as having internal linkage, an inline function definition can co-exist with a regular, external definition of a function with the same name in another translation unit. However, when you call the function from the file containing the inline definition, the compiler may choose *either* the inline version defined in the same file *or* the external version defined in another file for the call; your program should not rely on the inline version being called. In the following example, the call to `foo` from function `g` could return either 6 or 3:

```

// a.c

#include <stdio.h>

inline int foo(){
return 6;
}

void g() {
printf("foo called from g: return value = %d\n", foo());
}

// b.c

#include <stdio.h>

int foo(){
return 3;
}

void g();

int main() {
printf("foo called from main: return value = %d\n", foo());
g();
}

```

Similarly, if you define a function as `extern inline`, or redeclare an `inline` function as `extern`, the function simply becomes a regular, external function and is not inlined.

 If you specify the `__inline__` keyword, with the trailing underscores, the compiler uses the GNU C semantics for inline functions. In contrast to the C99 semantics, a function defined as `__inline__` provides an external definition only; a function defined as `static __inline__` provides an inline definition with internal linkage (as in C99); and a function defined as `extern __inline__`, when compiled with optimization enabled, allows the co-existence of an inline and external definition of the same function. For more information on the GNU C implementation of inline functions, see the GCC documentation, available at <http://gcc.gnu.org/onlinedocs/>.

C++ only

You must define an inline function in exactly the same way in each translation unit in which the function is used or called. Furthermore, if a function is defined as `inline`, but never used or called within the same translation unit, it is discarded by the compiler (unless you compile with the `-qkeepinlines` option).

Nevertheless, in C++, inline functions are treated by default as having *external* linkage, meaning that the program behaves as if there is only one copy of the function. The function will have the same address in all translation copy units and each translation unit will share any static locals and string literals. Therefore, compiling the previous example gives the following output:

```
foo called from main: return value = 3, address = 0x10000580
foo called from g: return value = 3, address = 0x10000580
```

Redefining an inline function with the same name but with a different function body is illegal; however, the compiler does not flag this as an error, but simply generates a function body for the version defined in the first file entered on the compilation command line, and discards the others. Therefore, the following example, in which inline function `foo` is defined differently in two different files, may not produce the expected results:

```
// a.C

#include <stdio.h>

inline int foo(){
return 6;
}

void g() {
printf("foo called from g: return value = %d, address = %p\n", foo(), &foo);
}

// b.C

#include <stdio.h>

inline int foo(){
return 3;
}

void g();

int main() {
printf("foo called from main: return value = %d, address = %p\n", foo(), &foo);
g();
}
```

When compiled with the command `ppuxlc++ a.C b.C` or `spuxlc++ a.C b.C`, the output is:

```
foo called from main: return value = 6, address = 0x10001640
foo called from g: return value = 6, address = 0x10001640
```

The call to `foo` from `main` does not use the inline definition provided in `b.C`, but rather calls `foo` as a regular external function defined in `a.C`. It is your

responsibility to ensure that inline function definitions with the same name match exactly across translation units, to avoid unexpected results.

Because inline functions are treated as having external linkage, any static local variables or string literals that are defined within the body of an inline function are treated as the same object across translation units. The following example demonstrates this:

```
// a.C

#include <stdio.h>

inline int foo(){
static int x = 23;
printf("address of x = %p\n", &x);
x++;
return x;
}

void g() {
printf("foo called from g: return value = %d\n", foo());
}

// b.C

#include <stdio.h>

inline int foo()
{
static int x=23;
printf("address of x = %p\n", &x);
x++;
return x;
}

void g();

int main() {
printf("foo called from main: return value = %d\n", foo());
g();
}
```

The output of this program shows that `x` in both definitions of `foo` is indeed the same object:

```
address of x = 0x10011d5c
foo called from main: return value = 24
address of x = 0x10011d5c
foo called from g: return value = 25
```

If you want to ensure that each instance of function defined as inline is treated as a separate function, you can use the `static` specifier in the function definition in each translation unit, or compile with the `-qstaticinline` option. Note, however, that static inline functions are removed from name lookup during template instantiation, and are not found.

Related information

- “The static storage class specifier” on page 206
- “The extern storage class specifier” on page 206
- `-qstaticinline` and `-qkeepinlines` in the *XL C/C++ Compiler Reference*

Function return type specifiers

The result of a function is called its *return value* and the data type of the return value is called the *return type*.

▶ **C++** Every function declaration and definition must specify a return type, whether or not it actually returns a value.

▶ **C** If a function declaration does not specify a return type, the compiler assumes an implicit return type of `int`. However, for conformance to C99, you should specify a return type for every function declaration and definition, whether or not the function returns `int`.

A function may be defined to return any type of value, except an array type or a function type; these exclusions must be handled by returning a pointer to the array or function. When a function does not return a value, `void` is the type specifier in the function declaration and definition.

A function cannot be declared as returning a data object having a `volatile` or `const` type, but it can return a pointer to a `volatile` or `const` object.

A function can have a return type that is a user-defined type. For example:

```
enum count {one, two, three};
enum count counter();
```

▶ **C** The user-defined type may also be defined within the function declaration. ▶ **C++** The user-defined type may not be defined within the function declaration.

```
enum count{one, two, three} counter(); // legal in C
enum count{one, two, three} counter(); // error in C++
```

▶ **C++** References can also be used as return types for functions. The reference returns the lvalue of the object to which it refers.

Related information

- “Type specifiers” on page 52

Function return values

▶ **C** If a function is defined as having a return type of `void`, it should not return a value. ▶ **C++** In C++, a function which is defined as having a return type of `void`, or is a constructor or destructor, *must* not return a value.

▶ **C** If a function is defined as having a return type other than `void`, it should return a value. Under compilation for strict C99 conformance, a function defined with a return type *must* include an expression containing the value to be returned.

▶ **C++** A function defined with a return type *must* include an expression containing the value to be returned.

When a function returns a value, the value is returned via a return statement to the caller of the function, after being implicitly converted to the return type of the function in which it is defined. The following code fragment shows a function definition including the return statement:

```
int add(int i, int j)
{
    return i + j; // return statement
}
```

The function `add()` can be called as shown in the following code fragment:

```
int a = 10,
    b = 20;
int answer = add(a, b); // answer is 30
```

In this example, the return statement initializes a variable of the returned type. The variable `answer` is initialized with the `int` value 30. The type of the returned expression is checked against the returned type. All standard and user-defined conversions are performed as necessary.

Each time a function is called, new copies of its variables with automatic storage are created. Because the storage for these automatic variables may be reused after the function has terminated, a pointer or reference to an automatic variable should not be returned. ▶ C++ If a class object is returned, a temporary object may be created if the class has copy constructors or a destructor.


Related information


- “The return statement” on page 191
- “Overloading assignments” on page 251
- “Overloading subscripting” on page 253
- “The auto storage class specifier” on page 46

Function declarators

Function declarators consist of the following elements:

- An *identifier*, or name
- Parameter declarations, which specify the parameters that can be passed to the function in a function call
- ▶ C++ Exception declarations, which include throw expressions; exception specifications are described in Chapter 16, “Exception handling (C++ only) (PPU only),” on page 371.
- ▶ C++ The type qualifiers `const` and `volatile`, which are used only in class member functions; they are described in “Constant and volatile member functions” on page 272.





Related information


- “Default arguments in C++ functions” on page 229

Parameter declarations


The function declarator includes the list of parameters that can be passed to the function when it is called by another function, or by itself.

C++ In C++, the parameter list of a function is referred to as its *signature*. The name and signature of a function uniquely identify it. As the word itself suggests, the function signature is used by the compiler to distinguish among the different instances of overloaded functions.

Function parameter declaration syntax



parameter



An empty argument list in a function declaration or definition indicates a function that takes no arguments. To explicitly indicate that a function does not take any arguments, you can declare the function in two ways: with an empty parameter list, or with the keyword `void`:

```
int f(void);
int f();
```



An empty argument list in a function *definition* indicates that a function that takes no arguments. An empty argument list in a function *declaration* indicates that a function may take any number or type of arguments. Thus,

```
int f()
{
...
}
```

indicates that function `f` takes no arguments. However,

```
int f();
```

simply indicates that the number and type of parameters is not known. To explicitly indicate that a function does not take any arguments, you should define the function with the keyword `void`.

End of C only

An ellipsis at the end of the parameter specifications is used to specify that a function has a variable number of parameters. The number of parameters is equal to, or greater than, the number of parameter specifications.

```
int f(int, ...);
```

C++ The comma before the ellipsis is optional. In addition, a parameter declaration is not required before the ellipsis.

C At least one parameter declaration, as well as a comma before the ellipsis, are both required in C.

Related information

- “The void type” on page 56
- “Type specifiers” on page 52
- “Type qualifiers” on page 72
- “Exception specifications” on page 382

Parameter types

In a function *declaration*, or prototype, the type of each parameter must be specified. **C++** In the function *definition*, the type of each parameter must also be specified. **C** In the function *definition*, if the type of a parameter is not specified, it is assumed to be `int`.

A variable of a user-defined type may be declared in a parameter declaration, as in the following example, in which `x` is declared for the first time:

```
struct X { int i; };
void print(struct X x);
```

C The user-defined type can also be defined within the parameter declaration. **C++** The user-defined type can not be defined within the parameter declaration.

```
void print(struct X { int i; } x); // legal in C
void print(struct X { int i; } x); // error in C++
```

Parameter names

In a function *definition*, each parameter must have an identifier. In a function *declaration*, or prototype, specifying an identifier is optional. Thus, the following example is legal in a function declaration:

```
int func(int, long);
```

The following constraints apply to the use of parameter names in function declarations:

- Two parameters cannot have the same name within a single declaration.
- If a parameter name is the same as a name outside the function, the name outside the function is hidden and cannot be used in the parameter declaration. In the following example, the third parameter name `intersects` is meant to have enumeration type `subway_line`, but this name is hidden by the name of the first parameter. The declaration of the function `subway()` causes a compile-time error because `subway_line` is not a valid type name because the first parameter name `subway_line` hides the namespace scope enum type and cannot be used again in the second parameter.

```
enum subway_line {yonge,
university, spadina, bloor};
int subway(char * subway_line, int stations,
            subway_line intersects);
```

Static array indices in function parameter declarations (C only)

Except in certain contexts, an unsubscripted array name (for example, `region` instead of `region[4]`) represents a pointer whose value is the address of the first element of the array, provided that the array has previously been declared. An array type in the parameter list of a function is also converted to the corresponding pointer type. Information about the size of the argument array is lost when the array is accessed from within the function body.


To preserve this information, which is useful for optimization, you may declare the index of the argument array using the `static` keyword. The constant expression specifies the minimum pointer size that can be used as an assumption for optimizations. This particular usage of the `static` keyword is highly prescribed. The keyword may only appear in the outermost array type derivation and only in function parameter declarations. If the caller of the function does not abide by these restrictions, the behavior is undefined.

The following examples show how the feature can be used.

```
void foo(int arr [static 10]);      /* arr points to the first of at least
                                10 ints                               */
void foo(int arr [const 10]);     /* arr is a const pointer                               */
void foo(int arr [static const i]); /* arr points to at least i ints;
                                i is computed at run time.           */
void foo(int arr [const static i]); /* alternate syntax to previous example */
void foo(int arr [const]);       /* const pointer to int                               */
```

Related information

- “The static storage class specifier” on page 46
- “Arrays” on page 90
- “Array subscripting operator []” on page 150



End of C++ only

The function attribute in a function declaration is always placed after the declarator, including the parenthesized parameter declaration:

```
/* Specify the attribute on a function prototype declaration */
void f(int i, int j) __attribute__((individual_attribute_name));
void f(int i, int j) { }
```

C++ only

In C++, the attribute specification must also follow any exception declaration that may be present for the function.

End of C++ only

C only

Due to ambiguities in parsing old-style parameter declarations, a function definition must have the attribute specification *precede* the declarator:

```
int __attribute__((individual_attribute_name)) foo(int i) { }
```

End of C only

A function attribute specification using the form `__attribute_name__` (that is, the attribute name with double underscore characters leading and trailing) reduces the likelihood of a name conflict with a macro of the same name.

The following function attributes are supported:

- “The alias function attribute” on page 220
- “The always_inline function attribute” on page 220
- “The const function attribute” on page 220
- “The constructor and destructor function attributes” on page 221
- “The format function attribute” on page 221
- “The format_arg function attribute” on page 222
- “The noinline function attribute” on page 222
- “The noreturn function attribute” on page 223
- “The pure function attribute” on page 223
- “The section function attribute” on page 223
- “The weak function attribute” on page 224

Related information

- “Variable attributes” on page 105

The alias function attribute

The alias function attribute causes the function declaration to appear in the object file as an alias for another symbol. This language feature provides a technique for coping with duplicate or cumbersome names.

alias function attribute syntax

```
►► __attribute__((alias ("original_function_name")))
```

C The aliased function can be defined after the specification of its alias with this function attribute. C also allows an alias specification in the absence of a definition of the aliased function in the same compilation unit.

C++ The *original_function_name* must be the mangled name.

The following declares bar to be an alias for `__foo`:

```
_____ C only _____  
void __foo(){ /* function body */ }  
void bar() __attribute__((alias("__foo")));  
_____ End of C only _____
```

```
_____ C++ only _____  
extern "C" __foo(){ /* function body */ }  
void bar() __attribute__((alias("__foo")));  
_____ End of C++ only _____
```

The compiler does not check for consistency between the declaration of bar and definition of `__foo`. Such consistency remains the responsibility of the programmer.

The always_inline function attribute

The `always_inline` function attribute instructs the compiler to inline an inline function, regardless of whether optimization was specified at compile time. However, the attribute has no effect if the program is compiled at no-opt levels. Specifying this attribute for a function without an `inline` specification also has no effect. The attribute takes precedence over inlining compiler options.

always_inline function attribute syntax

```
►► __attribute__((always_inline))
```

Related information


- “The inline function specifier” on page 208
- “The noinline function attribute” on page 222

The const function attribute

The const function attribute allows you to tell the compiler that the function can safely be called fewer times than indicated in the source code. The language

feature provides you with an explicit way to help the compiler optimize code by indicating that the function does not examine any values except its arguments and has no effects except for its return value.

const function attribute syntax



The following kinds of functions should not be declared const:

- A function with pointer arguments which examines the data pointed to.
- A function that calls a non-const function.

Related information


- `#pragma isolated_call` in the *XL C/C++ Compiler Reference*

The constructor and destructor function attributes

The constructor and destructor function attributes provide the ability to write a function that initializes data or releases storage that is used implicitly during program execution. A function to which the constructor function attribute has been applied is called automatically before execution enters `main`. Similarly, a function to which the destructor attribute has been applied is called automatically after calling `exit` or upon completion of `main`.

When the constructor or destructor function is called automatically, the return value of the function is ignored, and any parameters of the function are undefined.

constructor and destructor function attribute syntax




A function declaration containing a constructor or destructor function attribute must match all of its other declarations.

The format function attribute

The format function attribute provides a way to identify user-defined functions that take format strings as arguments so that calls to these functions will be type-checked against a format string, similar to the way the compiler checks calls to the functions `printf`, `scanf`, `strftime`, and `strfmon` for errors.

format function attribute syntax



where

noinline function attribute syntax

```
▶▶ __attribute__((noinline|_noinline_)) ▶▶
```

Other than preventing inlining, the attribute does not remove the semantics of inline functions.

The noreturn function attribute

The noreturn function attribute allows you to indicate to the compiler that the function is not intended to return. The language feature provides the programmer with another explicit way to help the compiler optimize code and to reduce false warnings for uninitialized variables.

The return type of the function should be void.

noreturn function attribute syntax

```
▶▶ __attribute__((noreturn|_noreturn_)) ▶▶
```

Registers saved by the calling function may not necessarily be restored before calling the nonreturning function.

Related information

- `#pragma leaves` in the *XL C/C++ Compiler Reference*

The pure function attribute

The pure function attribute allows you to declare a function that can be called fewer times than what is literally in the source code. Declaring a function with the attribute pure indicates that the function has no effect except a return value that depends only on the parameters, global variables, or both.

pure function attribute syntax

```
▶▶ __attribute__((pure|_pure_)) ▶▶
```

Related information

- `#pragma isolated_call` in the *XL C/C++ Compiler Reference*

The section function attribute

The section function attribute specifies the section in the object file in which the compiler should place its generated code. The language feature provides the ability to control the section in which a function should appear.

section function attribute syntax

```
▶▶ __attribute__((section("section_name")|_section_)) ▶▶
```

where *section_name* is a string literal.

Each defined function can reside in only one section. The section indicated in a function definition should match that in any previous declaration. The section indicated in a function definition cannot be overwritten, whereas one in a function declaration can be overwritten by a later specification. Moreover, if a section attribute is applied to a function declaration, the function will be placed in the specified section only if it is defined in the same compilation unit.

Related information

- “The section variable attribute” on page 109

The weak function attribute

The weak function attribute causes the symbol resulting from the function declaration to appear in the object file as a weak symbol, rather than a global one. The language feature provides the programmer writing library functions with a way to allow function definitions in user code to override the library function declaration without causing duplicate name errors.

weak function attribute syntax

```

▶▶ __attribute__((weak|__weak))

```

Related information

- “The alias function attribute” on page 220
- **#pragma weak** in the *XL C/C++ Compiler Reference*
- “The weak variable attribute” on page 110

End of IBM extension

The main() function

When a program begins running, the system calls the function `main`, which marks the entry point of the program. By default, `main` has the storage class `extern`. Every program must have one function named `main`, and the following constraints apply:

- No other function in the program can be called `main`.
- `main` cannot be defined as `inline` or `static`.
-  `main` cannot be called from within a program.
-  The address of `main` cannot be taken.
-  The `main` function cannot be overloaded.

The function `main` can be defined with or without parameters, using any of the following forms:

```

int main (void)
int main ( )
int main(int argc, char *argv[])
int main (int argc, char ** argv)

```

Although any name can be given to these parameters, they are usually referred to as `argc` and `argv`. The first parameter, `argc` (argument count) is an integer that indicates how many arguments were entered on the command line when the program was started. The second parameter, `argv` (argument vector), is an array of

pointers to arrays of character objects. The array objects are null-terminated strings, representing the arguments that were entered on the command line when the program was started.

The first element of the array, `argv[0]`, is a pointer to the character array that contains the program name or invocation name of the program that is being run from the command line. `argv[1]` indicates the first argument passed to the program, `argv[2]` the second argument, and so on. The following example program `backward` prints the arguments entered on a command line such that the last argument is printed first:

```
#include <stdio.h>
int main(int argc, char *argv[])
{
    while (--argc > 0)
        printf("%s ", argv[argc]);
    printf("\n");
}
```

Invoking this program from a command line with the following:

```
backward string1 string2
```

gives the following output:

```
string2 string1
```

The arguments `argc` and `argv` would contain the following values:

Object	Value
<code>argc</code>	3
<code>argv[0]</code>	pointer to string "backward"
<code>argv[1]</code>	pointer to string "string1"
<code>argv[2]</code>	pointer to string "string2"
<code>argv[3]</code>	NULL

Related information

- “The extern storage class specifier” on page 48
- “The inline function specifier” on page 208
- “The static storage class specifier” on page 46
- “Function calls”

Function calls

Once a function has been declared and defined, it can be *called* from anywhere within the program: from within the `main` function, from another function, and even from itself. Calling the function involves specifying the function name, followed by the function call operator and any data values the function expects to receive. These values are the *arguments* for the parameters defined for the function, and the process just described is called *passing arguments* to the function.

 A function may not be called if it has not already been declared.

Passing arguments can be done in two ways:

- Pass by value, which copies the *value* of an argument to the corresponding parameter in the called function

- Pass by reference, which passes the *address* of an argument to the corresponding parameter in the called function

C++ only

If a class has a destructor or a copy constructor that does more than a bitwise copy, passing a class object by value results in the construction of a temporary object that is actually passed by reference.

It is an error when a function argument is a class object and all of the following properties hold:

- The class needs a copy constructor.
- The class does not have a user-defined copy constructor.
- A copy constructor cannot be generated for that class.

End of C++ only

Related information

- “Function argument conversions” on page 119
- “Function call expressions” on page 127
- “Constructors” on page 319

Pass by value

When you use pass-by-*value*, the compiler copies the value of an argument in a calling function to a corresponding non-pointer or non-reference parameter in the called function definition. The parameter in the called function is initialized with the value of the passed argument. As long as the parameter has not been declared as constant, the value of the parameter can be changed, but the changes are only performed within the scope of the called function only; they have no effect on the value of the argument in the calling function.

In the following example, `main` passes `func` two values: 5 and 7. The function `func` receives copies of these values and accesses them by the identifiers `a` and `b`. The function `func` changes the value of `a`. When control passes back to `main`, the actual values of `x` and `y` are not changed.

```
/**
 ** This example illustrates calling a function by value
 **/

#include <stdio.h>

void func (int a, int b)
{
    a += b;
    printf("In func, a = %d    b = %d\n", a, b);
}

int main(void)
{
    int x = 5, y = 7;
    func(x, y);
    printf("In main, x = %d    y = %d\n", x, y);
    return 0;
}
```

The output of the program is:

In func, a = 12 b = 7
In main, x = 5 y = 7

Pass by reference

Passing by *reference* refers to a method of passing the address of an argument in the calling function to a corresponding parameter in the called function.

► **C** In C, the corresponding parameter in the called function must be declared as a pointer type. ► **C++** In C++, the corresponding parameter can be declared as any reference type, not just a pointer type.

In this way, the value of the argument in the calling function can be modified by the called function.

The following example shows how arguments are passed by reference. In C++, the reference parameters are initialized with the actual arguments when the function is called. In C, the pointer parameters are initialized with pointer values when the function is called.

```
_____ C++ only _____  
  
#include <stdio.h>  
  
void swapnum(int &i, int &j) {  
    int temp = i;  
    i = j;  
    j = temp;  
}  
  
int main(void) {  
    int a = 10;  
    int b = 20;  
  
    swapnum(a, b);  
    printf("A is %d and B is %d\n", a, b);  
    return 0;  
}  
_____ End of C++ only _____
```

```
_____ C only _____  
  
#include <stdio.h>  
  
void swapnum(int *i, int *j) {  
    int temp = *i;  
    *i = *j;  
    *j = temp;  
}  
  
int main(void) {  
    int a = 10;  
    int b = 20;  
  
    swapnum(&a, &b);  
    printf("A is %d and B is %d\n", a, b);  
    return 0;  
}  
_____ End of C only _____
```

When the function `swapnum()` is called, the actual values of the variables `a` and `b` are exchanged because they are passed by reference. The output is:

A is 20 and B is 10

C++ only

In order to modify a reference that is const-qualified, you must cast away its constness with the `const_cast` operator. The following example demonstrates this:

```
#include <iostream>
using namespace std;

void f(const int& x) {
    int* y = const_cast<int>(&x);
    (*y)++;
}

int main() {
    int a = 5;
    f(a);
    cout << a << endl;
}
```

This example outputs 6.

End of C++ only

Related information

- “References (C++ only)” on page 93
- “The `const_cast` operator (C++ only)” on page 161

Allocation and deallocation functions (C++ only)

You may define your own `new` operator or allocation function as a class member function or a global namespace function with the following restrictions:

- The first parameter must be of type `std::size_t`. It cannot have a default parameter.
- The return type must be of type `void*`.
- Your allocation function may be a template function. Neither the first parameter nor the return type may depend on a template parameter.
- If you declare your allocation function with the empty exception specification `throw()`, your allocation function must return a null pointer if your function fails. Otherwise, your function must throw an exception of type `std::bad_alloc` or a class derived from `std::bad_alloc` if your function fails.

You may define your own `delete` operator or deallocation function as a class member function or a global namespace function with the following restrictions:

- The first parameter must be of type `void*`.
- The return type must be of type `void`.
- Your deallocation function may be a template function. Neither the first parameter nor the return type may depend on a template parameter.

The following example defines replacement functions for global namespace `new` and `delete`:

```
#include <cstdio>
#include <cstdlib>

using namespace std;

void* operator new(size_t sz) {
```

```

    printf("operator new with %d bytes\n", sz);
    void* p = malloc(sz);
    if (p == 0) printf("Memory error\n");
    return p;
}

void operator delete(void* p) {
    if (p == 0) printf ("Deleting a null pointer\n");
    else {
        printf("delete object\n");
        free(p);
    }
}

struct A {
    const char* data;
    A() : data("Text String") { printf("Constructor of S\n"); }
    ~A() { printf("Destructor of S\n"); }
};

int main() {
    A* ap1 = new A;
    delete ap1;

    printf("Array of size 2:\n");
    A* ap2 = new A[2];
    delete[] ap2;
}

```

The following is the output of the above example:

```

operator
new with 40 bytes
operator new with 33 bytes
operator new with 4 bytes
Constructor of S
Destructor of S
delete object
Array of size 2:
operator new with 16 bytes
Constructor of S
Constructor of S
Destructor of S
Destructor of S
delete object

```

Related information

- “new expressions (C++ only)” on page 165

Default arguments in C++ functions

You can provide default values for function parameters. For example:

```

#include <iostream>
using namespace std;

int a = 1;
int f(int a) { return a; }
int g(int x = f(a)) { return x; }

int h() {
    a = 2;
    {
        int a = 3;
        return g();
    }
}

```

```

}

int main() {
    cout << h() << endl;
}

```

This example prints 2 to standard output, because the a referred to in the declaration of g() is the one at file scope, which has the value 2 when g() is called.

The default argument must be implicitly convertible to the parameter type.

A pointer to a function must have the same type as the function. Attempts to take the address of a function by reference without specifying the type of the function will produce an error. The type of a function is not affected by arguments with default values.

The following example shows that default arguments are not considered part of a function's type. The default argument allows you to call a function without specifying all of the arguments, it does not allow you to create a pointer to the function that does not specify the types of all the arguments. Function f can be called without an explicit argument, but the pointer badpointer cannot be defined without specifying the type of the argument:

```

int f(int = 0);
void g()
{
    int a = f(1);           // ok
    int b = f();           // ok, default argument used
}
int (*pointer)(int) = &f; // ok, type of f() specified (int)
int (*badpointer)() = &f; // error, badpointer and f have
                          // different types. badpointer must
                          // be initialized with a pointer to
                          // a function taking no arguments.

```

In this example, function f3 has a return type int, and takes an int argument with a default value that is the value returned from function f2:

```

const int j = 5;
int f3( int x = f2(j) );

```

Related information

- "Pointers to functions" on page 232

Restrictions on default arguments

Of the operators, only the function call operator and the operator new can have default arguments when they are overloaded.

Parameters with default arguments must be the trailing parameters in the function declaration parameter list. For example:

```

void f(int a, int b = 2, int c = 3); // trailing defaults
void g(int a = 1, int b = 2, int c); // error, leading defaults
void h(int a, int b = 3, int c);    // error, default in middle

```

Once a default argument has been given in a declaration or definition, you cannot redefine that argument, even to the same value. However, you can add default arguments not given in previous declarations. For example, the last declaration below attempts to redefine the default values for a and b:

```

void f(int a, int b, int c=1);    // valid
void f(int a, int b=1, int c);   // valid, add another default
void f(int a=1, int b, int c);   // valid, add another default
void f(int a=1, int b=1, int c=1); // error, redefined defaults

```

You can supply any default argument values in the function declaration or in the definition. Any parameters in the parameter list following a default argument value must have a default argument value specified in this or a previous declaration of the function.

You cannot use local variables in default argument expressions. For example, the compiler generates errors for both function `g()` and function `h()` below:

```

void f(int a)
{
    int b=4;
    void g(int c=a); // Local variable "a" cannot be used here
    void h(int d=b); // Local variable "b" cannot be used here
}

```

Related information

- “Function call expressions” on page 127
- “new expressions (C++ only)” on page 165

Evaluation of default arguments

When a function defined with default arguments is called with trailing arguments missing, the default expressions are evaluated. For example:

```

void f(int a, int b = 2, int c = 3); // declaration
// ...
int a = 1;
f(a);           // same as call f(a,2,3)
f(a,10);       // same as call f(a,10,3)
f(a,10,20);    // no default arguments

```

Default arguments are checked against the function declaration and evaluated when the function is called. The order of evaluation of default arguments is undefined. Default argument expressions cannot use other parameters of the function. For example:

```

int f(int q = 3, int r = q); // error

```

The argument `r` cannot be initialized with the value of the argument `q` because the value of `q` may not be known when it is assigned to `r`. If the above function declaration is rewritten:

```

int q=5;
int f(int q = 3, int r = q); // error

```

The value of `r` in the function declaration still produces an error because the variable `q` defined outside of the function is hidden by the argument `q` declared for the function. Similarly:

```

typedef double D;
int f(int D, int z = D(5.3) ); // error

```

Here the type `D` is interpreted within the function declaration as the name of an integer. The type `D` is hidden by the argument `D`. The cast `D(5.3)` is therefore not interpreted as a cast because `D` is the name of the argument not a type.

In the following example, the nonstatic member `a` cannot be used as an initializer because `a` does not exist until an object of class `X` is constructed. You can use the

static member `b` as an initializer because `b` is created independently of any objects of class `X`. You can declare the member `b` after its use as a default argument because the default values are not analyzed until after the final bracket `}` of the class declaration.

```
class X
{
    int a;
    f(int z = a) ; // error
    g(int z = b) ; // valid
    static int b;
};
```

Pointers to functions

A pointer to a function points to the address of the executable code of the function. You can use pointers to call functions and to pass functions as arguments to other functions. You cannot perform pointer arithmetic on pointers to functions.

The type of a pointer to a function is based on both the return type and parameter types of the function.

A declaration of a pointer to a function must have the pointer name in parentheses. The function call operator `()` has a higher precedence than the dereference operator `*`. Without them, the compiler interprets the statement as a function that returns a pointer to a specified return type. For example:

```
int *f(int a);          /* function f returning an int*          */
int (*g)(int a);       /* pointer g to a function returning an int          */
char (*h)(int, int)    /* h is a function
                        that takes two integer parameters and returns char */
```

In the first declaration, `f` is interpreted as a function that takes an `int` as argument, and returns a pointer to an `int`. In the second declaration, `g` is interpreted as a pointer to a function that takes an `int` argument and that returns an `int`.

Related information

- “Language linkage (C++ only)” on page 10
- “Pointers” on page 86
- “Pointer conversions” on page 117
- “The extern storage class specifier” on page 206

Nested functions

IBM extension

C only

A nested function is a function defined inside the definition of another function. It can be defined wherever a variable declaration is permitted, which allows nested functions within nested functions. Within the containing function, the nested function can be declared prior to being defined by using the `auto` keyword. Otherwise, a nested function has internal linkage. The language feature is an extension to C89 and C99, implemented to facilitate porting programs developed with GNU C.

A nested function can access all identifiers of the containing function that precede its definition.

A nested function must not be called after the containing function exits.

A nested function cannot use a `goto` statement to jump to a label in the containing function, or to a local label declared with the `__label__` keyword inherited from the containing function.

_____ **End of C only** _____

Related information

- “Locally declared labels” on page 176

_____ **End of IBM extension** _____

In this example, the IBM identifier is an alias for INTERNATIONAL_BUSINESS_MACHINES. This is useful for referring to long namespace identifiers.

If a namespace name or alias is declared as the name of any other entity in the same declarative region, a compiler error will result. Also, if a namespace name defined at global scope is declared as the name of any other entity in any global scope of the program, a compiler error will result.

Related information

- “File/global scope” on page 3

Creating an alias for a nested namespace

Namespace definitions hold declarations. Since a namespace definition is a declaration itself, namespace definitions can be nested.

An alias can also be applied to a nested namespace.

```
namespace INTERNATIONAL_BUSINESS_MACHINES {
    int j;
    namespace NESTED_IBM_PRODUCT {
        void a() { j++; }
        int j;
        void b() { j++; }
    }
}
namespace NIBM = INTERNATIONAL_BUSINESS_MACHINES::NESTED_IBM_PRODUCT
```

In this example, the NIBM identifier is an alias for the namespace NESTED_IBM_PRODUCT. This namespace is nested within the INTERNATIONAL_BUSINESS_MACHINES namespace.

Related information

- “Creating a namespace alias” on page 235

Extending namespaces

Namespaces are extensible. You can add subsequent declarations to a previously defined namespace. Extensions may appear in files separate from or attached to the original namespace definition. For example:

```
namespace X { // namespace definition
    int a;
    int b;
}

namespace X { // namespace extension
    int c;
    int d;
}

namespace Y { // equivalent to namespace X
    int a;
    int b;
    int c;
    int d;
}
```

In this example, namespace *X* is defined with *a* and *b* and later extended with *c* and *d*. namespace *X* now contains all four members. You may also declare all of the required members within one namespace. This method is represented by namespace *Y*. This namespace contains *a*, *b*, *c*, and *d*.

Namespaces and overloading

You can overload functions across namespaces. For example:

```
// Original X.h:
f(int);

// Original Y.h:
f(char);

// Original program.c:
#include "X.h"
#include "Y.h"

void z()
{
    f('a'); // calls f(char) from Y.h
}
```

Namespaces can be introduced to the previous example without drastically changing the source code.

```
// New X.h:
namespace X {
    f(int);
}

// New Y.h:
namespace Y {
    f(char);
}

// New program.c:
#include "X.h"
#include "Y.h"

using namespace X;
using namespace Y;

void z()
{
    f('a'); // calls f() from Y.h
}
```

In `program.c`, function `void z()` calls function `f()`, which is a member of namespace *Y*. If you place the `using` directives in the header files, the source code for `program.c` remains unchanged.

Related information

- Chapter 10, “Overloading (C++ only),” on page 243

Unnamed namespaces

A namespace with no identifier before an opening brace produces an *unnamed namespace*. Each translation unit may contain its own unique unnamed namespace. The following example demonstrates how unnamed namespaces are useful.

```

#include <iostream>

using namespace std;

namespace {
    const int i = 4;
    int variable;
}

int main()
{
    cout << i << endl;
    variable = 100;
    return 0;
}

```

In the previous example, the unnamed namespace permits access to `i` and `variable` without using a scope resolution operator.

The following example illustrates an improper use of unnamed namespaces.

```

#include <iostream>

using namespace std;

namespace {
    const int i = 4;
}

int i = 2;

int main()
{
    cout << i << endl; // error
    return 0;
}

```

Inside `main`, `i` causes an error because the compiler cannot distinguish between the global name and the unnamed namespace member with the same name. In order for the previous example to work, the namespace must be uniquely identified with an identifier and `i` must specify the namespace it is using.

You can extend an unnamed namespace within the same translation unit. For example:

```

#include <iostream>

using namespace std;

namespace {
    int variable;
    void funct (int);
}

namespace {
    void funct (int i) { cout << i << endl; }
}

int main()
{
    funct(variable);
    return 0;
}

```

both the prototype and definition for `funct` are members of the same unnamed namespace.

Note: Items defined in an unnamed namespace have internal linkage. Rather than using the keyword `static` to define items with internal linkage, define them in an unnamed namespace instead.

Related information

- “Program linkage” on page 7
- “Internal linkage” on page 8

Namespace member definitions

A namespace can define its own members within itself or externally using explicit qualification. The following is an example of a namespace defining a member internally:

```
namespace A {
    void b() { /* definition */ }
}
```

Within namespace `A` member `void b()` is defined internally.

A namespace can also define its members externally using explicit qualification on the name being defined. The entity being defined must already be declared in the namespace and the definition must appear after the point of declaration in a namespace that encloses the declaration’s namespace.

The following is an example of a namespace defining a member externally:

```
namespace A {
    namespace B {
        void f();
    }
    void B::f() { /* defined outside of B */ }
}
```

In this example, function `f()` is declared within namespace `B` and defined (outside `B`) in `A`.

Namespaces and friends

Every name first declared in a namespace is a member of that namespace. If a friend declaration in a non-local class first declares a class or function, the friend class or function is a member of the innermost enclosing namespace.

The following is an example of this structure:

```
// f has not yet been defined
void z(int);
namespace A {
    class X {
        friend void f(X); // A::f is a friend
    };
    // A::f is not visible here
    X x;
    void f(X) { /* definition */ } // f() is defined and known to be a friend
}

using A::x;
```

```

void z()
{
    A::f(x);    // OK
    A::X::f(x); // error: f is not a member of A::X
}

```

In this example, function `f()` can only be called through namespace `A` using the call `A::f(s)`. Attempting to call function `f()` through class `X` using the `A::X::f(x)`; call results in a compiler error. Since the friend declaration first occurs in a non-local class, the friend function is a member of the innermost enclosing namespace and may only be accessed through that namespace.

Related information

- “Friends” on page 286

The using directive

A `using` directive provides access to all namespace qualifiers and the scope operator. This is accomplished by applying the `using` keyword to a namespace identifier.

Using directive syntax

```

▶▶ using namespace name; ◀◀

```

The *name* must be a previously defined namespace. The `using` directive may be applied at the global and local scope but not the class scope. Local scope takes precedence over global scope by hiding similar declarations.

If a scope contains a `using` directive that nominates a second namespace and that second namespace contains another `using` directive, the `using` directive from the second namespace will act as if it resides within the first scope.

```

namespace A {
    int i;
}
namespace B {
    int i;
    using namespace A;
}
void f()
{
    using namespace B;
    i = 7; // error
}

```

In this example, attempting to initialize `i` within function `f()` causes a compiler error, because function `f()` cannot know which `i` to call; `i` from namespace `A`, or `i` from namespace `B`.

Related information

- “The using declaration and class members” on page 298

The using declaration and namespaces

A `using` declaration provides access to a specific namespace member. This is accomplished by applying the `using` keyword to a namespace name with its corresponding namespace member.

Using declaration syntax

►—using—namespace—::—member—◄

In this syntax diagram, the qualifier name follows the using declaration and the *member* follows the qualifier name. For the declaration to work, the member must be declared inside the given namespace. For example:

```
namespace A {  
    int i;  
    int k;  
    void f;  
    void g;  
}
```

```
using A::k
```

In this example, the using declaration is followed by *A*, the name of namespace *A*, which is then followed by the scope operator (`::`), and *k*. This format allows *k* to be accessed outside of namespace *A* through a using declaration. After issuing a using declaration, any extension made to that specific namespace will not be known at the point at which the using declaration occurs.

Overloaded versions of a given function must be included in the namespace prior to that given function's declaration. A using declaration may appear at namespace, block and class scope.

Related information

- “The using declaration and class members” on page 298

Explicit access

To explicitly qualify a member of a namespace, use the namespace identifier with a `::` scope resolution operator.

Explicit access qualification syntax

►—namespace_name—::—member—◄

For example:

```
namespace VENDITTI {  
    void j()  
};
```

```
VENDITTI::j();
```

In this example, the scope resolution operator provides access to the function *j* held within namespace *VENDITTI*. The scope resolution operator `::` is used to access identifiers in both global and local namespaces. Any identifier in an application can be accessed with sufficient qualification. Explicit access cannot be applied to an unnamed namespace.

Related information

- “Scope resolution operator `::` (C++ only)” on page 126

Chapter 10. Overloading (C++ only)

If you specify more than one definition for a function name or an operator in the same scope, you have *overloaded* that function name or operator. Overloaded functions and operators are described in “Overloading functions” and “Overloading operators” on page 245, respectively.

An *overloaded declaration* is a declaration that had been declared with the same name as a previously declared declaration in the same scope, except that both declarations have different types.

If you call an overloaded function name or operator, the compiler determines the most appropriate definition to use by comparing the argument types you used to call the function or operator with the parameter types specified in the definitions. The process of selecting the most appropriate overloaded function or operator is called *overload resolution*, as described in “Overload resolution” on page 255.

Overloading functions

You overload a function name *f* by declaring more than one function with the name *f* in the same scope. The declarations of *f* must differ from each other by the types and/or the number of arguments in the argument list. When you call an overloaded function named *f*, the correct function is selected by comparing the argument list of the function call with the parameter list of each of the overloaded candidate functions with the name *f*. A *candidate function* is a function that can be called based on the context of the call of the overloaded function name.

Consider a function `print`, which displays an `int`. As shown in the following example, you can overload the function `print` to display other types, for example, `double` and `char*`. You can have three functions with the same name, each performing a similar operation on a different data type:

```
#include <iostream>
using namespace std;

void print(int i) {
    cout << " Here is int " << i << endl;
}
void print(double f) {
    cout << " Here is float " << f << endl;
}

void print(char* c) {
    cout << " Here is char* " << c << endl;
}

int main() {
    print(10);
    print(10.10);
    print("ten");
}
```

The following is the output of the above example:

```
Here is int 10
Here is float 10.1
Here is char* ten
```

Function overloading based on vector parameter types is supported.

Related information

- “Restrictions on overloaded functions”
- “Derivation” on page 293

Restrictions on overloaded functions

You cannot overload the following function declarations if they appear in the same scope. Note that this list applies only to explicitly declared functions and those that have been introduced through using declarations:

- Function declarations that differ only by return type. For example, you cannot declare the following declarations:

```
int f();  
float f();
```

- Member function declarations that have the same name and the same parameter types, but one of these declarations is a static member function declaration. For example, you cannot declare the following two member function declarations of `f()`:

```
struct A {  
    static int f();  
    int f();  
};
```

- Member function template declarations that have the same name, the same parameter types, and the same template parameter lists, but one of these declarations is a static template member function declaration.
- Function declarations that have equivalent parameter declarations. These declarations are not allowed because they would be declaring the same function.
- Function declarations with parameters that differ only by the use of typedef names that represent the same type. Note that a typedef is a synonym for another type, not a separate type. For example, the following two declarations of `f()` are declarations of the same function:

```
typedef int I;  
void f(float, int);  
void f(float, I);
```

- Function declarations with parameters that differ only because one is a pointer and the other is an array. For example, the following are declarations of the same function:

```
f(char*);  
f(char[10]);
```

The first array dimension is insignificant when differentiating parameters; all other array dimensions are significant. For example, the following are declarations of the same function:

```
g(char*[20]);  
g(char[5][20]);
```

The following two declarations are *not* equivalent:

```
g(char*[20]);  
g(char*[40]);
```

- Function declarations with parameters that differ only because one is a function type and the other is a pointer to a function of the same type. For example, the following are declarations of the same function:

```
void f(int(float));  
void f(int (*)(float));
```

- Function declarations with parameters that differ only because of cv-qualifiers `const`, `volatile`, and `restrict`. This restriction only applies if any of these qualifiers appears at the outermost level of a parameter type specification. For example, the following are declarations of the same function:

```
int f(int);
int f(const int);
int f(volatile int);
```

Note that you can differentiate parameters with `const`, `volatile` and `restrict` qualifiers if you apply them *within* a parameter type specification. For example, the following declarations are *not* equivalent because `const` and `volatile` qualify `int`, rather than `*`, and thus are not at the outermost level of the parameter type specification.

```
void g(int*);
void g(const int*);
void g(volatile int*);
```

The following declarations are also not equivalent:

```
void g(float&);
void g(const float&);
void g(volatile float&);
```

- Function declarations with parameters that differ only because their default arguments differ. For example, the following are declarations of the same function:

```
void f(int);
void f(int i = 10);
```

- Multiple functions with extern "C" language-linkage and the same name, regardless of whether their parameter lists are different.

Related information

- “The using declaration and namespaces” on page 240
- “typedef definitions” on page 71
- “Type qualifiers” on page 72
- “Language linkage (C++ only)” on page 10

Overloading operators

You can redefine or overload the function of most built-in operators in C++. These operators can be overloaded globally or on a class-by-class basis. Overloaded operators are implemented as functions and can be member functions or global functions.

Operator overloading involving vector types is not supported.

An overloaded operator is called an *operator function*. You declare an operator function with the keyword `operator` preceding the operator. Overloaded operators are distinct from overloaded functions, but like overloaded functions, they are distinguished by the number and types of operands used with the operator.

Consider the standard `+` (plus) operator. When this operator is used with operands of different standard types, the operators have slightly different meanings. For example, the addition of two integers is not implemented in the same way as the addition of two floating-point numbers. C++ allows you to define your own meanings for the standard C++ operators when they are applied to class types. In the following example, a class called `complex` is defined to model complex numbers,

and the + (plus) operator is redefined in this class to add two complex numbers.

```

// This example illustrates overloading the plus (+) operator.

#include <iostream>
using namespace std;

class complx
{
    double real,
          imag;
public:
    complx( double real = 0., double imag = 0.); // constructor
    complx operator+(const complx& c) const;    // operator+()
};

// define constructor
complx::complx( double r, double i )
{
    real = r; imag = i;
}

// define overloaded + (plus) operator
complx complx::operator+ (const complx& c) const
{
    complx result;
    result.real = (this->real + c.real);
    result.imag = (this->imag + c.imag);
    return result;
}

int main()
{
    complx x(4,4);
    complx y(6,6);
    complx z = x + y; // calls complx::operator+()
}

```

You can overload any of the following operators:

+	-	*	/	%	^	&		~
!	=	<	>	+=	--	*=	/=	%=
^=	&=	=	<<	>>	<<=	>>=	==	!=
<=	>=	&&		++	--	,	->*	->
()	[]	new	delete	new[]	delete[]			

where () is the function call operator and [] is the subscript operator.

You can overload both the unary and binary forms of the following operators:

+	-	*	&
---	---	---	---

You cannot overload the following operators:

.	.*	::	?:
---	----	----	----

You cannot overload the preprocessor symbols # and ##.

An operator function can be either a nonstatic member function, or a nonmember function with at least one parameter that has class, reference to class, enumeration, or reference to enumeration type.

You cannot change the precedence, grouping, or the number of operands of an operator.

An overloaded operator (except for the function call operator) cannot have default arguments or an ellipsis in the argument list.

You must declare the overloaded `=`, `[]`, `()`, and `->` operators as nonstatic member functions to ensure that they receive lvalues as their first operands.

The operators `new`, `delete`, `new[]`, and `delete[]` do not follow the general rules described in this section.

All operators except the `=` operator are inherited.

Overloading unary operators

You overload a unary operator with either a nonstatic member function that has no parameters, or a nonmember function that has one parameter. Suppose a unary operator `@` is called with the statement `@t`, where `t` is an object of type `T`. A nonstatic member function that overloads this operator would have the following form:

```
return_type operator@()
```

A nonmember function that overloads the same operator would have the following form:

```
return_type operator@(T)
```

An overloaded unary operator may return any type.

The following example overloads the `!` operator:

```
#include <iostream>
using namespace std;

struct X { };

void operator!(X) {
    cout << "void operator!(X)" << endl;
}

struct Y {
    void operator!() {
        cout << "void Y::operator!()" << endl;
    }
};

struct Z { };

int main() {
    X ox; Y oy; Z oz;
    !ox;
    !oy;
    // !oz;
}
```

The following is the output of the above example:

```
void operator!(X)
void Y::operator!()
```

The operator function call `!ox` is interpreted as `operator!(X)`. The call `!oy` is interpreted as `Y::operator!()`. (The compiler would not allow `!oz` because the `!` operator has not been defined for class `Z`.)

Related information

- “Unary expressions” on page 128

Overloading increment and decrement operators

You overload the prefix increment operator `++` with either a nonmember function operator that has one argument of class type or a reference to class type, or with a member function operator that has no arguments.

In the following example, the increment operator is overloaded in both ways:

```
class X {
public:

    // member prefix ++x
    void operator++() { }
};

class Y { };

// non-member prefix ++y
void operator++(Y&) { }

int main() {
    X x;
    Y y;

    // calls x.operator++()
    ++x;

    // explicit call, like ++x
    x.operator++();

    // calls operator++(y)
    ++y;

    // explicit call, like ++y
    operator++(y);
}
```

The postfix increment operator `++` can be overloaded for a class type by declaring a nonmember function operator `operator++()` with two arguments, the first having class type and the second having type `int`. Alternatively, you can declare a member function operator `operator++()` with one argument having type `int`. The compiler uses the `int` argument to distinguish between the prefix and postfix increment operators. For implicit calls, the default value is zero.

For example:

```
class X {
public:

    // member postfix x++
    void operator++(int) { };
};

class Y { };

// nonmember postfix y++
void operator++(Y&, int) { };

int main() {
    X x;
    Y y;

    // calls x.operator++(0)
```

```

// default argument of zero is supplied by compiler
x++;
// explicit call to member postfix x++
x.operator++(0);

// calls operator++(y, 0)
y++;

// explicit call to non-member postfix y++
operator++(y, 0);
}

```

The prefix and postfix decrement operators follow the same rules as their increment counterparts.

Related information

- “Increment operator ++” on page 129
- “Decrement operator --” on page 130

Overloading binary operators

You overload a binary unary operator with either a nonstatic member function that has one parameter, or a nonmember function that has two parameters. Suppose a binary operator @ is called with the statement `t @ u`, where `t` is an object of type `T`, and `u` is an object of type `U`. A nonstatic member function that overloads this operator would have the following form:

```
return_type operator@(T)
```

A nonmember function that overloads the same operator would have the following form:

```
return_type operator@(T, U)
```

An overloaded binary operator may return any type.

The following example overloads the `*` operator:

```

struct X {

    // member binary operator
    void operator*(int) { }
};

// non-member binary operator
void operator*(X, float) { }

int main() {
    X x;
    int y = 10;
    float z = 10;

    x * y;
    x * z;
}

```

The call `x * y` is interpreted as `x.operator*(y)`. The call `x * z` is interpreted as `operator*(x, z)`.

Related information

- “Binary expressions” on page 139

Overloading assignments

You overload the assignment operator, `operator=`, with a nonstatic member function that has only one parameter. You cannot declare an overloaded assignment operator that is a nonmember function. The following example shows how you can overload the assignment operator for a particular class:

```
struct X {
    int data;
    X& operator=(X& a) { return a; }
    X& operator=(int a) {
        data = a;
        return *this;
    }
};

int main() {
    X x1, x2;
    x1 = x2;    // call x1.operator=(x2)
    x1 = 5;    // call x1.operator=(5)
}
```

The assignment `x1 = x2` calls the copy assignment operator `X& X::operator=(X&)`. The assignment `x1 = 5` calls the copy assignment operator `X& X::operator=(int)`. The compiler implicitly declares a copy assignment operator for a class if you do not define one yourself. Consequently, the copy assignment operator (`operator=`) of a derived class hides the copy assignment operator of its base class.

However, you can declare any copy assignment operator as virtual. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct A {
    A& operator=(char) {
        cout << "A& A::operator=(char)" << endl;
        return *this;
    }
    virtual A& operator=(const A&) {
        cout << "A& A::operator=(const A&)" << endl;
        return *this;
    }
};

struct B : A {
    B& operator=(char) {
        cout << "B& B::operator=(char)" << endl;
        return *this;
    }
    virtual B& operator=(const A&) {
        cout << "B& B::operator=(const A&)" << endl;
        return *this;
    }
};

struct C : B { };

int main() {
    B b1;
    B b2;
    A* ap1 = &b1;
    A* ap2 = &b1;
    *ap1 = 'z';
    *ap2 = b2;
}
```

```

    C c1;
    // c1 = 'z';
}

```

The following is the output of the above example:

```

A& A::operator=(char)
B& B::operator=(const A&)

```

The assignment `*ap1 = 'z'` calls `A& A::operator=(char)`. Because this operator has not been declared `virtual`, the compiler chooses the function based on the type of the pointer `ap1`. The assignment `*ap2 = b2` calls `B& B::operator=(const &A)`. Because this operator has been declared `virtual`, the compiler chooses the function based on the type of the object that the pointer `ap1` points to. The compiler would not allow the assignment `c1 = 'z'` because the implicitly declared copy assignment operator declared in class `C` hides `B& B::operator=(char)`.

Related information

- “Copy assignment operators” on page 335
- “Assignment operators” on page 140

Overloading function calls

The function call operator, when overloaded, does not modify how functions are called. Rather, it modifies how the operator is to be interpreted when applied to objects of a given type.

You overload the function call operator, `operator()`, with a nonstatic member function that has any number of parameters. If you overload a function call operator for a class its declaration will have the following form:

```

return_type operator()(parameter_list)

```

Unlike all other overloaded operators, you can provide default arguments and ellipses in the argument list for the function call operator.

The following example demonstrates how the compiler interprets function call operators:

```

struct A {
    void operator()(int a, char b, ...) { }
    void operator()(char c, int d = 20) { }
};

int main() {
    A a;
    a(5, 'z', 'a', 0);
    a('z');
    // a();
}

```

The function call `a(5, 'z', 'a', 0)` is interpreted as `a.operator()(5, 'z', 'a', 0)`. This calls `void A::operator()(int a, char b, ...)`. The function call `a('z')` is interpreted as `a.operator>('z')`. This calls `void A::operator()(char c, int d = 20)`. The compiler would not allow the function call `a()` because its argument list does not match any function call parameter list defined in class `A`.

The following example demonstrates an overloaded function call operator:

```

class Point {
private:
    int x, y;
public:
    Point() : x(0), y(0) { }
    Point& operator()(int dx, int dy) {
        x += dx;
        y += dy;
        return *this;
    }
};

int main() {
    Point pt;

    // Offset this coordinate x with 3 points
    // and coordinate y with 2 points.
    pt(3, 2);
}

```

The above example reinterprets the function call operator for objects of class `Point`. If you treat an object of `Point` like a function and pass it two integer arguments, the function call operator will add the values of the arguments you passed to `Point::x` and `Point::y` respectively.

Related information

- “Function call expressions” on page 127

Overloading subscripting

You overload `operator[]` with a nonstatic member function that has only one parameter. The following example is a simple array class that has an overloaded subscripting operator. The overloaded subscripting operator throws an exception if you try to access the array outside of its specified bounds:

```

#include <iostream>
using namespace std;

template <class T> class MyArray {
private:
    T* storage;
    int size;
public:
    MyArray(int arg = 10) {
        storage = new T[arg];
        size = arg;
    }

    ~MyArray() {
        delete[] storage;
        storage = 0;
    }

    T& operator[](const int location) throw (const char *);
};

template <class T> T& MyArray<T>::operator[](const int location)
    throw (const char *) {
    if (location < 0 || location >= size) throw "Invalid array access";
    else return storage[location];
}

int main() {
    try {
        MyArray<int> x(13);
    }
}

```

```

        x[0] = 45;
        x[1] = 2435;
        cout << x[0] << endl;
        cout << x[1] << endl;
        x[13] = 84;
    }
    catch (const char* e) {
        cout << e << endl;
    }
}

```

The following is the output of the above example:

```

45
2435
Invalid array access

```

The expression `x[1]` is interpreted as `x.operator[](1)` and calls `int& MyArray<int>::operator[](const int)`.

Related information

- “Array subscripting operator []” on page 150

Overloading class member access

You overload operator-`>` with a nonstatic member function that has no parameters. The following example demonstrates how the compiler interprets overloaded class member access operators:

```

struct Y {
    void f() { };
};

struct X {
    Y* ptr;
    Y* operator->() {
        return ptr;
    };
};

int main() {
    X x;
    x->f();
}

```

The statement `x->f()` is interpreted as `(x.operator->())->f()`.

The operator-`>` is used (often in conjunction with the pointer-dereference operator) to implement “smart pointers.” These pointers are objects that behave like normal pointers except they perform other tasks when you access an object through them, such as automatic object deletion (either when the pointer is destroyed, or the pointer is used to point to another object), or reference counting (counting the number of smart pointers that point to the same object, then automatically deleting the object when that count reaches zero).

One example of a smart pointer is included in the C++ Standard Library called `auto_ptr`. You can find it in the `<memory>` header. The `auto_ptr` class implements automatic object deletion.

Related information

- “Arrow operator ->” on page 128

Overload resolution

The process of selecting the most appropriate overloaded function or operator is called *overload resolution*.

Suppose that `f` is an overloaded function name. When you call the overloaded function `f()`, the compiler creates a set of *candidate functions*. This set of functions includes all of the functions named `f` that can be accessed from the point where you called `f()`. The compiler may include as a candidate function an alternative representation of one of those accessible functions named `f` to facilitate overload resolution.

After creating a set of candidate functions, the compiler creates a set of *viable functions*. This set of functions is a subset of the candidate functions. The number of parameters of each viable function agrees with the number of arguments you used to call `f()`.

The compiler chooses the *best viable function*, the function declaration that the C++ run-time environment will use when you call `f()`, from the set of viable functions. The compiler does this by *implicit conversion sequences*. An implicit conversion sequence is the sequence of conversions required to convert an argument in a function call to the type of the corresponding parameter in a function declaration. The implicit conversion sequences are ranked; some implicit conversion sequences are better than others. The best viable function is the one whose parameters all have either better or equal-ranked implicit conversion sequences than all of the other viable functions. The compiler will not allow a program in which the compiler was able to find more than one best viable function. Implicit conversion sequences are described in more detail in “Implicit conversion sequences” on page 256.

When a variable length array is a function parameter, the leftmost array dimension does not distinguish functions among candidate functions. In the following, the second definition of `f` is not allowed because `void f(int [])` has already been defined.

```
void f(int a[*]) {}  
void f(int a[5]) {} // illegal
```

However, array dimensions other than the leftmost in a variable length array do differentiate candidate functions when the variable length array is a function parameter. For example, the overload set for function `f` might comprise the following:

```
void f(int a[][5]) {}  
void f(int a[][4]) {}  
void f(int a[][g]) {} // assume g is a global int
```

but cannot include

```
void f(int a[][g2]) {} // illegal, assuming g2 is a global int
```

because having candidate functions with second-level array dimensions `g` and `g2` creates ambiguity about which function `f` should be called: neither `g` nor `g2` is known at compile time.

You can override an exact match by using an explicit cast. In the following example, the second call to `f()` matches with `f(void*)`:

```

void f(int) { };
void f(void*) { };

int main() {
    f(0xaabb);           // matches f(int);
    f((void*) 0xaabb);  // matches f(void*)
}

```

Implicit conversion sequences

An *implicit conversion sequence* is the sequence of conversions required to convert an argument in a function call to the type of the corresponding parameter in a function declaration.

The compiler will try to determine an implicit conversion sequence for each argument. It will then categorize each implicit conversion sequence in one of three categories and rank them depending on the category. The compiler will not allow any program in which it cannot find an implicit conversion sequence for an argument.

The following are the three categories of conversion sequences in order from best to worst:

- Standard conversion sequences
- User-defined conversion sequences
- Ellipsis conversion sequences

Note: Two standard conversion sequences or two user-defined conversion sequences may have different ranks.

Standard conversion sequences

Standard conversion sequences are categorized in one of three ranks. The ranks are listed in order from best to worst:

- Exact match: This rank includes the following conversions:
 - Identity conversions
 - Lvalue-to-rvalue conversions
 - Array-to-pointer conversions
 - Qualification conversions
- Promotion: This rank includes integral and floating point promotions.
- Conversion: This rank includes the following conversions:
 - Integral and floating-point conversions
 - Floating-integral conversions
 - Pointer conversions
 - Pointer-to-member conversions
 - Boolean conversions

The compiler ranks a standard conversion sequence by its worst-ranked standard conversion. For example, if a standard conversion sequence has a floating-point conversion, then that sequence has conversion rank.

Related information

- “Lvalue-to-rvalue conversions” on page 116
- “Pointer conversions” on page 117
- “Qualification conversions (C++ only)” on page 119

- “Integral conversions” on page 114
- “Floating-point conversions” on page 114
- “Boolean conversions” on page 114

User-defined conversion sequences

A *user-defined conversion sequence* consists of the following:

- A standard conversion sequence
- A user-defined conversion
- A second standard conversion sequence

A user-defined conversion sequence A is better than a user-defined conversion sequence B if the both have the same user-defined conversion function or constructor, and the second standard conversion sequence of A is better than the second standard conversion sequence of B.

Ellipsis conversion sequences

An *ellipsis conversion sequence* occurs when the compiler matches an argument in a function call with a corresponding ellipsis parameter.

Resolving addresses of overloaded functions

If you use an overloaded function name *f* without any arguments, that name can refer to a function, a pointer to a function, a pointer to member function, or a specialization of a function template. Because you did not provide any arguments, the compiler cannot perform overload resolution the same way it would for a function call or for the use of an operator. Instead, the compiler will try to choose the best viable function that matches the type of one of the following expressions, depending on where you have used *f*:

- An object or reference you are initializing
- The left side of an assignment
- A parameter of a function or a user-defined operator
- The return value of a function, operator, or conversion
- An explicit type conversion

If the compiler chose a declaration of a nonmember function or a static member function when you used *f*, the compiler matched the declaration with an expression of type pointer-to-function or reference-to-function. If the compiler chose a declaration of a nonstatic member function, the compiler matched that declaration with an expression of type pointer-to-member function. The following example demonstrates this:

```
struct X {
    int f(int) { return 0; }
    static int f(char) { return 0; }
};

int main() {
    int (X::*a)(int) = &X::f;
    // int (*b)(int) = &X::f;
}
```

The compiler will not allow the initialization of the function pointer *b*. No nonmember function or static function of type `int(int)` has been declared.

If *f* is a template function, the compiler will perform template argument deduction to determine which template function to use. If successful, it will add that function to the list of viable functions. If there is more than one function in this set,

including a non-template function, the compiler will eliminate all template functions from the set and choose the non-template function. If there are only template functions in this set, the compiler will choose the most specialized template function. The following example demonstrates this:

```
template<class T> int f(T) { return 0; }
template<> int f(int) { return 0; }
int f(int) { return 0; }

int main() {
    int (*a)(int) = f;
    a(1);
}
```

The function call `a(1)` calls `int f(int)`.

Related information

- “Pointers to functions” on page 232
- “Pointers to members” on page 274
- “Function templates” on page 348
- “Explicit specialization” on page 359

Chapter 11. Classes (C++ only)

A *class* is a mechanism for creating user-defined data types. It is similar to the C language structure data type. In C, a structure is composed of a set of data members. In C++, a class type is like a C structure, except that a class is composed of a set of data members and a set of operations that can be performed on the class.

In C++, a class type can be declared with the keywords `union`, `struct`, or `class`. A union object can hold any one of a set of named members. Structure and class objects hold a complete set of members. Each class type represents a unique set of class members including data members, member functions, and other type names. The default access for members depends on the class key:

- The members of a class declared with the keyword `class` are private by default. A class is inherited privately by default.
- The members of a class declared with the keyword `struct` are public by default. A structure is inherited publicly by default.
- The members of a union (declared with the keyword `union`) are public by default. A union cannot be used as a base class in derivation.

Once you create a class type, you can declare one or more objects of that class type. For example:

```
class X
{
    /* define class members here */
};
int main()
{
    X xobject1;    // create an object of class type X
    X xobject2;    // create another object of class type X
}
```

You may have *polymorphic* classes in C++. Polymorphism is the ability to use a function name that appears in different classes (related by inheritance), without knowing exactly the class the function belongs to at compile time.

C++ allows you to redefine standard operators and functions through the concept of overloading. Operator overloading facilitates data abstraction by allowing you to use classes as easily as built-in types.

Related information


- “Structures and unions” on page 59
- Chapter 12, “Class members and friends (C++ only),” on page 269
- Chapter 13, “Inheritance (C++ only),” on page 291
- Chapter 10, “Overloading (C++ only),” on page 243
- “Virtual functions” on page 309

Declaring class types

A class declaration creates a unique type class name.

A *class specifier* is a type specifier used to declare a class. Once a class specifier has been seen and its members declared, a class is considered to be defined even if the member functions of that class are not yet defined.

Class specifier syntax



The *class_name* is a unique identifier that becomes a reserved word within its scope. Once a class name is declared, it hides other declarations of the same name within the enclosing scope.

The *member_list* specifies the class members, both data and functions, of the class *class_name*. If the *member_list* of a class is empty, objects of that class have a nonzero size. You can use a *class_name* within the *member_list* of the class specifier itself as long as the size of the class is not required.

The *base_clause* specifies the base class or classes from which the class *class_name* inherits members. If the *base_clause* is not empty, the class *class_name* is called a *derived class*.

A *structure* is a class declared with the *class_key* struct. The members and base classes of a structure are public by default. A *union* is a class declared with the *class_key* union. The members of a union are public by default; a union holds only one data member at a time.

An *aggregate class* is a class that has no user-defined constructors, no private or protected non-static data members, no base classes, and no virtual functions.

Related information

- “Class member lists” on page 269
- “Derivation” on page 293

Using class objects

You can use a class type to create instances or *objects* of that class type. For example, you can declare a class, structure, and union with class names X, Y, and Z respectively:

```
class X {  
    // members of class X  
};  
  
struct Y {  
    // members of struct Y  
};  
  
union Z {  
    // members of union Z  
};
```

You can then declare objects of each of these class types. Remember that classes, structures, and unions are all types of C++ classes.

```
int main()
{
    X xobj;    // declare a class object of class type X
    Y yobj;    // declare a struct object of class type Y
    Z zobj;    // declare a union object of class type Z
}
```

In C++, unlike C, you do not need to precede declarations of class objects with the keywords `union`, `struct`, and `class` unless the name of the class is hidden. For example:

```
struct Y { /* ... */ };
class X { /* ... */ };
int main ()
{
    int X;        // hides the class name X
    Y yobj;       // valid
    X xobj;       // error, class name X is hidden
    class X xobj; // valid
}
```

When you declare more than one class object in a declaration, the declarators are treated as if declared individually. For example, if you declare two objects of class `S` in a single declaration:

```
class S { /* ... */ };
int main()
{
    S S,T; // declare two objects of class type S
}
```

this declaration is equivalent to:

```
class S { /* ... */ };
int main()
{
    S S;
    class S T; // keyword class is required
               // since variable S hides class type S
}
```

but is not equivalent to:

```
class S { /* ... */ };
int main()
{
    S S;
    S T; // error, S class type is hidden
}
```

You can also declare references to classes, pointers to classes, and arrays of classes. For example:

```
class X { /* ... */ };
struct Y { /* ... */ };
union Z { /* ... */ };
int main()
{
    X xobj;
    X &xref = xobj; // reference to class object of type X
    Y *yptr;       // pointer to struct object of type Y
    Z zarray[10];  // array of 10 union objects of type Z
}
```

You can initialize classes in external, static, and automatic definitions. The initializer contains an = (equal sign) followed by a brace-enclosed, comma-separated list of values. You do not need to initialize all members of a class.

Objects of class types that are not copy restricted can be assigned, passed as arguments to functions, and returned by functions.

Related information

- “Structures and unions” on page 59
- “References (C++ only)” on page 93
- “Scope of class names” on page 263

Classes and structures

The C++ class is an extension of the C language structure. Because the only difference between a structure and a class is that structure members have public access by default and class members have private access by default, you can use the keywords `class` or `struct` to define equivalent classes.

For example, in the following code fragment, the class `X` is equivalent to the structure `Y`:

```
class X {  
  
    // private by default  
    int a;  
  
public:  
  
    // public member function  
    int f() { return a = 5; };  
};  
  
struct Y {  
  
    // public by default  
    int f() { return a = 5; };  
  
private:  
  
    // private data member  
    int a;  
};
```

If you define a structure and then declare an object of that structure using the keyword `class`, the members of the object are still public by default. In the following example, `main()` has access to the members of `obj_X` even though `obj_X` has been declared using an elaborated type specifier that uses the class key `class`:

```
#include <iostream>  
using namespace std;  
  
struct X {  
    int a;  
    int b;  
};  
  
class X obj_X;  
  
int main() {
```

```

obj_X.a = 0;
obj_X.b = 1;
cout << "Here are a and b: " << obj_X.a << " " << obj_X.b << endl;
}

```

The following is the output of the above example:

Here are a and b: 0 1

Related information


- “Structures and unions” on page 59

Scope of class names


A class declaration introduces the class name into the scope where it is declared. Any class, object, function or other declaration of that name in an enclosing scope is hidden.

If a class name is declared in the same scope as a function, enumerator, or object with the same name, you must refer to that class using an *elaborated type specifier*:

Elaborated type specifier syntax



Nested name specifier:



The following example must use an elaborated type specifier to refer to class A because this class is hidden by the definition of the function A():

```

class A { };

void A (class A*) { };

int main()
{
    class A* x;
    A(x);
}

```

The declaration `class A* x` is an elaborated type specifier. Declaring a class with the same name of another function, enumerator, or object as demonstrated above is not recommended.

An elaborated type specifier can also be used in the incomplete declaration of a class type to reserve the name for a class type within the current scope.

Related information

- “Class scope (C++ only)” on page 5

- “Incomplete class declarations”

Incomplete class declarations

An *incomplete class declaration* is a class declaration that does not define any class members. You cannot declare any objects of the class type or refer to the members of a class until the declaration is complete. However, an incomplete declaration allows you to make specific references to a class prior to its definition as long as the size of the class is not required.

For example, you can define a pointer to the structure first in the definition of the structure second. Structure first is declared in an incomplete class declaration prior to the definition of second, and the definition of oneptr in structure second does not require the size of first:

```
struct first;           // incomplete declaration of struct first

struct second          // complete declaration of struct second
{
    first* oneptr;     // pointer to struct first refers to
                     // struct first prior to its complete
                     // declaration

    first one;        // error, you cannot declare an object of
                     // an incompletely declared class type

    int x, y;
};

struct first           // complete declaration of struct first
{
    second two;       // define an object of class type second
    int z;
};
```

However, if you declare a class with an empty member list, it is a complete class declaration. For example:

```
class X;               // incomplete class declaration
class Z {};           // empty member list
class Y
{
public:
    X yobj;           // error, cannot create an object of an
                     // incomplete class type
    Z zobj;           // valid
};
```

Related information

- “Class member lists” on page 269

Nested classes

A *nested class* is declared within the scope of another class. The name of a nested class is local to its enclosing class. Unless you use explicit pointers, references, or object names, declarations in a nested class can only use visible constructs, including type names, static members, and enumerators from the enclosing class and global variables.

Member functions of a nested class follow regular access rules and have no special access privileges to members of their enclosing classes. Member functions of the enclosing class have no special access to members of a nested class. The following example demonstrates this:

```

class A {
    int x;

    class B { };

    class C {

        // The compiler cannot allow the following
        // declaration because A::B is private:
        //   B b;

        int y;
        void f(A* p, int i) {

            // The compiler cannot allow the following
            // statement because A::x is private:
            //   p->x = i;

        }
    };

    void g(C* p) {

        // The compiler cannot allow the following
        // statement because C::y is private:
        //   int z = p->y;
    }
};

int main() { }

```

The compiler would not allow the declaration of object `b` because class `A::B` is private. The compiler would not allow the statement `p->x = i` because `A::x` is private. The compiler would not allow the statement `int z = p->y` because `C::y` is private.

You can define member functions and static data members of a nested class in namespace scope. For example, in the following code fragment, you can access the static members `x` and `y` and member functions `f()` and `g()` of the nested class nested by using a qualified type name. Qualified type names allow you to define a typedef to represent a qualified class name. You can then use the typedef with the `::` (scope resolution) operator to refer to a nested class or class member, as shown in the following example:

```

class outside
{
public:
    class nested
    {
    public:
        static int x;
        static int y;
        int f();
        int g();
    };
};

int outside::nested::x = 5;
int outside::nested::f() { return 0; };

typedef outside::nested outnest;    // define a typedef
int outnest::y = 10;                // use typedef with ::
int outnest::g() { return 0; };

```

However, using a typedef to represent a nested class name hides information and may make the code harder to understand.

You cannot use a typedef name in an elaborated type specifier. To illustrate, you cannot use the following declaration in the above example:

```
class outnest obj;
```

A nested class may inherit from private members of its enclosing class. The following example demonstrates this:

```
class A {
private:
    class B { };
    B *z;

    class C : private B {
private:
        B y;
//      A::B y2;
        C *x;
//      A::C *x2;
    };
};
```

The nested class `A::C` inherits from `A::B`. The compiler does not allow the declarations `A::B y2` and `A::C *x2` because both `A::B` and `A::C` are private.

Related information

- “Class scope (C++ only)” on page 5
- “Scope of class names” on page 263
- “Member access” on page 283
- “Static members” on page 279

Local classes

A *local class* is declared within a function definition. Declarations in a local class can only use type names, enumerations, static variables from the enclosing scope, as well as external variables and functions.

For example:

```
int x; // global variable
void f() // function definition
{
    static int y; // static variable y can be used by
                // local class
    int x; // auto variable x cannot be used by
          // local class
    extern int g(); // extern function g can be used by
                  // local class

    class local // local class
    {
        int g() { return x; } // error, local variable x
                            // cannot be used by g
        int h() { return y; } // valid, static variable y
        int k() { return ::x; } // valid, global x
        int l() { return g(); } // valid, extern function g
    };
}

int main()
{
    local* z; // error: the class local is not visible
    // ...}
```

Member functions of a local class have to be defined within their class definition, if they are defined at all. As a result, member functions of a local class are inline functions. Like all member functions, those defined within the scope of a local class do not need the keyword `inline`.

A local class cannot have static data members. In the following example, an attempt to define a static member of a local class causes an error:

```
void f()
{
    class local
    {
        int f();           // error, local class has noninline
                          // member function
        int g() {return 0;} // valid, inline member function
        static int a;      // error, static is not allowed for
                          // local class
        int b;            // valid, nonstatic variable
    };
}
// . . .
```

An enclosing function has no special access to members of the local class.

Related information

- “Member functions” on page 271
- “The inline function specifier” on page 208

Local type names

Local type names follow the same scope rules as other names. Type names defined within a class declaration have class scope and cannot be used outside their class without qualification.

If you use a class name, typedef name, or a constant name that is used in a type name, in a class declaration, you cannot redefine that name after it is used in the class declaration.

For example:

```
int main ()
{
    typedef double db;
    struct st
    {
        db x;
        typedef int db; // error
        db y;
    };
}
```

The following declarations are valid:

```
typedef float T;
class s {
    typedef int T;
    void f(const T);
};
```

Here, function `f()` takes an argument of type `s::T`. However, the following declarations, where the order of the members of `s` has been reversed, cause an error:

```
typedef float T;
class s {
    void f(const T);
    typedef int T;
};
```

In a class declaration, you cannot redefine a name that is not a class name, or a typedef name to a class name or typedef name once you have used that name in the class declaration.

Related information

- “Scope” on page 2
- “typedef definitions” on page 71

Chapter 12. Class members and friends (C++ only)

This section discusses the declaration of class members with respect to the information hiding mechanism and how a class can grant functions and classes access to its nonpublic members by the use of the friend mechanism. C++ expands the concept of information hiding to include the notion of having a public class interface but a private implementation. It is the mechanism for limiting direct access to the internal representation of a class type by functions in a program.


Related information

- “Member access” on page 283
- “Inherited member access” on page 296

Class member lists

An optional *member list* declares subobjects called *class members*. Class members can be data, functions, nested types, and enumerators.

Class member list syntax



The member list follows the class name and is placed between braces. The following applies to member lists, and members of member lists:

- A *member_declaration* or a *member_definition* may be a declaration or definition of a data member, member function, nested type, or enumeration. (The enumerators of an enumeration defined in a class member list are also members of the class.)
- A member list is the only place where you can declare class members.
- Friend declarations are not class members but must appear in member lists.
- The member list in a class definition declares all the members of a class; you cannot add members elsewhere.
- You cannot declare a member twice in a member list.
- You may declare a data member or member function as `static` but not `auto`, `extern`, or `register`.
- You may declare a nested class, a member class template, or a member function, and define it outside the class.
- You must define static data members outside the class.
- Nonstatic members that are class objects must be objects of previously defined classes; a class A cannot contain an object of class A, but it can contain a pointer or reference to an object of class A.
- You must specify all dimensions of a nonstatic array member.

A *constant initializer* (= *constant_expression*) may only appear in a class member of integral or enumeration type that has been declared `static`.

A *pure specifier* (= `0`) indicates that a function has no definition. It is only used with member functions declared as `virtual` and replaces the function definition of a member function in the member list.

An *access specifier* is one of `public`, `private`, or `protected`.

A *member declaration* declares a class member for the class containing the declaration.

The order of allocation of nonstatic class members separated by an *access_specifier* is implementation-dependent.

The order of allocation of nonstatic class members separated by an *access_specifier* is implementation-dependent. The compiler allocates class members in the order in which they are declared.

Suppose `A` is a name of a class. The following class members of `A` must have a name different from `A`:

- All data members
- All type members
- All enumerators of enumerated type members
- All members of all anonymous union members

Related information

- “Declaring class types” on page 260
- “Member access” on page 283
- “Inherited member access” on page 296
- “Static members” on page 279

Data members

Data members include members that are declared with any of the fundamental types, as well as other types, including pointer, reference, array types, bit fields, and user-defined types. You can declare a data member the same way as a variable, except that explicit initializers are not allowed inside the class definition. However, a `const static` data member of integral or enumeration type may have an explicit initializer.

If an array is declared as a nonstatic class member, you must specify all of the dimensions of the array.

A class can have members that are of a class type or are pointers or references to a class type. Members that are of a class type must be of a class type that has been previously declared. An incomplete class type can be used in a member declaration as long as the size of the class is not needed. For example, a member can be declared that is a pointer to an incomplete class type.

A class `X` cannot have a member that is of type `X`, but it can contain pointers to `X`, references to `X`, and static objects of `X`. Member functions of `X` can take arguments of type `X` and have a return type of `X`. For example:

```

class X
{
    X();
    X *xptr;
    X &xref;
    static X xcount;
    X xfunc(X);
};

```

Related information

- “Member access” on page 283
- “Inherited member access” on page 296
- “Static members” on page 279

Member functions

Member functions are operators and functions that are declared as members of a class. Member functions do not include operators and functions declared with the friend specifier. These are called *friends* of a class. You can declare a member function as `static`; this is called a *static member function*. A member function that is not declared as `static` is called a *nonstatic member function*.

The definition of a member function is within the scope of its enclosing class. The body of a member function is analyzed after the class declaration so that members of that class can be used in the member function body, even if the member function definition appears before the declaration of that member in the class member list. When the function `add()` is called in the following example, the data variables `a`, `b`, and `c` can be used in the body of `add()`.

```

class x
{
public:
    int add()           // inline member function add
    {return a+b+c;};
private:
    int a,b,c;
};

```

Inline member functions

You may either define a member function inside its class definition, or you may define it outside if you have already declared (but not defined) the member function in the class definition.

A member function that is defined inside its class member list is called an *inline member function*. Member functions containing a few lines of code are usually declared inline. In the above example, `add()` is an inline member function. If you define a member function outside of its class definition, it must appear in a namespace scope enclosing the class definition. You must also qualify the member function name using the scope resolution (`::`) operator.

An equivalent way to declare an inline member function is to either declare it in the class with the `inline` keyword (and define the function outside of its class) or to define it outside of the class declaration using the `inline` keyword.

In the following example, member function `Y::f()` is an inline member function:

```

struct Y {
private:
    char* a;
public:
    char* f() { return a; }
};

```

The following example is equivalent to the previous example; `Y::f()` is an inline member function:

```

struct Y {
private:
    char* a;
public:
    char* f();
};

inline char* Y::f() { return a; }

```

The `inline` specifier does not affect the linkage of a member or nonmember function: linkage is external by default.

Member functions of a local class must be defined within their class definition. As a result, member functions of a local class are implicitly inline functions. These inline member functions have no linkage.

Related information

- “Friends” on page 286
- “Static member functions” on page 282
- “The inline function specifier” on page 208
- “Local classes” on page 266

Constant and volatile member functions

A member function declared with the `const` qualifier can be called for constant and nonconstant objects. A nonconstant member function can only be called for a nonconstant object. Similarly, a member function declared with the `volatile` qualifier can be called for volatile and nonvolatile objects. A nonvolatile member function can only be called for a nonvolatile object.

Related information

- “Type qualifiers” on page 72
- “The `this` pointer” on page 276

Virtual member functions

Virtual member functions are declared with the keyword `virtual`. They allow dynamic binding of member functions. Because all virtual functions must be member functions, virtual member functions are simply called *virtual functions*.

If the definition of a virtual function is replaced by a pure specifier in the declaration of the function, the function is said to be declared pure. A class that has at least one pure virtual function is called an *abstract class*.

Related information

- “Virtual functions” on page 309
- “Abstract classes” on page 314

Special member functions

Special member functions are used to create, destroy, initialize, convert, and copy class objects. These include the following:

- Constructors
- Destructors
- Conversion constructors
- Conversion functions
- Copy constructors

For full descriptions of these functions, see Chapter 14, “Special member functions (C++ only),” on page 317.

Member scope

Member functions and static members can be defined outside their class declaration if they have already been declared, but not defined, in the class member list. Nonstatic data members are defined when an object of their class is created. The declaration of a static data member is not a definition. The declaration of a member function is a definition if the body of the function is also given.

Whenever the definition of a class member appears outside of the class declaration, the member name must be qualified by the class name using the `::` (scope resolution) operator.

The following example defines a member function outside of its class declaration.

```
#include <iostream>
using namespace std;

struct X {
    int a, b ;

    // member function declaration only
    int add();
};

// global variable
int a = 10;

// define member function outside its class declaration
int X::add() { return a + b; }

int main() {
    int answer;
    X xobject;
    xobject.a = 1;
    xobject.b = 2;
    answer = xobject.add();
    cout << xobject.a << " + " << xobject.b << " = " << answer << endl;
}
```

The output for this example is: 1 + 2 = 3

All member functions are in class scope even if they are defined outside their class declaration. In the above example, the member function `add()` returns the data member `a`, not the global variable `a`.

The name of a class member is local to its class. Unless you use one of the class access operators, `.` (dot), or `->` (arrow), or `::` (scope resolution) operator, you can

only use a class member in a member function of its class and in nested classes. You can only use types, enumerations and static members in a nested class without qualification with the `::` operator.

The order of search for a name in a member function body is:

1. Within the member function body itself
2. Within all the enclosing classes, including inherited members of those classes
3. Within the lexical scope of the body declaration

The search of the enclosing classes, including inherited members, is demonstrated in the following example:

```
class A { /* ... */ };
class B { /* ... */ };
class C { /* ... */ };
class Z : A {
    class Y : B {
        class X : C { int f(); /* ... */ };
    };
};
int Z::Y::X f()
{
    char j;
    return 0;
}
```

In this example, the search for the name `j` in the definition of the function `f` follows this order:

1. In the body of the function `f`
2. In `X` and in its base class `C`
3. In `Y` and in its base class `B`
4. In `Z` and in its base class `A`
5. In the lexical scope of the body of `f`. In this case, this is global scope.

Note that when the containing classes are being searched, only the definitions of the containing classes and their base classes are searched. The scope containing the base class definitions (global scope, in this example) is not searched.

Related information

- “Class scope (C++ only)” on page 5

Pointers to members

Pointers to members allow you to refer to nonstatic members of class objects. You cannot use a pointer to member to point to a static class member because the address of a static member is not associated with any particular object. To point to a static class member, you must use a normal pointer.

You can use pointers to member functions in the same manner as pointers to functions. You can compare pointers to member functions, assign values to them, and use them to call member functions. Note that a member function does not have the same type as a nonmember function that has the same number and type of arguments and the same return type.

Pointers to members can be declared and used as shown in the following example:

```
#include <iostream>
using namespace std;

class X {
```

```

public:
    int a;
    void f(int b) {
        cout << "The value of b is " << b << endl;
    }
};

int main() {

    // declare pointer to data member
    int X::*ptiptr = &X::a;

    // declare a pointer to member function
    void (X::* ptfptr) (int) = &X::f;

    // create an object of class type X
    X xobject;

    // initialize data member
    xobject.*ptiptr = 10;

    cout << "The value of a is " << xobject.*ptiptr << endl;

    // call member function
    (xobject.*ptfptr) (20);
}

```

The output for this example is:

```

The value of a is 10
The value of b is 20

```

To reduce complex syntax, you can declare a typedef to be a pointer to a member. A pointer to a member can be declared and used as shown in the following code fragment:

```

typedef int X::*my_pointer_to_member;
typedef void (X::*my_pointer_to_function) (int);

int main() {
    my_pointer_to_member ptiptr = &X::a;
    my_pointer_to_function ptfptr = &X::f;
    X xobject;
    xobject.*ptiptr = 10;
    cout << "The value of a is " << xobject.*ptiptr << endl;
    (xobject.*ptfptr) (20);
}

```

The pointer to member operators `.*` and `->*` are used to bind a pointer to a member of a specific class object. Because the precedence of `()` (function call operator) is higher than `.*` and `->*`, you must use parentheses to call the function pointed to by `ptf`.

Pointer-to-member conversion can occur when pointers to members are initialized, assigned, or compared. Note that pointer to a member is not the same as a pointer to an object or a pointer to a function.

Related information

- “Pointer to member operators `.*` `->*` (C++ only)” on page 153

The this pointer

The keyword `this` identifies a special type of pointer. Suppose that you create an object named `x` of class `A`, and class `A` has a nonstatic member function `f()`. If you call the function `x.f()`, the keyword `this` in the body of `f()` stores the address of `x`. You cannot declare the `this` pointer or make assignments to it.

A static member function does not have a `this` pointer.

The type of the `this` pointer for a member function of a class type `X`, is `X* const`. If the member function is declared with the `const` qualifier, the type of the `this` pointer for that member function for class `X`, is `const X* const`.

A `const this` pointer can be used only with `const` member functions. Data members of the class will be constant within that function. The function is still able to change the value, but requires a `const_cast` to do so:

```
void foo::p() const{
    member = 1;                // illegal
    const_cast <int&> (member) = 1; // a bad practice but legal
}
```

A better technique would be to declare member mutable.

If the member function is declared with the `volatile` qualifier, the type of the `this` pointer for that member function for class `X` is `volatile X* const`. For example, the compiler will not allow the following:

```
struct A {
    int a;
    int f() const { return a++; }
};
```

The compiler will not allow the statement `a++` in the body of function `f()`. In the function `f()`, the `this` pointer is of type `A* const`. The function `f()` is trying to modify part of the object to which `this` points.

The `this` pointer is passed as a hidden argument to all nonstatic member function calls and is available as a local variable within the body of all nonstatic functions.

For example, you can refer to the particular class object that a member function is called for by using the `this` pointer in the body of the member function. The following code example produces the output `a = 5`:

```
#include <iostream>
using namespace std;

struct X {
private:
    int a;
public:
    void Set_a(int a) {

        // The 'this' pointer is used to retrieve 'xobj.a'
        // hidden by the automatic variable 'a'
        this->a = a;
    }
    void Print_a() { cout << "a = " << a << endl; }
};

int main() {
    X xobj;
```

```
int a = 5;
xobj.Set_a(a);
xobj.Print_a();
}
```

In the member function `Set_a()`, the statement `this->a = a` uses the `this` pointer to retrieve `xobj.a` hidden by the automatic variable `a`.

Unless a class member name is hidden, using the class member name is equivalent to using the class member name with the `this` pointer and the class member access operator (`->`).

The example in the first column of the following table shows code that uses class members without the `this` pointer. The code in the second column uses the variable `THIS` to simulate the first column's hidden use of the `this` pointer:

Code without using this pointer	Equivalent code, the THIS variable simulating the hidden use of the this pointer
<pre> #include <string> #include <iostream> using namespace std; struct X { private: int len; char *ptr; public: int GetLen() { return len; } char * GetPtr() { return ptr; } X& Set(char *); X& Cat(char *); X& Copy(X&); void Print(); }; X& X::Set(char *pc) { len = strlen(pc); ptr = new char[len]; strcpy(ptr, pc); return *this; } X& X::Cat(char *pc) { len += strlen(pc); strcat(ptr,pc); return *this; } X& X::Copy(X& x) { Set(x.GetPtr()); return *this; } void X::Print() { cout << ptr << endl; } int main() { X xobj1; xobj1.Set("abcd") .Cat("efgh"); xobj1.Print(); X xobj2; xobj2.Copy(xobj1) .Cat("ijkl"); xobj2.Print(); } </pre>	<pre> #include <string> #include <iostream> using namespace std; struct X { private: int len; char *ptr; public: int GetLen (X* const THIS) { return THIS->len; } char * GetPtr (X* const THIS) { return THIS->ptr; } X& Set(X* const, char *); X& Cat(X* const, char *); X& Copy(X* const, X&); void Print(X* const); }; X& X::Set(X* const THIS, char *pc) { THIS->len = strlen(pc); THIS->ptr = new char[THIS->len]; strcpy(THIS->ptr, pc); return *THIS; } X& X::Cat(X* const THIS, char *pc) { THIS->len += strlen(pc); strcat(THIS->ptr, pc); return *THIS; } X& X::Copy(X* const THIS, X& x) { THIS->Set(THIS, x.GetPtr(&x)); return *THIS; } void X::Print(X* const THIS) { cout << THIS->ptr << endl; } int main() { X xobj1; xobj1.Set(&xobj1 , "abcd") .Cat(&xobj1 , "efgh"); xobj1.Print(&xobj1); X xobj2; xobj2.Copy(&xobj2 , xobj1) .Cat(&xobj2 , "ijkl"); xobj2.Print(&xobj2); } </pre>

Both examples produce the following output:

```

abcdefgh
abcdefghijkl

```

Related information

- “Overloading assignments” on page 251
- “Copy constructors” on page 333

Static members

Class members can be declared using the storage class specifier `static` in the class member list. Only one copy of the static member is shared by all objects of a class in a program. When you declare an object of a class having a static member, the static member is not part of the class object.

A typical use of static members is for recording data common to all objects of a class. For example, you can use a static data member as a counter to store the number of objects of a particular class type that are created. Each time a new object is created, this static data member can be incremented to keep track of the total number of objects.

You access a static member by qualifying the class name using the `::` (scope resolution) operator. In the following example, you can refer to the static member `f()` of class type `X` as `X::f()` even if no object of type `X` is ever declared:

```
struct X {
    static int f();
};

int main() {
    X::f();
}
```

Related information

- “Constant and volatile member functions” on page 272
- “The static storage class specifier” on page 46
- “Class member lists” on page 269

Using the class access operators with static members

You do not have to use the class member access syntax to refer to a static member; to access a static member `s` of class `X`, you could use the expression `X::s`. The following example demonstrates accessing a static member:

```
#include <iostream>
using namespace std;

struct A {
    static void f() { cout << "In static function A::f()" << endl; }
};

int main() {

    // no object required for static member
    A::f();

    A a;
    A* ap = &a;
    a.f();
    ap->f();
}
```

The three statements `A::f()`, `a.f()`, and `ap->f()` all call the same static member function `A::f()`.

You can directly refer to a static member in the same scope of its class, or in the scope of a class derived from the static member's class. The following example demonstrates the latter case (directly referring to a static member in the scope of a class derived from the static member's class):

```
#include <iostream>
using namespace std;

int g() {
    cout << "In function g()" << endl;
    return 0;
}

class X {
public:
    static int g() {
        cout << "In static member function X::g()" << endl;
        return 1;
    }
};

class Y: public X {
public:
    static int i;
};

int Y::i = g();

int main() { }
```

The following is the output of the above code:

```
In static member function X::g()
```

The initialization `int Y::i = g()` calls `X::g()`, not the function `g()` declared in the global namespace.

Related information

- “The static storage class specifier” on page 46
- “Scope resolution operator `::` (C++ only)” on page 126
- “Dot operator `.`” on page 128
- “Arrow operator `->`” on page 128

Static data members

The declaration of a static data member in the member list of a class is not a definition. You must define the static member outside of the class declaration, in namespace scope. For example:

```
class X
{
public:
    static int i;
};
int X::i = 0; // definition outside class declaration
```

Once you define a static data member, it exists even though no objects of the static data member's class exist. In the above example, no objects of class `X` exist even though the static data member `X::i` has been defined.

Static data members of a class in namespace scope have external linkage. The initializer for a static data member is in the scope of the class declaring the member.

A static data member can be of any type except for `void` or `void` qualified with `const` or `volatile`. You cannot declare a static data member as `mutable`.

You can only have one definition of a static member in a program. Unnamed classes, classes contained within unnamed classes, and local classes cannot have static data members.

Static data members and their initializers can access other static private and protected members of their class. The following example shows how you can initialize static members using other static members, even though these members are private:

```
class C {
    static int i;
    static int j;
    static int k;
    static int l;
    static int m;
    static int n;
    static int p;
    static int q;
    static int r;
    static int s;
    static int f() { return 0; }
    int a;
public:
    C() { a = 0; }
};

C c;
int C::i = C::f();    // initialize with static member function
int C::j = C::i;     // initialize with another static data member
int C::k = c.f();    // initialize with member function from an object
int C::l = c.j;     // initialize with data member from an object
int C::s = c.a;     // initialize with nonstatic data member
int C::r = 1;       // initialize with a constant value

class Y : private C { } y;

int C::m = Y::f();
int C::n = Y::r;
int C::p = y.r;     // error
int C::q = y.f();   // error
```

The initializations of `C::p` and `C::q` cause errors because `y` is an object of a class that is derived privately from `C`, and its members are not accessible to members of `C`.

If a static data member is of `const` integral or `const` enumeration type, you may specify a *constant initializer* in the static data member's declaration. This constant initializer must be an integral constant expression. Note that the constant initializer is not a definition. You still need to define the static member in an enclosing namespace. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct X {
    static const int a = 76;
};

const int X::a;
```

```
int main() {
    cout << X::a << endl;
}
```

The tokens = 76 at the end of the declaration of static data member a is a constant initializer.

Related information

- “External linkage” on page 8
- “Member access” on page 283
- “Local classes” on page 266

Static member functions

You cannot have static and nonstatic member functions with the same names and the same number and type of arguments.

Like static data members, you may access a static member function `f()` of a class `A` without using an object of class `A`.

A static member function does not have a `this` pointer. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct X {
private:
    int i;
    static int si;
public:
    void set_i(int arg) { i = arg; }
    static void set_si(int arg) { si = arg; }

    void print_i() {
        cout << "Value of i = " << i << endl;
        cout << "Again, value of i = " << this->i << endl;
    }

    static void print_si() {
        cout << "Value of si = " << si << endl;
        // cout << "Again, value of si = " << this->si << endl;
    }
};

int X::si = 77;    // Initialize static data member

int main() {
    X xobj;
    xobj.set_i(11);
    xobj.print_i();

    // static data members and functions belong to the class and
    // can be accessed without using an object of class X
    X::print_si();
    X::set_si(22);
    X::print_si();
}
```

The following is the output of the above example:

```
Value of i = 11
Again, value of i = 11
Value of si = 77
Value of si = 22
```

The compiler does not allow the member access operation `this->si` in function `A::print_si()` because this member function has been declared as static, and therefore does not have a `this` pointer.

You can call a static member function using the `this` pointer of a nonstatic member function. In the following example, the nonstatic member function `printall()` calls the static member function `f()` using the `this` pointer:

```
#include <iostream>
using namespace std;

class C {
    static void f() {
        cout << "Here is i: " << i << endl;
    }
    static int i;
    int j;
public:
    C(int firstj): j(firstj) { }
    void printall();
};

void C::printall() {
    cout << "Here is j: " << this->j << endl;
    this->f();
}

int C::i = 3;

int main() {
    C obj_C(0);
    obj_C.printall();
}
```

The following is the output of the above example:

```
Here is j: 0
Here is i: 3
```

A static member function cannot be declared with the keywords `virtual`, `const`, `volatile`, or `const volatile`.

A static member function can access only the names of static members, enumerators, and nested types of the class in which it is declared. Suppose a static member function `f()` is a member of class `X`. The static member function `f()` cannot access the nonstatic members `X` or the nonstatic members of a base class of `X`.

Related information

- “The `this` pointer” on page 276

Member access

Member access determines if a class member is accessible in an expression or declaration. Suppose `x` is a member of class `A`. Class member `x` can be declared to have one of the following levels of accessibility:

- `public`: `x` can be used anywhere without the access restrictions defined by `private` or `protected`.
- `private`: `x` can be used only by the members and friends of class `A`.
- `protected`: `x` can be used only by the members and friends of class `A`, and the members and friends of classes derived from class `A`.

Members of classes declared with the keyword `class` are `private` by default. Members of classes declared with the keyword `struct` or `union` are `public` by default.

To control the access of a class member, you use one of the *access specifiers* `public`, `private`, or `protected` as a label in a class member list. The following example demonstrates these access specifiers:

```
struct A {
    friend class C;
private:
    int a;
public:
    int b;
protected:
    int c;
};

struct B : A {
    void f() {
        // a = 1;
        b = 2;
        c = 3;
    }
};

struct C {
    void f(A x) {
        x.a = 4;
        x.b = 5;
        x.c = 6;
    }
};

int main() {
    A y;
    // y.a = 7;
    y.b = 8;
    // y.c = 9;

    B z;
    // z.a = 10;
    z.b = 11;
    // z.c = 12;
}
```

The following table lists the access of data members `A::a`, `A::b`, and `A::c` in various scopes of the above example.

Scope	<code>A::a</code>	<code>A::b</code>	<code>A::c</code>
function <code>B::f()</code>	No access. Member <code>A::a</code> is private.	Access. Member <code>A::b</code> is public.	Access. Class <code>B</code> inherits from <code>A</code> .
function <code>C::f()</code>	Access. Class <code>C</code> is a friend of <code>A</code> .	Access. Member <code>A::b</code> is public.	Access. Class <code>C</code> is a friend of <code>A</code> .

Scope	A::a	A::b	A::c
object y in main()	No access. Member y.a is private.	Access. Member y.a is public.	No access. Member y.c is protected.
object z in main()	No access. Member z.a is private.	Access. Member z.a is public.	No access. Member z.c is protected.

An access specifier specifies the accessibility of members that follow it until the next access specifier or until the end of the class definition. You can use any number of access specifiers in any order. If you later define a class member within its class definition, its access specification must be the same as its declaration. The following example demonstrates this:

```
class A {
    class B;
    public:
        class B { };
};
```

The compiler will not allow the definition of class B because this class has already been declared as private.

A class member has the same access control regardless whether it has been defined within its class or outside its class.

Access control applies to names. In particular, if you add access control to a typedef name, it affects only the typedef name. The following example demonstrates this:

```
class A {
    class B { };
    public:
        typedef B C;
};

int main() {
    A::C x;
    // A::B y;
}
```

The compiler will allow the declaration `A::C x` because the typedef name `A::C` is public. The compiler would not allow the declaration `A::B y` because `A::B` is private.

Note that accessibility and visibility are independent. Visibility is based on the scoping rules of C++. A class member can be visible and inaccessible at the same time.

Related information

- “Scope” on page 2
- “Class member lists” on page 269
- “Inherited member access” on page 296

Friends

A friend of a class X is a function or class that is not a member of X , but is granted the same access to X as the members of X . Functions declared with the friend specifier in a class member list are called *friend functions* of that class. Classes declared with the friend specifier in the member list of another class are called *friend classes* of that class.

A class Y must be defined before any member of Y can be declared a friend of another class.

In the following example, the friend function `print` is a member of class Y and accesses the private data members `a` and `b` of class X .

```
#include <iostream>
using namespace std;

class X;

class Y {
public:
    void print(X& x);
};

class X {
    int a, b;
    friend void Y::print(X& x);
public:
    X() : a(1), b(2) { }
};

void Y::print(X& x) {
    cout << "a is " << x.a << endl;
    cout << "b is " << x.b << endl;
}

int main() {
    X xobj;
    Y yobj;
    yobj.print(xobj);
}
```

The following is the output of the above example:

```
a is 1
b is 2
```

You can declare an entire class as a friend. Suppose class F is a friend of class A . This means that every member function and static data member definition of class F has access to class A .

In the following example, the friend class F has a member function `print` that accesses the private data members `a` and `b` of class X and performs the same task as the friend function `print` in the above example. Any other members declared in class F also have access to all members of class X .

```
#include <iostream>
using namespace std;

class X {
    int a, b;
    friend class F;
public:
    X() : a(1), b(2) { }
```

```

};

class F {
public:
    void print(X& x) {
        cout << "a is " << x.a << endl;
        cout << "b is " << x.b << endl;
    }
};

int main() {
    X xobj;
    F fobj;
    fobj.print(xobj);
}

```

The following is the output of the above example:

```

a is 1
b is 2

```

You must use an elaborated type specifier when you declare a class as a friend. The following example demonstrates this:

```

class F;
class G;
class X {
    friend class F;
    friend G;
};

```

The compiler will warn you that the friend declaration of G must be an elaborated class name.

You cannot define a class in a friend declaration. For example, the compiler will not allow the following:

```

class F;
class X {
    friend class F { };
};

```

However, you can define a function in a friend declaration. The class must be a non-local class, function, the function name must be unqualified, and the function has namespace scope. The following example demonstrates this:

```

class A {
    void g();
};

void z() {
    class B {
    // friend void f() { };
    };
}

class C {
// friend void A::g() { }
    friend void h() { }
};

```

The compiler would not allow the function definition of f() or g(). The compiler will allow the definition of h().

You cannot declare a friend with a storage class specifier.

Related information

- “Member access” on page 283
- “Inherited member access” on page 296

Friend scope

The name of a friend function or class first introduced in a friend declaration is not in the scope of the class granting friendship (also called the *enclosing class*) and is not a member of the class granting friendship.

The name of a function first introduced in a friend declaration is in the scope of the first nonclass scope that contains the enclosing class. The body of a function provided in a friend declaration is handled in the same way as a member function defined within a class. Processing of the definition does not start until the end of the outermost enclosing class. In addition, unqualified names in the body of the function definition are searched for starting from the class containing the function definition.

If the name of a friend class has been introduced before the friend declaration, the compiler searches for a class name that matches the name of the friend class beginning at the scope of the friend declaration. If the declaration of a nested class is followed by the declaration of a friend class with the same name, the nested class is a friend of the enclosing class.

The scope of a friend class name is the first nonclass enclosing scope. For example:

```
class A {
    class B { // arbitrary nested class definitions
        friend class C;
    };
};
```

is equivalent to:

```
class C;
class A {
    class B { // arbitrary nested class definitions
        friend class C;
    };
};
```

If the friend function is a member of another class, you need to use the scope resolution operator (`::`). For example:

```
class A {
public:
    int f() { }
};

class B {
    friend int A::f();
};
```

Friends of a base class are not inherited by any classes derived from that base class. The following example demonstrates this:

```
class A {
    friend class B;
    int a;
};

class B { };
```

```

class C : public B {
    void f(A* p) {
        //    p->a = 2;
    }
};

```

The compiler would not allow the statement `p->a = 2` because class `C` is not a friend of class `A`, although `C` inherits from a friend of `A`.

Friendship is not transitive. The following example demonstrates this:

```

class A {
    friend class B;
    int a;
};

class B {
    friend class C;
};

class C {
    void f(A* p) {
        //    p->a = 2;
    }
};

```

The compiler would not allow the statement `p->a = 2` because class `C` is not a friend of class `A`, although `C` is a friend of a friend of `A`.

If you declare a friend in a local class, and the friend's name is unqualified, the compiler will look for the name only within the innermost enclosing nonclass scope. You must declare a function before declaring it as a friend of a local scope. You do not have to do so with classes. However, a declaration of a friend class will hide a class in an enclosing scope with the same name. The following example demonstrates this:

```

class X { };
void a();

void f() {
    class Y { };
    void b();
    class A {
        friend class X;
        friend class Y;
        friend class Z;
        //    friend void a();
        friend void b();
        //    friend void c();
    };
    ::X moocow;
    // X moocow2;
}

```

In the above example, the compiler will allow the following statements:

- `friend class X`: This statement does not declare `::X` as a friend of `A`, but the local class `X` as a friend, even though this class is not otherwise declared.
- `friend class Y`: Local class `Y` has been declared in the scope of `f()`.
- `friend class Z`: This statement declares the local class `Z` as a friend of `A` even though `Z` is not otherwise declared.
- `friend void b()`: Function `b()` has been declared in the scope of `f()`.
- `::X moocow`: This declaration creates an object of the nonlocal class `::X`.

The compiler would not allow the following statements:

- `friend void a();` This statement does not consider function `a()` declared in namespace scope. Since function `a()` has not been declared in the scope of `f()`, the compiler would not allow this statement.
- `friend void c();` Since function `c()` has not been declared in the scope of `f()`, the compiler would not allow this statement.
- `X moocow2;` This declaration tries to create an object of the local class `X`, not the nonlocal class `::X`. Since local class `X` has not been defined, the compiler would not allow this statement.

Related information

- “Scope of class names” on page 263
- “Nested classes” on page 264
- “Local classes” on page 266

Friend access

A friend of a class can access the private and protected members of that class. Normally, you can only access the private members of a class through member functions of that class, and you can only access the protected members of a class through member functions of a class or classes derived from that class.

Friend declarations are not affected by access specifiers.

Related information

- “Member access” on page 283

Chapter 13. Inheritance (C++ only)

Inheritance is a mechanism of reusing and extending existing classes without modifying them, thus producing hierarchical relationships between them.

Inheritance is almost like embedding an object into a class. Suppose that you declare an object *x* of class *A* in the class definition of *B*. As a result, class *B* will have access to all the public data members and member functions of class *A*. However, in class *B*, you have to access the data members and member functions of class *A* through object *x*. The following example demonstrates this:

```
#include <iostream>
using namespace std;

class A {
    int data;
public:
    void f(int arg) { data = arg; }
    int g() { return data; }
};

class B {
public:
    A x;
};

int main() {
    B obj;
    obj.x.f(20);
    cout << obj.x.g() << endl;
    // cout << obj.g() << endl;
}
```

In the main function, object *obj* accesses function *A::f()* through its data member *B::x* with the statement *obj.x.f(20)*. Object *obj* accesses *A::g()* in a similar manner with the statement *obj.x.g()*. The compiler would not allow the statement *obj.g()* because *g()* is a member function of class *A*, not class *B*.

The inheritance mechanism lets you use a statement like *obj.g()* in the above example. In order for that statement to be legal, *g()* must be a member function of class *B*.

Inheritance lets you include the names and definitions of another class's members as part of a new class. The class whose members you want to include in your new class is called a *base class*. Your new class is *derived* from the base class. The new class contains a *subobject* of the type of the base class. The following example is the same as the previous example except it uses the inheritance mechanism to give class *B* access to the members of class *A*:

```
#include <iostream>
using namespace std;

class A {
    int data;
public:
    void f(int arg) { data = arg; }
    int g() { return data; }
};
```

```

class B : public A { };

int main() {
    B obj;
    obj.f(20);
    cout << obj.g() << endl;
}

```

Class A is a base class of class B. The names and definitions of the members of class A are included in the definition of class B; class B inherits the members of class A. Class B is derived from class A. Class B contains a subobject of type A.

You can also add new data members and member functions to the derived class. You can modify the implementation of existing member functions or data by overriding base class member functions or data in the newly derived class.

You may derive classes from other derived classes, thereby creating another level of inheritance. The following example demonstrates this:

```

struct A { };
struct B : A { };
struct C : B { };

```

Class B is a derived class of A, but is also a base class of C. The number of levels of inheritance is only limited by resources.

Multiple inheritance allows you to create a derived class that inherits properties from more than one base class. Because a derived class inherits members from all its base classes, ambiguities can result. For example, if two base classes have a member with the same name, the derived class cannot implicitly differentiate between the two members. Note that, when you are using multiple inheritance, the access to names of base classes may be ambiguous. See “Multiple inheritance” on page 302 for more detailed information.

A *direct base class* is a base class that appears directly as a base specifier in the declaration of its derived class.

An *indirect base class* is a base class that does not appear directly in the declaration of the derived class but is available to the derived class through one of its base classes. For a given class, all base classes that are not direct base classes are indirect base classes. The following example demonstrates direct and indirect base classes:

```

class A {
public:
    int x;
};
class B : public A {
public:
    int y;
};
class C : public B { };

```

Class B is a direct base class of C. Class A is a direct base class of B. Class A is an indirect base class of C. (Class C has x and y as its data members.)

Polymorphic functions are functions that can be applied to objects of more than one type. In C++, polymorphic functions are implemented in two ways:


- Overloaded functions are statically bound at compile time.

- C++ provides virtual functions. A *virtual function* is a function that can be called for a number of different user-defined types that are related through derivation. Virtual functions are bound dynamically at run time. They are described in more detail in “Virtual functions” on page 309.

Derivation

Inheritance is implemented in C++ through the mechanism of derivation. Derivation allows you to derive a class, called a *derived class*, from another class, called a *base class*.

Derived class syntax



In the declaration of a derived class, you list the base classes of the derived class. The derived class inherits its members from these base classes.

The *qualified_class_specifier* must be a class that has been previously declared in a class declaration.

An *access specifier* is one of `public`, `private`, or `protected`.

The `virtual` keyword can be used to declare virtual base classes.

The following example shows the declaration of the derived class `D` and the base classes `V`, `B1`, and `B2`. The class `B1` is both a base class and a derived class because it is derived from class `V` and is a base class for `D`:

```
class V { /* ... */ };
class B1 : virtual public V { /* ... */ };
class B2 { /* ... */ };
class D : public B1, private B2 { /* ... */ };
```

Classes that are declared but not defined are not allowed in base lists.

For example:

```
class X;

// error
class Y: public X { };
```

The compiler will not allow the declaration of class `Y` because `X` has not been defined.

When you derive a class, the derived class inherits class members of the base class. You can refer to inherited members (base class members) as if they were members of the derived class. For example:

```
class Base {
public:
    int a,b;
};

class Derived : public Base {
public:
    int c;
};

int main() {
    Derived d;
    d.a = 1;    // Base::a
    d.b = 2;    // Base::b
    d.c = 3;    // Derived::c
}
```

The derived class can also add new class members and redefine existing base class members. In the above example, the two inherited members, a and b, of the derived class d, in addition to the derived class member c, are assigned values. If you redefine base class members in the derived class, you can still refer to the base class members by using the :: (scope resolution) operator. For example:

```
#include <iostream>
using namespace std;

class Base {
public:
    char* name;
    void display() {
        cout << name << endl;
    }
};

class Derived: public Base {
public:
    char* name;
    void display() {
        cout << name << ", " << Base::name << endl;
    }
};

int main() {
    Derived d;
    d.name = "Derived Class";
    d.Base::name = "Base Class";

    // call Derived::display()
    d.display();

    // call Base::display()
    d.Base::display();
}
```

The following is the output of the above example:

```
Derived Class, Base Class
Base Class
```

You can manipulate a derived class object as if it were a base class object. You can use a pointer or a reference to a derived class object in place of a pointer or reference to its base class. For example, you can pass a pointer or reference to a

derived class object D to a function expecting a pointer or reference to the base class of D. You do not need to use an explicit cast to achieve this; a standard conversion is performed. You can implicitly convert a pointer to a derived class to point to an accessible unambiguous base class. You can also implicitly convert a reference to a derived class to a reference to a base class.

The following example demonstrates a standard conversion from a pointer to a derived class to a pointer to a base class:

```
#include <iostream>
using namespace std;

class Base {
public:
    char* name;
    void display() {
        cout << name << endl;
    }
};

class Derived: public Base {
public:
    char* name;
    void display() {
        cout << name << ", " << Base::name << endl;
    }
};

int main() {
    Derived d;
    d.name = "Derived Class";
    d.Base::name = "Base Class";

    Derived* dptr = &d;

    // standard conversion from Derived* to Base*
    Base* bptr = dptr;

    // call Base::display()
    bptr->display();
}
```

The following is the output of the above example:

```
Base Class
```

The statement `Base* bptr = dptr` converts a pointer of type `Derived` to a pointer of type `Base`.

The reverse case is not allowed. You cannot implicitly convert a pointer or a reference to a base class object to a pointer or reference to a derived class. For example, the compiler will not allow the following code if the classes `Base` and `Class` are defined as in the above example:

```
int main() {
    Base b;
    b.name = "Base class";

    Derived* dptr = &b;
}
```

The compiler will not allow the statement `Derived* dptr = &b` because the statement is trying to implicitly convert a pointer of type `Base` to a pointer of type `Derived`.

If a member of a derived class and a member of a base class have the same name, the base class member is hidden in the derived class. If a member of a derived class has the same name as a base class, the base class name is hidden in the derived class.

Related information

- “Virtual base classes” on page 303
- “Incomplete class declarations” on page 264
- “Scope resolution operator :: (C++ only)” on page 126

Inherited member access

The following sections discuss the access rules affecting a protected nonstatic base class member and how to declare a derived class using an access specifier:

- “Protected members”
- “Access control of base class members” on page 297

Related information

- “Member access” on page 283

Protected members

A protected nonstatic base class member can be accessed by members and friends of any classes derived from that base class by using one of the following:

- A pointer to a directly or indirectly derived class
- A reference to a directly or indirectly derived class
- An object of a directly or indirectly derived class

If a class is derived privately from a base class, all protected base class members become private members of the derived class.

If you reference a protected nonstatic member *x* of a base class *A* in a friend or a member function of a derived class *B*, you must access *x* through a pointer to, reference to, or object of a class derived from *A*. However, if you are accessing *x* to create a pointer to member, you must qualify *x* with a nested name specifier that names the derived class *B*. The following example demonstrates this:

```
class A {
public:
protected:
    int i;
};

class B : public A {
    friend void f(A*, B*);
    void g(A*);
};

void f(A* pa, B* pb) {
    // pa->i = 1;
    pb->i = 2;

    // int A::* point_i = &A::i;
    int A::* point_i2 = &B::i;
}

void B::g(A* pa) {
    // pa->i = 1;
```

```

    i = 2;

    // int A::* point_i = &A::i;
    int A::* point_i2 = &B::i;
}

void h(A* pa, B* pb) {
    // pa->i = 1;
    // pb->i = 2;
}

int main() { }

```

Class A contains one protected data member, an integer `i`. Because B derives from A, the members of B have access to the protected member of A. Function `f()` is a friend of class B:

- The compiler would not allow `pa->i = 1` because `pa` is not a pointer to the derived class B.
- The compiler would not allow `int A::* point_i = &A::i` because `i` has not been qualified with the name of the derived class B.

Function `g()` is a member function of class B. The previous list of remarks about which statements the compiler would and would not allow apply for `g()` except for the following:

- The compiler allows `i = 2` because it is equivalent to `this->i = 2`.

Function `h()` cannot access any of the protected members of A because `h()` is neither a friend or a member of a derived class of A.

Access control of base class members

When you declare a derived class, an access specifier can precede each base class in the base list of the derived class. This does not alter the access attributes of the individual members of a base class as seen by the base class, but allows the derived class to restrict the access control of the members of a base class.

You can derive classes using any of the three access specifiers:

- In a `public` base class, `public` and `protected` members of the base class remain `public` and `protected` members of the derived class.
- In a `protected` base class, `public` and `protected` members of the base class are `protected` members of the derived class.
- In a `private` base class, `public` and `protected` members of the base class become `private` members of the derived class.

In all cases, `private` members of the base class remain `private`. `Private` members of the base class cannot be used by the derived class unless friend declarations within the base class explicitly grant access to them.

In the following example, class `d` is derived publicly from class `b`. Class `b` is declared a `public` base class by this declaration.

```

class b { };
class d : public b // public derivation
{ };

```

You can use both a structure and a class as base classes in the base list of a derived class declaration:

- If the derived class is declared with the keyword `class`, the default access specifier in its base list specifiers is `private`.

- If the derived class is declared with the keyword `struct`, the default access specifier in its base list specifiers is `public`.

In the following example, private derivation is used by default because no access specifier is used in the base list and the derived class is declared with the keyword `class`:

```
struct B
{ };
class D : B // private derivation
{ };
```

Members and friends of a class can implicitly convert a pointer to an object of that class to a pointer to either:

- A direct private base class
- A protected base class (either direct or indirect)

Related information

- “Member access” on page 283
- “Member scope” on page 273

The using declaration and class members

A using declaration in a definition of a class *A* allows you to introduce a *name* of a data member or member function from a base class of *A* into the scope of *A*.

You would need a using declaration in a class definition if you want to create a set of overload a member functions from base and derived classes, or you want to change the access of a class member.

using declaration syntax

```
using _typename _::_ nested_name_specifier unqualified_id ;
_::_ unqualified_id ;
```

A using declaration in a class *A* may name one of the following:

- A member of a base class of *A*
- A member of an anonymous union that is a member of a base class of *A*
- An enumerator for an enumeration type that is a member of a base class of *A*

The following example demonstrates this:

```
struct Z {
    int g();
};

struct A {
    void f();
    enum E { e };
    union { int u; };
};

struct B : A {
    using A::f;
    using A::e;
    using A::u;
    // using Z::g;
};
```

The compiler would not allow the using declaration `using Z::g` because `Z` is not a base class of `A`.

A using declaration cannot name a template. For example, the compiler will not allow the following:

```
struct A {
    template<class T> void f(T);
};

struct B : A {
    using A::f<int>;
};
```

Every instance of the name mentioned in a using declaration must be accessible. The following example demonstrates this:

```
struct A {
private:
    void f(int);
public:
    int f();
protected:
    void g();
};

struct B : A {
    // using A::f;
    using A::g;
};
```

The compiler would not allow the using declaration `using A::f` because `void A::f(int)` is not accessible from `B` even though `int A::f()` is accessible.

Related information

- “Scope of class names” on page 263
- “The using declaration and namespaces” on page 240

Overloading member functions from base and derived classes

A member function named `f` in a class `A` will hide all other members named `f` in the base classes of `A`, regardless of return types or arguments. The following example demonstrates this:

```
struct A {
    void f() { }
};

struct B : A {
    void f(int) { }
};

int main() {
    B obj_B;
    obj_B.f(3);
    // obj_B.f();
}
```

The compiler would not allow the function call `obj_B.f()` because the declaration of `void B::f(int)` has hidden `A::f()`.

To overload, rather than hide, a function of a base class A in a derived class B, you introduce the name of the function into the scope of B with a using declaration. The following example is the same as the previous example except for the using declaration using A::f:

```
struct A {
    void f() { }
};

struct B : A {
    using A::f;
    void f(int) { }
};

int main() {
    B obj_B;
    obj_B.f(3);
    obj_B.f();
}
```

Because of the using declaration in class B, the name f is overloaded with two functions. The compiler will now allow the function call obj_B.f().

You can overload virtual functions in the same way. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct A {
    virtual void f() { cout << "void A::f()" << endl; }
    virtual void f(int) { cout << "void A::f(int)" << endl; }
};

struct B : A {
    using A::f;
    void f(int) { cout << "void B::f(int)" << endl; }
};

int main() {
    B obj_B;
    B* pb = &obj_B;
    pb->f(3);
    pb->f();
}
```

The following is the output of the above example:

```
void B::f(int)
void A::f()
```

Suppose that you introduce a function f from a base class A a derived class B with a using declaration, and there exists a function named B::f that has the same parameter types as A::f. Function B::f will hide, rather than conflict with, function A::f. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct A {
    void f() { }
    void f(int) { cout << "void A::f(int)" << endl; }
};

struct B : A {
    using A::f;
}
```

```

    void f(int) { cout << "void B::f(int)" << endl; }
};

int main() {
    B obj_B;
    obj_B.f(3);
}

```

The following is the output of the above example:
void B::f(int)

Related information

- Chapter 10, “Overloading (C++ only),” on page 243
- “Name hiding (C++ only)” on page 7
- “The using declaration and class members” on page 298

Changing the access of a class member

Suppose class B is a direct base class of class A. To restrict access of class B to the members of class A, derive B from A using either the access specifiers `protected` or `private`.

To increase the access of a member `x` of class A inherited from class B, use a `using` declaration. You cannot restrict the access to `x` with a `using` declaration. You may increase the access of the following members:

- A member inherited as `private`. (You cannot increase the access of a member declared as `private` because a `using` declaration must have access to the member’s name.)
- A member either inherited or declared as `protected`

The following example demonstrates this:

```

struct A {
protected:
    int y;
public:
    int z;
};

struct B : private A { };

struct C : private A {
public:
    using A::y;
    using A::z;
};

struct D : private A {
protected:
    using A::y;
    using A::z;
};

struct E : D {
    void f() {
        y = 1;
        z = 2;
    }
};

struct F : A {
public:
    using A::y;
};

```

```

private:
    using A::z;
};

int main() {
    B obj_B;
    // obj_B.y = 3;
    // obj_B.z = 4;

    C obj_C;
    obj_C.y = 5;
    obj_C.z = 6;

    D obj_D;
    // obj_D.y = 7;
    // obj_D.z = 8;

    F obj_F;
    obj_F.y = 9;
    obj_F.z = 10;
}

```

The compiler would not allow the following assignments from the above example:

- `obj_B.y = 3` and `obj_B.z = 4`: Members `y` and `z` have been inherited as private.
- `obj_D.y = 7` and `obj_D.z = 8`: Members `y` and `z` have been inherited as private, but their access have been changed to protected.

The compiler allows the following statements from the above example:

- `y = 1` and `z = 2` in `D::f()`: Members `y` and `z` have been inherited as private, but their access have been changed to protected.
- `obj_C.y = 5` and `obj_C.z = 6`: Members `y` and `z` have been inherited as private, but their access have been changed to public.
- `obj_F.y = 9`: The access of member `y` has been changed from protected to public.
- `obj_F.z = 10`: The access of member `z` is still public. The private using declaration using `A::z` has no effect on the access of `z`.

Related information

- “Member access” on page 283
- “Inherited member access” on page 296

Multiple inheritance

You can derive a class from any number of base classes. Deriving a class from more than one direct base class is called *multiple inheritance*.


In the following example, classes `A`, `B`, and `C` are direct base classes for the derived class `X`:

```

class A { /* ... */ };
class B { /* ... */ };
class C { /* ... */ };
class X : public A, private B, public C { /* ... */ };

```

The following *inheritance graph* describes the inheritance relationships of the above example. An arrow points to the direct base class of the class at the tail of the arrow:




The order of derivation is relevant only to determine the order of default initialization by constructors and cleanup by destructors.

A direct base class cannot appear in the base list of a derived class more than once:

```
class B1 { /* ... */ }; // direct base class
class D : public B1, private B1 { /* ... */ }; // error
```

However, a derived class can inherit an indirect base class more than once, as shown in the following example:



```
class L { /* ... */ }; // indirect base class
class B2 : public L { /* ... */ };
class B3 : public L { /* ... */ };
class D : public B2, public B3 { /* ... */ }; // valid
```

In the above example, class D inherits the indirect base class L once through class B2 and once through class B3. However, this may lead to ambiguities because two subobjects of class L exist, and both are accessible through class D. You can avoid this ambiguity by referring to class L using a qualified class name. For example:

B2::L

or

B3::L.


You can also avoid this ambiguity by using the base specifier `virtual` to declare a base class, as described in “Derivation” on page 293.

Virtual base classes

Suppose you have two derived classes B and C that have a common base class A, and you also have another class D that inherits from B and C. You can declare the base class A as *virtual* to ensure that B and C share the same subobject of A.

In the following example, an object of class D has two distinct subobjects of class L, one through class B1 and another through class B2. You can use the keyword `virtual` in front of the base class specifiers in the *base lists* of classes B1 and B2 to indicate that only one subobject of type L, shared by class B1 and class B2, exists.


For example:



```
class L { /* ... */ }; // indirect base class
class B1 : virtual public L { /* ... */ };
class B2 : virtual public L { /* ... */ };
class D : public B1, public B2 { /* ... */ }; // valid
```

Using the keyword `virtual` in this example ensures that an object of class `D` inherits only one subobject of class `L`.

A derived class can have both virtual and nonvirtual base classes. For example:



```
class V { /* ... */ };
class B1 : virtual public V { /* ... */ };
class B2 : virtual public V { /* ... */ };
class B3 : public V { /* ... */ };
class X : public B1, public B2, public B3 { /* ... */ };
```

In the above example, class `X` has two subobjects of class `V`, one that is shared by classes `B1` and `B2` and one through class `B3`.

Related information

- “Derivation” on page 293

Multiple access

In an inheritance graph containing virtual base classes, a name that can be reached through more than one path is accessed through the path that gives the most access.

For example:

```
class L {
public:
    void f();
};
```

```

class B1 : private virtual L { };

class B2 : public virtual L { };

class D : public B1, public B2 {
public:
    void f() {
        // L::f() is accessed through B2
        // and is public
        L::f();
    }
};

```

In the above example, the function `f()` is accessed through class `B2`. Because class `B2` is inherited publicly and class `B1` is inherited privately, class `B2` offers more access.

Related information

- “Member access” on page 283
- “Protected members” on page 296
- “Access control of base class members” on page 297

Ambiguous base classes

When you derive classes, ambiguities can result if base and derived classes have members with the same names. Access to a base class member is ambiguous if you use a name or qualified name that does not refer to a unique function or object. The declaration of a member with an ambiguous name in a derived class is not an error. The ambiguity is only flagged as an error if you use the ambiguous member name.

For example, suppose that two classes named `A` and `B` both have a member named `x`, and a class named `C` inherits from both `A` and `B`. An attempt to access `x` from class `C` would be ambiguous. You can resolve ambiguity by qualifying a member with its class name using the scope resolution (`::`) operator.

```

class B1 {
public:
    int i;
    int j;
    void g(int) { }
};

class B2 {
public:
    int j;
    void g() { }
};

class D : public B1, public B2 {
public:
    int i;
};

int main() {
    D dobj;
    D *dptr = &dobj;
    dptr->i = 5;
    // dptr->j = 10;
    dptr->B1::j = 10;
    // dobj.g();
    dobj.B2::g();
}

```

The statement `dptr->j = 10` is ambiguous because the name `j` appears both in `B1` and `B2`. The statement `dobj.g()` is ambiguous because the name `g` appears both in `B1` and `B2`, even though `B1::g(int)` and `B2::g()` have different parameters.

The compiler checks for ambiguities at compile time. Because ambiguity checking occurs before access control or type checking, ambiguities may result even if only one of several members with the same name is accessible from the derived class.

Name hiding

Suppose two subobjects named `A` and `B` both have a member name `x`. The member name `x` of subobject `B` *hides* the member name `x` of subobject `A` if `A` is a base class of `B`. The following example demonstrates this:

```
struct A {
    int x;
};

struct B: A {
    int x;
};

struct C: A, B {
    void f() { x = 0; }
};

int main() {
    C i;
    i.f();
}
```

The assignment `x = 0` in function `C::f()` is not ambiguous because the declaration `B::x` has hidden `A::x`. However, the compiler will warn you that deriving `C` from `A` is redundant because you already have access to the subobject `A` through `B`.

A base class declaration can be hidden along one path in the inheritance graph and not hidden along another path. The following example demonstrates this:

```
struct A { int x; };
struct B { int y; };
struct C: A, virtual B { };
struct D: A, virtual B {
    int x;
    int y;
};
struct E: C, D { };

int main() {
    E e;
    // e.x = 1;
    e.y = 2;
}
```

The assignment `e.x = 1` is ambiguous. The declaration `D::x` hides `A::x` along the path `D::A::x`, but it does not hide `A::x` along the path `C::A::x`. Therefore the variable `x` could refer to either `D::x` or `A::x`. The assignment `e.y = 2` is not ambiguous. The declaration `D::y` hides `B::y` along both paths `D::B::y` and `C::B::y` because `B` is a virtual base class.

Ambiguity and using declarations

Suppose you have a class named `C` that inherits from a class named `A`, and `x` is a member name of `A`. If you use a using declaration to declare `A::x` in `C`, then `x` is

also a member of C; C::x does not hide A::x. Therefore using declarations cannot resolve ambiguities due to inherited members. The following example demonstrates this:

```
struct A {
    int x;
};

struct B: A { };

struct C: A {
    using A::x;
};

struct D: B, C {
    void f() { x = 0; }
};

int main() {
    D i;
    i.f();
}
```

The compiler will not allow the assignment `x = 0` in function `D::f()` because it is ambiguous. The compiler can find `x` in two ways: as `B::x` or as `C::x`.

Unambiguous class members

The compiler can unambiguously find static members, nested types, and enumerators defined in a base class A regardless of the number of subobjects of type A an object has. The following example demonstrates this:

```
struct A {
    int x;
    static int s;
    typedef A* Pointer_A;
    enum { e };
};

int A::s;

struct B: A { };

struct C: A { };

struct D: B, C {
    void f() {
        s = 1;
        Pointer_A pa;
        int i = e;
        // x = 1;
    }
};

int main() {
    D i;
    i.f();
}
```

The compiler allows the assignment `s = 1`, the declaration `Pointer_A pa`, and the statement `int i = e`. There is only one static variable `s`, only one typedef `Pointer_A`, and only one enumerator `e`. The compiler would not allow the assignment `x = 1` because `x` can be reached either from class B or class C.

Pointer conversions

Conversions (either implicit or explicit) from a derived class pointer or reference to a base class pointer or reference must refer unambiguously to the same accessible base class object. (An *accessible base class* is a publicly derived base class that is neither hidden nor ambiguous in the inheritance hierarchy.) For example:

```
class W { /* ... */ };
class X : public W { /* ... */ };
class Y : public W { /* ... */ };
class Z : public X, public Y { /* ... */ };
int main ()
{
    Z z;
    X* xptr = &z;           // valid
    Y* yptr = &z;           // valid
    W* wptr = &z;           // error, ambiguous reference to class W
                             // X's W or Y's W ?
}
```

You can use virtual base classes to avoid ambiguous reference. For example:

```
class W { /* ... */ };
class X : public virtual W { /* ... */ };
class Y : public virtual W { /* ... */ };
class Z : public X, public Y { /* ... */ };
int main ()
{
    Z z;
    X* xptr = &z;           // valid
    Y* yptr = &z;           // valid
    W* wptr = &z;           // valid, W is virtual therefore only one
                             // W subobject exists
}
```

A pointer to a member of a base class can be converted to a pointer to a member of a derived class if the following conditions are true:

- The conversion is not ambiguous. The conversion is ambiguous if multiple instances of the base class are in the derived class.
- A pointer to the derived class can be converted to a pointer to the base class. If this is the case, the base class is said to be *accessible*.
- Member types must match. For example suppose class A is a base class of class B. You cannot convert a pointer to member of A of type int to a pointer to member of type B of type float.
- The base class cannot be virtual.

Overload resolution

Overload resolution takes place *after* the compiler unambiguously finds a given function name. The following example demonstrates this:

```
struct A {
    int f() { return 1; }
};

struct B {
    int f(int arg) { return arg; }
};

struct C: A, B {
    int g() { return f(); }
};
```

The compiler will not allow the function call to `f()` in `C::g()` because the name `f` has been declared both in `A` and `B`. The compiler detects the ambiguity error before overload resolution can select the base match `A::f()`.

Related information

- “Scope resolution operator `::` (C++ only)” on page 126
- “Virtual base classes” on page 303

Virtual functions

By default, C++ matches a function call with the correct function definition at compile time. This is called *static binding*. You can specify that the compiler match a function call with the correct function definition at run time; this is called *dynamic binding*. You declare a function with the keyword `virtual` if you want the compiler to use dynamic binding for that specific function.

The following examples demonstrate the differences between static and dynamic binding. The first example demonstrates static binding:

```
#include <iostream>
using namespace std;

struct A {
    void f() { cout << "Class A" << endl; }
};

struct B: A {
    void f() { cout << "Class B" << endl; }
};

void g(A& arg) {
    arg.f();
}

int main() {
    B x;
    g(x);
}
```

The following is the output of the above example:

```
Class A
```

When function `g()` is called, function `A::f()` is called, although the argument refers to an object of type `B`. At compile time, the compiler knows only that the argument of function `g()` will be a reference to an object derived from `A`; it cannot determine whether the argument will be a reference to an object of type `A` or type `B`. However, this can be determined at run time. The following example is the same as the previous example, except that `A::f()` is declared with the `virtual` keyword:

```
#include <iostream>
using namespace std;

struct A {
    virtual void f() { cout << "Class A" << endl; }
};

struct B: A {
    void f() { cout << "Class B" << endl; }
};

void g(A& arg) {
    arg.f();
}
```

```

}

int main() {
    B x;
    g(x);
}

```

The following is the output of the above example:

```
Class B
```

The `virtual` keyword indicates to the compiler that it should choose the appropriate definition of `f()` not by the type of reference, but by the type of object that the reference refers to.

Therefore, a *virtual function* is a member function you may redefine for other derived classes, and can ensure that the compiler will call the redefined virtual function for an object of the corresponding derived class, even if you call that function with a pointer or reference to a base class of the object.

A class that declares or inherits a virtual function is called a *polymorphic class*.

You redefine a virtual member function, like any member function, in any derived class. Suppose you declare a virtual function named `f` in a class `A`, and you derive directly or indirectly from `A` a class named `B`. If you declare a function named `f` in class `B` with the same name and same parameter list as `A::f`, then `B::f` is also virtual (regardless whether or not you declare `B::f` with the `virtual` keyword) and it *overrides* `A::f`. However, if the parameter lists of `A::f` and `B::f` are different, `A::f` and `B::f` are considered different, `B::f` does not override `A::f`, and `B::f` is not virtual (unless you have declared it with the `virtual` keyword). Instead `B::f` *hides* `A::f`. The following example demonstrates this:

```

#include <iostream>
using namespace std;

struct A {
    virtual void f() { cout << "Class A" << endl; }
};

struct B: A {
    void f(int) { cout << "Class B" << endl; }
};

struct C: B {
    void f() { cout << "Class C" << endl; }
};

int main() {
    B b; C c;
    A* pa1 = &b;
    A* pa2 = &c;
    // b.f();
    pa1->f();
    pa2->f();
}

```

The following is the output of the above example:

```
Class A
Class C
```

The function `B::f` is not virtual. It hides `A::f`. Thus the compiler will not allow the function call `b.f()`. The function `C::f` is virtual; it overrides `A::f` even though `A::f` is not visible in `C`.

If you declare a base class destructor as virtual, a derived class destructor will override that base class destructor, even though destructors are not inherited.

The return type of an overriding virtual function may differ from the return type of the overridden virtual function. This overriding function would then be called a *covariant virtual function*. Suppose that `B::f` overrides the virtual function `A::f`. The return types of `A::f` and `B::f` may differ if all the following conditions are met:

- The function `B::f` returns a reference or pointer to a class of type `T`, and `A::f` returns a pointer or a reference to an unambiguous direct or indirect base class of `T`.
- The `const` or `volatile` qualification of the pointer or reference returned by `B::f` has the same or less `const` or `volatile` qualification of the pointer or reference returned by `A::f`.
- The return type of `B::f` must be complete at the point of declaration of `B::f`, or it can be of type `B`.

The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct A { };

class B : private A {
    friend class D;
    friend class F;
};

A global_A;
B global_B;

struct C {
    virtual A* f() {
        cout << "A* C::f()" << endl;
        return &global_A;
    }
};

struct D : C {
    B* f() {
        cout << "B* D::f()" << endl;
        return &global_B;
    }
};

struct E;

struct F : C {

    // Error:
    // E is incomplete
    // E* f();
};

struct G : C {

    // Error:
    // A is an inaccessible base class of B
    // B* f();
};
```

```

};

int main() {
    D d;
    C* cp = &d;
    D* dp = &d;

    A* ap = cp->f();
    B* bp = dp->f();
};

```

The following is the output of the above example:

```

B* D::f()
B* D::f()

```

The statement `A* ap = cp->f()` calls `D::f()` and converts the pointer returned to type `A*`. The statement `B* bp = dp->f()` calls `D::f()` as well but does not convert the pointer returned; the type returned is `B*`. The compiler would not allow the declaration of the virtual function `F::f()` because `E` is not a complete class. The compiler would not allow the declaration of the virtual function `G::f()` because class `A` is not an accessible base class of `B` (unlike friend classes `D` and `F`, the definition of `B` does not give access to its members for class `G`).

A virtual function cannot be global or static because, by definition, a virtual function is a member function of a base class and relies on a specific object to determine which implementation of the function is called. You can declare a virtual function to be a friend of another class.

If a function is declared virtual in its base class, you can still access it directly using the scope resolution (`::`) operator. In this case, the virtual function call mechanism is suppressed and the function implementation defined in the base class is used. In addition, if you do not override a virtual member function in a derived class, a call to that function uses the function implementation defined in the base class.

A virtual function must be one of the following:

- Defined
- Declared pure
- Defined and declared pure

A base class containing one or more pure virtual member functions is called an *abstract class*.

Related information

- “Abstract classes” on page 314

Ambiguous virtual function calls

You cannot override one virtual function with two or more ambiguous virtual functions. This can happen in a derived class that inherits from two nonvirtual bases that are derived from a virtual base class.

For example:

```

class V {
public:
    virtual void f() { }
};

```

```

class A : virtual public V {
    void f() { }
};

class B : virtual public V {
    void f() { }
};

// Error:
// Both A::f() and B::f() try to override V::f()
class D : public A, public B { };

int main() {
    D d;
    V* vptr = &d;

    // which f(), A::f() or B::f()?
    vptr->f();
}

```

The compiler will not allow the definition of class D. In class A, only A::f() will override V::f(). Similarly, in class B, only B::f() will override V::f(). However, in class D, both A::f() and B::f() will try to override V::f(). This attempt is not allowed because it is not possible to decide which function to call if a D object is referenced with a pointer to class V, as shown in the above example. Only one function can override a virtual function.

A special case occurs when the ambiguous overriding virtual functions come from separate instances of the same class type. In the following example, class D has two separate subobjects of class A:

```

#include <iostream>
using namespace std;

struct A {
    virtual void f() { cout << "A::f()" << endl; };
};

struct B : A {
    void f() { cout << "B::f()" << endl; };
};

struct C : A {
    void f() { cout << "C::f()" << endl; };
};

struct D : B, C { };

int main() {
    D d;

    B* bp = &d;
    A* ap = bp;
    D* dp = &d;

    ap->f();
    // dp->f();
}

```

Class D has two occurrences of class A, one inherited from B, and another inherited from C. Therefore there are also two occurrences of the virtual function A::f. The statement ap->f() calls D::B::f. However the compiler would not allow the statement dp->f() because it could either call D::B::f or D::C::f.

Virtual function access

The access for a virtual function is specified when it is declared. The access rules for a virtual function are not affected by the access rules for the function that later overrides the virtual function. In general, the access of the overriding member function is not known.

If a virtual function is called with a pointer or reference to a class object, the type of the class object is not used to determine the access of the virtual function. Instead, the type of the pointer or reference to the class object is used.

In the following example, when the function `f()` is called using a pointer having type `B*`, `bptr` is used to determine the access to the function `f()`. Although the definition of `f()` defined in class `D` is executed, the access of the member function `f()` in class `B` is used. When the function `f()` is called using a pointer having type `D*`, `dptr` is used to determine the access to the function `f()`. This call produces an error because `f()` is declared private in class `D`.

```
class B {
public:
    virtual void f();
};

class D : public B {
private:
    void f();
};

int main() {
    D dobj;
    B* bptr = &dobj;
    D* dptr = &dobj;

    // valid, virtual B::f() is public,
    // D::f() is called
    bptr->f();

    // error, D::f() is private
    dptr->f();
}
```

Abstract classes

An *abstract class* is a class that is designed to be specifically used as a base class. An abstract class contains at least one *pure virtual function*. You declare a pure virtual function by using a *pure specifier* (`= 0`) in the declaration of a virtual member function in the class declaration.

The following is an example of an abstract class:

```
class AB {
public:
    virtual void f() = 0;
};
```

Function `AB::f` is a pure virtual function. A function declaration cannot have both a pure specifier and a definition. For example, the compiler will not allow the following:

```
struct A {
    virtual void g() { } = 0;
};
```

You cannot use an abstract class as a parameter type, a function return type, or the type of an explicit conversion, nor can you declare an object of an abstract class. You can, however, declare pointers and references to an abstract class. The following example demonstrates this:

```
struct A {
    virtual void f() = 0;
};

struct B : A {
    virtual void f() { }
};

// Error:
// Class A is an abstract class
// A g();

// Error:
// Class A is an abstract class
// void h(A);
A& i(A&);

int main() {

// Error:
// Class A is an abstract class
// A a;

    A* pa;
    B b;

// Error:
// Class A is an abstract class
// static_cast<A>(b);
}
```

Class A is an abstract class. The compiler would not allow the function declarations `A g()` or `void h(A)`, declaration of object `a`, nor the static cast of `b` to type `A`.

Virtual member functions are inherited. A class derived from an abstract base class will also be abstract unless you override each pure virtual function in the derived class.

For example:

```
class AB {
public:
    virtual void f() = 0;
};

class D2 : public AB {
    void g();
};

int main() {
    D2 d;
}
```

The compiler will not allow the declaration of object `d` because `D2` is an abstract class; it inherited the pure virtual function `f()` from `AB`. The compiler will allow the declaration of object `d` if you define function `D2::g()`.

Note that you can derive an abstract class from a nonabstract class, and you can override a non-pure virtual function with a pure virtual function.

You can call member functions from a constructor or destructor of an abstract class. However, the results of calling (directly or indirectly) a pure virtual function from its constructor are undefined. The following example demonstrates this:

```
struct A {
    A() {
        direct();
        indirect();
    }
    virtual void direct() = 0;
    virtual void indirect() { direct(); }
};
```

The default constructor of A calls the pure virtual function `direct()` both directly and indirectly (through `indirect()`).

The compiler issues a warning for the direct call to the pure virtual function, but not for the indirect call.

Related information

- “Virtual functions” on page 309
- “Virtual function access” on page 314

Chapter 14. Special member functions (C++ only)

The default constructor, destructor, copy constructor, and copy assignment operator are *special member functions*. These functions create, destroy, convert, initialize, and copy class objects, and are discussed in the following sections:

- “Overview of constructors and destructors”
- “Constructors” on page 319
- “Destructors” on page 326
- “Conversion constructors” on page 330
- “Conversion functions” on page 332
- “Copy constructors” on page 333

Overview of constructors and destructors

Because classes have complicated internal structures, including data and functions, object initialization and cleanup for classes is much more complicated than it is for simple data structures. Constructors and destructors are special member functions of classes that are used to construct and destroy class objects. Construction may involve memory allocation and initialization for objects. Destruction may involve cleanup and deallocation of memory for objects.

Like other member functions, constructors and destructors are declared within a class declaration. They can be defined inline or external to the class declaration. Constructors can have default arguments. Unlike other member functions, constructors can have member initialization lists. The following restrictions apply to constructors and destructors:

- Constructors and destructors do not have return types nor can they return values.
- References and pointers cannot be used on constructors and destructors because their addresses cannot be taken.
- Constructors cannot be declared with the keyword `virtual`.
- Constructors and destructors cannot be declared `static`, `const`, or `volatile`.
- Unions cannot contain class objects that have constructors or destructors.

Constructors and destructors obey the same access rules as member functions. For example, if you declare a constructor with protected access, only derived classes and friends can use it to create class objects.

The compiler automatically calls constructors when defining class objects and calls destructors when class objects go out of scope. A constructor does not allocate memory for the class object its `this` pointer refers to, but may allocate storage for more objects than its class object refers to. If memory allocation is required for objects, constructors can explicitly call the `new` operator. During cleanup, a destructor may release objects allocated by the corresponding constructor. To release objects, use the `delete` operator.

Derived classes do not inherit or overload constructors or destructors from their base classes, but they do call the constructor and destructor of base classes. Destructors can be declared with the keyword `virtual`.

Constructors are also called when local or temporary class objects are created, and destructors are called when local or temporary objects go out of scope.

You can call member functions from constructors or destructors. You can call a virtual function, either directly or indirectly, from a constructor or destructor of a class A. In this case, the function called is the one defined in A or a base class of A, but not a function overridden in any class derived from A. This avoids the possibility of accessing an unconstructed object from a constructor or destructor. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct A {
    virtual void f() { cout << "void A::f()" << endl; }
    virtual void g() { cout << "void A::g()" << endl; }
    virtual void h() { cout << "void A::h()" << endl; }
};

struct B : A {
    virtual void f() { cout << "void B::f()" << endl; }
    B() {
        f();
        g();
        h();
    }
};

struct C : B {
    virtual void f() { cout << "void C::f()" << endl; }
    virtual void g() { cout << "void C::g()" << endl; }
    virtual void h() { cout << "void C::h()" << endl; }
};

int main() {
    C obj;
}
```

The following is the output of the above example:

```
void B::f()
void A::g()
void A::h()
```

The constructor of B does not call any of the functions overridden in C because C has been derived from B, although the example creates an object of type C named obj.

You can use the typeid or the dynamic_cast operator in constructors or destructors, as well as member initializers of constructors.

Related information

- “The constructor and destructor function attributes” on page 221
- “new expressions (C++ only)” on page 165
- “delete expressions (C++ only)” on page 169

Constructors

A *constructor* is a member function with the same name as its class. For example:

```
class X {
public:
    X();      // constructor for class X
};
```

Constructors are used to create, and can initialize, objects of their class type.

You cannot declare a constructor as `virtual` or `static`, nor can you declare a constructor as `const`, `volatile`, or `const volatile`.

You do not specify a return type for a constructor. A return statement in the body of a constructor cannot have a return value.

Related information

- “The constructor and destructor function attributes” on page 221

Default constructors

A *default constructor* is a constructor that either has no parameters, or if it has parameters, *all* the parameters have default values.

If no user-defined constructor exists for a class *A* and one is needed, the compiler implicitly *declares* a default parameterless constructor `A::A()`. This constructor is an inline public member of its class. The compiler will implicitly *define* `A::A()` when the compiler uses this constructor to create an object of type *A*. The constructor will have no constructor initializer and a null body.

The compiler first implicitly defines the implicitly declared constructors of the base classes and nonstatic data members of a class *A* before defining the implicitly declared constructor of *A*. No default constructor is created for a class that has any constant or reference type members.

A constructor of a class *A* is *trivial* if all the following are true:

- It is implicitly defined
- *A* has no virtual functions and no virtual base classes
- All the direct base classes of *A* have trivial constructors
- The classes of all the nonstatic data members of *A* have trivial constructors

If any of the above are false, then the constructor is *nontrivial*.

A union member cannot be of a class type that has a nontrivial constructor.

Like all functions, a constructor can have default arguments. They are used to initialize member objects. If default values are supplied, the trailing arguments can be omitted in the expression list of the constructor. Note that if a constructor has any arguments that do not have default values, it is not a default constructor.

A *copy constructor* for a class *A* is a constructor whose first parameter is of type `A&`, `const A&`, `volatile A&`, or `const volatile A&`. Copy constructors are used to make a copy of one class object from another class object of the same class type. You cannot use a copy constructor with an argument of the same type as its class; you must use a reference. You can provide copy constructors with additional

parameters as long as they all have default arguments. If a user-defined copy constructor does not exist for a class and one is needed, the compiler implicitly creates a copy constructor, with public access, for that class. A copy constructor is not created for a class if any of its members or base classes have an inaccessible copy constructor.

The following code fragment shows two classes with constructors, default constructors, and copy constructors:

```
class X {
public:

    // default constructor, no arguments
    X();

    // constructor
    X(int, int , int = 0);

    // copy constructor
    X(const X&);

    // error, incorrect argument type
    X(X);
};

class Y {
public:

    // default constructor with one
    // default argument
    Y( int = 0);

    // default argument
    // copy constructor
    Y(const Y&, int = 0);
};
```

Related information

- “The constructor and destructor function attributes” on page 221
- “Copy constructors” on page 333

Explicit initialization with constructors

A class object with a constructor must be explicitly initialized or have a default constructor. Except for aggregate initialization, explicit initialization using a constructor is the only way to initialize non-static constant and reference class members.

A class object that has no user-declared constructors, no virtual functions, no private or protected non-static data members, and no base classes is called an *aggregate*. Examples of aggregates are C-style structures and unions.

You explicitly initialize a class object when you create that object. There are two ways to initialize a class object:

- Using a parenthesized expression list. The compiler calls the constructor of the class using this list as the constructor’s argument list.
- Using a single initialization value and the = operator. Because this type of expression is an initialization, not an assignment, the assignment operator function, if one exists, is not called. The type of the single argument must match the type of the first argument to the constructor. If the constructor has remaining arguments, these arguments must have default values.


```

re = 1 im = 0
re = 1 im = 0
re = 3 im = 4
re = 0 im = 0
re = 5 im = 0

```

Related information

- “The constructor and destructor function attributes” on page 221
- “Initializers” on page 93

Initialization of base classes and members

Constructors can initialize their members in two different ways. A constructor can use the arguments passed to it to initialize member variables in the constructor definition:

```

complx(double r, double i = 0.0) { re = r; im = i; }

```

Or a constructor can have an *initializer list* within the definition but prior to the constructor body:


```

complx(double r, double i = 0) : re(r), im(i) { /* ... */ }

```

Both methods assign the argument values to the appropriate data members of the class.

Initializer list syntax



Include the initialization list as part of the constructor definition, not as part of the constructor declaration. For example:

```

#include <iostream>
using namespace std;

class B1 {
    int b;
public:
    B1() { cout << "B1::B1()" << endl; };

    // inline constructor
    B1(int i) : b(i) { cout << "B1::B1(int)" << endl; }
};

class B2 {
    int b;
protected:
    B2() { cout << "B2::B2()" << endl; }

    // noninline constructor
    B2(int i);
};

// B2 constructor definition including initialization list
B2::B2(int i) : b(i) { cout << "B2::B2(int)" << endl; }

class D : public B1, public B2 {
    int d1, d2;
public:

```

```

    D(int i, int j) : B1(i+1), B2(), d1(i) {
        cout << "D1::D1(int, int)" << endl;
        d2 = j;}
};

int main() {
    D obj(1, 2);
}

```

The following is the output of the above example:

```

B1::B1(int)
B1::B1()
D1::D1(int, int)

```

If you do not explicitly initialize a base class or member that has constructors by calling a constructor, the compiler automatically initializes the base class or member with a default constructor. In the above example, if you leave out the call `B2()` in the constructor of class `D` (as shown below), a constructor initializer with an empty expression list is automatically created to initialize `B2`. The constructors for class `D`, shown above and below, result in the same construction of an object of class `D`:

```

class D : public B1, public B2 {
    int d1, d2;
public:

    // call B2() generated by compiler
    D(int i, int j) : B1(i+1), d1(i) {
        cout << "D1::D1(int, int)" << endl;
        d2 = j;}
};

```

In the above example, the compiler will automatically call the default constructor for `B2()`.

Note that you must declare constructors as public or protected to enable a derived class to call them. For example:

```

class B {
    B() { }
};

class D : public B {

    // error: implicit call to private B() not allowed
    D() { }
};

```

The compiler does not allow the definition of `D::D()` because this constructor cannot access the private constructor `B::B()`.

You must initialize the following with an initializer list: base classes with no default constructors, reference data members, non-static const data members, or a class type which contains a constant data member. The following example demonstrates this:

```

class A {
public:
    A(int) { }
};

class B : public A {
    static const int i;
    const int j;
};

```

```

    int &k;
public:
    B(int& arg) : A(0), j(1), k(arg) { }
};

int main() {
    int x = 0;
    B obj(x);
};

```

The data members `j` and `k`, as well as the base class `A` must be initialized in the initializer list of the constructor of `B`.

You can use data members when initializing members of a class. The following example demonstrate this:

```

struct A {
    int k;
    A(int i) : k(i) { }
};
struct B: A {
    int x;
    int i;
    int j;
    int& r;
    B(int i): r(x), A(i), j(this->i), i(i) { }
};

```

The constructor `B(int i)` initializes the following:

- `B::r` to refer to `B::x`
- Class `A` with the value of the argument to `B(int i)`
- `B::j` with the value of `B::i`
- `B::i` with the value of the argument to `B(int i)`

You can also call member functions (including virtual member functions) or use the operators `typeid` or `dynamic_cast` when initializing members of a class. However if you perform any of these operations in a member initialization list before all base classes have been initialized, the behavior is undefined. The following example demonstrates this:

```

#include <iostream>
using namespace std;

struct A {
    int i;
    A(int arg) : i(arg) {
        cout << "Value of i: " << i << endl;
    }
};

struct B : A {
    int j;
    int f() { return i; }
    B();
};

B::B() : A(f()), j(1234) {
    cout << "Value of j: " << j << endl;
}

int main() {
    B obj;
}

```

The output of the above example would be similar to the following:

```
Value of i: 8
Value of j: 1234
```

The behavior of the initializer `A(f())` in the constructor of `B` is undefined. The run time will call `B::f()` and try to access `A::i` even though the base `A` has not been initialized.

The following example is the same as the previous example except that the initializers of `B::B()` have different arguments:

```
#include <iostream>
using namespace std;

struct A {
    int i;
    A(int arg) : i(arg) {
        cout << "Value of i: " << i << endl;
    }
};

struct B : A {
    int j;
    int f() { return i; }
    B();
};

B::B() : A(5678), j(f()) {
    cout << "Value of j: " << j << endl;
}

int main() {
    B obj;
}
```

The following is the output of the above example:

```
Value of i: 5678
Value of j: 5678
```

The behavior of the initializer `j(f())` in the constructor of `B` is well-defined. The base class `A` is already initialized when `B::j` is initialized.

Related information

- “The typeid operator (C++ only) (PPU only)” on page 133
- “The dynamic_cast operator (C++ only) (PPU only)” on page 162

Construction order of derived class objects

When a derived class object is created using constructors, it is created in the following order:

1. Virtual base classes are initialized, in the order they appear in the base list.
2. Nonvirtual base classes are initialized, in declaration order.
3. Class members are initialized in declaration order (regardless of their order in the initialization list).
4. The body of the constructor is executed.

The following example demonstrates this:

```
#include <iostream>
using namespace std;
struct V {
```

```

    V() { cout << "V()" << endl; }
};
struct V2 {
    V2() { cout << "V2()" << endl; }
};
struct A {
    A() { cout << "A()" << endl; }
};
struct B : virtual V {
    B() { cout << "B()" << endl; }
};
struct C : B, virtual V2 {
    C() { cout << "C()" << endl; }
};
struct D : C, virtual V {
    A obj_A;
    D() { cout << "D()" << endl; }
};
int main() {
    D c;
}

```

The following is the output of the above example:

```

V()
V2()
B()
C()
A()
D()

```

The above output lists the order in which the C++ run time calls the constructors to create an object of type D.

Related information

- “Virtual base classes” on page 303

Destructors

Destructors are usually used to deallocate memory and do other cleanup for a class object and its class members when the object is destroyed. A destructor is called for a class object when that object passes out of scope or is explicitly deleted.

A destructor is a member function with the same name as its class prefixed by a ~ (tilde). For example:

```

class X {
public:
    // Constructor for class X
    X();
    // Destructor for class X
    ~X();
};

```

A destructor takes no arguments and has no return type. Its address cannot be taken. Destructors cannot be declared `const`, `volatile`, `const volatile` or `static`. A destructor can be declared `virtual` or `pure virtual`.

If no user-defined destructor exists for a class and one is needed, the compiler implicitly declares a destructor. This implicitly declared destructor is an inline public member of its class.

The compiler will implicitly define an implicitly declared destructor when the compiler uses the destructor to destroy an object of the destructor's class type. Suppose a class A has an implicitly declared destructor. The following is equivalent to the function the compiler would implicitly define for A:

```
A::~~A() { }
```

The compiler first implicitly defines the implicitly declared destructors of the base classes and nonstatic data members of a class A before defining the implicitly declared destructor of A

A destructor of a class A is *trivial* if all the following are true:

- It is implicitly defined
- All the direct base classes of A have trivial destructors
- The classes of all the nonstatic data members of A have trivial destructors

If any of the above are false, then the destructor is *nontrivial*.

A union member cannot be of a class type that has a nontrivial destructor.

Class members that are class types can have their own destructors. Both base and derived classes can have destructors, although destructors are not inherited. If a base class A or a member of A has a destructor, and a class derived from A does not declare a destructor, a default destructor is generated.

The default destructor calls the destructors of the base class and members of the derived class.

The destructors of base classes and members are called in the reverse order of the completion of their constructor:

1. The destructor for a class object is called before destructors for members and bases are called.
2. Destructors for nonstatic members are called before destructors for base classes are called.
3. Destructors for nonvirtual base classes are called before destructors for virtual base classes are called.

When an exception is thrown for a class object with a destructor, the destructor for the temporary object thrown is not called until control passes out of the catch block.

Destructors are implicitly called when an automatic object (a local object that has been declared `auto` or `register`, or not declared as `static` or `extern`) or temporary object passes out of scope. They are implicitly called at program termination for constructed external and static objects. Destructors are invoked when you use the `delete` operator for objects created with the `new` operator.

For example:

```
#include <string>

class Y {
private:
    char * string;
    int number;
public:
    // Constructor
```

```

    Y(const char*, int);
    // Destructor
    ~Y() { delete[] string; }
};

// Define class Y constructor
Y::Y(const char* n, int a) {
    string = strcpy(new char[strlen(n) + 1 ], n);
    number = a;
}

int main () {
    // Create and initialize
    // object of class Y
    Y yobj = Y("somestring", 10);

    // ...

    // Destructor ~Y is called before
    // control returns from main()
}

```

You can use a destructor explicitly to destroy objects, although this practice is not recommended. However to destroy an object created with the placement new operator, you can explicitly call the object's destructor. The following example demonstrates this:

```

#include <new>
#include <iostream>
using namespace std;
class A {
public:
    A() { cout << "A::A()" << endl; }
    ~A() { cout << "A::~A()" << endl; }
};
int main () {
    char* p = new char[sizeof(A)];
    A* ap = new (p) A;
    ap->A::~A();
    delete [] p;
}

```

The statement `A* ap = new (p) A` dynamically creates a new object of type A not in the free store but in the memory allocated by p. The statement `delete [] p` will delete the storage allocated by p, but the run time will still believe that the object pointed to by ap still exists until you explicitly call the destructor of A (with the statement `ap->A::~A()`).


Related information

- “The constructor and destructor function attributes” on page 221

Pseudo-destructors

A *pseudo-destructor* is a destructor of a nonclass type.

Pseudo-destructor syntax



The following example calls the pseudo destructor for an integer type:

```
typedef int I;
int main() {
    I x = 10;
    x.I::~~I();
    x = 20;
}
```

The call to the pseudo destructor, `x.I::~~I()`, has no effect at all. Object `x` has not been destroyed; the assignment `x = 20` is still valid. Because pseudo destructors require the syntax for explicitly calling a destructor for a nonclass type to be valid, you can write code without having to know whether or not a destructor exists for a given type.

Related information

- Chapter 12, “Class members and friends (C++ only),” on page 269
- “Scope of class names” on page 263

User-defined conversions

User-defined conversions allow you to specify object conversions with constructors or with conversion functions. User-defined conversions are implicitly used in addition to standard conversions for conversion of initializers, functions arguments, function return values, expression operands, expressions controlling iteration, selection statements, and explicit type conversions.

There are two types of user-defined conversions:

- Conversion constructors
- Conversion functions

The compiler can use only one user-defined conversion (either a conversion constructor or a conversion function) when implicitly converting a single value. The following example demonstrates this:

```
class A {
    int x;
public:
    operator int() { return x; };
};

class B {
    A y;
public:
    operator A() { return y; };
};

int main () {
    B b_obj;
    // int i = b_obj;
    int j = A(b_obj);
}
```

The compiler would not allow the statement `int i = b_obj`. The compiler would have to implicitly convert `b_obj` into an object of type `A` (with `B::operator A()`), then implicitly convert that object to an integer (with `A::operator int()`). The statement `int j = A(b_obj)` explicitly converts `b_obj` into an object of type `A`, then implicitly converts that object to an integer.

User-defined conversions must be unambiguous, or they are not called. A conversion function in a derived class does not hide another conversion function in a base class unless both conversion functions convert to the same type. Function overload resolution selects the most appropriate conversion function. The following example demonstrates this:

```
class A {
    int a_int;
    char* a_carp;
public:
    operator int() { return a_int; }
    operator char*() { return a_carp; }
};

class B : public A {
    float b_float;
    char* b_carp;
public:
    operator float() { return b_float; }
    operator char*() { return b_carp; }
};

int main () {
    B b_obj;
    // long a = b_obj;
    char* c_p = b_obj;
}
```

The compiler would not allow the statement `long a = b_obj`. The compiler could either use `A::operator int()` or `B::operator float()` to convert `b_obj` into a `long`. The statement `char* c_p = b_obj` uses `B::operator char*()` to convert `b_obj` into a `char*` because `B::operator char*()` hides `A::operator char*()`.

When you call a constructor with an argument and you have not defined a constructor accepting that argument type, only standard conversions are used to convert the argument to another argument type acceptable to a constructor for that class. No other constructors or conversions functions are called to convert the argument to a type acceptable to a constructor defined for that class. The following example demonstrates this:

```
class A {
public:
    A() { }
    A(int) { }
};

int main() {
    A a1 = 1.234;
    // A moocow = "text string";
}
```

The compiler allows the statement `A a1 = 1.234`. The compiler uses the standard conversion of converting 1.234 into an `int`, then implicitly calls the converting constructor `A(int)`. The compiler would not allow the statement `A moocow = "text string"`; converting a text string to an integer is not a standard conversion.

Related information

- Chapter 5, “Type conversions,” on page 113

Conversion constructors

A *conversion constructor* is a single-parameter constructor that is declared without the function specifier `explicit`. The compiler uses conversion constructors to

convert objects from the type of the first parameter to the type of the conversion constructor's class. The following example demonstrates this:

```
class Y {
    int a, b;
    char* name;
public:
    Y(int i) { };
    Y(const char* n, int j = 0) { };
};

void add(Y) { };

int main() {

    // equivalent to
    // obj1 = Y(2)
    Y obj1 = 2;

    // equivalent to
    // obj2 = Y("somestring",0)
    Y obj2 = "somestring";

    // equivalent to
    // obj1 = Y(10)
    obj1 = 10;

    // equivalent to
    // add(Y(5))
    add(5);
}
```

The above example has the following two conversion constructors:

- `Y(int i)` which is used to convert integers to objects of class `Y`.
- `Y(const char* n, int j = 0)` which is used to convert pointers to strings to objects of class `Y`.

The compiler will not implicitly convert types as demonstrated above with constructors declared with the `explicit` keyword. The compiler will only use explicitly declared constructors in new expressions, the `static_cast` expressions and explicit casts, and the initialization of bases and members. The following example demonstrates this:

```
class A {
public:
    explicit A() { };
    explicit A(int) { };
};

int main() {
    A z;
    // A y = 1;
    A x = A(1);
    A w(1);
    A* v = new A(1);
    A u = (A)1;
    A t = static_cast<A>(1);
}
```

The compiler would not allow the statement `A y = 1` because this is an implicit conversion; class `A` has no conversion constructors.

A copy constructor is a conversion constructor.

Related information

- “new expressions (C++ only)” on page 165
- “The static_cast operator (C++ only)” on page 159

The explicit specifier

The `explicit` function specifier controls unwanted implicit type conversions. It can only be used in declarations of constructors within a class declaration. For example, except for the default constructor, the constructors in the following class are converting constructors.

```
class A
{ public:
  A();
  A(int);
  A(const char*, int = 0);
};
```

The following declarations are legal.

```
A c = 1;
A d = "Venditti";
```

The first declaration is equivalent to `A c = A(1)`.

If you declare the constructor of the class with the `explicit` keyword, the previous declarations would be illegal.

For example, if you declare the class as:

```
class A
{ public:
  explicit A();
  explicit A(int);
  explicit A(const char*, int = 0);
};
```

You can only assign values that match the values of the class type.

For example, the following statements will be legal:

```
A a1;
A a2 = A(1);
A a3(1);
A a4 = A("Venditti");
A* p = new A(1);
A a5 = (A)1;
A a6 = static_cast<A>(1);
```


Related information


- “Conversion constructors” on page 330

Conversion functions

You can define a member function of a class, called a *conversion function*, that converts from the type of its class to another specified type.

Conversion function syntax





A conversion function that belongs to a class X specifies a conversion from the class type X to the type specified by the *conversion_type*. The following code fragment shows a conversion function called `operator int()`:

```
class Y {
    int b;
public:
    operator int();
};
Y::operator int() {
    return b;
}
void f(Y obj) {
    int i = int(obj);
    int j = (int)obj;
    int k = i + obj;
}
```

All three statements in function `f(Y)` use the conversion function `Y::operator int()`.

Classes, enumerations, typedef names, function types, or array types cannot be declared or defined in the *conversion_type*. You cannot use a conversion function to convert an object of type A to type A , to a base class of A , or to `void`.

Conversion functions have no arguments, and the return type is implicitly the conversion type. Conversion functions can be inherited. You can have virtual conversion functions but not static ones.

Related information

- Chapter 5, “Type conversions,” on page 113
- “User-defined conversions” on page 329
- “Conversion constructors” on page 330
- Chapter 5, “Type conversions,” on page 113

Copy constructors

The *copy constructor* lets you create a new object from an existing one by initialization. A copy constructor of a class A is a non-template constructor in which the first parameter is of type `A&`, `const A&`, `volatile A&`, or `const volatile A&`, and the rest of its parameters (if there are any) have default values.

If you do not declare a copy constructor for a class A , the compiler will implicitly declare one for you, which will be an inline public member.

The following example demonstrates implicitly defined and user-defined copy constructors:

```
#include <iostream>
using namespace std;

struct A {
    int i;
    A() : i(10) { }
```

```

};

struct B {
    int j;
    B() : j(20) {
        cout << "Constructor B(), j = " << j << endl;
    }

    B(B& arg) : j(arg.j) {
        cout << "Copy constructor B(B&), j = " << j << endl;
    }

    B(const B&, int val = 30) : j(val) {
        cout << "Copy constructor B(const B&, int), j = " << j << endl;
    }
};

struct C {
    C() { }
    C(C&) { }
};

int main() {
    A a;
    A a1(a);
    B b;
    const B b_const;
    B b1(b);
    B b2(b_const);
    const C c_const;
    // C c1(c_const);
}

```

The following is the output of the above example:

```

Constructor B(), j = 20
Constructor B(), j = 20
Copy constructor B(B&), j = 20
Copy constructor B(const B&, int), j = 30

```

The statement `A a1(a)` creates a new object from `a` with an implicitly defined copy constructor. The statement `B b1(b)` creates a new object from `b` with the user-defined copy constructor `B::B(B&)`. The statement `B b2(b_const)` creates a new object with the copy constructor `B::B(const B&, int)`. The compiler would not allow the statement `C c1(c_const)` because a copy constructor that takes as its first parameter an object of type `const C&` has not been defined.

The implicitly declared copy constructor of a class `A` will have the form `A::A(const A&)` if the following are true:

- The direct and virtual bases of `A` have copy constructors whose first parameters have been qualified with `const` or `const volatile`
- The nonstatic class type or array of class type data members of `A` have copy constructors whose first parameters have been qualified with `const` or `const volatile`

If the above are not true for a class `A`, the compiler will implicitly declare a copy constructor with the form `A::A(A&)`.

The compiler cannot allow a program in which the compiler must implicitly define a copy constructor for a class `A` and one or more of the following are true:

- Class `A` has a nonstatic data member of a type which has an inaccessible or ambiguous copy constructor.

- Class A is derived from a class which has an inaccessible or ambiguous copy constructor.

The compiler will implicitly define an implicitly declared constructor of a class A if you initialize an object of type A or an object derived from class A.

An implicitly defined copy constructor will copy the bases and members of an object in the same order that a constructor would initialize the bases and members of the object.

Related information

- “Overview of constructors and destructors” on page 317

Copy assignment operators

The *copy assignment operator* lets you create a new object from an existing one by initialization. A copy assignment operator of a class A is a nonstatic non-template member function that has one of the following forms:

- `A::operator=(A)`
- `A::operator=(A&)`
- `A::operator=(const A&)`
- `A::operator=(volatile A&)`
- `A::operator=(const volatile A&)`

If you do not declare a copy assignment operator for a class A, the compiler will implicitly declare one for you which will be inline public.

The following example demonstrates implicitly defined and user-defined copy assignment operators:

```
#include <iostream>
using namespace std;

struct A {
    A& operator=(const A&) {
        cout << "A::operator=(const A&)" << endl;
        return *this;
    }

    A& operator=(A&) {
        cout << "A::operator=(A&)" << endl;
        return *this;
    }
};

class B {
    A a;
};

struct C {
    C& operator=(C&) {
        cout << "C::operator=(C&)" << endl;
        return *this;
    }
    C() { }
};

int main() {
    B x, y;
    x = y;
}
```

```

A w, z;
w = z;

C i;
const C j();
// i = j;
}

```

The following is the output of the above example:

```

A::operator=(const A&)
A::operator=(A&)

```

The assignment `x = y` calls the implicitly defined copy assignment operator of `B`, which calls the user-defined copy assignment operator `A::operator=(const A&)`. The assignment `w = z` calls the user-defined operator `A::operator=(A&)`. The compiler will not allow the assignment `i = j` because an operator `C::operator=(const C&)` has not been defined.

The implicitly declared copy assignment operator of a class `A` will have the form `A& A::operator=(const A&)` if the following are true:

- A direct or virtual base `B` of class `A` has a copy assignment operator whose parameter is of type `const B&`, `const volatile B&`, or `B`.
- A non-static class type data member of type `X` that belongs to class `A` has a copy constructor whose parameter is of type `const X&`, `const volatile X&`, or `X`.

If the above are not true for a class `A`, the compiler will implicitly declare a copy assignment operator with the form `A& A::operator=(A&)`.

The implicitly declared copy assignment operator returns a reference to the operator's argument.

The copy assignment operator of a derived class hides the copy assignment operator of its base class.

The compiler cannot allow a program in which the compiler must implicitly define a copy assignment operator for a class `A` and one or more of the following are true:

- Class `A` has a nonstatic data member of a `const` type or a reference type
- Class `A` has a nonstatic data member of a type which has an inaccessible copy assignment operator
- Class `A` is derived from a base class with an inaccessible copy assignment operator.

An implicitly defined copy assignment operator of a class `A` will first assign the direct base classes of `A` in the order that they appear in the definition of `A`. Next, the implicitly defined copy assignment operator will assign the nonstatic data members of `A` in the order of their declaration in the definition of `A`.

Related information

- "Assignment operators" on page 140

Chapter 15. Templates (C++ only)

A *template* describes a set of related classes or set of related functions in which a list of parameters in the declaration describe how the members of the set vary. The compiler generates new classes or functions when you supply arguments for these parameters; this process is called *template instantiation*, and is described in detail in “Template instantiation” on page 356. This class or function definition generated from a template and a set of template parameters is called a *specialization*, as described in “Template specialization” on page 359.

Template declaration syntax

► `[export] template <—template_parameter_list—> —declaration—` ►

The compiler accepts and silently ignores the `export` keyword on a template.

The *template_parameter_list* is a comma-separated list of template parameters, which are described in “Template parameters” on page 338.

The *declaration* is one of the following::

- a declaration or definition of a function or a class
- a definition of a member function or a member class of a class template
- a definition of a static data member of a class template
- a definition of a static data member of a class nested within a class template
- a definition of a member template of a class or class template

The *identifier* of a *type* is defined to be a *type_name* in the scope of the template declaration. A template declaration can appear as a namespace scope or class scope declaration.

The following example demonstrates the use of a class template:

```
template<class T> class Key
{
    T k;
    T* kptr;
    int length;
public:
    Key(T);
    // ...
};
```

Suppose the following declarations appear later:

```
Key<int> i;
Key<char*> c;
Key<mytype> m;
```

The compiler would create three instances of class `Key`. The following table shows the definitions of these three class instances if they were written out in source form as regular classes, not as templates:

<code>class Key<int> i;</code>	<code>class Key<char*> c;</code>	<code>class Key<mytype> m;</code>
<pre>class Key { int k; int * kptr; int length; public: Key(int); // ... };</pre>	<pre>class Key { char* k; char** kptr; int length; public: Key(char*); // ... };</pre>	<pre>class Key { mytype k; mytype* kptr; int length; public: Key(mytype); // ... };</pre>

Note that these three classes have different names. The arguments contained within the angle braces are not just the arguments to the class names, but part of the class names themselves. `Key<int>` and `Key<char*>` are class names.

Template parameters

There are three kinds of template parameters:

- Type template parameters
- Non-type template parameters
- Template template parameters

You can interchange the keywords `class` and `typename` in a template parameter declaration. You cannot use storage class specifiers (`static` and `auto`) in a template parameter declaration.

Type template parameters

Type template parameter declaration syntax



The *identifier* is the name of a type.

Related information

- “The `typename` keyword” on page 367

Non-type template parameters

The syntax of a non-type template parameter is the same as a declaration of one of the following types:

- integral or enumeration
- pointer to object or pointer to function
- reference to object or reference to function
- pointer to member

Non-type template parameters that are declared as arrays or functions are converted to pointers or pointers to functions, respectively. The following example demonstrates this:

```
template<int a[4]> struct A { };
template<int f(int)> struct B { };

int i;
```

```
int g(int) { return 0;}

A<&i> x;
B<&g> y;
```

The type deduced from `&i` is `int *`, and the type deduced from `&g` is `int (*)(int)`.

You may qualify a non-type template parameter with `const` or `volatile`.

You cannot declare a non-type template parameter as a floating point, class, or void type.

Non-type template parameters are not lvalues.

Related information

- “Type qualifiers” on page 72
- “Lvalues and rvalues” on page 121

Template template parameters

Template template parameter declaration syntax

```
▶▶—template—◀◀—template-parameter-list—▶▶—class—          —          —◀◀
                                          |                                          |
                                          |_identifier_|                          |_id-expression_|
```

The following example demonstrates a declaration and use of a template template parameter:

```
template<template <class T> class X> class A { };
template<class T> class B { };

A<B> a;
```

Related information

- “Template parameters” on page 338

Default arguments for template parameters

Template parameters may have default arguments. The set of default template arguments accumulates over all declarations of a given template. The following example demonstrates this:

```
template<class T, class U = int> class A;
template<class T = float, class U> class A;

template<class T, class U> class A {
public:
    T x;
    U y;
};

A<> a;
```

The type of member `a.x` is `float`, and the type of `a.y` is `int`.

You cannot give default arguments to the same template parameters in different declarations in the same scope. For example, the compiler will not allow the following:

```
template<class T = char> class X;  
template<class T = char> class X { };
```

If one template parameter has a default argument, then all template parameters following it must also have default arguments. For example, the compiler will not allow the following:

```
template<class T = char, class U, class V = int> class X { };
```

Template parameter U needs a default argument or the default for T must be removed.

The scope of a template parameter starts from the point of its declaration to the end of its template definition. This implies that you may use the name of a template parameter in other template parameter declarations and their default arguments. The following example demonstrates this:

```
template<class T = int> class A;  
template<class T = float> class B;  
template<class V, V obj> class C;  
// a template parameter (T) used as the default argument  
// to another template parameter (U)  
template<class T, class U = T> class D { };
```

Related information

- “Template parameters” on page 338

Template arguments

There are three kinds of template arguments corresponding to the three types of template parameters:

- Template type arguments
- Template non-type arguments
- Template template arguments

A template argument must match the type and form specified by the corresponding parameter declared in the template.

To use the default value of a template parameter, you omit the corresponding template argument. However, even if all template parameters have defaults, you still must use the <> brackets. For example, the following will yield a syntax error:

```
template<class T = int> class X { };  
X<> a;  
X b;
```

The last declaration, X b, will yield an error.

Template type arguments

You cannot use one of the following as a template argument for a type template parameter:

- a local type
- a type with no linkage
- an unnamed type
- a type compounded from any of the above types

If it is ambiguous whether a template argument is a type or an expression, the template argument is considered to be a type. The following example demonstrates this:

```
template<class T> void f() { };
template<int i> void f() { };

int main() {
    f<int>();
}
```

The function call `f<int>()` calls the function with `T` as a template argument – the compiler considers `int()` as a type – and therefore implicitly instantiates and calls the first `f()`.

Related information

- “Block/local scope” on page 2
- “No linkage” on page 9
- “Bit field members” on page 63
- “typedef definitions” on page 71

Template non-type arguments

A non-type template argument provided within a template argument list is an expression whose value can be determined at compile time. Such arguments must be constant expressions, addresses of functions or objects with external linkage, or addresses of static class members. Non-type template arguments are normally used to initialize a class or to specify the sizes of class members.

For non-type integral arguments, the instance argument matches the corresponding template parameter as long as the instance argument has a value and sign appropriate to the parameter type.

For non-type address arguments, the type of the instance argument must be of the form *identifier* or *&identifier*, and the type of the instance argument must match the template parameter exactly, except that a function name is changed to a pointer to function type before matching.

The resulting values of non-type template arguments within a template argument list form part of the template class type. If two template class names have the same template name and if their arguments have identical values, they are the same class.

In the following example, a class template is defined that requires a non-type template `int` argument as well as the type argument:

```
template<class T, int size> class Myfilebuf
{
    T* filepos;
    static int array[size];
public:
    Myfilebuf() { /* ... */ }
    ~Myfilebuf();
    advance(); // function defined elsewhere in program
};
```

In this example, the template argument `size` becomes a part of the template class name. An object of such a template class is created with both the type argument `T` of the class and the value of the non-type template argument `size`.

An object `x`, and its corresponding template class with arguments `double` and `size=200`, can be created from this template with a value as its second template argument:

```
Myfilebuf<double,200> x;
```

`x` can also be created using an arithmetic expression:

```
Myfilebuf<double,10*20> x;
```

The objects created by these expressions are identical because the template arguments evaluate identically. The value `200` in the first expression could have been represented by an expression whose result at compile time is known to be equal to `200`, as shown in the second construction.

Note: Arguments that contain the `<` symbol or the `>` symbol must be enclosed in parentheses to prevent either symbol from being parsed as a template argument list delimiter when it is in fact being used as a relational operator. For example, the arguments in the following definition are valid:

```
Myfilebuf<double, (75>25)> x;          // valid
```

The following definition, however, is not valid because the greater than operator (`>`) is interpreted as the closing delimiter of the template argument list:

```
Myfilebuf<double, 75>25> x;          // error
```

If the template arguments do not evaluate identically, the objects created are of different types:

```
Myfilebuf<double,200> x;              // create object x of class
// Myfilebuf<double,200>
Myfilebuf<double,200.0> y;            // error, 200.0 is a double,
// not an int
```

The instantiation of `y` fails because the value `200.0` is of type `double`, and the template argument is of type `int`.

The following two objects:

```
Myfilebuf<double, 128> x
Myfilebuf<double, 512> y
```

are objects of separate template specializations. Referring either of these objects later with `Myfilebuf<double>` is an error.

A class template does not need to have a type argument if it has non-type arguments. For example, the following template is a valid class template:

```
template<int i> class C
{
    public:
        int k;
        C() { k = i; }
};
```

This class template can be instantiated by declarations such as:

```
class C<100>;
class C<200>;
```

Again, these two declarations refer to distinct classes because the values of their non-type arguments differ.

Related information

- “Integer constant expressions” on page 124
- “References (C++ only)” on page 93
- “External linkage” on page 8
- “Static members” on page 279

Template template arguments

A template argument for a template template parameter is the name of a class template.

When the compiler tries to find a template to match the template template argument, it only considers primary class templates. (A *primary template* is the template that is being specialized.) The compiler will not consider any partial specialization even if their parameter lists match that of the template template parameter. For example, the compiler will not allow the following code:

```
template<class T, int i> class A {
    int x;
};

template<class T> class A<T, 5> {
    short x;
};

template<template<class T> class U> class B1 { };

B1<A> c;
```

The compiler will not allow the declaration `B1<A> c`. Although the partial specialization of `A` seems to match the template template parameter `U` of `B1`, the compiler considers only the primary template of `A`, which has different template parameters than `U`.

The compiler considers the partial specializations based on a template template argument once you have instantiated a specialization based on the corresponding template template parameter. The following example demonstrates this:

```
#include <iostream>
using namespace std;

template<class T, class U> class A {
    int x;
};

template<class U> class A<int, U> {
    short x;
};

template<template<class T, class U> class V> class B {
    V<int, char> i;
    V<char, char> j;
};

B<A> c;

int main() {
    cout << typeid(c.i.x).name() << endl;
    cout << typeid(c.j.x).name() << endl;
}
```

The following is the output of the above example:

```
short
int
```

The declaration `V<int, char> i` uses the partial specialization while the declaration `V<char, char> j` uses the primary template.

Related information

- “Partial specialization” on page 364
- “Template instantiation” on page 356

Class templates

The relationship between a class template and an individual class is like the relationship between a class and an individual object. An individual class defines how a group of objects can be constructed, while a class template defines how a group of classes can be generated.

Note the distinction between the terms *class template* and *template class*:

Class template

is a template used to generate template classes. You cannot declare an object of a class template.

Template class

is an instance of a class template.

A template definition is identical to any valid class definition that the template might generate, except for the following:

- The class template definition is preceded by
`template< template-parameter-list >`

where *template-parameter-list* is a comma-separated list of one or more of the following kinds of template parameters:

- type
- non-type
- template

- Types, variables, constants and objects within the class template can be declared using the template parameters as well as explicit types (for example, `int` or `char`).

A class template can be declared without being defined by using an elaborated type specifier. For example:

```
template<class L, class T> class Key;
```

This reserves the name as a class template name. All template declarations for a class template must have the same types and number of template arguments. Only one template declaration containing the class definition is allowed.

Note: When you have nested template argument lists, you must have a separating space between the `>` at the end of the inner list and the `>` at the end of the outer list. Otherwise, there is an ambiguity between the extraction operator `>>` and two template list delimiters `>`.

```
template<class L, class T> class Key { /* ... */};
template<class L> class Vector { /* ... */};
```

```

int main ()
{
    class Key <int, Vector<int> > my_key_vector;
    // implicitly instantiates template
}

```

Objects and function members of individual template classes can be accessed by any of the techniques used to access ordinary class member objects and functions. Given a class template:

```

template<class T> class Vehicle
{
public:
    Vehicle() { /* ... */ }    // constructor
    ~Vehicle() {};           // destructor
    T kind[16];
    T* drive();
    static void roadmap();
    // ...
};

```

and the declaration:

```
Vehicle<char> bicycle; // instantiates the template
```

the constructor, the constructed object, and the member function drive() can be accessed with any of the following (assuming the standard header file string.h is included in the program file):

constructor	Vehicle<char> bicycle; // constructor called automatically, // object bicycle created
object bicycle	strcpy (bicycle.kind, "10 speed"); bicycle.kind[0] = '2';
function drive()	char* n = bicycle.drive();
function roadmap()	Vehicle<char>::roadmap();

Related information

- “Declaring class types” on page 260
- “Scope of class names” on page 263

Class template declarations and definitions

A class template must be declared before any instantiation of a corresponding template class. A class template definition can only appear once in any single translation unit. A class template must be defined before any use of a template class that requires the size of the class or refers to members of the class.

In the following example, the class template Key is declared before it is defined. The declaration of the pointer keyiptr is valid because the size of the class is not needed. The declaration of keyi, however, causes an error.

```

template <class L> class Key;    // class template declared,
                                // not defined yet
                                //
class Key<int> *keyiptr;        // declaration of pointer
                                //
class Key<int> keyi;           // error, cannot declare keyi

```

```

// without knowing size
//
template <class L> class Key // now class template defined
{ /* ... */ };

```

If a template class is used before the corresponding class template is defined, the compiler issues an error. A class name with the appearance of a template class name is considered to be a template class. In other words, angle brackets are valid in a class name only if that class is a template class.

The previous example uses the elaborated type specifier `class` to declare the class template `key` and the pointer `keyiptr`. The declaration of `keyiptr` can also be made without the elaborated type specifier.

```

template <class L> class Key; // class template declared,
// not defined yet
//
Key<int> *keyiptr; // declaration of pointer
//
Key<int> keyi; // error, cannot declare keyi
// without knowing size
//
template <class L> class Key // now class template defined
{ /* ... */ };

```

Related information

- “Class templates” on page 344

Static data members and templates

Each class template instantiation has its own copy of any static data members. The static declaration can be of template argument type or of any defined type.

You must separately define static members. The following example demonstrates this:

```

template <class T> class K
{
public:
    static T x;
};
template <class T> T K<T> ::x;

int main()
{
    K<int>::x = 0;
}

```

The statement `template T K::x` defines the static member of class `K`, while the statement in the `main()` function assigns a value to the data member for `K <int>`.

Related information

- “Static members” on page 279

Member functions of class templates

You may define a template member function outside of its class template definition.

When you call a member function of a class template specialization, the compiler will use the template arguments that you used to generate the class template. The following example demonstrates this:

```

template<class T> class X {
public:
    T operator+(T);
};

template<class T> T X<T>::operator+(T arg1) {
    return arg1;
};

int main() {
    X<char> a;
    X<int> b;
    a + 'z';
    b + 4;
}

```

The overloaded addition operator has been defined outside of class X. The statement `a + 'z'` is equivalent to `a.operator+'z')`. The statement `b + 4` is equivalent to `b.operator+(4)`.

Related information

- “Member functions” on page 271

Friends and templates

There are four kinds of relationships between classes and their friends when templates are involved:

- *One-to-many*: A non-template function may be a friend to all template class instantiations.
- *Many-to-one*: All instantiations of a template function may be friends to a regular non-template class.
- *One-to-one*: A template function instantiated with one set of template arguments may be a friend to one template class instantiated with the same set of template arguments. This is also the relationship between a regular non-template class and a regular non-template friend function.
- *Many-to-many*: All instantiations of a template function may be a friend to all instantiations of the template class.

The following example demonstrates these relationships:

```

class B{
    template<class V> friend int j();
}

template<class S> g();

template<class T> class A {
    friend int e();
    friend int f(T);
    friend int g<T>();
    template<class U> friend int h();
};

```

- Function `e()` has a one-to-many relationship with class A. Function `e()` is a friend to all instantiations of class A.
- Function `f()` has a one-to-one relationship with class A. The compiler will give you a warning for this kind of declaration similar to the following:

The friend function declaration "f" will cause an error when the enclosing template class is instantiated with arguments that declare a friend function that does not match an existing definition. The function declares only one function because it is not a template but the function type depends on one or more template parameters.

- Function `g()` has a one-to-one relationship with class A. Function `g()` is a function template. It must be declared before here or else the compiler will not recognize `g<T>` as a template name. For each instantiation of A there is one matching instantiation of `g()`. For example, `g<int>` is a friend of `A<int>`.
- Function `h()` has a many-to-many relationship with class A. Function `h()` is a function template. For all instantiations of A all instantiations of `h()` are friends.
- Function `j()` has a many-to-one relationship with class B.

These relationships also apply to friend classes.

Related information

- “Friends” on page 286

Function templates

A *function template* defines how a group of functions can be generated.

A non-template function is not related to a function template, even though the non-template function may have the same name and parameter profile as those of a specialization generated from a template. A non-template function is never considered to be a specialization of a function template.

The following example implements the quicksort algorithm with a function template named `quicksort`:

```
#include <iostream>
#include <cstdlib>
using namespace std;

template<class T> void quicksort(T a[], const int& leftarg, const int& rightarg)
{
    if (leftarg < rightarg) {

        T pivotvalue = a[leftarg];
        int left = leftarg - 1;
        int right = rightarg + 1;

        for(;;) {

            while (a[--right] > pivotvalue);
            while (a[++left] < pivotvalue);

            if (left >= right) break;

            T temp = a[right];
            a[right] = a[left];
            a[left] = temp;
        }

        int pivot = right;
        quicksort(a, leftarg, pivot);
        quicksort(a, pivot + 1, rightarg);
    }
}

int main(void) {
    int sortme[10];

    for (int i = 0; i < 10; i++) {
        sortme[i] = rand();
        cout << sortme[i] << " ";
    };
    cout << endl;
}
```

```

quicksort<int>(sortme, 0, 10 - 1);

for (int i = 0; i < 10; i++) cout << sortme[i] << "
";
cout << endl;
return 0;
}

```

The above example will have output similar to the following:

```

16838 5758 10113 17515 31051 5627 23010 7419 16212 4086
4086 5627 5758 7419 10113 16212 16838 17515 23010 31051

```

This quicksort algorithm will sort an array of type `T` (whose relational and assignment operators have been defined). The template function takes one template argument and three function arguments:

- the type of the array to be sorted, `T`
- the name of the array to be sorted, `a`
- the lower bound of the array, `leftarg`
- the upper bound of the array, `rightarg`

In the above example, you can also call the `quicksort()` template function with the following statement:

```
quicksort(sortme, 0, 10 - 1);
```

You may omit any template argument if the compiler can deduce it by the usage and context of the template function call. In this case, the compiler deduces that `sortme` is an array of type `int`.

Template argument deduction

When you call a template function, you may omit any template argument that the compiler can determine or *deduce* by the usage and context of that template function call.

The compiler tries to deduce a template argument by comparing the type of the corresponding template parameter with the type of the argument used in the function call. The two types that the compiler compares (the template parameter and the argument used in the function call) must be of a certain structure in order for template argument deduction to work. The following lists these type structures:

```

T
const T
volatile T
T&
T*
T[10]
A<T>
C(*) (T)
T(*) ()
T(*) (U)
T C::*
C T::*
T U::*
T (C::*) ()
C (T::*) ()
D (C::*) (T)
C (T::*) (U)
T (C::*) (U)
T (U::*) ()
T (U::*) (V)

```

```
E[10][i]
B<i>
TT<T>
TT<i>
TT<C>
```

- T, U, and V represent a template type argument
- 10 represents any integer constant
- i represents a template non-type argument
- [i] represents an array bound of a reference or pointer type, or a non-major array bound of a normal array.
- TT represents a template template argument
- (T), (U), and (V) represents an argument list that has at least one template type argument
- () represents an argument list that has no template arguments
- <T> represents a template argument list that has at least one template type argument
- <i> represents a template argument list that has at least one template non-type argument
- <C> represents a template argument list that has no template arguments dependent on a template parameter

The following example demonstrates the use of each of these type structures. The example declares a template function using each of the above structures as an argument. These functions are then called (without template arguments) in order of declaration. The example outputs the same list of type structures:

```
#include <iostream>
using namespace std;

template<class T> class A { };
template<int i> class B { };

class C {
public:
    int x;
};

class D {
public:
    C y;
    int z;
};

template<class T> void f (T)           { cout << "T" << endl; };
template<class T> void f1(const T)    { cout << "const T" << endl; };
template<class T> void f2(volatile T) { cout << "volatile T" << endl; };
template<class T> void g (T*)        { cout << "T*" << endl; };
template<class T> void g (T&)        { cout << "T&" << endl; };
template<class T> void g1(T[10])     { cout << "T[10]" << endl; };
template<class T> void h1(A<T>)      { cout << "A<T>" << endl; };

void test_1() {
    A<char> a;
    C c;

    f(c);   f1(c);   f2(c);
    g(c);   g(&c);   g1(&c);
    h1(a);
}

template<class T> void j(C(*) (T)) { cout << "C(*) (T)" << endl; };
```

```

template<class T>          void j(T*()) { cout << "T* ()" << endl; }
template<class T, class U> void j(T*(U)) { cout << "T* (U)" << endl; };

void test_2() {
    C (*c_pfunct1)(int);
    C (*c_pfunct2)(void);
    int (*c_pfunct3)(int);
    j(c_pfunct1);
    j(c_pfunct2);
    j(c_pfunct3);
}

template<class T>          void k(T C::*) { cout << "T C::*" << endl; };
template<class T>          void k(C T::*) { cout << "C T::*" << endl; };
template<class T, class U> void k(T U::*) { cout << "T U::*" << endl; };

void test_3() {
    k(&C::x);
    k(&D::y);
    k(&D::z);
}

template<class T>          void m(T (C::*)() )
    { cout << "T (C::*)()" << endl; };
template<class T>          void m(C (T::*)() )
    { cout << "C (T::*)()" << endl; };
template<class T>          void m(D (C::*)(T))
    { cout << "D (C::*)(T)" << endl; };
template<class T, class U> void m(C (T::*)(U))
    { cout << "C (T::*)(U)" << endl; };
template<class T, class U> void m(T (C::*)(U))
    { cout << "T (C::*)(U)" << endl; };
template<class T, class U> void m(T (U::*)() )
    { cout << "T (U::*)()" << endl; };
template<class T, class U, class V> void m(T (U::*)(V))
    { cout << "T (U::*)(V)" << endl; };

void test_4() {
    int (C::*f_membp1)(void);
    C (D::*f_membp2)(void);
    D (C::*f_membp3)(int);
    m(f_membp1);
    m(f_membp2);
    m(f_membp3);

    C (D::*f_membp4)(int);
    int (C::*f_membp5)(int);
    int (D::*f_membp6)(void);
    m(f_membp4);
    m(f_membp5);
    m(f_membp6);

    int (D::*f_membp7)(int);
    m(f_membp7);
}

template<int i> void n(C[10][i]) { cout << "E[10][i]" << endl; };
template<int i> void n(B<i>) { cout << "B<i>" << endl; };

void test_5() {
    C array[10][20];
    n(array);
    B<20> b;
    n(b);
}

template<template<class> class TT, class T> void p1(TT<T>)

```

```

    { cout << "TT<T>" << endl; };
template<template<int> class TT, int i>      void p2(TT<i>)
    { cout << "TT<i>" << endl; };
template<template<class> class TT>         void p3(TT<C>)
    { cout << "TT<C>" << endl; };

void test_6() {
    A<char> a;
    B<20> b;
    A<C> c;
    p1(a);
    p2(b);
    p3(c);
}

int main() { test_1(); test_2(); test_3(); test_4(); test_5(); test_6(); }

```

Deducing type template arguments

The compiler can deduce template arguments from a type composed of several of the listed type structures. The following example demonstrates template argument deduction for a type composed of several type structures:

```

template<class T> class Y { };

template<class T, int i> class X {
public:
    Y<T> f(char[20][i]) { return x; };
    Y<T> x;
};

template<template<class> class T, class U, class V, class W, int i>
    void g( T<U> (V::*)(W[20][i]) ) { };

int main()
{
    Y<int> (X<int, 20>::*p)(char[20][20]) = &X<int, 20>::f;
    g(p);
}

```

The type `Y<int> (X<int, 20>::*p)(char[20][20])T<U> (V::*)(W[20][i])` is based on the type structure `T (U::*)(V)`:

- T is `Y<int>`
- U is `X<int, 20>`
- V is `char[20][20]`

If you qualify a type with the class to which that type belongs, and that class (a *nested name specifier*) depends on a template parameter, the compiler will not deduce a template argument for that parameter. If a type contains a template argument that cannot be deduced for this reason, all template arguments in that type will not be deduced. The following example demonstrates this:

```

template<class T, class U, class V>
    void h(typename Y<T>::template Z<U>, Y<T>, Y<V>) { };

int main() {
    Y<int>::Z<char> a;
    Y<int> b;
    Y<float> c;

    h<int, char, float>(a, b, c);
    h<int, char>(a, b, c);
    // h<int>(a, b, c);
}

```

The compiler will not deduce the template arguments T and U in typename Y<T>::template Z<U> (but it will deduce the T in Y<T>). The compiler would not allow the template function call h<int>(a, b, c) because U is not deduced by the compiler.

The compiler can deduce a function template argument from a pointer to function or pointer to member function argument given several overloaded function names. However, none of the overloaded functions may be function templates, nor can more than one overloaded function match the required type. The following example demonstrates this:

```
template<class T> void f(void(*) (T,int)) { };

template<class T> void g1(T, int) { };

void g2(int, int) { };
void g2(char, int) { };

void g3(int, int, int) { };
void g3(float, int) { };

int main() {
//  f(&g1);
//  f(&g2);
    f(&g3);
}
```

The compiler would not allow the call f(&g1) because g1() is a function template. The compiler would not allow the call f(&g2) because both functions named g2() match the type required by f().

The compiler cannot deduce a template argument from the type of a default argument. The following example demonstrates this:

```
template<class T> void f(T = 2, T = 3) { };

int main() {
    f(6);
//  f();
    f<int>();
}
```

The compiler allows the call f(6) because the compiler deduces the template argument (int) by the value of the function call's argument. The compiler would not allow the call f() because the compiler cannot deduce template argument from the default arguments of f().

The compiler cannot deduce a template type argument from the type of a non-type template argument. For example, the compiler will not allow the following:

```
template<class T, T i> void f(int[20][i]) { };

int main() {
    int a[20][30];
    f(a);
}
```

The compiler cannot deduce the type of template parameter T.

Deducing non-type template arguments

The compiler cannot deduce the value of a major array bound unless the bound refers to a reference or pointer type. Major array bounds are not part of function parameter types. The following code demonstrates this:

```

template<int i> void f(int a[10][i]) { };
template<int i> void g(int a[i]) { };
template<int i> void h(int (&a)[i]) { };

int main () {
    int b[10][20];
    int c[10];
    f(b);
    // g(c);
    h(c);
}

```

The compiler would not allow the call `g(c)`; the compiler cannot deduce template argument `i`.

The compiler cannot deduce the value of a non-type template argument used in an expression in the template function's parameter list. The following example demonstrates this:

```

template<int i> class X { };

template<int i> void f(X<i - 1>) { };

int main () {
    X<0> a;
    f<1>(a);
    // f(a);
}

```

In order to call function `f()` with object `a`, the function must accept an argument of type `X<0>`. However, the compiler cannot deduce that the template argument `i` must be equal to 1 in order for the function template argument type `X<i - 1>` to be equivalent to `X<0>`. Therefore the compiler would not allow the function call `f(a)`.

If you want the compiler to deduce a non-type template argument, the type of the parameter must match exactly the type of value used in the function call. For example, the compiler will not allow the following:

```

template<int i> class A { };
template<short d> void f(A<d>) { };

int main() {
    A<1> a;
    f(a);
}

```

The compiler will not convert `int` to `short` when the example calls `f()`.

However, deduced array bounds may be of any integral type.

Overloading function templates

You may overload a function template either by a non-template function or by another function template.

If you call the name of an overloaded function template, the compiler will try to deduce its template arguments and check its explicitly declared template arguments. If successful, it will instantiate a function template specialization, then add this specialization to the set of *candidate functions* used in overload resolution. The compiler proceeds with overload resolution, choosing the most appropriate function from the set of candidate functions. Non-template functions take precedence over template functions. The following example describes this:

```

#include <iostream>
using namespace std;

template<class T> void f(T x, T y) { cout << "Template" << endl; }

void f(int w, int z) { cout << "Non-template" << endl; }

int main() {
    f( 1, 2 );
    f('a', 'b');
    f( 1, 'b');
}

```

The following is the output of the above example:

```

Non-template
Template
Non-template

```

The function call `f(1, 2)` could match the argument types of both the template function and the non-template function. The non-template function is called because a non-template function takes precedence in overload resolution.

The function call `f('a', 'b')` can only match the argument types of the template function. The template function is called.

Argument deduction fails for the function call `f(1, 'b')`; the compiler does not generate any template function specialization and overload resolution does not take place. The non-template function resolves this function call after using the standard conversion from `char` to `int` for the function argument `'b'`.

Related information

- “Overload resolution” on page 255

Partial ordering of function templates

A function template specialization might be ambiguous because template argument deduction might associate the specialization with more than one of the overloaded definitions. The compiler will then choose the definition that is the most specialized. This process of selecting a function template definition is called *partial ordering*.

A template `X` is more specialized than a template `Y` if every argument list that matches the one specified by `X` also matches the one specified by `Y`, but not the other way around. The following example demonstrates partial ordering:

```

template<class T> void f(T) { }
template<class T> void f(T*) { }
template<class T> void f(const T*) { }

template<class T> void g(T) { }
template<class T> void g(T&) { }

template<class T> void h(T) { }
template<class T> void h(T, ...) { }

int main() {
    const int *p;
    f(p);
}

```

```

    int q;
    // g(q);
    // h(q);
}

```

The declaration `template<class T> void f(const T*)` is more specialized than `template<class T> void f(T*)`. Therefore, the function call `f(p)` calls `template<class T> void f(const T*)`. However, neither `void g(T)` nor `void g(T&)` is more specialized than the other. Therefore, the function call `g(q)` would be ambiguous.

Ellipses do not affect partial ordering. Therefore, the function call `h(q)` would also be ambiguous.

The compiler uses partial ordering in the following cases:

- Calling a function template specialization that requires overload resolution.
- Taking the address of a function template specialization.
- When a friend function declaration, an explicit instantiation, or explicit specialization refers to a function template specialization.
- Determining the appropriate deallocation function that is also a function template for a given placement operator `new`.

Related information

- “Template specialization” on page 359
- “new expressions (C++ only)” on page 165

Template instantiation

The act of creating a new definition of a function, class, or member of a class from a template declaration and one or more template arguments is called *template instantiation*. The definition created from a template instantiation to handle a specific set of template arguments is called a *specialization*.

IBM extension

A forward declaration of a template instantiation has the form of an explicit template instantiation preceded by the `extern` keyword.

End of IBM extension

Template instantiation declaration syntax

►►—`extern—template—template_declaration`—◄◄

The language feature is an orthogonal extension to Standard C++ for compatibility with GNU C++, and is described further in “Explicit instantiation” on page 358.

Related information

- “Template specialization” on page 359

Implicit instantiation

Unless a template specialization has been explicitly instantiated or explicitly specialized, the compiler will generate a specialization for the template only when it needs the definition. This is called *implicit instantiation*.

If the compiler must instantiate a class template specialization and the template is declared, you must also define the template.

For example, if you declare a pointer to a class, the definition of that class is not needed and the class will not be implicitly instantiated. The following example demonstrates when the compiler instantiates a template class:

```
template<class T> class X {
public:
    X* p;
    void f();
    void g();
};

X<int>* q;
X<int> r;
X<float>* s;
r.f();
s->g();
```

The compiler requires the instantiation of the following classes and functions:

- `X<int>` when the object `r` is declared
- `X<int>::f()` at the member function call `r.f()`
- `X<float>` and `X<float>::g()` at the class member access function call `s->g()`

Therefore, the functions `X<T>::f()` and `X<T>::g()` must be defined in order for the above example to compile. (The compiler will use the default constructor of class `X` when it creates object `r`.) The compiler does not require the instantiation of the following definitions:

- class `X` when the pointer `p` is declared
- `X<int>` when the pointer `q` is declared
- `X<float>` when the pointer `s` is declared

The compiler will implicitly instantiate a class template specialization if it is involved in pointer conversion or pointer to member conversion. The following example demonstrates this:

```
template<class T> class B { };
template<class T> class D : public B<T> { };

void g(D<double>* p, D<int>* q)
{
    B<double>* r = p;
    delete q;
}
```

The assignment `B<double>* r = p` converts `p` of type `D<double>*` to a type of `B<double>*`; the compiler must instantiate `D<double>`. The compiler must instantiate `D<int>` when it tries to delete `q`.

If the compiler implicitly instantiates a class template that contains static members, those static members are not implicitly instantiated. The compiler will instantiate a static member only when the compiler needs the static member's definition. Every instantiated class template specialization has its own copy of static members. The following example demonstrates this:

```
template<class T> class X {
public:
    static T v;
};

template<class T> T X<T>::v = 0;
```

```
X<char*> a;
X<float> b;
X<float> c;
```

Object a has a static member variable v of type char*. Object b has a static variable v of type float. Objects b and c share the single static data member v.

An implicitly instantiated template is in the same namespace where you defined the template.

If a function template or a member function template specialization is involved with overload resolution, the compiler implicitly instantiates a declaration of the specialization.

Related information

- “Template instantiation” on page 356

Explicit instantiation

You can explicitly tell the compiler when it should generate a definition from a template. This is called *explicit instantiation*.

Explicit instantiation declaration syntax

►—template—*template_declaration*—◄

The following are examples of explicit instantiations:

```
template<class T> class Array { void mf(); };
template class Array<char>; // explicit instantiation
template void Array<int>::mf(); // explicit instantiation

template<class T> void sort(Array<T>& v) { }
template void sort(Array<char>&); // explicit instantiation

namespace N {
    template<class T> void f(T&) { }
}

template void N::f<int>(int&);
// The explicit instantiation is in namespace N.

int* p = 0;
template<class T> T g(T = &p);
template char g(char); // explicit instantiation

template <class T> class X {
    private:
        T v(T arg) { return arg; };
};

template int X<int>::v(int); // explicit instantiation

template<class T> T g(T val) { return val;}
template<class T> void Array<T>::mf() { }
```

A template declaration must be in scope at the point of instantiation of the template’s explicit instantiation. An explicit instantiation of a template specialization is in the same namespace where you defined the template.

Access checking rules do not apply to names in explicit instantiations. Template arguments and names in a declaration of an explicit instantiation may be private types or objects. In the above example, the compiler allows the explicit instantiation `template int X<int>::v(int)` even though the member function is declared private.

The compiler does not use default arguments when you explicitly instantiate a template. In the above example the compiler allows the explicit instantiation `template char g(char)` even though the default argument is an address of type `int`.

IBM extension

An extern-qualified template declaration does not instantiate the class or function. For both classes and functions, the extern template instantiation prevents the instantiation of parts of the template, provided that the instantiation has not already been triggered by code prior to the extern template instantiation. For classes, the members (both static and nonstatic) are not instantiated. The class itself is instantiated if required to map the class. For functions, the prototype is instantiated, but the body of the template function is not instantiated.

The following examples show template instantiation using extern:

```
template<class T>class C {
    static int i;
    void f(T) { }
};
template<class U>int C<U>::i = 0;
extern template C<int>; // extern explicit template instantiation
C<int>c; // does not cause instantiation of C<int>::i
        // or C<int>::f(int) in this file,
        // but the class is instantiated for mapping
C<char>d; // normal instantiations
template<class C> C foo(C c) { return c; }
extern template int foo<int>(int); // extern explicit template instantiation
int i = foo(1); // does not cause instantiation of the body of foo<int>
```

End of IBM extension

Template specialization

The act of creating a new definition of a function, class, or member of a class from a template declaration and one or more template arguments is called *template instantiation*. The definition created from a template instantiation is called a *specialization*. A *primary template* is the template that is being specialized.

Related information

- “Template instantiation” on page 356

Explicit specialization

When you instantiate a template with a given set of template arguments the compiler generates a new definition based on those template arguments. You can override this behavior of definition generation. You can instead specify the definition the compiler uses for a given set of template arguments. This is called *explicit specialization*. You can explicitly specialize any of the following:

- Function or class template

- Member function of a class template
- Static data member of a class template
- Member class of a class template
- Member function template of a class template
- Member class template of a class template

Explicit specialization declaration syntax

►—template—<—>—*declaration_name*—<—*template_argument_list*—>—*declaration_body*—►

The `template<>` prefix indicates that the following template declaration takes no template parameters. The *declaration_name* is the name of a previously declared template. Note that you can forward-declare an explicit specialization so the *declaration_body* is optional, at least until the specialization is referenced.

The following example demonstrates explicit specialization:

```
using namespace std;

template<class T = float, int i = 5> class A
{
    public:
        A();
        int value;
};

template<> class A<> { public: A(); };
template<> class A<double, 10> { public: A(); };

template<class T, int i> A<T, i>::A() : value(i) {
    cout << "Primary template, "
         << "non-type argument is " << value << endl;
}

A<>::A() {
    cout << "Explicit specialization "
         << "default arguments" << endl;
}

A<double, 10>::A() {
    cout << "Explicit specialization "
         << "<double, 10>" << endl;
}

int main() {
    A<int,6> x;
    A<> y;
    A<double, 10> z;
}
```

The following is the output of the above example:

```
Primary template non-type argument is: 6
Explicit specialization default arguments
Explicit specialization <double, 10>
```

This example declared two explicit specializations for the *primary template* (the template which is being specialized) class A. Object x uses the constructor of the primary template. Object y uses the explicit specialization `A<>::A()`. Object z uses the explicit specialization `A<double, 10>::A()`.

Related information

- “Function templates” on page 348
- “Class templates” on page 344
- “Member functions of class templates” on page 346
- “Static data members and templates” on page 346

Definition and declaration of explicit specializations

The definition of an explicitly specialized class is unrelated to the definition of the primary template. You do not have to define the primary template in order to define the specialization (nor do you have to define the specialization in order to define the primary template). For example, the compiler will allow the following:

```
template<class T> class A;
template<> class A<int>;

template<> class A<int> { /* ... */ };
```

The primary template is not defined, but the explicit specialization is.

You can use the name of an explicit specialization that has been declared but not defined the same way as an incompletely defined class. The following example demonstrates this:

```
template<class T> class X { };
template<> class X<char>;
X<char>* p;
X<int> i;
// X<char> j;
```

The compiler does not allow the declaration `X<char> j` because the explicit specialization of `X<char>` is not defined.

Explicit specialization and scope

A declaration of a primary template must be in scope at the *point of declaration* of the explicit specialization. In other words, an explicit specialization declaration must appear after the declaration of the primary template. For example, the compiler will not allow the following:

```
template<> class A<int>;
template<class T> class A;
```

An explicit specialization is in the same namespace as the definition of the primary template.

Class members of explicit specializations

A member of an explicitly specialized class is not implicitly instantiated from the member declaration of the primary template. You have to explicitly define members of a class template specialization. You define members of an explicitly specialized template class as you would normal classes, without the `template<>` prefix. In addition, you can define the members of an explicit specialization inline; no special template syntax is used in this case. The following example demonstrates a class template specialization:

```
template<class T> class A {
    public:
        void f(T);
};

template<> class A<int> {
    public:
        int g(int);
};
```

```

};

int A<int>::g(int arg) { return 0; }

int main() {
    A<int> a;
    a.g(1234);
}

```

The explicit specialization `A<int>` contains the member function `g()`, which the primary template does not.

If you explicitly specialize a template, a member template, or the member of a class template, then you must declare this specialization before that specialization is implicitly instantiated. For example, the compiler will not allow the following code:

```

template<class T> class A { };

void f() { A<int> x; }
template<> class A<int> { };

int main() { f(); }

```

The compiler will not allow the explicit specialization `template<> class A<int> { };` because function `f()` uses this specialization (in the construction of `x`) before the specialization.

Explicit specialization of function templates

In a function template specialization, a template argument is optional if the compiler can deduce it from the type of the function arguments. The following example demonstrates this:

```

template<class T> class X { };
template<class T> void f(X<T>);
template<> void f(X<int>);

```

The explicit specialization `template<> void f(X<int>)` is equivalent to `template<> void f<int>(X<int>)`.

You cannot specify default function arguments in a declaration or a definition for any of the following:

- Explicit specialization of a function template
- Explicit specialization of a member function template

For example, the compiler will not allow the following code:

```

template<class T> void f(T a) { };
template<> void f<int>(int a = 5) { };

template<class T> class X {
    void f(T a) { }
};
template<> void X<int>::f(int a = 10) { };

```

Related information

- “Function templates” on page 348

Explicit specialization of members of class templates

Each instantiated class template specialization has its own copy of any static members. You may explicitly specialize static members. The following example demonstrates this:

```

template<class T> class X {
public:
    static T v;
    static void f(T);
};

template<class T> T X<T>::v = 0;
template<class T> void X<T>::f(T arg) { v = arg; }

template<> char* X<char*>::v = "Hello";
template<> void X<float>::f(float arg) { v = arg * 2; }

int main() {
    X<char*> a, b;
    X<float> c;
    c.f(10);
}

```

This code explicitly specializes the initialization of static data member `X::v` to point to the string "Hello" for the template argument `char*`. The function `X::f()` is explicitly specialized for the template argument `float`. The static data member `v` in objects `a` and `b` point to the same string, "Hello". The value of `c.v` is equal to 20 after the call function call `c.f(10)`.

You can nest member templates within many enclosing class templates. If you explicitly specialize a template nested within several enclosing class templates, you must prefix the declaration with `template<>` for every enclosing class template you specialize. You may leave some enclosing class templates unspecialized, however you cannot explicitly specialize a class template unless its enclosing class templates are also explicitly specialized. The following example demonstrates explicit specialization of nested member templates:

```

#include <iostream>
using namespace std;

template<class T> class X {
public:
    template<class U> class Y {
    public:
        template<class V> void f(U,V);
        void g(U);
    };
};

template<class T> template<class U> template<class V>
void X<T>::Y<U>::f(U, V) { cout << "Template 1" << endl; }

template<class T> template<class U>
void X<T>::Y<U>::g(U) { cout << "Template 2" << endl; }

template<> template<>
void X<int>::Y<int>::g(int) { cout << "Template 3" << endl; }

template<> template<> template<class V>
void X<int>::Y<int>::f(int, V) { cout << "Template 4" << endl; }

template<> template<> template<>
void X<int>::Y<int>::f<int>(int, int) { cout << "Template 5" << endl; }

// template<> template<class U> template<class V>
// void X<char>::Y<U>::f(U, V) { cout << "Template 6" << endl; }

// template<class T> template<>
// void X<T>::Y<float>::g(float) { cout << "Template 7" << endl; }

```

```

int main() {
    X<int>::Y<int> a;
    X<char>::Y<char> b;
    a.f(1, 2);
    a.f(3, 'x');
    a.g(3);
    b.f('x', 'y');
    b.g('z');
}

```

The following is the output of the above program:

```

Template 5
Template 4
Template 3
Template 1
Template 2

```

- The compiler would not allow the template specialization definition that would output "Template 6" because it is attempting to specialize a member (function `f()`) without specialization of its containing class (`Y`).
- The compiler would not allow the template specialization definition that would output "Template 7" because the enclosing class of class `Y` (which is class `X`) is not explicitly specialized.

A friend declaration cannot declare an explicit specialization.

Related information

- "Static data members and templates" on page 346

Partial specialization

When you instantiate a class template, the compiler creates a definition based on the template arguments you have passed. Alternatively, if all those template arguments match those of an explicit specialization, the compiler uses the definition defined by the explicit specialization.

A *partial specialization* is a generalization of explicit specialization. An explicit specialization only has a template argument list. A partial specialization has both a template argument list and a template parameter list. The compiler uses the partial specialization if its template argument list matches a subset of the template arguments of a template instantiation. The compiler will then generate a new definition from the partial specialization with the rest of the unmatched template arguments of the template instantiation.

You cannot partially specialize function templates.

Partial specialization syntax

```

▶▶—template—<template_parameter_list>—declaration_name—————▶
▶—<template_argument_list>—declaration_body—————▶▶

```

The *declaration_name* is a name of a previously declared template. Note that you can forward-declare a partial specialization so that the *declaration_body* is optional.

The following demonstrates the use of partial specializations:

```

#include <iostream>
using namespace std;

```

```

template<class T, class U, int I> struct X
    { void f() { cout << "Primary template" << endl; } };

template<class T, int I> struct X<T, T*, I>
    { void f() { cout << "Partial specialization 1" << endl;
    } };

template<class T, class U, int I> struct X<T*, U, I>
    { void f() { cout << "Partial specialization 2" << endl;
    } };

template<class T> struct X<int, T*, 10>
    { void f() { cout << "Partial specialization 3" << endl;
    } };

template<class T, class U, int I> struct X<T, U*, I>
    { void f() { cout << "Partial specialization 4" << endl;
    } };

int main() {
    X<int, int, 10> a;
    X<int, int*, 5> b;
    X<int*, float, 10> c;
    X<int, char*, 10> d;
    X<float, int*, 10> e;
    // X<int, int*, 10> f;
    a.f(); b.f(); c.f(); d.f(); e.f();
}

```

The following is the output of the above example:

```

Primary template
Partial specialization 1
Partial specialization 2
Partial specialization 3
Partial specialization 4

```

The compiler would not allow the declaration `X<int, int*, 10> f` because it can match template struct `X<T, T*, I>`, template struct `X<int, T*, 10>`, or template struct `X<T, U*, I>`, and none of these declarations are a better match than the others.

Each class template partial specialization is a separate template. You must provide definitions for each member of a class template partial specialization.

Template parameter and argument lists of partial specializations

Primary templates do not have template argument lists; this list is implied in the template parameter list.

Template parameters specified in a primary template but not used in a partial specialization are omitted from the template parameter list of the partial specialization. The order of a partial specialization's argument list is the same as the order of the primary template's implied argument list.

In a template argument list of a partial template parameter, you cannot have an expression that involves non-type arguments unless that expression is only an identifier. In the following example, the compiler will not allow the first partial specialization, but will allow the second one:

```

template<int I, int J> class X { };

// Invalid partial specialization

```

```
template<int I> class X <I * 4, I + 3> { };

// Valid partial specialization
template <int I> class X <I, I> { };
```

The type of a non-type template argument cannot depend on a template parameter of a partial specialization. The compiler will not allow the following partial specialization:

```
template<class T, T i> class X { };

// Invalid partial specialization
template<class T> class X<T, 25> { };
```

A partial specialization's template argument list cannot be the same as the list implied by the primary template.

You cannot have default values in the template parameter list of a partial specialization.

Related information

- "Template parameters" on page 338
- "Template arguments" on page 340

Matching of class template partial specializations

The compiler determines whether to use the primary template or one of its partial specializations by matching the template arguments of the class template specialization with the template argument lists of the primary template and the partial specializations:

- If the compiler finds only one specialization, then the compiler generates a definition from that specialization.
- If the compiler finds more than one specialization, then the compiler tries to determine which of the specializations is the most specialized. A template *X* is more specialized than a template *Y* if every argument list that matches the one specified by *X* also matches the one specified by *Y*, but not the other way around. If the compiler cannot find the most specialized specialization, then the use of the class template is ambiguous; the compiler will not allow the program.
- If the compiler does not find any matches, then the compiler generates a definition from the primary template.

Name binding and dependent names

Name binding is the process of finding the declaration for each name that is explicitly or implicitly used in a template. The compiler may bind a name in the definition of a template, or it may bind a name at the instantiation of a template.

A *dependent name* is a name that depends on the type or the value of a template parameter. For example:

```
template<class T> class U : A<T>
{
    typename T::B x;
    void f(A<T>& y)
    {
        *y++;
    }
};
```

The dependent names in this example are the base class `A<T>`, the type name `T::B`, and the variable `y`.

The compiler binds dependent names when a template is instantiated. The compiler binds non-dependent names when a template is defined. For example:

```
void f(double) { cout << "Function f(double)" << endl; }

template<class T> void g(T a) {
    f(123);
    h(a);
}

void f(int) { cout << "Function f(int)" << endl; }
void h(double) { cout << "Function h(double)" << endl; }

void i() {
    extern void h(int);
    g<int>(234);
}

void h(int) { cout << "Function h(int)" << endl; }
```

The following is the output if you call function `i()`:

```
Function f(double)
Function h(double)
```

The *point of definition* of a template is located immediately before its definition. In this example, the point of definition of the function template `g(T)` is located immediately before the keyword `template`. Because the function call `f(123)` does not depend on a template argument, the compiler will consider names declared before the definition of function template `g(T)`. Therefore, the call `f(123)` will call `f(double)`. Although `f(int)` is a better match, it is not in scope at the point of definition of `g(T)`.

The *point of instantiation* of a template is located immediately before the declaration that encloses its use. In this example, the point of instantiation of the call to `g<int>(234)` is located immediately before `i()`. Because the function call `h(a)` depends on a template argument, the compiler will consider names declared before the instantiation of function template `g(T)`. Therefore, the call `h(a)` will call `h(double)`. It will not consider `h(int)`, because this function was not in scope at the point of instantiation of `g<int>(234)`.

Point of instantiation binding implies the following:

- A template parameter cannot depend on any local name or class member.
- An unqualified name in a template cannot depend on a local name or class member.

Related information

- “Template instantiation” on page 356

The typename keyword

Use the keyword `typename` if you have a qualified name that refers to a type and depends on a template parameter. Only use the keyword `typename` in template declarations and definitions. The following example illustrates the use of the keyword `typename`:

```

template<class T> class A
{
    T::x(y);
    typedef char C;
    A::C d;
}

```

The statement `T::x(y)` is ambiguous. It could be a call to function `x()` with a nonlocal argument `y`, or it could be a declaration of variable `y` with type `T::x`. C++ will interpret this statement as a function call. In order for the compiler to interpret this statement as a declaration, you would add the keyword `typename` to the beginning of it. The statement `A::C d;` is ill-formed. The class `A` also refers to `A<T>` and thus depends on a template parameter. You must add the keyword `typename` to the beginning of this declaration:

```

typename A::C d;

```

You can also use the keyword `typename` in place of the keyword `class` in template parameter declarations.

Related information

- “Template parameters” on page 338

The template keyword as qualifier

Use the keyword `template` as a qualifier to distinguish member templates from other names. The following example illustrates when you must use `template` as a qualifier:

```

class A
{
public:
    template<class T> T function_m() { };
};

template<class U> void function_n(U argument)
{
    char object_x = argument.function_m<char>();
}

```

The declaration `char object_x = argument.function_m<char>();` is ill-formed. The compiler assumes that the `<` is a less-than operator. In order for the compiler to recognize the function template call, you must add the `template` quantifier:

```

char object_x = argument.template function_m<char>();

```

If the name of a member template specialization appears after a `.`, `->`, or `::` operator, and that name has explicitly qualified template parameters, prefix the member template name with the keyword `template`. The following example demonstrates this use of the keyword `template`:

```

#include <iostream>
using namespace std;

class X {
public:
    template <int j> struct S {
        void h() {
            cout << "member template's member function: " << j << endl;
        }
    };
    template <int i> void f() {
        cout << "Primary: " << i << endl;
    }
};

```

```

    }
};

template<> void X::f<20>() {
    cout << "Specialized, non-type argument = 20" << endl;
}

template<class T> void g(T* p) {
    p->template f<100>();
    p->template f<20>();
    typename T::template S<40> s; // use of scope operator on a member template
    s.h();
}

int main()
{
    X temp;
    g(&temp);
}

```

The following is the output of the above example:

```

Primary: 100
Specialized, non-type argument = 20
member template's member function: 40

```

If you do not use the keyword `template` in these cases, the compiler will interpret the `<` as a less-than operator. For example, the following line of code is ill-formed:

```
p->f<100>();
```

The compiler interprets `f` as a non-template member, and the `<` as a less-than operator.

Chapter 16. Exception handling (C++ only) (PPU only)

Exception handling is a mechanism that separates code that detects and handles exceptional circumstances from the rest of your program. Note that an exceptional circumstance is not necessarily an error.

When a function detects an exceptional situation, you represent this with an object. This object is called an *exception object*. In order to deal with the exceptional situation you *throw the exception*. This passes control, as well as the exception, to a designated block of code in a direct or indirect caller of the function that threw the exception. This block of code is called a *handler*. In a handler, you specify the types of exceptions that it may process. The C++ run time, together with the generated code, will pass control to the first appropriate handler that is able to process the exception thrown. When this happens, an exception is *caught*. A handler may *rethrow* an exception so it can be caught by another handler.

The exception handling mechanism is made up of the following elements:

- try blocks
- catch blocks
- throw expressions
- Exception specifications


try blocks

You use a *try block* to indicate which areas in your program that might throw exceptions you want to handle immediately. You use a *function try block* to indicate that you want to detect exceptions in the entire body of a function.

try block syntax



Function try block syntax



The following is an example of a function try block with a member initializer, a function try block and a try block:

```
#include <iostream>
using namespace std;

class E {
public:
    const char* error;
    E(const char* arg) : error(arg) { }
};
```

```

class A {
public:
    int i;

    // A function try block with a member
    // initializer
    A() try : i(0) {
        throw E("Exception thrown in A()");
    }
    catch (E& e) {
        cout << e.error << endl;
    }
};

// A function try block
void f() try {
    throw E("Exception thrown in f()");
}
catch (E& e) {
    cout << e.error << endl;
}

void g() {
    throw E("Exception thrown in g()");
}

int main() {
    f();

    // A try block
    try {
        g();
    }
    catch (E& e) {
        cout << e.error << endl;
    }
    try {
        A x;
    }
    catch(...) { }
}

```

The following is the output of the above example:

```

Exception thrown in f()
Exception thrown in g()
Exception thrown in A()

```

The constructor of class A has a function try block with a member initializer. Function f() has a function try block. The main() function contains a try block.

Related information

- “Initialization of base classes and members” on page 322

Nested try blocks

When try blocks are nested and a throw occurs in a function called by an inner try block, control is transferred outward through the nested try blocks until the first catch block is found whose argument matches the argument of the throw expression.

For example:

```
try
{
    func1();
    try
    {
        func2();
    }
    catch (spec_err) { /* ... */ }
    func3();
}
catch (type_err) { /* ... */ }
// if no throw is issued, control resumes here.
```

In the above example, if `spec_err` is thrown within the inner try block (in this case, from `func2()`), the exception is caught by the inner catch block, and, assuming this catch block does not transfer control, `func3()` is called. If `spec_err` is thrown after the inner try block (for instance, by `func3()`), it is not caught and the function `terminate()` is called. If the exception thrown from `func2()` in the inner try block is `type_err`, the program skips out of both try blocks to the second catch block without invoking `func3()`, because no appropriate catch block exists following the inner try block.

You can also nest a try block within a catch block.

catch blocks

catch block syntax

►►—catch—(—*exception_declaration*—)—{—*statements*—}—►►

You can declare a handler to catch many types of exceptions. The allowable objects that a function can catch are declared in the parentheses following the catch keyword (the *exception_declaration*). You can catch objects of the fundamental types, base and derived class objects, references, and pointers to all of these types. You can also catch `const` and `volatile` types. The *exception_declaration* cannot be an incomplete type, or a reference or pointer to an incomplete type other than one of the following:

- `void*`
- `const void*`
- `volatile void*`
- `const volatile void*`

You cannot define a type in an *exception_declaration*.

You can also use the `catch(...)` form of the handler to catch all thrown exceptions that have not been caught by a previous catch block. The ellipsis in the catch argument indicates that any exception thrown can be handled by this handler.

If an exception is caught by a `catch(...)` block, there is no direct way to access the object thrown. Information about an exception caught by `catch(...)` is very limited.

You can declare an optional variable name if you want to access the thrown object in the catch block.

A catch block can only catch accessible objects. The object caught must have an accessible copy constructor.

Related information

- “Type qualifiers” on page 72
- “Member access” on page 283

Function try block handlers

The scope and lifetime of the parameters of a function or constructor extend into the handlers of a function try block. The following example demonstrates this:

```
void f(int &x) try {
    throw 10;
}
catch (const int &i)
{
    x = i;
}

int main() {
    int v = 0;
    f(v);
}
```

The value of `v` after `f()` is called is 10.

A function try block on `main()` does not catch exceptions thrown in destructors of objects with static storage duration, or constructors of namespace scope objects.

The following example throws an exception from a destructor of a static object:

```
#include <iostream>
using namespace std;

class E {
public:
    const char* error;
    E(const char* arg) : error(arg) { }
};

class A {
public: ~A() { throw E("Exception in ~A()"); }
};

class B {
public: ~B() { throw E("Exception in ~B()"); }
};

int main() try {
    cout << "In main" << endl;
    static A cow;
    B bull;
}
catch (E& e) {
    cout << e.error << endl;
}
```

The following is the output of the above example:

```
In main
Exception in ~B()
```

The run time will not catch the exception thrown when object `cow` is destroyed at the end of the program.

The following example throws an exception from a constructor of a namespace scope object:

```
#include <iostream>
using namespace std;

class E {
public:
    const char* error;
    E(const char* arg) : error(arg) { }
};

namespace N {
    class C {
    public:
        C() {
            cout << "In C()" << endl;
            throw E("Exception in C()");
        }
    };

    C calf;
};

int main() try {
    cout << "In main" << endl;
}
catch (E& e) {
    cout << e.error << endl;
}
```

The following is the output of the above example:

```
In C()
```

The compiler will not catch the exception thrown when object `calf` is created.

In a function try block's handler, you cannot have a jump into the body of a constructor or destructor.

A return statement cannot appear in a function try block's handler of a constructor.

When the function try block's handler of an object's constructor or destructor is entered, fully constructed base classes and members of that object are destroyed. The following example demonstrates this:

```
#include <iostream>
using namespace std;

class E {
public:
    const char* error;
    E(const char* arg) : error(arg) { };
};

class B {
public:
    B() { };
    ~B() { cout << "~B() called" << endl; };
};

class D : public B {
public:
```

```

        D();
        ~D() { cout << "~D() called" << endl; };
};

D::~D() try : B() {
    throw E("Exception in D()");
}
catch(E& e) {
    cout << "Handler of function try block of D(): " << e.error << endl;
};

int main() {
    try {
        D val;
    }
    catch(...) { }
}

```

The following is the output of the above example:

```

~B() called
Handler of function try block of D(): Exception in D()

```

When the function try block's handler of D() is entered, the run time first calls the destructor of the base class of D, which is B. The destructor of D is not called because val is not fully constructed.

The run time will rethrow an exception at the end of a function try block's handler of a constructor or destructor. All other functions will return once they have reached the end of their function try block's handler. The following example demonstrates this:

```

#include <iostream>
using namespace std;

class E {
public:
    const char* error;
    E(const char* arg) : error(arg) { };
};

class A {
public:
    A() try { throw E("Exception in A()"); }
    catch(E& e) { cout << "Handler in A(): " << e.error << endl; }
};

int f() try {
    throw E("Exception in f()");
    return 0;
}
catch(E& e) {
    cout << "Handler in f(): " << e.error << endl;
    return 1;
}

int main() {
    int i = 0;
    try { A a; }
    catch(E& e) {
        cout << "Handler in main(): " << e.error << endl;
    }

    try { i = f(); }
    catch(E& e) {
        cout << "Another handler in main(): " << e.error << endl;
    }
}

```

```

    }
    cout << "Returned value of f(): " << i << endl;
}

```

The following is the output of the above example:

```

Handler in A(): Exception in A()
Handler in main(): Exception in A()
Handler in f(): Exception in f()
Returned value of f(): 1

```

Related information

- “The main() function” on page 224
- “The static storage class specifier” on page 46
- Chapter 9, “Namespaces (C++ only),” on page 235
- “Destructors” on page 326

Arguments of catch blocks

If you specify a class type for the argument of a catch block (the *exception_declaration*), the compiler uses a copy constructor to initialize that argument. If that argument does not have a name, the compiler initializes a temporary object and destroys it when the handler exits.

The ISO C++ specifications do not require the compiler to construct temporary objects in cases where they are redundant. The compiler takes advantage of this rule to create more efficient, optimized code. Take this into consideration when debugging your programs, especially for memory problems.

Matching between exceptions thrown and caught

An argument in the catch argument of a handler matches an argument in the *assignment_expression* of the throw expression (throw argument) if any of the following conditions is met:

- The catch argument type matches the type of the thrown object.
- The catch argument is a public base class of the thrown class object.
- The catch specifies a pointer type, and the thrown object is a pointer type that can be converted to the pointer type of the catch argument by standard pointer conversion.

Note: If the type of the thrown object is `const` or `volatile`, the catch argument must also be a `const` or `volatile` for a match to occur. However, a `const`, `volatile`, or reference type catch argument can match a nonconstant, nonvolatile, or nonreference object type. A nonreference catch argument type matches a reference to an object of the same type.

Related information

- “Pointer conversions” on page 117
- “Type qualifiers” on page 72
- “References (C++ only)” on page 93

Order of catching

If the compiler encounters an exception in a try block, it will try each handler in order of appearance.

If a catch block for objects of a base class precedes a catch block for objects of a class derived from that base class, the compiler issues a warning and continues to compile the program despite the unreachable code in the derived class handler.

A catch block of the form `catch(...)` must be the last catch block following a try block or an error occurs. This placement ensures that the `catch(...)` block does not prevent more specific catch blocks from catching exceptions intended for them.

If the run time cannot find a matching handler in the current scope, the run time will continue to find a matching handler in a dynamically surrounding try block. The following example demonstrates this:

```
#include <iostream>
using namespace std;

class E {
public:
    const char* error;
    E(const char* arg) : error(arg) { };
};

class F : public E {
public:
    F(const char* arg) : E(arg) { };
};

void f() {
    try {
        cout << "In try block of f()" << endl;
        throw E("Class E exception");
    }
    catch (F& e) {
        cout << "In handler of f()";
        cout << e.error << endl;
    }
};

int main() {
    try {
        cout << "In main" << endl;
        f();
    }
    catch (E& e) {
        cout << "In handler of main: ";
        cout << e.error << endl;
    };
    cout << "Resume execution in main" << endl;
}
```

The following is the output of the above example:

```
In main
In try block of f()
In handler of main: Class E exception
Resume execution in main
```

In function `f()`, the run time could not find a handler to handle the exception of type `E` thrown. The run time finds a matching handler in a dynamically surrounding try block: the try block in the `main()` function.

If the run time cannot find a matching handler in the program, it calls the `terminate()` function.

Related information

- “try blocks” on page 371

throw expressions

You use a *throw expression* to indicate that your program has encountered an exception.

throw expression syntax

```

▶▶ throw _____ ▶▶
      └──assignment_expression──┘
  
```

The type of *assignment_expression* cannot be an incomplete type, or a pointer to an incomplete type other than one of the following:

- void*
- const void*
- volatile void*
- const volatile void*

The *assignment_expression* is treated the same way as a function argument in a call or the operand of a return statement.

If the *assignment_expression* is a class object, the copy constructor and destructor of that object must be accessible. For example, you cannot throw a class object that has its copy constructor declared as private.

Related information

- “Incomplete types” on page 42

Rethrowing an exception

If a catch block cannot handle the particular exception it has caught, you can rethrow the exception. The rethrow expression (`throw` without *assignment_expression*) causes the originally thrown object to be rethrown.

Because the exception has already been caught at the scope in which the rethrow expression occurs, it is rethrown out to the next dynamically enclosing try block. Therefore, it cannot be handled by catch blocks at the scope in which the rethrow expression occurred. Any catch blocks for the dynamically enclosing try block have an opportunity to catch the exception.

The following example demonstrates rethrowing an exception:

```

#include <iostream>
using namespace std;

struct E {
    const char* message;
    E() : message("Class E") { }
};

struct E1 : E {
    const char* message;
    E1() : message("Class E1") { }
};

struct E2 : E {
    const char* message;
    E2() : message("Class E2") { }
}
  
```

```

};

void f() {
    try {
        cout << "In try block of f()" << endl;
        cout << "Throwing exception of type E1" << endl;
        E1 myException;
        throw myException;
    }
    catch (E2& e) {
        cout << "In handler of f(), catch (E2& e)" << endl;
        cout << "Exception: " << e.message << endl;
        throw;
    }
    catch (E1& e) {
        cout << "In handler of f(), catch (E1& e)" << endl;
        cout << "Exception: " << e.message << endl;
        throw;
    }
    catch (E& e) {
        cout << "In handler of f(), catch (E& e)" << endl;
        cout << "Exception: " << e.message << endl;
        throw;
    }
}

int main() {
    try {
        cout << "In try block of main()" << endl;
        f();
    }
    catch (E2& e) {
        cout << "In handler of main(), catch (E2& e)" << endl;
        cout << "Exception: " << e.message << endl;
    }
    catch (...) {
        cout << "In handler of main(), catch (...)" << endl;
    }
}

```

The following is the output of the above example:

```

In try block of main()
In try block of f()
Throwing exception of type E1
In handler of f(), catch (E1& e)
Exception: Class E1
In handler of main(), catch (...)

```

The try block in the main() function calls function f(). The try block in function f() throws an object of type E1 named myException. The handler catch (E1 &e) catches myException. The handler then rethrows myException with the statement throw to the next dynamically enclosing try block: the try block in the main() function. The handler catch(...) catches myException.

Stack unwinding

When an exception is thrown and control passes from a try block to a handler, the C++ run time calls destructors for all automatic objects constructed since the beginning of the try block. This process is called *stack unwinding*. The automatic objects are destroyed in reverse order of their construction. (Automatic objects are local objects that have been declared auto or register, or not declared static or extern. An automatic object x is deleted whenever the program exits the block in which x is declared.)

If an exception is thrown during construction of an object consisting of subobjects or array elements, destructors are only called for those subobjects or array elements successfully constructed before the exception was thrown. A destructor for a local static object will only be called if the object was successfully constructed.

If during stack unwinding a destructor throws an exception and that exception is not handled, the `terminate()` function is called. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct E {
    const char* message;
    E(const char* arg) : message(arg) { }
};

void my_terminate() {
    cout << "Call to my_terminate" << endl;
};

struct A {
    A() { cout << "In constructor of A" << endl; }
    ~A() {
        cout << "In destructor of A" << endl;
        throw E("Exception thrown in ~A()");
    }
};

struct B {
    B() { cout << "In constructor of B" << endl; }
    ~B() { cout << "In destructor of B" << endl; }
};

int main() {
    set_terminate(my_terminate);

    try {
        cout << "In try block" << endl;
        A a;
        B b;
        throw("Exception thrown in try block of main()");
    }
    catch (const char* e) {
        cout << "Exception: " << e << endl;
    }
    catch (...) {
        cout << "Some exception caught in main()" << endl;
    }

    cout << "Resume execution of main()" << endl;
}
```

The following is the output of the above example:

```
In try block
In constructor of A
In constructor of B
In destructor of B
In destructor of A
Call to my_terminate
```

In the try block, two automatic objects are created: `a` and `b`. The try block throws an exception of type `const char*`. The handler `catch (const char* e)` catches this

exception. The C++ run time unwinds the stack, calling the destructors for `a` and `b` in reverse order of their construction. The destructor for `a` throws an exception. Since there is no handler in the program that can handle this exception, the C++ run time calls `terminate()`. (The function `terminate()` calls the function specified as the argument to `set_terminate()`. In this example, `terminate()` has been specified to call `my_terminate()`.)

Exception specifications

C++ provides a mechanism to ensure that a given function is limited to throwing only a specified list of exceptions. An exception specification at the beginning of any function acts as a guarantee to the function's caller that the function will throw only the exceptions contained in the exception specification.

For example, a function:

```
void translate() throw(unknown_word,bad_grammar) { /* ... */ }
```

explicitly states that it will only throw exception objects whose types are `unknown_word` or `bad_grammar`, or any type derived from `unknown_word` or `bad_grammar`.

Exception specification syntax

►► throw (type_id_list) ◀◀

The *type_id_list* is a comma-separated list of types. In this list you cannot specify an incomplete type, a pointer or a reference to an incomplete type, other than a pointer to `void`, optionally qualified with `const` and/or `volatile`. You cannot define a type in an exception specification.

A function with no exception specification allows all exceptions. A function with an exception specification that has an empty *type_id_list*, `throw()`, does not allow any exceptions to be thrown.

An exception specification is not part of a function's type.

An exception specification may only appear at the end of a function declarator of a function, pointer to function, reference to function, pointer to member function declaration, or pointer to member function definition. An exception specification cannot appear in a typedef declaration. The following declarations demonstrate this:

```
void f() throw(int);
void (*g)() throw(int);
void h(void i() throw(int));
// typedef int (*j)() throw(int); This is an error.
```

The compiler would not allow the last declaration, `typedef int (*j)() throw(int)`.

Suppose that class `A` is one of the types in the *type_id_list* of an exception specification of a function. That function may throw exception objects of class `A`, or any class publicly derived from class `A`. The following example demonstrates this:

```
class A { };
class B : public A { };
class C { };
```

```

void f(int i) throw (A) {
    switch (i) {
        case 0: throw A();
        case 1: throw B();
        default: throw C();
    }
}

void g(int i) throw (A*) {
    A* a = new A();
    B* b = new B();
    C* c = new C();
    switch (i) {
        case 0: throw a;
        case 1: throw b;
        default: throw c;
    }
}

```

Function `f()` can throw objects of types `A` or `B`. If the function tries to throw an object of type `C`, the compiler will call `unexpected()` because type `C` has not been specified in the function's exception specification, nor does it derive publicly from `A`. Similarly, function `g()` cannot throw pointers to objects of type `C`; the function may throw pointers of type `A` or pointers of objects that derive publicly from `A`.

A function that overrides a virtual function can only throw exceptions specified by the virtual function. The following example demonstrates this:

```

class A {
public:
    virtual void f() throw (int, char);
};

class B : public A{
public: void f() throw (int) { }
};

/* The following is not allowed. */
/*
class C : public A {
public: void f() { }
};

class D : public A {
public: void f() throw (int, char, double) { }
};
*/

```

The compiler allows `B::f()` because the member function may throw only exceptions of type `int`. The compiler would not allow `C::f()` because the member function may throw any kind of exception. The compiler would not allow `D::f()` because the member function can throw more types of exceptions (`int`, `char`, and `double`) than `A::f()`.

Suppose that you assign or initialize a pointer to function named `x` with a function or pointer to function named `y`. The pointer to function `x` can only throw exceptions specified by the exception specifications of `y`. The following example demonstrates this:

```

void (*f)();
void (*g)();
void (*h)() throw (int);

```

```

void i() {
    f = h;
    // h = g; This is an error.
}

```

The compiler allows the assignment `f = h` because `f` can throw any kind of exception. The compiler would not allow the assignment `h = g` because `h` can only throw objects of type `int`, while `g` can throw any kind of exception.

Implicitly declared special member functions (default constructors, copy constructors, destructors, and copy assignment operators) have exception specifications. An implicitly declared special member function will have in its exception specification the types declared in the functions' exception specifications that the special function invokes. If any function that a special function invokes allows all exceptions, then that special function allows all exceptions. If all the functions that a special function invokes allow no exceptions, then that special function will allow no exceptions. The following example demonstrates this:

```

class A {
public:
    A() throw (int);
    A(const A&) throw (float);
    ~A() throw();
};

class B {
public:
    B() throw (char);
    B(const A&);
    ~B() throw();
};

class C : public B, public A { };

```

The following special functions in the above example have been implicitly declared:

```

C::C() throw (int, char);
C::C(const C&); // Can throw any type of exception, including float
C::~~C() throw();

```

The default constructor of `C` can throw exceptions of type `int` or `char`. The copy constructor of `C` can throw any kind of exception. The destructor of `C` cannot throw any exceptions.

Related information

- “Incomplete types” on page 42
- “Function declarations and definitions” on page 201
- “Pointers to functions” on page 232
- Chapter 14, “Special member functions (C++ only),” on page 317

Special exception handling functions

Not all thrown errors can be caught and successfully dealt with by a catch block. In some situations, the best way to handle an exception is to terminate the program. Two special library functions are implemented in C++ to process exceptions not properly handled by catch blocks or exceptions thrown outside of a valid try block. These functions are:

- The `unexpected()` function

- The `terminate()` function

The `unexpected()` function

When a function with an exception specification throws an exception that is not listed in its exception specification, the C++ run time does the following:

1. The `unexpected()` function is called.
2. The `unexpected()` function calls the function pointed to by `unexpected_handler`. By default, `unexpected_handler` points to the function `terminate()`.

You can replace the default value of `unexpected_handler` with the function `set_unexpected()`.

Although `unexpected()` cannot return, it may throw (or rethrow) an exception. Suppose the exception specification of a function `f()` has been violated. If `unexpected()` throws an exception allowed by the exception specification of `f()`, then the C++ run time will search for another handler at the call of `f()`. The following example demonstrates this:

```
#include <iostream>
using namespace std;

struct E {
    const char* message;
    E(const char* arg) : message(arg) { }
};

void my_unexpected() {
    cout << "Call to my_unexpected" << endl;
    throw E("Exception thrown from my_unexpected");
}

void f() throw(E) {
    cout << "In function f(), throw const char* object" << endl;
    throw("Exception, type const char*, thrown from f()");
}

int main() {
    set_unexpected(my_unexpected);
    try {
        f();
    }
    catch (E& e) {
        cout << "Exception in main(): " << e.message << endl;
    }
}
```

The following is the output of the above example:

```
In function f(), throw const char* object
Call to my_unexpected
Exception in main(): Exception thrown from my_unexpected
```

The `main()` function's `try` block calls function `f()`. Function `f()` throws an object of type `const char*`. However the exception specification of `f()` allows only objects of type `E` to be thrown. The function `unexpected()` is called. The function `unexpected()` calls `my_unexpected()`. The function `my_unexpected()` throws an object of type `E`. Since `unexpected()` throws an object allowed by the exception specification of `f()`, the handler in the `main()` function may handle the exception.

If `unexpected()` did not throw (or rethrow) an object allowed by the exception specification of `f()`, then the C++ run time does one of two things:

- If the exception specification of `f()` included the class `std::bad_exception`, `unexpected()` will throw an object of type `std::bad_exception`, and the C++ run time will search for another handler at the call of `f()`.
- If the exception specification of `f()` did not include the class `std::bad_exception`, the function `terminate()` is called.

Related information

- “Special exception handling functions” on page 384
- “The `set_unexpected()` and `set_terminate()` functions” on page 387

The `terminate()` function

In some cases, the exception handling mechanism fails and a call to `void terminate()` is made. This `terminate()` call occurs in any of the following situations:

- The exception handling mechanism cannot find a handler for a thrown exception. The following are more specific cases of this:
 - During stack unwinding, a destructor throws an exception and that exception is not handled.
 - The expression that is thrown also throws an exception, and that exception is not handled.
 - The constructor or destructor of a nonlocal static object throws an exception, and the exception is not handled.
 - A function registered with `atexit()` throws an exception, and the exception is not handled. The following demonstrates this:

```
extern "C" printf(char* ...);
#include <exception>
#include <cstdlib>
using namespace std;

void f() {
    printf("Function f()\n");
    throw "Exception thrown from f()";
}

void g() { printf("Function g()\n"); }
void h() { printf("Function h()\n"); }

void my_terminate() {
    printf("Call to my_terminate\n");
    abort();
}

int main() {
    set_terminate(my_terminate);
    atexit(f);
    atexit(g);
    atexit(h);
    printf("In main\n");
}
```

The following is the output of the above example:

```
In main
Function h()
Function g()
Function f()
Call to my_terminate
```

To register a function with `atexit()`, you pass a parameter to `atexit()` a pointer to the function you want to register. At normal program termination,

`atexit()` calls the functions you have registered with no arguments in reverse order. The `atexit()` function is in the `<cstdlib>` library.

- A throw expression without an operand tries to rethrow an exception, and no exception is presently being handled.
- A function `f()` throws an exception that violates its exception specification. The `unexpected()` function then throws an exception which violates the exception specification of `f()`, and the exception specification of `f()` did not include the class `std::bad_exception`.
- The default value of `unexpected_handler` is called.

The `terminate()` function calls the function pointed to by `terminate_handler`. By default, `terminate_handler` points to the function `abort()`, which exits from the program. You can replace the default value of `terminate_handler` with the function `set_terminate()`.

A terminate function cannot return to its caller, either by using `return` or by throwing an exception.

Related information

- “The `set_unexpected()` and `set_terminate()` functions”

The `set_unexpected()` and `set_terminate()` functions

The function `unexpected()`, when invoked, calls the function most recently supplied as an argument to `set_unexpected()`. If `set_unexpected()` has not yet been called, `unexpected()` calls `terminate()`.

The function `terminate()`, when invoked, calls the function most recently supplied as an argument to `set_terminate()`. If `set_terminate()` has not yet been called, `terminate()` calls `abort()`, which ends the program.

You can use `set_unexpected()` and `set_terminate()` to register functions you define to be called by `unexpected()` and `terminate()`. The functions `set_unexpected()` and `set_terminate()` are included in the standard header files. Each of these functions has as its return type and its argument type a pointer to function with a `void` return type and no arguments. The pointer to function you supply as the argument becomes the function called by the corresponding special function: the argument to `set_unexpected()` becomes the function called by `unexpected()`, and the argument to `set_terminate()` becomes the function called by `terminate()`.

Both `set_unexpected()` and `set_terminate()` return a pointer to the function that was previously called by their respective special functions (`unexpected()` and `terminate()`). By saving the return values, you can restore the original special functions later so that `unexpected()` and `terminate()` will once again call `terminate()` and `abort()`.

If you use `set_terminate()` to register your own function, the function should not return to its caller but terminate execution of the program.

Example using the exception handling functions

The following example shows the flow of control and special functions used in exception handling:

```
#include <iostream>
#include <exception>
using namespace std;
```

```

class X { };
class Y { };
class A { };

// pfv type is pointer to function returning void
typedef void (*pfv)();

void my_terminate() {
    cout << "Call to my terminate" << endl;
    abort();
}

void my_unexpected() {
    cout << "Call to my_unexpected()" << endl;
    throw;
}

void f() throw(X,Y, bad_exception) {
    throw A();
}

void g() throw(X,Y) {
    throw A();
}

int main()
{
    pfv old_term = set_terminate(my_terminate);
    pfv old_unex = set_unexpected(my_unexpected);
    try {
        cout << "In first try block" << endl;
        f();
    }
    catch(X) {
        cout << "Caught X" << endl;
    }
    catch(Y) {
        cout << "Caught Y" << endl;
    }
    catch (bad_exception& e1) {
        cout << "Caught bad_exception" << endl;
    }
    catch (...) {
        cout << "Caught some exception" << endl;
    }

    cout << endl;

    try {
        cout << "In second try block" << endl;
        g();
    }
    catch(X) {
        cout << "Caught X" << endl;
    }
    catch(Y) {
        cout << "Caught Y" << endl;
    }
    catch (bad_exception& e2) {
        cout << "Caught bad_exception" << endl;
    }
    catch (...) {
        cout << "Caught some exception" << endl;
    }
}

```

The following is the output of the above example:

```
In first try block  
Call to my_unexpected()  
Caught bad_exception
```

```
In second try block  
Call to my_unexpected()  
Call to my_terminate
```

At run time, this program behaves as follows:

1. The call to `set_terminate()` assigns to `old_term` the address of the function last passed to `set_terminate()` when `set_terminate()` was previously called.
2. The call to `set_unexpected()` assigns to `old_unex` the address of the function last passed to `set_unexpected()` when `set_unexpected()` was previously called.
3. Within the first try block, function `f()` is called. Because `f()` throws an `unexpected` exception, a call to `unexpected()` is made. `unexpected()` in turn calls `my_unexpected()`, which prints a message to standard output. The function `my_unexpected()` tries to rethrow the exception of type `A`. Because class `A` has not been specified in the exception specification of function `f()`, `my_unexpected()` throws an exception of type `bad_exception`.
4. Because `bad_exception` has been specified in the exception specification of function `f()`, the handler `catch (bad_exception& e1)` is able to handle the exception.
5. Within the second try block, function `g()` is called. Because `g()` throws an `unexpected` exception, a call to `unexpected()` is made. The `unexpected()` throws an exception of type `bad_exception`. Because `bad_exception` has not been specified in the exception specification of `g()`, `unexpected()` calls `terminate()`, which calls the function `my_terminate()`.
6. `my_terminate()` displays a message then calls `abort()`, which terminates the program.

Note that the catch blocks following the second try block are not entered, because the exception was handled by `my_unexpected()` as an `unexpected` throw, not as a valid exception.

Chapter 17. Preprocessor directives

The preprocessor is a program that is invoked by the compiler to process code before compilation. Commands for that program, known as *directives*, are lines of the source file beginning with the character #, which distinguishes them from lines of source program text. The effect of each preprocessor directive is a change to the text of the source code, and the result is a new source code file, which does not contain the directives. The preprocessed source code, an intermediate file, must be a valid C or C++ program, because it becomes the input to the compiler.

Preprocessor directives consist of the following:

- Macro definition directives, which replace tokens in the current file with specified replacement tokens
- File inclusion directives, which imbed files within the current file
- Conditional compilation directives, which conditionally compile sections of the current file
- Message generation directives, which control the generation of diagnostic messages
-  Assertion directives, which specify attributes of the system the program is to run on
- The null directive (#), which performs no action
- Pragma directives, which apply compiler-specific rules to specified sections of code

Preprocessor directives begin with the # token followed by a preprocessor keyword. The # token must appear as the first character that is not white space on a line. The # is not part of the directive name and can be separated from the name with white spaces.

A preprocessor directive ends at the new-line character unless the last character of the line is the \ (backslash) character. If the \ character appears as the last character in the preprocessor line, the preprocessor interprets the \ and the new-line character as a continuation marker. The preprocessor deletes the \ (and the following new-line character) and splices the physical source lines into continuous logical lines. White space is allowed between backslash and the end of line character or the physical end of record. However, this white space is usually not visible during editing.

Except for some #pragma directives, preprocessor directives can appear anywhere in a program.

Macro definition directives

Macro definition directives include the following directives and operators:


- The #define directive, which defines a macro
- The #undef directive, which removes a macro definition

“Standard predefined macro names” on page 398 describes the macros that are predefined by the ISO C standard.  Macros that are predefined for XL C/C++ are described in “Predefined macros” in the *XL C/C++ Compiler Reference*.

The #define directive

A *preprocessor define directive* directs the preprocessor to replace all subsequent occurrences of a macro with specified replacement tokens.

#define directive syntax



The #define directive can contain:

- Object-like macros
- Function-like macros

The following are some differences between #define and the const type qualifier:

- The #define directive can be used to create a name for a numerical, character, or string constant, whereas a const object of any type can be declared.
- A const object is subject to the scoping rules for variables, whereas a constant created using #define is not.
- Unlike a const object, the value of a macro does not appear in the intermediate source code used by the compiler because they are expanded inline. The inline expansion makes the macro value unavailable to the debugger.
- A macro can be used in a constant expression, such as an array bound, whereas a const object cannot.
-  C++ The compiler does not type-check a macro, including macro arguments.

Related information

- “The const type qualifier” on page 76

Object-like macros

An *object-like macro definition* replaces a single identifier with the specified replacement tokens. The following object-like definition causes the preprocessor to replace all subsequent instances of the identifier COUNT with the constant 1000 :

```
#define COUNT 1000
```

If the statement

```
int arry[COUNT];
```

appears after this definition and in the same file as the definition, the preprocessor would change the statement to

```
int arry[1000];
```

in the output of the preprocessor.

Other definitions can make reference to the identifier COUNT:

```
#define MAX_COUNT COUNT + 100
```

The preprocessor replaces each subsequent occurrence of `MAX_COUNT` with `COUNT + 100`, which the preprocessor then replaces with `1000 + 100`.

If a number that is partially built by a macro expansion is produced, the preprocessor does not consider the result to be a single value. For example, the following will not result in the value 10.2 but in a syntax error.

```
#define a 10
a.2
```

Identifiers that are partially built from a macro expansion may not be produced. Therefore, the following example contains two identifiers and results in a syntax error:

```
#define d efg
abcd
```

Function-like macros

More complex than object-like macros, a function-like macro definition declares the names of formal parameters within parentheses, separated by commas. An empty formal parameter list is legal: such a macro can be used to simulate a function that takes no arguments. C99 adds support for function-like macros with a variable number of arguments. XL C++ supports function-like macros with a variable number of arguments, as a language extension for compatibility with C.

Function-like macro definition:

An identifier followed by a parameter list in parentheses and the replacement tokens. The parameters are imbedded in the replacement code. White space cannot separate the identifier (which is the name of the macro) and the left parenthesis of the parameter list. A comma must separate each parameter.

For portability, you should not have more than 31 parameters for a macro. The parameter list may end with an ellipsis (...). In this case, the identifier `__VA_ARGS__` may appear in the replacement list.

Function-like macro invocation:

An identifier followed by a comma-separated list of arguments in parentheses. The number of arguments should match the number of parameters in the macro definition, unless the parameter list in the definition ends with an ellipsis. In this latter case, the number of arguments in the invocation should exceed the number of parameters in the definition. The excess are called *trailing arguments*. Once the preprocessor identifies a function-like macro invocation, argument substitution takes place. A parameter in the replacement code is replaced by the corresponding argument. If trailing arguments are permitted by the macro definition, they are merged with the intervening commas to replace the identifier `__VA_ARGS__`, as if they were a single argument. Any macro invocations contained in the argument itself are completely replaced before the argument replaces its corresponding parameter in the replacement code.

A macro argument can be empty (consisting of zero preprocessing tokens). For example,

```
#define SUM(a,b,c) a + b + c
SUM(1,,3) /* No error message.
           1 is substituted for a, 3 is substituted for c. */
```

If the identifier list does not end with an ellipsis, the number of arguments in a macro invocation must be the same as the number of parameters in the

corresponding macro definition. During parameter substitution, any arguments remaining after all specified arguments have been substituted (including any separating commas) are combined into one argument called the variable argument. The variable argument will replace any occurrence of the identifier `__VA_ARGS__` in the replacement list. The following example illustrates this:

```
#define debug(...) fprintf(stderr, __VA_ARGS__)

debug("flag");    /* Becomes fprintf(stderr, "flag"); */
```

Commas in the macro invocation argument list do not act as argument separators when they are:

- In character constants
- In string literals
- Surrounded by parentheses

The following line defines the macro `SUM` as having two parameters `a` and `b` and the replacement tokens `(a + b)`:

```
#define SUM(a,b) (a + b)
```

This definition would cause the preprocessor to change the following statements (if the statements appear after the previous definition):

```
c = SUM(x,y);
c = d * SUM(x,y);
```

In the output of the preprocessor, these statements would appear as:

```
c = (x + y);
c = d * (x + y);
```

Use parentheses to ensure correct evaluation of replacement text. For example, the definition:

```
#define SQR(c) ((c) * (c))
```

requires parentheses around each parameter `c` in the definition in order to correctly evaluate an expression like:

```
y = SQR(a + b);
```

The preprocessor expands this statement to:

```
y = ((a + b) * (a + b));
```

Without parentheses in the definition, the correct order of evaluation is not preserved, and the preprocessor output is:

```
y = (a + b * a + b);
```

Arguments of the `#` and `##` operators are converted *before* replacement of parameters in a function-like macro.

Once defined, a preprocessor identifier remains defined and in scope independent of the scoping rules of the language. The scope of a macro definition begins at the definition and does not end until a corresponding `#undef` directive is encountered. If there is no corresponding `#undef` directive, the scope of the macro definition lasts until the end of the translation unit.

A recursive macro is not fully expanded. For example, the definition

```
#define x(a,b) x(a+1,b+1) + 4
```

expands

```
x(20,10)
```

to

```
x(20+1,10+1) + 4
```

rather than trying to expand the macro `x` over and over within itself. After the macro `x` is expanded, it is a call to function `x()`.

A definition is not required to specify replacement tokens. The following definition removes all instances of the token `debug` from subsequent lines in the current file:

```
#define debug
```

You can change the definition of a defined identifier or macro with a second preprocessor `#define` directive only if the second preprocessor `#define` directive is preceded by a preprocessor `#undef` directive. The `#undef` directive nullifies the first definition so that the same identifier can be used in a redefinition.

Within the text of the program, the preprocessor does not scan character constants or string constants for macro invocations.

The following example program contains two macro definitions and a macro invocation that refers to both of the defined macros:

```
/**
 ** This example illustrates #define directives.
 **/

#include <stdio.h>

#define SQR(s) ((s) * (s))
#define PRINT(a,b) \
    printf("value 1 = %d\n", a); \
    printf("value 2 = %d\n", b) ;

int main(void)
{
    int x = 2;
    int y = 3;

    PRINT(SQR(x),y);

    return(0);
}
```

After being interpreted by the preprocessor, this program is replaced by code equivalent to the following:

```
#include <stdio.h>

int main(void)
{
    int x = 2;
    int y = 3;

    printf("value 1 = %d\n", ( x ) * ( x ) );
    printf("value 2 = %d\n", y);

    return(0);
}
```

This program produces the following output:

```
value 1 = 4
value 2 = 3
```

Related information

- “Operator precedence and associativity” on page 170
- “Parenthesized expressions ()” on page 125

Variadic macro extensions:

IBM extension

Variadic macro extensions refer to two extensions to C99 and Standard C++ related to macros with variable number of arguments. One extension is a mechanism for renaming the variable argument identifier from `__VA_ARGS__` to a user-defined identifier. The other extension provides a way to remove the dangling comma in a variadic macro when no variable arguments are specified. Both extensions have been implemented to facilitate porting programs developed with GNU C and C++.

The following examples demonstrate the use of an identifier in place of `__VA_ARGS__`. The first definition of the macro `debug` exemplifies the usual usage of `__VA_ARGS__`. The second definition shows the use of the identifier `args` in place of `__VA_ARGS__`.

```
#define debug1(format, ...) printf(format, ## __VA_ARGS__)
#define debug2(format, args ...) printf(format, ## args)
```

Invocation

```
debug1("Hello %s\n", "World");
debug2("Hello %s\n", "World");
```

Result of macro expansion

```
printf("Hello %s\n", "World");
printf("Hello %s\n", "World");
```

The preprocessor removes the trailing comma if the variable arguments to a function macro are omitted or empty and the comma followed by `##` precedes the variable argument identifier in the function macro definition.

End of IBM extension

The `#undef` directive

A *preprocessor undef directive* causes the preprocessor to end the scope of a preprocessor definition.

#undef directive syntax

```
▶▶ #undef identifier ◀◀
```

If the identifier is not currently defined as a macro, `#undef` is ignored.

The following directives define `BUFFER` and `SQR`:

```
#define BUFFER 512
#define SQR(x) ((x) * (x))
```

The following directives nullify these definitions:

```
#undef BUFFER
#undef SQR
```

Any occurrences of the identifiers `BUFFER` and `SQR` that follow these `#undef` directives are not replaced with any replacement tokens. Once the definition of a macro has been removed by an `#undef` directive, the identifier can be used in a new `#define` directive.

The # operator

The `#` (single number sign) operator converts a parameter of a function-like macro into a character string literal. For example, if macro `ABC` is defined using the following directive:

```
#define ABC(x)  #x
```

all subsequent invocations of the macro `ABC` would be expanded into a character string literal containing the argument passed to `ABC`. For example:

Invocation	Result of macro expansion
<code>ABC(1)</code>	<code>"1"</code>
<code>ABC>Hello there)</code>	<code>"Hello there"</code>

The `#` operator should not be confused with the null directive.

Use the `#` operator in a function-like macro definition according to the following rules:

- A parameter following `#` operator in a function-like macro is converted into a character string literal containing the argument passed to the macro.
- White-space characters that appear before or after the argument passed to the macro are deleted.
- Multiple white-space characters imbedded within the argument passed to the macro are replaced by a single space character.
- If the argument passed to the macro contains a string literal and if a `\` (backslash) character appears within the literal, a second `\` character is inserted before the original `\` when the macro is expanded.
- If the argument passed to the macro contains a `"` (double quotation mark) character, a `\` character is inserted before the `"` when the macro is expanded.
- The conversion of an argument into a string literal occurs before macro expansion on that argument.
- If more than one `##` operator or `#` operator appears in the replacement list of a macro definition, the order of evaluation of the operators is not defined.
- If the result of the macro expansion is not a valid character string literal, the behavior is undefined.

The following examples demonstrate the use of the `#` operator:

```
#define STR(x)      #x
#define XSTR(x)    STR(x)
#define ONE        1
```

Invocation	Result of macro expansion
<code>STR(\n "\n" '\n')</code>	<code>"\n \\\"n\" '\n'"</code>
<code>STR(ONE)</code>	<code>"ONE"</code>
<code>XSTR(ONE)</code>	<code>"1"</code>
<code>XSTR("hello")</code>	<code>"\"hello\""</code>

Related information

- “The null directive (#)” on page 410

The ## operator

The ## (double number sign) operator concatenates two tokens in a macro invocation (text and/or arguments) given in a macro definition.

If a macro XY was defined using the following directive:

```
#define XY(x,y) x##y
```

the last token of the argument for x is concatenated with the first token of the argument for y.

Use the ## operator according to the following rules:

- The ## operator cannot be the very first or very last item in the replacement list of a macro definition.
- The last token of the item in front of the ## operator is concatenated with first token of the item following the ## operator.
- Concatenation takes place before any macros in arguments are expanded.
- If the result of a concatenation is a valid macro name, it is available for further replacement even if it appears in a context in which it would not normally be available.
- If more than one ## operator and/or # operator appears in the replacement list of a macro definition, the order of evaluation of the operators is not defined.

The following examples demonstrate the use of the ## operator:

```
#define ArgArg(x, y) x##y
#define ArgText(x) x##TEXT
#define TextArg(x) TEXT##x
#define TextText TEXT##text
#define Jitter 1
#define bug 2
#define Jitterbug 3
```

Invocation	Result of macro expansion
------------	---------------------------

ArgArg(lady, bug)	"\ladybug"
ArgText(con)	"conTEXT"
TextArg(book)	"TEXTbook"
TextText	"TEXTtext"
ArgArg(Jitter, bug)	3

Standard predefined macro names

Both C and C++ provide the following predefined macro names as specified in the ISO C language standard. Except for `__FILE__` and `__LINE__`, the value of the predefined macros remains constant throughout the translation unit.

`__DATE__` A character string literal containing the date when the source file was compiled.

The value of `__DATE__` changes as the compiler processes any include files that are part of your source program. The date is in the form:

```
"Mmm dd yyyy"
```

where:

Mmm Represents the month in an abbreviated form (Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, or Dec).

dd Represents the day. If the day is less than 10, the first d is a blank character.

yyyy Represents the year.

__FILE__ A character string literal containing the name of the source file.

The value of `__FILE__` changes as the compiler processes include files that are part of your source program. It can be set with the `#line` directive.

__LINE__ An integer representing the current source line number.

The value of `__LINE__` changes during compilation as the compiler processes subsequent lines of your source program. It can be set with the `#line` directive.

__STDC__ For C, the integer 1 (one) indicates that the C compiler supports the ISO standard. If you set the language level to **classic**, this macro is undefined. (When a macro is undefined, it behaves as if it had the integer value 0 when used in a `#if` statement.)

For C++, this macro is predefined to have the value 0 (zero). This indicates that the C++ language is not a proper superset of C, and that the compiler does not conform to ISO C.

C only

__STDC_HOSTED__

The value of this C99 macro is 1, indicating that the C compiler is a hosted implementation. Note that this macro is only defined if `__STDC__` is also defined.

End of C only

C only

__STDC_VERSION__

The integer constant of type `long int`: 199409L for the C89 language level, 199901L for C99. Note that this macro is only defined if `__STDC__` is also defined.

End of C only

__TIME__

A character string literal containing the time when the source file was compiled.

The value of `__TIME__` changes as the compiler processes any include files that are part of your source program. The time is in the form:

"hh:mm:ss"

where:

hh Represents the hour.

mm Represents the minutes.

ss Represents the seconds.

C++ only

__cplusplus For C++ programs, this macro expands to the long integer literal 199711L, indicating that the compiler is a C++ compiler. For C programs, this macro is not defined. Note that this macro name has no trailing underscores.

End of C++ only

Related information

- “The #line directive” on page 407
- “Object-like macros” on page 392

File inclusion directives


File inclusion directives consist of:

- The #include directive, which inserts text from another source file
-  The #include_next directive, which causes the compiler to omit the directory of the including file from the search path when searching for include files

The #include directive

A *preprocessor include directive* causes the preprocessor to replace the directive with the contents of the specified file.

#include directive syntax



If the *file_name* is enclosed in double quotation marks, for example:

```
#include "payroll.h"
```

it is treated as a user-defined file, and may represent a header or source file.

If the *file_name* is enclosed in angle brackets, for example:

```
#include <stdio.h>
```

it is treated as a system-defined file, and must represent a header file.

The new-line and > characters cannot appear in a file name delimited by < and >. The new-line and " (double quotation marks) characters cannot appear in a file name delimited by " and ", although > can.

The *file_path* can be an absolute or relative path. If the double quotation marks are used, and *file_path* is a relative path, or is not specified, the preprocessor adds the directory of the including file to the list of paths to be searched for the included file. If the double angle brackets are used, and *file_path* is a relative path, or is not

specified, the preprocessor does *not* add the directory of the including file to the list of paths to be searched for the included file.

The preprocessor resolves macros contained in an `#include` directive. After macro replacement, the resulting token sequence consists of a file name enclosed in either double quotation marks or the characters `<` and `>`. For example:

```
#define MONTH <july.h>
#include MONTH
```

Declarations that are used by several files can be placed in one file and included with `#include` in each file that uses them. For example, the following file `defs.h` contains several definitions and an inclusion of an additional file of declarations:

```
/* defs.h */
#define TRUE 1
#define FALSE 0
#define BUFFERSIZE 512
#define MAX_ROW 66
#define MAX_COLUMN 80
int hour;
int min;
int sec;
#include "mydefs.h"
```

You can embed the definitions that appear in `defs.h` with the following directive:

```
#include "defs.h"
```

In the following example, a `#define` combines several preprocessor macros to define a macro that represents the name of the C standard I/O header file. A `#include` makes the header file available to the program.

```
#define C_IO_HEADER <stdio.h>

/* The following is equivalent to:
 * #include <stdio.h>
 */


#include C_IO_HEADER
```

The `#include_next` directive

IBM extension

The preprocessor directive `#include_next` behaves like the `#include` directive, except that it specifically excludes the directory of the including file from the paths to be searched for the named file. All search paths up to and including the directory of the including file are omitted from the list of paths to be searched for the included file. This allows you to include multiple versions of a file with the same name in different parts of an application; or to include one header file in another header file with the same name (without the header including itself recursively). Provided that the different file versions are stored in different directories, the directive ensures you can access each version of the file, without requiring that you use absolute paths to specify the file name.

#include_next directive syntax



The directive must only be used in header files, and the file specified by the *file_name* must be a header file. There is no distinction between the use of double quotation marks and angle brackets to enclose the file name.

As an example of how search paths are resolved with the `#include_next` directive, assume that there are two versions of the file `t.h`: the first one, which is included in the source file `t.c`, is located in the subdirectory `path1`; the second one, which is included in the first one, is located in the subdirectory `path2`. Both directories are specified as include file search paths when `t.c` is compiled.

```
/* t.c */

#include "t.h"

int main()
{
    printf("%d", ret_val);
}

/* t.h in path1 */

#include_next "t.h"

int ret_val = RET;

/* t.h in path2 */

#define RET 55;
```

The `#include_next` directive instructs the preprocessor to skip the `path1` directory and start the search for the included file from the `path2` directory. This directive allows you to use two different versions of `t.h` and it prevents `t.h` from being included recursively.

End of IBM extension

Conditional compilation directives

A *preprocessor conditional compilation directive* causes the preprocessor to conditionally suppress the compilation of portions of source code. These directives test a constant expression or an identifier to determine which tokens the preprocessor should pass on to the compiler and which tokens should be bypassed during preprocessing. The directives are:

- The `#if` and `#elif` directives, which conditionally include or suppress portions of source code, depending on the result of a constant expression
- The `#ifdef` directive, which conditionally includes source text if a macro name is defined
- The `#ifndef` directive, which conditionally includes source text if a macro name is not defined
- The `#else` directive, which conditionally includes source text if the previous `#if`, `#ifdef`, `#ifndef`, or `#elif` test fails
- The `#endif` directive, which ends conditional text

The preprocessor conditional compilation directive spans several lines:

- The condition specification line (beginning with `#if`, `#ifdef`, or `#ifndef`)
- Lines containing code that the preprocessor passes on to the compiler if the condition evaluates to a nonzero value (optional)
- The `#elif` line (optional)
- Lines containing code that the preprocessor passes on to the compiler if the condition evaluates to a nonzero value (optional)
- The `#else` line (optional)
- Lines containing code that the preprocessor passes on to the compiler if the condition evaluates to zero (optional)
- The preprocessor `#endif` directive

For each `#if`, `#ifdef`, and `#ifndef` directive, there are zero or more `#elif` directives, zero or one `#else` directive, and one matching `#endif` directive. All the matching directives are considered to be at the same nesting level.

You can nest conditional compilation directives. In the following directives, the first `#else` is matched with the `#if` directive.

```
#ifdef MACNAME
/* tokens added if MACNAME is defined */
#   if TEST <=10
/* tokens added if MACNAME is defined and TEST <= 10 */
#   else
/* tokens added if MACNAME is defined and TEST > 10 */
#   endif
#else
/* tokens added if MACNAME is not defined */
#endif
```

Each directive controls the block immediately following it. A block consists of all the tokens starting on the line following the directive and ending at the next conditional compilation directive at the same nesting level.

Each directive is processed in the order in which it is encountered. If an expression evaluates to zero, the block following the directive is ignored.


When a block following a preprocessor directive is to be ignored, the tokens are examined only to identify preprocessor directives within that block so that the conditional nesting level can be determined. All tokens other than the name of the directive are ignored.

Only the first block whose expression is nonzero is processed. The remaining blocks at that nesting level are ignored. If none of the blocks at that nesting level has been processed and there is a `#else` directive, the block following the `#else` directive is processed. If none of the blocks at that nesting level has been processed and there is no `#else` directive, the entire nesting level is ignored.

The `#if` and `#elif` directives

The `#if` and `#elif` directives compare the value of *constant_expression* to zero:

#if and #elif directive syntax



If the constant expression evaluates to a nonzero value, the lines of code that immediately follow the condition are passed on to the compiler.

If the expression evaluates to zero and the conditional compilation directive contains a preprocessor `#elif` directive, the source text located between the `#elif` and the next `#elif` or preprocessor `#else` directive is selected by the preprocessor to be passed on to the compiler. The `#elif` directive cannot appear after the preprocessor `#else` directive.

All macros are expanded, any `defined()` expressions are processed and all remaining identifiers are replaced with the token `0`.

The *constant_expression* that is tested must be integer constant expressions with the following properties:

- No casts are performed.
- Arithmetic is performed using long int values.
- The *constant_expression* can contain defined macros. No other identifiers can appear in the expression.
- The *constant_expression* can contain the unary operator `defined`. This operator can be used only with the preprocessor keyword `#if` or `#elif`. The following expressions evaluate to 1 if the *identifier* is defined in the preprocessor, otherwise to 0:

```
defined identifier
defined(identifier)
```

For example:

```
#if defined(TEST1) || defined(TEST2)
```

Note: If a macro is not defined, a value of 0 (zero) is assigned to it. In the following example, `TEST` must be a macro identifier:


```
#if TEST >= 1
    printf("i = %d\n", i);
    printf("array[i] = %d\n", array[i]);
#elif TEST < 0
    printf("array subscript out of bounds \n");
#endif
```

The #ifdef directive

The `#ifdef` directive checks for the existence of macro definitions.

If the identifier specified is defined as a macro, the lines of code that immediately follow the condition are passed on to the compiler.

#ifdef directive syntax



The following example defines MAX_LEN to be 75 if EXTENDED is defined for the preprocessor. Otherwise, MAX_LEN is defined to be 50.

```
#ifdef EXTENDED
#   define MAX_LEN 75
#else
#   define MAX_LEN 50
#endif
```

The #ifndef directive

The #ifndef directive checks whether a macro is not defined.

If the identifier specified is not defined as a macro, the lines of code immediately follow the condition are passed on to the compiler.

#ifndef directive syntax




An identifier must follow the #ifndef keyword. The following example defines MAX_LEN to be 50 if EXTENDED is not defined for the preprocessor. Otherwise, MAX_LEN is defined to be 75.

```
#ifndef EXTENDED
#   define MAX_LEN 50
#else
#   define MAX_LEN 75
#endif
```

The #else directive

If the condition specified in the #if, #ifdef, or #ifndef directive evaluates to 0, and the conditional compilation directive contains a preprocessor #else directive, the lines of code located between the preprocessor #else directive and the preprocessor #endif directive is selected by the preprocessor to be passed on to the compiler.

#else directive syntax



The #endif directive

The preprocessor #endif directive ends the conditional compilation directive.

#endif directive syntax

▶ `#endif` *newline_character* ▶

Examples of conditional compilation directives

The following example shows how you can nest preprocessor conditional compilation directives:

```
#if defined(TARGET1)
#   define SIZEOF_INT 16
#   ifdef PHASE2
#       define MAX_PHASE 2
#   else
#       define MAX_PHASE 8
#   endif
#elif defined(TARGET2)
#   define SIZEOF_INT 32
#   define MAX_PHASE 16
#else
#   define SIZEOF_INT 32
#   define MAX_PHASE 32
#endif
```

The following program contains preprocessor conditional compilation directives:

```
/**
** This example contains preprocessor
** conditional compilation directives.
**/

#include <stdio.h>

int main(void)
{
    static int array[ ] = { 1, 2, 3, 4, 5 };
    int i;

    for (i = 0; i <= 4; i++)
    {
        array[i] *= 2;
    }

    #if TEST >= 1
        printf("i = %d\n", i);
        printf("array[i] = %d\n",
            array[i]);
    #endif

    }
    return(0);
}
```

Message generation directives

Message generation directives include the following:

- The `#error` directive, which defines text for a compile-time error message
- The `#warning` directive, which defines text for a compile-time warning message
- The `#line` directive, which supplies a line number for compiler messages

Related information

- “Conditional compilation directives” on page 402

The #error directive

A *preprocessor error directive* causes the preprocessor to generate an error message and causes the compilation to fail.

#error directive syntax



```
▶▶ #error [preprocessor_token]
```

The argument *preprocessor_token* is not subject to macro expansion.

The #error directive is often used in the #else portion of a #if-#elif-#else construct, as a safety check during compilation. For example, #error directives in the source file can prevent code generation if a section of the program is reached that should be bypassed.

For example, the directive

```
#define BUFFER_SIZE 255

#if BUFFER_SIZE < 256
#error "BUFFER_SIZE is too small."
#endif
```

generates the error message:

```
BUFFER_SIZE is too small.
```

The #warning directive

IBM extension

A *preprocessor warning directive* causes the preprocessor to generate a warning message but allows compilation to continue. The argument to #warning is not subject to macro expansion.

#warning directive syntax



```
▶▶ #warning [preprocessor_token]
```


The preprocessor #warning directive is a language extension provided to facilitate handling programs developed with GNU C. The IBM implementation preserves multiple white spaces.

End of IBM extension

The #line directive

A *preprocessor line control directive* supplies line numbers for compiler messages. It causes the compiler to view the line number of the next source line as the specified number.

#line directive syntax



In order for the compiler to produce meaningful references to line numbers in preprocessed source, the preprocessor inserts #line directives where necessary (for example, at the beginning and after the end of included text).

A file name specification enclosed in double quotation marks can follow the line number. If you specify a file name, the compiler views the next line as part of the specified file. If you do not specify a file name, the compiler views the next line as part of the current source file.

At the C99 language level, the maximum value of the #line preprocessing directive is 2147483647.

In all C and C++ implementations, the token sequence on a #line directive is subject to macro replacement. After macro replacement, the resulting character sequence must consist of a decimal constant, optionally followed by a file name enclosed in double quotation marks.

You can use #line control directives to make the compiler provide more meaningful error messages. The following example program uses #line control directives to give each function an easily recognizable line number:

```
/**
 ** This example illustrates #line directives.
 **/

#include <stdio.h>
#define LINE200 200

int main(void)
{
    func_1();
    func_2();
}

#line 100
func_1()
{
    printf("Func_1 - the current line number is %d\n",_ _LINE_ _);
}

#line LINE200
func_2()
{
    printf("Func_2 - the current line number is %d\n",_ _LINE_ _);
}
```

This program produces the following output:

```
Func_1 - the current line number is 102
Func_2 - the current line number is 202
```

Assertion directives

IBM extension

An *assertion directive* is an alternative to a macro definition, used to define the computer or system the compiled program will run on. Assertions are usually predefined, but you can define them with the `#assert` preprocessor directive.

#assert directive syntax

```
▶▶ #assert predicate (—answer—) ▶▶
```

The *predicate* represents the assertion entity you are defining. The *answer* represents a value you are assigning to the assertion. You can make several assertions using the same predicate and different answers. All the answers for any given predicate are simultaneously true. For example, the following directives create assertions regarding font properties:

```
#assert font(arial)
#assert font(blue)
```

Once an assertion has been defined, the assertion predicate can be used in conditional directives to test the current system. The following directive tests whether `arial` or `blue` is asserted for `font`:

```
#if #font(arial) || #font(blue)
```

You can test whether any answer is asserted for a predicate by omitting the answer in the conditional:

```
#if #font
```

Assertions can be cancelled with the `#unassert` directive. If you use the same syntax as the `#assert` directive, the directive cancels only the answer you specify. For example, the following directive cancels the `arial` answer for the `font` predicate:

```
#unassert font(arial)
```

An entire predicate is cancelled by omitting the answer from the `#unassert` directive. The following directive cancels the `font` directive altogether:

```
#unassert font
```

Related information

- “Conditional compilation directives” on page 402

Predefined assertions

The following assertions are predefined for the Cell Broadband Engine platform:

Table 30. Predefined assertions for Cell Broadband Engine

PPU	SPU
#machine(powerpc)	#machine(bigendian)
#machine(bigendian)	#cpu(spu)
#cpu(powerpc)	

End of IBM extension

The null directive (#)

The *null directive* performs no action. It consists of a single # on a line of its own.

The null directive should not be confused with the # operator or the character that starts a preprocessor directive.

In the following example, if MINVAL is a defined macro name, no action is performed. If MINVAL is not a defined identifier, it is defined 1.

```
#ifndef MINVAL
#
#else
#define MINVAL 1
#endif
```


Related information

- “The # operator” on page 397

Pragma directives

A *pragma* is an implementation-defined instruction to the compiler. It has the general form:

#pragma directive syntax



The *character_sequence* is a series of characters giving a specific compiler instruction and arguments, if any. The token STDC indicates a standard pragma; consequently, no macro substitution takes place on the directive. The *new-line* character must terminate a pragma directive.

The *character_sequence* on a pragma is subject to macro substitutions. For example,

```
#define
XX_ISO_DATA
isolated_call(LG_ISO_DATA)
// ...
#pragma XX_ISO_DATA
```

Note: You can also use the `_Pragma` operator syntax to specify a pragma directive; for details, see “The `_Pragma` preprocessing operator.”

More than one pragma construct can be specified on a single pragma directive. The compiler ignores unrecognized pragmas.

Standard C pragmas are described in “Standard pragmas” on page 411.  Pragmas available for XL C/C++ are described in “General purpose pragmas” in the *XL C/C++ Compiler Reference*.

The `_Pragma` preprocessing operator

The unary operator `_Pragma`, which is a C99 feature, allows a preprocessor macro to be contained in a pragma directive.

`_Pragma` operator syntax

►► `_Pragma` (`"` `string_literal` `"`) ◀◀

The `string_literal` may be prefixed with `L`, making it a wide-string literal.

The string literal is dstringized and tokenized. The resulting sequence of tokens is processed as if it appeared in a pragma directive. For example:

```
_Pragma ( "pack(full)" )
```

would be equivalent to

```
#pragma pack(full)
```

Standard pragmas

A *standard pragma* is a pragma preprocessor directive for which the C Standard defines the syntax and semantics and for which no macro replacement is performed. A standard pragma must be one of the following:

►► `#pragma` `STDC`

<code>FP_CONTRACT</code>	<code>ON</code>
<code>FENV_ACCESS</code>	<code>OFF</code>
<code>CX_LIMITED_RANGE</code>	<code>DEFAULT</code>

`new-line` ◀◀

C The `FP_CONTRACT` and `FENV_ACCESS` pragmas are recognized and ignored.

`CX_LIMITED_RANGE` is described below.

`pragma STDC CX_LIMITED_RANGE`

The usual mathematical formulas for complex multiplication, division, and absolute value are problematic because of their treatment of infinities and because of undue overflow and underflow. The usual formulas are as follows:

$$(x + iy) \times (u + iv) = (xu - yv) + i(yu + xv)$$

$$(x + iy)/(u + iv) = [(xu + yv) + i(yu - xv)]/(u^2 + v^2)$$

$$|x + iy| = \text{sqrt}(x^2 + y^2)$$

By default, the compiler uses slightly more complex but mathematically safer algorithms to implement these calculations. Where you determine that the usual mathematical formulas are safe, you can use the `STDC CX_LIMITED_RANGE` pragma to inform the compiler that, when the state is "on", the formulas are acceptable. In doing so, you allow the compiler to generate faster code for these computations. When the state is "off", the compiler will continue to use the safer algorithms. For details on the implementation of this pragma, see **`#pragma STDC CX_LIMITED_RANGE`** in the *XL C/C++ Compiler Reference*.

Appendix A. The IBM XL C language extensions

This appendix presents the IBM XL C extensions in the following categories:

- “C99 features as extensions to C89”
- “Extensions for Unicode support” on page 415
- “Extensions for GNU C compatibility” on page 415
- “Extensions for vector processing support” on page 417

C99 features as extensions to C89

The following features are enabled by default when you compile with any of the following:

- the **ppuxlc** invocation command
- the **spuxlc** invocation command
- the **ppuc99** invocation command
- the **-qlanglvl=extc99 | stcd99 | extc89 | extended** options

For more information on these options, see the **-qlanglvl** option in the *XL C/C++ Compiler Reference*.

Table 31. Default C99 features as extensions to C89

Language feature	Discussed in:
Hexadecimal floating-point constants	“Hexadecimal floating-point literals” on page 23
<code>__func__</code> predefined identifier	“The <code>__func__</code> predefined identifier” on page 16
Concatenation of wide and non-wide character strings	“String concatenation” on page 30
Mixed declarations and code	“Overview of data declarations and definitions” on page 43
long long data type	“Integral types” on page 53
Complex data type	“Complex floating-point types ” on page 55
<code>_Bool</code> data type	“Boolean types” on page 54
Trailing comma allowed in enum declaration	“Enumeration type definition” on page 67
Duplicate type qualifiers	“Type qualifiers” on page 72
Variable length arrays	“Variable length arrays ” on page 91
Non-lvalue array subscripts	“Array subscripting operator []” on page 150
Flexible array members at the end of a structure or union	“Flexible array members ” on page 61
Non-constant expression in initializer for structure or union	“Initialization of structures and unions” on page 98
Designated initializers	“Designated initializers for aggregate types (C only)” on page 95
Removal of implicit function declaration	“Function declarations” on page 202

Table 31. Default C99 features as extensions to C89 (continued)

Language feature	Discussed in:
Removal of implicit int return type in function declarations	"Function return type specifiers" on page 213
Static arrays as function parameters	"Static array indices in function parameter declarations (C only)" on page 217
Variable arguments in function-like macros	"Function-like macros" on page 393
Empty arguments in function-like macros	"Function-like macros" on page 393
Additional predefined macro names	"Standard predefined macro names" on page 398
Compound literals	"Compound literal expressions" on page 164
_Pragma operator	"The _Pragma preprocessing operator " on page 410
Standard pragmas	"Standard pragmas " on page 411
New limit for #line directive	"The #line directive" on page 407

The following features are enabled by default when you compile with any of the following:

- the **ppuxlc** invocation command
- the **spuxlc** invocation command
- the **ppuc99** invocation command
- the **-qlanglvl=extc99 | stcd99 | extc89 | extended** options

They are also enabled or disabled by specific compiler options, which are listed in the table below.

Table 32. Default C99 features as extensions to C89, with individual option controls

Language feature	Discussed in:	Individual option control
Digraphs	"Digraph characters" on page 36	-q[no]digraph
C++ style comments	"Comments" on page 37	-q[no]cpluscmt
The inline function specifier	"The inline function specifier" on page 208	-qkeyword=inline

The following feature is enabled by default when you compile with any of the following:

- the **ppuxlc** invocation command
- the **spuxlc** invocation command
- the **ppuc99** invocation command
- the **-qlanglvl=extc99 | stcd99** options

Language feature	Discussed in:
Unsuffixed long long integer literals	"Decimal integer literals" on page 20

The following features are enabled by default when you compile with any of the following:

- the **ppuxlc** invocation command

- the **spuxlc** invocation command
- the **ppuc99** invocation command
- the **-qlanglvl=extc99 | stcd99** options

They are also enabled or disabled by specific compiler options, which are listed in the table below.

Table 33. Strict C99 features as extensions to C89, with individual option controls

Language feature	Discussed in:	Individual option control
Universal character names	"The Unicode standard" on page 34	-qlanglvl=[no]ucs
The restrict type qualifier	"The restrict type qualifier" on page 77	-qkeyword=restrict

Related information

- **-qcplusmt**, **-qkeyword**, and **-qdigraph** in the *XL C/C++ Compiler Reference*
- "Invoking the compiler" in the *XL C/C++ Compiler Reference*

Extensions for Unicode support

The following feature requires compilation with the use of an additional option.

Language feature	Discussed in:	Required compilation option
UTF-16, UTF-32 literals	"UTF literals" on page 35	-qutf

Related information

- **-qutf** in the *XL C/C++ Compiler Reference*

Extensions for GNU C compatibility

The following feature is enabled by default at all language levels.

Table 34. Default IBM XL C extensions for GNU C compatibility

Language feature	Discussed in:
#include_next preprocessor directive	"The #include_next directive" on page 401

The following features are enabled by default when you compile with any of the following:

- the **ppuxlc** invocation command
- the **spuxlc** invocation command
- the **-qlanglvl=extc99 | extc89 | extended** options

Table 35. Default IBM XL C extensions for GNU C compatibility

Language feature	Discussed in:
Alternate keywords	"Keywords for language extensions" on page 14
__extension__ keyword	"Keywords for language extensions" on page 14

Table 35. Default IBM XL C extensions for GNU C compatibility (continued)

Language feature	Discussed in:
asm labels	"Assembly labels" on page 17
Complex literal suffixes	"Complex literals" on page 24
Global register variables	"Variables in specified registers" on page 50
Placement of flexible array members anywhere in structure or union	"Flexible array members " on page 61
Static initialization of flexible array members of aggregates	"Flexible array members " on page 61
Zero-extent arrays	"Zero-extent array members" on page 62
Type attributes	"Type attributes" on page 78
Variable attributes	"Variable attributes" on page 105
Locally declared labels	"Locally declared labels" on page 176
Labels as values	"Labels as values" on page 176
<code>__alignof__</code> operator	"The <code>__alignof__</code> operator" on page 135
<code>__typeof__</code> operator	"The <code>typeof</code> operator" on page 137
Generalized lvalues	"Lvalues and rvalues" on page 121
Complex type arguments to unary operators	"Unary expressions" on page 128
Initialization of static variables by compound literals	"Compound literal expressions" on page 164
<code>__imag__</code> and <code>__real__</code> complex type operators	"The <code>__real__</code> and <code>__imag__</code> operators " on page 138
Cast to a union type	"Cast to union type (C only)" on page 158
Computed goto statements	"Computed goto statement" on page 193
Statements and declarations in expressions	"Statement expressions" on page 179
Function attributes	"Function attributes" on page 218
<code>__inline__</code> function specifier	"The inline function specifier" on page 208
Nested functions	"Nested functions" on page 232
Variadic macro extensions	"Variadic macro extensions" on page 396
<code>#warning</code> preprocessor directive	"The <code>#warning</code> directive" on page 407
<code>#assert</code> , <code>#unassert</code> , <code>#cpu</code> , <code>#machine</code> , <code>#system</code> preprocessor directives	"Assertion directives" on page 409

The following features are enabled by default when you compile with any of the following:

- the **ppuxlc** invocation command
- the **spuxlc** invocation command
- the **-qlanglvl=extc99 | extc89 | extended** options

They are also enabled or disabled by specific compiler options, which are listed in the table below.

Table 36. IBM XL C extensions for GNU C compatibility with individual option controls

Language feature	Discussed in:	Individual option controls
asm, and __asm keywords	“Assembly labels” on page 17, “Inline assembly statements” on page 194	-qkeyword=asm, -qasm
asm inline assembly-language statements	“Inline assembly statements” on page 194	-qasm
The __thread storage class specifier	“The __thread storage class specifier (PPU only)” on page 52	-qtls

The following features require compilation with the use of an additional option.

Table 37. IBM XL C extensions for GNU C compatibility, requiring additional compiler options

Language feature	Discussed in:	Required compilation option
Dollar signs in identifiers	“Characters in identifiers” on page 16	-qdollar
The typeof keyword	“The typeof operator” on page 137	-qkeyword=typeof

Related information

- **-qkeyword, -qasm,-qtls** and **-qdollar** in the *XL C/C++ Compiler Reference*
- “Invoking the compiler” in the *XL C/C++ Compiler Reference*

Extensions for vector processing support

Table 38. IBM XL C extensions to support the AltiVec Application Programming Interface specification and the Cell Broadband Processor

Language feature	Discussed in:
Vector programming language extensions	“Vector types” on page 57, “Vector literals” on page 25

The following features are IBM extensions to the AltiVec Application Programming Interface specification. These extensions are supported for both AltiVec types on the PPU and vector types on the SPU.

Table 39. IBM XL C extensions to the AltiVec Application Programming Interface specification

Language extension	Discussed in:
Initializer lists for vector constants	“Initialization of vectors” on page 97
typedef definitions for vector types	“typedef definitions” on page 71
compound literals as initializers for static vector variables	“Compound literal expressions” on page 164
vector types as arguments to the __alignof__ and typeof operators	“The __alignof__ operator” on page 135, “The typeof operator” on page 137

Appendix B. The IBM XL C++ language extensions

This appendix presents the IBM XL C++ extensions to Standard C++ in the following categories:

- “General IBM extensions”
- “Extensions for C99 compatibility”
- “Extensions for Unicode support” on page 420
- “Extensions for GNU C compatibility” on page 420
- “Extensions for GNU C++ compatibility” on page 422
- “Extensions for vector processing support” on page 423

General IBM extensions

The following feature is enabled by default. It can also be enabled or disabled by a specific compiler option, listed in the table below.

Language feature	Discussed in:	Individual option controls
long long data type	“Integral types” on page 53	-q[no]longlong

Related information

- -qlonglong in the *XL C/C++ Compiler Reference*

Extensions for C99 compatibility

IBM XL C++ adds support for the following C99 language features. All of these features are enabled by default.

Table 40. Default C99 features as extensions to Standard C++

Language feature	Discussed in:
Duplicate type qualifiers	“Type qualifiers” on page 72
Flexible array members at the end of a structure or union	“Flexible array members ” on page 61
_Pragma operator	“The _Pragma preprocessing operator ” on page 410
Additional predefined macro names	“Standard predefined macro names” on page 398
Empty arguments in function-like macros	“Function-like macros” on page 393
C standard pragmas	“Standard pragmas ” on page 411

The following features are enabled by default. They can also be enabled or disabled by specific compiler options, which are listed in the table below.

Table 41. Default C99 features as extensions to Standard C++, with individual option controls

Language feature	Discussed in:	Individual option controls
<code>__func__</code> predefined identifier	“The <code>__func__</code> predefined identifier” on page 16	<code>-qlanglvl=[no]c99__func__</code>
Hexadecimal floating-point literals	“Hexadecimal floating-point literals” on page 23	<code>-qlanglvl=[no]c99hexfloat</code>
Complex data type	“Complex floating-point types ” on page 55	<code>-qlanglvl=[no]c99complex</code>
Trailing comma allowed in enum declaration	“Enumeration type definition” on page 67	<code>-qlanglvl=[no]trailenum</code>
The restrict type qualifier	“The restrict type qualifier” on page 77	<code>-q[no]keyword=restrict</code>
Variable length arrays	“Variable length arrays ” on page 91	<code>-qlanglvl=[no]c99vla</code>
Compound literals	“Compound literal expressions” on page 164	<code>-qlanglvl=[no]c99compoundliteral</code>
Variable arguments in function-like macros	“Function-like macros” on page 393	<code>-qlanglvl=[no]varargmacros</code>

The following feature is only enabled by a specific compiler option, listed in the table below.

Table 42. C99 features as extensions to Standard C++, with individual option controls

Language feature	Discussed in:	Individual option control
Universal character names	“The Unicode standard” on page 34	<code>-qlanglvl=ucs</code>

Related information

- `-qlanglvl` and `-qkeyword` in the *XL C/C++ Compiler Reference*

Extensions for Unicode support

The following feature is enabled by default. It can also be enabled or disabled by a specific compiler option, which is listed in the table below.

Language feature	Discussed in:	Required compilation option
UTF-16, UTF-32 literals	“UTF literals” on page 35	<code>-qutf</code>

Related information

- `-qutf` in the *XL C/C++ Compiler Reference*

Extensions for GNU C compatibility

The following subset of the GNU C language extensions is enabled by default.

Table 43. Default IBM XL C++ extensions for compatibility with GNU C

Language feature	Discussed in:
Assembly labels	"Assembly labels" on page 17
Placement of flexible array members anywhere in structure or union	"Flexible array members " on page 61
Static initialization of flexible array members of aggregates	"Flexible array members " on page 61
Statements and declarations in expressions	"Statement expressions" on page 179
<code>__alignof__</code> operator	"The <code>__alignof__</code> operator" on page 135
<code>__typeof__</code> operator	"The <code>typeof</code> operator" on page 137
Generalized lvalues	"Lvalues and rvalues" on page 121
Complex type arguments to unary operators	"Unary expressions" on page 128
<code>__imag__</code> and <code>__real__</code> complex type operators	"The <code>__real__</code> and <code>__imag__</code> operators " on page 138
Function attributes	"Function attributes" on page 218

The following subset of the GNU C language extensions is enabled by default. These extensions can also be enabled or disabled by specific compiler options, which are listed in the table below.

Table 44. Default IBM XL C++ extensions for compatibility with GNU C, with individual option controls

Language feature	Discussed in:	Individual option controls
Alternate keywords	"Keywords for language extensions" on page 14	<code>-q[no]keyword=token</code>
<code>__extension__</code> keyword	"Keywords for language extensions" on page 14	<code>-q[no]keyword=__extension__</code>
Complex literal suffixes	"Complex literals" on page 24	<code>-qlanglvl=[no]gnu_suffixij</code>
Type attributes	"Type attributes" on page 78	<code>-q[no]keyword=__attribute__</code>
Variable attributes	"Variable attributes" on page 105	<code>-q[no]keyword=__attribute__</code>
Zero-extent arrays	"Zero-extent array members" on page 62	<code>-qlanglvl=[no]zeroextarray</code>
<code>typeof</code> operator	"The <code>typeof</code> operator" on page 137	<code>-q[no]keyword=typeof</code>
Locally declared labels	"Locally declared labels" on page 176	<code>-qlanglvl=[no]gnu_locallabel</code>
Labels as values	"Labels as values" on page 176	<code>-qlanglvl=[no]gnu_labelvalue</code>
Computed goto statements	"Computed goto statement" on page 193	<code>-qlanglvl=[no]gnu_computedgoto</code>
inline assembly-language statements	"Inline assembly statements" on page 194	<code>-qasm</code>
Variadic macro extensions	"Variadic macro extensions" on page 396	<code>-qlanglvl=[no]gnu_varargmacros</code>

Table 44. Default IBM XL C++ extensions for compatibility with GNU C, with individual option controls (continued)

Language feature	Discussed in:	Individual option controls
#warning preprocessor directive	“The #warning directive” on page 407	-qlanglvl=[no]gnu_warning
#assert, #unassert, #cpu, #machine, #system preprocessor directives	“Assertion directives” on page 409	-qlanglvl=[no]gnu_assert

The following feature requires compilation with the use of an additional option, listed in the table below.

Table 45. IBM XL C++ extensions for GNU C compatibility, requiring additional compiler options

Language feature	Discussed in:	Required compilation option
Dollar signs in identifiers	“Characters in identifiers” on page 16	-qdollar

Related information

- **-qkeyword**, **-qasm**, **-qtls** and **-qdollar** in the *XL C/C++ Compiler Reference*

Extensions for GNU C++ compatibility

The following GNU C++ language extension is enabled by default.

Table 46. IBM XL C++ language extensions for compatibility with GNU C++

Language feature	Discussed in:
init_priority variable attribute	“The init_priority variable attribute (C++ only)” on page 109

The following GNU C++ language extensions are enabled by default. They can also be enabled or disabled by specific compiler options, listed in the table below.

Table 47. IBM XL C++ language extensions for compatibility with GNU C++, with individual option controls

Language feature	Discussed in:	Individual option control
Template instantiations declared as extern	“Template instantiation” on page 356	-qlanglvl=[no]gnu_externtemplate
The __thread storage class specifier	“The __thread storage class specifier (PPU only)” on page 52	-qtls

Related information

- **-qlanglvl** and **-qtls** in the *XL C/C++ Compiler Reference*

Extensions for vector processing support

Table 48. IBM XL C++ extensions to support the AltiVec Application Programming Interface specification and the Cell Broadband Processor

Language feature	Discussed in:
Vector programming language extensions	"Vector types" on page 57, "Vector literals" on page 25

The following features are IBM extensions to the AltiVec Application Programming Interface specification. These extensions are supported for both AltiVec types on the PPU and vector types on the SPU.

Table 49. IBM XL C++ extensions to the AltiVec Application Programming Interface specification

Language extension	Discussed in:
Initializer lists for vector constants	"Initialization of vectors" on page 97
typedef definitions for vector types	"typedef definitions" on page 71
compound literals as initializers for static vector variables	"Compound literal expressions" on page 164
vector types as arguments to the <code>__alignof__</code> and <code>typeof</code> operators	"The <code>__alignof__</code> operator" on page 135, "The <code>typeof</code> operator" on page 137

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

Lab Director
IBM Canada Ltd. Laboratory
8200 Warden Avenue
Markham, Ontario L6G 1C7
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 1998, 2007. All rights reserved.

Trademarks and service marks

Company, product, or service names identified in the text may be trademarks or service marks of IBM or other companies. Information on the trademarks of International Business Machines Corporation in the United States, other countries, or both is located at <http://www.ibm.com/legal/copytrade.shtml>.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Cell Broadband Engine is a trademark of the Sony Corporation and/or the Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

Industry standards

The following standards are supported:

- The C language is consistent with the International Standard for Information Systems-Programming Language C (ISO/IEC 9899-1990).
- The C language is also consistent with the International Standard for Information Systems-Programming Language C (ISO/IEC 9899-1999 (E)).
- The C++ language is consistent with the International Standard for Information Systems-Programming Language C++ (ISO/IEC 14882:1998).
- The C++ language is also consistent with the International Standard for Information Systems-Programming Language C++ (ISO/IEC 14882:2003 (E)).

Index

Special characters

`__align` 74
`__cdecl` 232
`__func__` 16
`__VA_ARGS__` 393, 396
`_Pragma` 410
- (subtraction operator) 143
- (unary minus operator) 130
-- (decrement operator) 130
-> (arrow operator) 128
, (comma operator) 151
:: (scope resolution operator) 126
! (logical negation operator) 131
!= (not equal to operator) 146
?: (conditional operators) 153
/ (division operator) 142
/= (compound assignment operator) 141
. (dot operator) 128
\$ 16, 32
* (indirection operator) 133
* (multiplication operator) 142
*= (compound assignment operator) 141
\
 continuation character 29, 391
\
 escape character 33
[] (array subscript operator) 150
% (remainder) 143
> (greater than operator) 144
>> (right-shift operator) 144
>>= (compound assignment operator) 141
>= (greater than or equal to operator) 144
< (less than operator) 144
<< (left-shift operator) 144
<<= (compound assignment operator) 141
<= (less than or equal to operator) 144
| (bitwise inclusive OR operator) 148
|| (logical OR operator) 149
& (address operator) 132
& (bitwise AND operator) 147
& (reference declarator) 93
&& (label value operator) 170, 176
&& (logical AND operator) 148
&= (compound assignment operator) 141
preprocessor directive character 391
preprocessor operator 397
(macro concatenation) 398
+ (addition operator) 143
+ (unary plus operator) 130
++ (increment operator) 129
+= (compound assignment operator) 141
= (simple assignment operator) 140
== (equal to operator) 146
^ (bitwise exclusive OR operator) 147
^= (compound assignment operator) 141
~ (bitwise negation operator) 131

A

aborting functions 387
abstract classes 312, 314
access rules
 base classes 297
 class types 259, 284
 friends 290
 members 283
 multiple access 304
 protected members 296
 virtual functions 314
access specifiers 270, 283, 293, 301
 in class derivations 297
accessibility 283, 304
addition operator (+) 143
address operator (&) 132
 GNU C extension 176
aggregate types 41, 320
 initialization 98, 320
alias 93
 function 17
 type-based aliasing 88
alignment 107, 108
 bit fields 64
 structure members 60
 structures 107
 structures and unions 74
alignof operator 135
allocation
 expressions 165
 functions 228
 uninitialized global variables 108
ambiguities
 base and derived member names 305
 base classes 303
 resolving 177, 306
 virtual function calls 312
AND operator, bitwise (&) 147
AND operator, logical (&&) 148
argc (argument count) 224
 example 225
arguments
 default 229
 evaluation 231
 macro 393
 main function 224
 of catch blocks 377
 passing 201, 225
 passing by reference 227
 passing by value 226
 trailing 393, 396
argv (argument vector) 224
 example 225
arithmetic conversions 113
arithmetic types
 type compatibility 57
arrays
 array-to-pointer conversions 117
 as function parameter 47, 217
 declaration 47, 217, 270
 description 90

arrays (*continued*)
 flexible array member 60, 61
 initialization 95
 initializing 101
 multidimensional 91
 subscripting operator 150
 type compatibility 92
 variable length 86, 91
 zero-extent 60, 62
ASCII character codes 34
asm 13
 keyword 17, 50
 labels 17
 statements 194
assembly
 labels 17
 statements 194
assignment operator (=)
 compound 141
 pointers 89
 simple 140
associativity of operators 170
atexit function 386
auto storage class specifier 46

B

base classes
 abstract 314
 access rules 297
 ambiguities 303, 305
 direct 302
 indirect 292, 303
 initialization 322
 multiple access 304
 pointers to 295
 virtual 303, 308
base list 303
best viable function 255
binary expressions and operators 139
binding 104
 direct 105
 dynamic 309
 static 309
 virtual functions 309
bit fields 63
 as structure member 60
 type name 138
bitwise negation operator (~) 131
block statement 178
block visibility 3
bool 57
Boolean
 conversions 114
 data types 54
 literals 21
break statement 189
built-in data types 41

C

- candidate functions 243, 255
- case label 182
- cast expressions 27, 59, 156
 - union type 158
- catch blocks 371, 373
 - argument matching 377
 - order of catching 377
- char type specifier 55
- character
 - data types 55
 - literals 28
 - multibyte 30, 32
- character set
 - extended 32
 - source 32
- class members
 - access operators 128
 - access rules 283
 - class member list 269
 - declaration 270
 - initialization 322
 - order of allocation 270
- class templates
 - declaration and definition 345
 - distinction from *template class* 344
 - explicit specialization 362
 - member functions 346
 - static data members 346
- classes 262
 - abstract 314
 - access rules 283
 - aggregate 260
 - base 293
 - base list 293
 - class objects 41
 - class specifiers 260
 - class templates 344
 - declarations 260
 - incomplete 264, 270
 - derived 293
 - friends 286
 - inheritance 291
 - keywords 259
 - local 266
 - member functions 271
 - member lists 269
 - member scope 273
 - nested 264, 288
 - overview 259
 - polymorphic 259
 - scope of names 263
 - static members 279
 - this pointer 276
 - using declaration 298
 - virtual 303, 309
- comma 151
 - in enumerator list 67
- comments 37
- compatibility
 - C89 and C99 413
 - data types 43
 - user-defined types 70
 - XL C and GCC 415
 - XL C/C++ and GCC viii
 - XL C++ and C99 viii, 419
 - XL C++ and GCC 420, 422
- compatible types
 - across source files 70
 - arithmetic types 57
 - arrays 92
 - in conditional expressions 153
- complex types 55
- composite types 43
 - across source files 70
- compound
 - assignment 141
 - expression 142
 - literal 164
 - statement 178
 - types 41
- computed goto 170, 176, 193
- concatenation
 - macros 398
 - multibyte characters 30
 - u-literals, U-literals 36
- conditional compilation directives 402
 - elif preprocessor directive 403
 - else preprocessor directive 405
 - endif preprocessor directive 405
 - examples 406
 - if preprocessor directive 403
 - ifdef preprocessor directive 404
 - ifndef preprocessor directive 405
- conditional expression (? :) 142, 153
- const 57, 76
 - casting away constness 228
 - function attribute 220
 - member functions 272
 - object 121
 - placement in type name 86
 - qualifier 72
 - vs. #define 392
- const_cast 161, 228
- constant expressions 67, 124
- constant initializers 270
- constructors 319
 - converting 330, 332
 - copy 333
 - exception handling 380
 - exception thrown in constructor 374
 - initialization
 - explicit 320
 - nontrivial 319, 327
 - overview 317
 - trivial 319, 327
- continuation character 29, 391
- continue statement 190
- conversion
 - constructors 330
 - function 332
 - implicit conversion sequences 256
- conversion sequence
 - ellipsis 257
 - implicit 256
 - standard 256
 - user-defined 257
- conversions
 - arithmetic 113
 - array-to-pointer 117
 - Boolean 114
 - cast 156
 - complex to real 115
 - explicit keyword 332
- conversions (*continued*)
 - floating-point 114
 - function arguments 119
 - function-to-pointer 117
 - integral 114
 - lvalue-to-rvalue 116, 121, 256
 - pointer 117
 - pointer to derived class 308
 - pointer to member 275
 - qualification 119
 - references 118
 - standard 113
 - user-defined 329
 - void pointer 118
- copy assignment operators 335
- copy constructors 333
- covariant virtual functions 311
- CPLUSPLUS macro 400
- cv-qualifier 72, 84
 - in parameter type specification 245
 - syntax 73

D

- data members
 - description 270
 - scope 273
 - static 280
- data types
 - aggregates 41
 - Boolean 54
 - built-in 41
 - character 55
 - compatible 43
 - complex 55
 - composite 43
 - compound 41
 - enumerated 67
 - floating 54
 - incomplete 42
 - integral 53
 - scalar 41
 - user-defined 41, 59
 - vector 57
 - void 56
- DATE macro 398
- deallocation
 - expressions 169
 - functions 228
- decimal integer literals 20
- declaration 201
- declarations
 - classes 260, 264
 - description 43
 - duplicate type qualifiers 73
 - friend specifier in member list 286
 - friends 290
 - pointers to members 274
 - resolving ambiguous statements 177
 - syntax 44, 85, 202
 - unsubscripted arrays 91
 - vector types 57
- declarative region 2
- declarators
 - description 83
 - reference 93
- decrement operator (--) 130

- default
 - clause 182
 - label 182
- default constructor 319
- define preprocessor directive 392
- defined unary operator 404
- definitions
 - description 43
 - macro 392
 - member function 271
 - tentative 45
- delete operator 169
- dependent names 366
- dereferencing operator 133
- derivation 293
 - array type 90
 - public, protected, private 297
- derived classes
 - catch block 377
 - construction order 325
 - pointers to 295
- designated initializer
 - aggregate types 95
 - union 98
- designator 95
 - designation 95
 - designator list 95
 - union 98
- destructors 326
 - exception handling 380
 - exception thrown in destructor 374
 - overview 317
 - pseudo 328, 329
- digraph characters 36
- direct base class 302
- division operator (/) 142
- do statement 186
- dollar sign 16, 32
- dot operator 128
- double type specifier 54
- downcast 163
- dynamic binding 309
- dynamic_cast 162

E

- EBCDIC character codes 34
- elaborated type specifier 263
- elif preprocessor directive 403
- ellipsis
 - conversion sequence 257
 - in function declaration 216
 - in function definition 216
 - in macro argument list 393
- else
 - preprocessor directive 405
 - statement 180
- enclosing class 271, 288
- endif preprocessor directive 405
- entry point
 - program 224
- enum
 - keyword 67
- enumerations 67
 - compatibility 70
 - declaration 67
 - initialization 100

- enumerations (*continued*)
 - trailing comma 67
- enumerator 68
- equal to operator (==) 146
- error preprocessor directive 407
- escape character \ 33
- escape sequence 33
 - alarm \a 34
 - backslash \\ 34
 - backspace \b 34
 - carriage return \r 34
 - double quotation mark \" 34
 - form feed \f 34
 - horizontal tab \t 34
 - new-line \n 34
 - question mark \? 34
 - single quotation mark \' 34
 - vertical tab \v 34
- exception handling 371
 - argument matching 377
 - catch blocks 373
 - arguments 377
 - constructors 380
 - destructors 380
 - example, C++ 387
 - exception objects 371
 - function try blocks 371
 - handlers 371, 373
 - order of catching 377
 - rethrowing exceptions 379
 - set_terminate 387
 - set_unexpected 387
 - special functions 384
 - stack unwinding 380
 - terminate function 386
 - throw expressions 372, 379
 - try blocks 371
 - try exceptions 374
 - unexpected function 385
- exceptions
 - declaration 373
 - function try block handlers 374
 - specification 382
- exclusive OR operator, bitwise (^) 147
- explicit
 - instantiation, templates 358
 - keyword 330, 332
 - specializations, templates 359, 361
 - type conversions 156
- exponent 22
- expressions
 - allocation 165
 - assignment 140
 - binary 139
 - cast 156
 - comma 151
 - conditional 153
 - deallocation 169
 - description 121
 - integer constant 124
 - new initializer 168
 - parenthesized 125
 - pointer to member 153
 - primary 123
 - resolving ambiguous statements 177
 - statement 177
 - throw 170, 379

- expressions (*continued*)
 - unary 128
- extern storage class specifier 8, 11, 48, 206
 - with function pointers 232
 - with template declaration 356, 359
 - with variable length arrays 91

F

- file inclusion 400, 401
- FILE macro 399
- file scope data declarations
 - unsubscripted arrays 91
- flexible array member 61
- float type specifier 54
- floating-point
 - constant 23
 - conversions 114
 - literal 22
 - promotion 115
- floating-point types 54
- for statement 187
- free store
 - delete operator 169
 - new operator 165
- friend
 - access rules 290
 - implicit conversion of pointers 298
 - member functions 271
 - nested classes 288
 - relationships with classes when templates are involved 347
 - scope 288
 - specifier 286
 - virtual functions 312
- function
 - aliases 17
- function attributes 218
- function designator 122
- function specifier
 - explicit 330, 332
- function templates
 - explicit specialization 362
- function try blocks 371
 - handlers 374
- function-like macro 393
- functions 201
 - allocation 228
 - arguments 201, 226
 - conversions 119
 - block 201
 - body 201
 - calling 225
 - calls 127
 - as lvalue 122
 - class templates 346
 - conversion function 332
 - deallocation 228
 - declaration 201
 - C++ 272
 - examples 203
 - multiple 205
 - parameter names 216
 - default arguments 229
 - evaluation 231
 - restrictions 230

- functions (*continued*)
 - definition 201
 - examples 204
 - exception handling 384
 - exception specification 382
 - friends 286
 - function call operator 201
 - function templates 348
 - function-to-pointer conversions 117
 - inline 208, 271
 - library functions 201
 - main 224
 - name 201
 - diagnostic 16
 - overloading 243
 - parameters 225, 226
 - pointers to 232
 - polymorphic 292
 - predefined identifier 16
 - prototype 201
 - return statements 191
 - return type 201, 213
 - return value 201, 213
 - signature 215
 - specifiable attributes 218
 - specifiers 208
 - template function
 - template argument deduction 349
 - type name 86
 - virtual 272, 309, 312

G

- global register variables 50
- global variable 3, 8
 - uninitialized 95, 108
- goto statement 192
 - computed goto 193
 - restrictions 193
- greater than operator (>) 144
- greater than or equal to operator (>=) 144

H

- handlers 373
- hexadecimal
 - floating constants 23
- hexadecimal integer literals 21
- hidden names 261, 263

I

- identifiers 15, 123
 - case sensitivity 16
 - id-expression 84, 124
 - labels 175
 - linkage 8
 - namespace 6
 - predefined 16
 - reserved 13, 14, 16
 - special characters 16, 32
- if
 - preprocessor directive 403
 - statement 180
- ifndef preprocessor directive 404

- ifndef preprocessor directive 405
- implicit conversion 113
 - Boolean 114
 - floating-point 114
 - integral 114
 - lvalue 121
 - pointer to derived class 295, 298
 - pointers to base class 295
 - types 113
- implicit conversions
 - complex to real 115
- implicit instantiation
 - templates 356
- include preprocessor directive 400
- include_next preprocessor directive 401
- inclusive OR operator, bitwise (|) 148
- incomplete type 90
 - as structure member 60, 61
 - class declaration 264
- incomplete types 42
- increment operator (++) 129
- indentation of code 391
- indirect base class 292, 303
- indirection operator (*) 59, 133
- information hiding 2, 3, 269, 296
- inheritance
 - multiple 292, 302
 - overview 291
- initialization
 - aggregate types 98
 - auto object 94
 - base classes 322
 - class members 322
 - extern object 95
 - order of 109
 - references 118
 - register object 95
 - static data members 282
 - static object 94, 165
 - union member 100
 - vector types 97
- initializer lists 93, 97, 165, 322
- initializers 93
 - aggregate types 95, 98
 - enumerations 100
 - unions 100
 - vector types 97
- inline
 - assembly statements 194
 - function specifier 208
 - functions 208, 271
- integer
 - constant expressions 67, 124
 - data types 53
 - literals 19
 - promotion 115
- integral
 - conversions 114

K

- keywords 13
 - exception handling 371
 - language extension 14
 - template 337, 367, 368
 - underscore characters 14

L

- label
 - as values 176
 - implicit declaration 3
 - in switch statement 182
 - locally declared 176
 - statement 175
- language extension vii, 14
 - C99 413
 - GNU C viii, 413
- left-shift operator (<<) 144
- less than operator (<) 144
- less than or equal to operator (<=) 144
- LINE macro 399
- line preprocessor directive 407
- linkage 1, 7
 - auto storage class specifier 46
 - const cv-qualifier 76
 - extern storage class specifier 11, 48
 - external 8
 - in function definition 206
 - inline member functions 272
 - internal 8, 46, 206
 - language 10
 - multiple function declarations 205
 - none 9
 - register storage class specifier 50
 - specifications 10
 - static storage class specifier 47
 - weak symbols 110
 - with function pointers 232
- linking to non-C++ programs 10
- literals 19
 - Boolean 21
 - character 28
 - compound 164
 - floating-point 22
 - integer 19, 20
 - decimal 20
 - hexadecimal 21
 - octal 21
 - string 28
 - Unicode 35
- local
 - classes 266
 - type names 267
- logical operators
 - ! (logical negation) 131
 - || (logical OR) 149
 - && (logical AND) 148
- long double type specifier 54
- long long type specifier 53, 57
- long type specifier 53, 57
- lvalues 72, 121, 123
 - casting 156
 - conversions 116, 121, 256

M

- macro
 - definition 392
 - typeof operator 138
 - function-like 393
 - invocation 393
 - object-like 392
 - variable argument 393, 396

- main function 224
 - arguments 224
 - example 225
- member functions
 - const and volatile 272
 - definition 271
 - friend 271
 - special 273
 - static 282
 - this pointer 276, 313
- member lists 260, 269
- members
 - access 283
 - access control 301
 - class member access operators 128
 - data 270
 - pointers to 153, 274
 - protected 296
 - scope 273
 - static 265, 279
 - virtual functions 272
- modifiable lvalue 121, 140
- modulo operator (%) 143
- multibyte character 32
 - concatenation 30
- multicharacter literal 28
- multidimensional arrays 91
- multiple
 - access 304
 - inheritance 292, 302
- multiplication operator (*) 142
- mutable storage class specifier 49

N

- name binding 366
- name hiding 7, 126
 - accessible base class 308
 - ambiguities 306
- names
 - conflicts 6
 - hidden 126, 261, 263
 - local type 267
 - mangling 10
 - resolution 3, 298, 306
- namespace
 - class names 263
 - context 6
 - of identifiers 6
- namespaces 235
 - alias 235, 236
 - declaring 235
 - defining 235
 - explicit access 241
 - extending 236
 - friends 239
 - member definitions 239
 - namespace scope object
 - exception thrown in
 - constructor 374
 - overloading 237
 - unnamed 237
 - user-defined 3
 - using declaration 240
 - using directive 240
- narrow character literal 28

- nested classes
 - friend scope 288
 - scope 264
- new operator
 - default arguments 230
 - description 165
 - initializer expression 168
 - placement syntax 166
 - set_new_handler function 168
- not equal to operator (!=) 146
- null
 - character \0 29
 - pointer 101
 - pointer constants 117
 - preprocessor directive 410
 - statement 194
- number sign (#)
 - preprocessor directive character 391
 - preprocessor operator 397

O

- object-like macro 392
- objects 121
 - class
 - declarations 260
 - description 41
 - lifetime 1
 - namespace scope
 - exception thrown in
 - constructor 374
 - restrict-qualified pointer 77
 - static
 - exception thrown in
 - destructor 374
- octal integer literals 21
- one's complement operator (~) 131
- operator functions 245
- operators 30
 - (subtraction) 143
 - (decrement) 130
 - > (arrow) 128
 - >* (pointer to member) 153
 - , (comma) 151
 - :: (scope resolution) 126
 - ! (logical negation) 131
 - != (not equal to) 146
 - ? : (conditional) 153
 - / (division) 142
 - . (dot) 128
 - .* (pointer to member) 153
 - () (function call) 127, 201
 - * (indirection) 133
 - * (multiplication) 142
 - (unary minus) 130
 - [] (array subscripting) 150
 - % (remainder) 143
 - > (greater than) 144
 - >> (right- shift) 144
 - >= (greater than or equal to) 144
 - < (less than) 144
 - << (left- shift) 144
 - <= (less than or equal to) 144
 - | (bitwise inclusive OR) 148
 - || (logical OR) 149
 - & (address) 132
 - & (bitwise AND) 147

- operators (*continued*)
 - && (logical AND) 148
 - + (addition) 143
 - ++ (increment) 129
 - = (simple assignment) 140
 - == (equal to) 146
 - ^ (bitwise exclusive OR) 147
 - alternative representations 31
 - assignment 140
 - copy assignment 335
 - associativity 170
 - binary 139
 - bitwise negation operator (~) 131
 - compound assignment 141
 - const_cast 161
 - defined 404
 - delete 169
 - dynamic_cast 162
 - equality 146
 - new 165
 - overloading 245, 271
 - binary 250
 - unary 248
 - pointer to member 153, 275
 - precedence 170
 - examples 173
 - type names 85
 - preprocessor
 - # 397
 - ## 398
 - pragma 410
 - reinterpret_cast 160
 - relational 144
 - scope resolution 294, 305, 312
 - sizeof 135
 - static_cast 159
 - typeid 133
 - typeof 137
 - unary 128
 - unary plus operator (+) 130
- OR operator, logical (||) 149
- overload resolution 255, 308
 - resolving addresses of overloaded
 - functions 257
- overloading
 - description 243
 - function templates 354
 - functions 243, 299
 - restrictions 244
 - operators 245, 259
 - assignment 251
 - binary 250
 - class member access 254
 - decrement 249
 - function call 252
 - increment 249
 - subscripting 253
 - unary 248
- overriding virtual functions 313
 - covariant virtual function 311

P

- packed
 - structure member 63
 - variable attribute 108
- parenthesized expressions 85, 125

- pass by reference 93, 227
- pass by value 226
- pixel 57
- placement syntax 166
- pointer to member
 - conversions 275
 - declarations 274
 - operators 153, 275
- pointers
 - conversions 117, 308
 - cv-qualified 86
 - dereferencing 88
 - description 86
 - generic 118
 - null 101
 - pointer arithmetic 59, 87
 - restrict-qualified 77
 - this 276
 - to functions 232
 - to members 153, 274
 - type-qualified 86
 - vector types 59
 - void* 117
- polymorphism
 - polymorphic classes 259, 310
 - polymorphic functions 292
- postfix
 - ++ and -- 129, 130
- pound sign (#)
 - preprocessor directive character 391
 - preprocessor operator 397
- pragma operator 410
- pragmas
 - _Pragma 410
 - preprocessor directive 410
 - standard 411
- precedence of operators 170
- predefined identifier 16
- predefined macros
 - CPLUSPLUS 400
 - DATE 398
 - FILE 399
 - LINE 399
 - STDC 399
 - STDC_HOSTED 399
 - STDC_VERSION 399
 - TIME 399
- prefix
 - ++ and -- 129, 130
 - hexadecimal floating constants 23
 - hexadecimal integer literals 21
 - octal integer literals 21
- preprocessor directives 391
 - conditional compilation 402
 - preprocessing overview 391
 - special character 391
 - warning 407
- preprocessor operator
 - _Pragma 410
 - # 397
 - ## 398
- primary expressions 123
- promotions
 - integral and floating-point 115
- pseudo-destructors 328
- punctuators 30
 - alternative representations 31

- pure specifier 270, 272, 312, 314
- pure virtual functions 314

Q

- qualification conversions 119
- qualified name 126, 265
- qualifiers
 - const 72
 - in parameter type specification 245
 - restrict 77
 - volatile 72, 78

R

- references
 - as return types 213
 - binding 104
 - conversions 118
 - declarator 132
 - description 93
 - initialization 104
- register storage class specifier 49
- register variables 50
- reinterpret_cast 160
- remainder operator (%) 143
- restrict 77
 - in parameter type specification 245
- return statement 191, 213
- return type
 - reference as 213
 - size_t 135
- right-shift operator (>>) 144
- RTTI support 133, 163
- rvalues 121

S

- scalar types 41, 86
- scope 1
 - class 5
 - class names 263
 - description 2
 - enclosing and nested 3
 - friends 288
 - function 3
 - function prototype 3
 - global 3
 - global namespace 3
 - identifiers 6
 - local (block) 2
 - local classes 266
 - macro names 396
 - member 273
 - nested classes 264
- scope resolution operator
 - ambiguous base classes 305
 - description 126
 - inheritance 294
 - virtual functions 312
- sequence point 152
- set_new_handler function 168
- set_terminate function 387
- set_unexpected function 385, 387
- shift operators << and >> 144
- short type specifier 53

- side effect 78
- signed type specifiers
 - char 55
 - int 53
 - long 53
 - long long 53
- size_t 135
- sizeof operator 135
 - with variable length arrays 92
- space character 391
- special characters 32
- special member functions 273
- specifiers
 - access control 297
 - inline 208
 - pure 272
 - storage class 45
- splice preprocessor directive ## 398
- stack unwinding 380
- standard type conversions 113
- statement expression 179
- statements 175
 - block 178
 - break 189
 - compound 179
 - continue 190
 - do 186
 - expressions 177
 - for 187
 - goto 192
 - if 180
 - labels 175
 - null 194
 - resolving ambiguities 177
 - return 191, 213
 - selection 180, 181
 - switch 181
 - while 185
- static 57
 - binding 309
 - data members 280
 - in array declaration 47, 217
 - initialization of data members 282
 - member functions 282
 - members 265, 279
 - storage class specifier 46, 206
 - linkage 47
 - with variable length arrays 91
- static storage class specifier 8
- static_cast 159
- STDC macro 399
- STDC_HOSTED macro 399
- STDC_VERSION macro 399
- storage class specifiers 45
 - auto 46
 - extern 48, 206
 - mutable 49
 - register 49
 - static 46, 206
- storage duration 1
 - auto storage class specifier 46
 - extern storage class specifier 48
 - register storage class specifier 50
 - static 46, 206
 - exception thrown in
 - destructor 374

- string
 - literal 28
 - terminator 29
- stringize preprocessor directive # 397
- struct type specifier 60
- structures 59, 262
 - alignment 74
 - as base class 297
 - as class type 259, 260
 - compatibility 70
 - flexible array member 60, 61
 - identifier (tag) 60
 - initialization 98
 - members 60
 - alignment 60
 - incomplete types 61
 - layout in memory 60, 98
 - packed 63
 - padding 60
 - zero-extent array 60
 - namespaces within 6
 - packed 60
 - unnamed members 98
- subscript declarator
 - in arrays 90
- subscripting operator 90, 150
 - in type name 86
- subtraction operator (-) 143
- suffix
 - floating-point literals 22
 - hexadecimal floating constants 23
 - integer literal constants 19
- switch statement 181

T

- tags
 - enumeration 67
 - structure 60
 - union 60
- template arguments 340
 - deduction 349
 - deduction, non-type 353
 - deduction, type 352
 - non-type 341
 - template 343
 - type 340
- template keyword 368
- templates
 - arguments
 - non-type 341
 - type 340
 - class
 - declaration and definition 345
 - distinction from *template class* 344
 - explicit specialization 362
 - member functions 346
 - static data members 346
 - declaration 337
 - dependent names 366
 - explicit specializations 359, 361
 - class members 361
 - declaration 360
 - definition and declaration 361
 - function templates 362
 - function
 - argument deduction 353

- templates (*continued*)
 - function (*continued*)
 - overloading 354
 - partial ordering 355
 - function templates 348
 - type template argument deduction 352
 - instantiation 337, 356, 359
 - explicit 358
 - forward declaration 356, 359
 - implicit 356
 - name binding 366
 - parameters 338
 - default arguments 339
 - non-type 338
 - template 339
 - type 338
 - partial specialization 364
 - matching 366
 - parameter and argument lists 365
 - point of definition 367
 - point of instantiation 367
 - relationship between classes and their friends 347
 - scope 361
 - specialization 337, 356, 359
- temporary objects 377
- tentative definition 45
- terminate function 371, 372, 377, 380, 384, 386
 - set_terminate 387
- this pointer 276, 313
- throw expressions 170, 371, 379
 - argument matching 377
 - rethrowing exceptions 379
 - within nested try blocks 372
- TIME macro 399
- tokens 13, 391
 - alternative representations for operators and punctuators 31
- translation unit 1
- trigraph sequences 37
- truncation
 - integer division 142
- try blocks 371
 - nested 372
- try keyword 371
- type attributes 78
- type name 85
 - local 267
 - qualified 126, 265
 - typename keyword 367
 - typeof operator 137
- type qualifiers
 - const 72, 76
 - const and volatile 84
 - duplicate 73
 - in function parameters 245
 - restrict 73, 77
 - volatile 72
- type specifier
 - class types 260
 - elaborated 263
 - overriding 109
- type specifiers 52
 - _Bool 54
 - bool 54

- type specifiers (*continued*)
 - char 55
 - complex 55
 - double 54
 - enumeration 67
 - float 54
 - int 53
 - long 53
 - long double 54
 - long long 53
 - short 53
 - unsigned 53
 - vector data types 57
 - void 56
 - wchar_t 53, 55
- typedef specifier 71
 - class declaration 267
 - local type names 267
 - pointers to members 275
 - qualified type name 265
 - with variable length arrays 92
- typeid operator 133
- typename keyword 367
- typeof operator 137
- types
 - class 260
 - conversions 156
 - type-based aliasing 88
 - variably modified 90

U

- u-literal, U-literal 35
- unary expressions 128
- unary operators 128
 - label value 170
 - minus (-) 130
 - plus (+) 130
- undef preprocessor directive 396
- underscore character 14, 16
- unexpected function 371, 384, 385, 386
 - set_unexpected 387
- Unicode 34
- unions 59, 60
 - as class type 259, 260
 - cast to union type 158
 - compatibility 70
 - designated initializer 95
 - initialization 100
 - specifier 60
 - unnamed members 98
- universal character name 16, 28, 34
- unnamed namespaces 237
- unsigned type specifiers
 - char 55
 - int 53
 - long 53
 - long long 53
 - short 53
- unsubscripted arrays
 - description 91, 217
- user-defined conversions 329
- user-defined data types 41, 59
- using declarations 240, 298, 306
 - changing member access 301
 - overloading member functions 299
- using directive 240

V

- variable
 - in specified registers 50
- variable attributes 105
- variable length array 42, 91, 193
 - as function parameter 92, 226, 255
 - sizeof 124
 - type name 86
- variably modified types 90, 91, 183
- vector data types 57
- vector processing support ix, 417, 423
- vector types 137
 - in typedef declarations 71
 - literals 25
- virtual
 - base classes 293, 303, 308
- virtual functions 272, 309
 - access 314
 - ambiguous calls 312
 - overriding 313
 - pure specifier 314
- visibility 1, 7
 - block 3
 - class members 285
- void 56
 - in function definition 213, 216
 - pointer 117, 118
- volatile
 - member functions 272
 - qualifier 72, 78

W

- warning preprocessor directive 407
- wchar_t type specifier 28, 53, 55
- weak symbol 110
- while statement 185
- white space 13, 37, 391, 397
- wide characters
 - literals 28
- wide string literal 30

Z

- zero-extent array 62



Program Number: 5724-T42, 5724-T43

SC23-8519-00

